From Satellite Fire Detections to Agricultural Statistics: Measuring Crop-Residue Burning in U.S. Agriculture

Thi Hoang Hieu Nguyen*

Abstract

U.S. agricultural statistics provide detailed measures of acreage, production, yields, prices, and farm income, but they contain limited information on crop-specific residue-management practices. This paper develops a framework for extending agricultural measurement to crop-residue burning, a field-level practice that is not consistently observed in standard county-level agricultural data. I construct a county-crop-year measure of detected crop-residue burning intensity by linking VIIRS active fire detections, the USDA Cropland Data Layer, and USDA QuickStats harvested acreage for corn, soybeans, wheat, cotton, and rice from 2012 to 2024. The measure converts satellite fire detections into an agricultural statistical object through three steps: assigning detections to crop pixels, normalizing crop-specific detections by harvested acreage, and separating annual detections into crop-specific peak-season and off-peak components.

The results show that measurement choice substantially changes the empirical picture of agricultural burning. Raw crop-assigned fire counts emphasize large crop systems, especially soybeans and corn. After acreage normalization, rice, cotton, and wheat emerge as the highest-intensity crop systems. Raw-count and intensity rankings also identify different county-crop units, indicating that unnormalized fire counts primarily capture the scale of crop production, while acreage-normalized intensity identifies systems with high detected burning relative to harvested acreage. Construct-validity checks show that same-crop fire detections are positively related to same-crop harvested acreage within counties over time, and crop-by-crop placebo comparisons clarify both the specificity and the limits of the measure. Seasonal decomposition further shows that detected burning is crop-specific: rice has a concentrated peak-season signal, whereas soybeans and wheat display more diffuse seasonal patterns. Descriptive linkages to NOAA HMS smoke and QCEW agricultural labor-market data illustrate how the measure can be connected to external environmental and economic datasets without estimating causal damages. The paper contributes a transparent statistical bridge between satellite fire detections and the county-crop-year structure of U.S. agricultural data, converting physical fire signals into crop-specific, acreage-normalized, and seasonally interpretable variables for research on the agricultural economy.

Keywords: Agricultural measurement; crop-residue burning; satellite data; VIIRS; Cropland Data Layer; USDA QuickStats
JEL Codes: Q10, Q15, Q53, C80

*CERGE-EI (The Center for Economic Research and Graduate Education – Economics Institute), Czech Republic. hieuthihoang.nguyen@cerge-ei.cz.

1 Introduction

U.S. agricultural statistics provide detailed county-level information on acreage, production, yields, prices, and farm income. These data are central to measuring the agricultural economy, but they are less complete for field-level management practices that shape production decisions, residue management, smoke exposure, and the environmental footprint of agriculture. Crop-residue burning is one such practice. Farmers may burn residues to clear fields, manage pests and weeds, or prepare land for the next planting cycle. Yet residue burning is not consistently measured as a crop-specific county-year variable in standard agricultural data. This paper develops such a measure. Satellite data offer a way to observe agricultural practices that are difficult to measure through surveys or administrative records. As emphasized in the broader economics literature on remote sensing, the value of satellite data lies not only in their spatial coverage, but also in their ability to measure otherwise unobserved economic activity and environmental conditions (Donaldson and Storeygard, 2016). For agricultural and environmental applications, however, satellite observations must be translated into variables that correspond to the economic object of interest (Wuepper et al., 2025; Kubitza et al., 2020). This translation is especially important for crop-residue burning. A satellite active-fire detection records a thermal anomaly at the time of overpass. It does not, by itself, identify the crop system associated with the fire, adjust for the scale of the crop system, or distinguish crop-plausible burning windows from more diffuse fire activity. Raw fire counts are therefore not yet measures of agricultural burning practices.

I develop a framework for converting satellite fire detections into an agricultural economic measurement object. I combine VIIRS active fire detections, the USDA Cropland Data Layer, and USDA QuickStats harvested acreage to construct county-crop-year measures of detected burning intensity for corn, soybeans, wheat, cotton, and rice from 2012 to 2024. The main measure is crop-assigned fire detections per 1,000 harvested acres. The construction has three steps. First, I assign VIIRS detections to crop systems using the Cropland Data Layer. This converts generic fire detections into crop-specific fire counts. Second, I normalize these counts by crop-specific harvested acreage from USDA QuickStats. This separates detected burning intensity from crop scale. Third, I use crop-specific seasonal timing to divide annual detections into peak-season and off-peak components. These steps produce a county-crop-year measure designed for economic analysis, rather than a new fire-detection algorithm, burned-area product, or emissions inventory. The distinction matters because different measurement choices lead to different empirical pictures of agricultural burning. Raw fire counts identify places with many detected fires. Crop-assigned fire counts identify detections that occur on target crop pixels. Acreage-normalized intensity identifies crop

systems where detected burning is high relative to harvested acreage. If these measures ranked crops and counties similarly, the transformation would mainly change units. They do not. Raw crop-assigned fire counts emphasize large crop systems, especially soybeans and corn. After acreage normalization, rice has the highest detected burning intensity, followed by cotton and wheat, while soybeans and corn rank lower. The same pattern appears within crops across counties: the county-crop units with the most raw fire detections are often not the county-crop units with the highest detections per 1,000 harvested acres. Raw counts therefore do not simply measure the same object with noise; they often emphasize a different object, crop scale rather than burning intensity.

This paper uses these differences to study the consequences of measurement choice. I compare raw VIIRS fire counts, crop-assigned fire counts, and acreage-normalized detected burning intensity across crops, county-crop units, seasonal windows, and external linkages. The results show that measurement choice changes which crops appear most burning-intensive, which county-crop systems appear highest burning, and which county-years are linked to external environmental and economic datasets. Descriptive comparisons with NOAA HMS smoke and QCEW agricultural labor-market data illustrate the linkability of the constructed measure. They are not used to estimate causal smoke, health, or employment effects. Instead, they show that the acreage-normalized burning-intensity measure identifies a different agricultural economic geography than raw fire counts or unnormalized crop-pixel fire counts. The validation exercises evaluate the measure at the level at which it is designed to be used: the county-crop-year. Same-crop fire detections are positively related to same-crop harvested acreage within counties over time. Crop-by-crop placebo tests show that same-crop relationships are consistently positive and stronger than off-crop relationships on average. Geographic plausibility checks show that high-intensity crop systems appear in plausible regions, especially for rice, cotton, and southern row-crop systems. Denominator-robustness checks show that the crop-level intensity ranking is not driven by very small harvested-acreage denominators. These exercises do not turn satellite detections into field-level audit records. Rather, they assess whether the aggregate measure behaves like a crop-specific agricultural practice variable. Seasonal patterns provide an additional diagnostic. Rice has a sharply concentrated burning calendar, with most detections occurring in peak burning months. Corn also shows a stronger peak-season than off-peak relationship. By contrast, soybean and wheat detections are more diffuse across the year. Seasonal concentration is therefore useful, but crop-specific. It should not be treated as a universal replacement for annual detected burning intensity. This result is consistent with the idea that agricultural fire detections must be interpreted in relation to crop calendars, while also showing that a single national burning season would be too restrictive for a multi-crop measurement framework.

This paper contributes to economic measurement in the agricultural economy by constructing a crop-specific measure of residue-burning practice that is not directly observed in standard agricultural statistics. Existing county-level data measure agricultural production in detail, but they provide limited information on residue-management behavior. The framework developed here uses public satellite, crop-map, and acreage data to fill part of this gap. It produces a transparent measurement layer that can be linked to standard agricultural, environmental, and economic datasets. The paper also contributes to three connected literatures. First, it contributes to work using satellite data to expand economic measurement when conventional data do not observe the relevant behavior (Donaldson and Storeygard, 2016; Kubitza et al., 2020). Second, it builds on crop-mapping research that uses operational land-cover products, especially the Cropland Data Layer, to identify crop systems across space and time (Lark et al., 2021; Zhang et al., 2025). Third, it relates to active-fire and burned-area remote sensing. Active-fire products capture fires burning at satellite overpass times, while burned-area products attempt to map post-fire surface changes. For crop-residue burning, this distinction is important because agricultural fires are often small, short-lived, and followed quickly by field operations that can remove the burn scar before the next optical overpass (Wooster et al., 2021; Hall et al., 2021). I do not attempt to improve fire detection or estimate burned acreage. Instead, I construct an intermediate measurement layer that links satellite fire detections, crop-specific land-cover data, and USDA acreage statistics. The paper is not an emissions inventory, burned-area validation study, or causal analysis of smoke, health, or labor-market effects. Its purpose is to construct and evaluate a crop-specific agricultural measurement object from satellite active-fire detections. This distinction is important because the central question is not whether crop-residue burning causes a particular downstream outcome, but how satellite detections should be transformed into an economically interpretable measure of agricultural practice.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data sources and construction of the crop-residue burning panel. Section 4 presents the measurement framework. Section 5 reports construct-validity results for the crop-specific detected burning-intensity measure. Section 6 shows how measurement choice changes crop rankings, county-crop rankings, geographic classification, seasonal interpretation, and linkages to external environmental and labor-market data. Section 7 develops the seasonal concentration measures and evaluates how their performance differs across crops. Section 8 concludes.

2 Literature Review

This paper connects four strands of research: economic measurement using satellite data, remote-sensing studies of agricultural and crop-residue burning, crop-specific land-cover products, and economic studies that use satellite-derived burning measures for policy and impact analysis. The common lesson across these literatures is that satellite observations become useful for economic research only after the measurement object has been defined. A satellite fire detection is a physical signal: it records a thermal anomaly at a particular place and time. It is not, by itself, a crop-specific agricultural practice, a burned-area estimate, an emissions measure, or an economic exposure variable. The measurement problem is therefore one of translation. Which detections belong to which crop systems? How should detected fire activity be scaled by the relevant agricultural production base? How should seasonal timing be interpreted? I build on these literatures by constructing and evaluating a crop-specific, acreage-normalized, and seasonally decomposed measure of detected crop-residue burning intensity.

First, the paper contributes to the growing literature on satellite data and economic measurement. Satellite data are valuable to economists because they provide repeated, spatially explicit observations of outcomes and behaviors that are difficult to observe through surveys or administrative records (Donaldson and Storeygard, 2016). This advantage is especially important in agriculture, where many production practices occur at fine spatial scales and are not consistently recorded in standard administrative datasets. At the same time, satellite products are not economic variables by construction. Researchers must choose the spatial unit, temporal aggregation, processing rule, and validation strategy that connect the remotely sensed signal to the economic object of interest. Recent methodological work in agricultural and environmental economics makes this point explicitly: remote-sensing data must be converted into research-ready variables before they can be used for economic analysis (Wuepper et al., 2025). This paper applies that measurement logic to crop-residue burning, a field-management practice that is not measured consistently as a crop-specific county-year variable in U.S. agricultural statistics. The relevant economic object is therefore not the raw satellite detection, but a measure of detected burning intensity that can be compared across crops, counties, and years and linked to agricultural data. This distinction also relates to the broader review by Kubitza et al. (2020), who emphasize that satellite data can improve the measurement of agricultural practices and their impacts when survey measures are costly, incomplete, or subject to reporting error. The setting here is different from many developing-country applications, but the measurement principle is similar. I use satellite data not only to observe an environmental signal, but also to construct a variable for an

agricultural management practice. The contribution is therefore an intermediate economic measurement layer between the physical satellite product and the agricultural data system.

Second, the paper builds on remote-sensing studies of agricultural burning and crop-residue burning. This literature has established that cropland fires can be observed from space and that agricultural burning is spatially and seasonally structured. Early evidence from the southeastern United States shows that MODIS detections can identify agricultural fires and that these fires vary across crops, locations, and management calendars (McCarty et al., 2007). Subsequent work extends the analysis to the contiguous United States and provides a satellite-based baseline for the spatial and temporal distribution of crop-residue burning (McCarty et al., 2009). Related emissions-inventory work combines satellite-derived burned area, crop information, fuel loads, combustion completeness, and emission factors to estimate annual and seasonal emissions from U.S. crop-residue burning (McCarty, 2011). More recent local studies use higher-resolution imagery to identify burned cropland and examine $PM_{2.5}$ dispersion from agricultural burning episodes (Zamanialaei et al., 2023), while atmospheric evidence shows that emission factors vary across residues and combustion conditions (Travis et al., 2023). These studies motivate crop-specific measurement because the environmental and agricultural interpretation of a fire detection depends partly on the crop system and residue context. My paper takes a complementary step. I do not estimate burned area, fuel consumption, or emissions. Those are natural objects for physical science and emissions accounting. Agricultural economics often requires a different object: a panel measure defined at a unit that can be linked to USDA acreage, production, labor-market, environmental, and policy datasets. The county-crop-year measure constructed here is designed for that purpose. It retains the crop specificity emphasized in the remote-sensing and emissions literature, but shifts the focus from physical accounting to economic measurement. This distinction matters because a measure designed to estimate total emissions is not necessarily the measure best suited for comparing detected burning practices across crops of very different scale.

Third, the paper relates to work on active-fire products, burned-area products, and the validation challenges created by small agricultural fires. Active-fire products detect thermal anomalies while combustion is occurring, whereas burned-area products identify post-fire changes in surface reflectance. This distinction is especially important for crop-residue burning. Agricultural fires are often small, short-lived, and followed quickly by plowing or other field operations, so a field may lose its burn-scar signature before the next suitable optical image is observed. Validation evidence from cropland settings shows that coarse burned-area products can have high omission and commission errors for crop-residue fires (Hall et al., 2021). For some questions, especially those involving short-lived field fires, active-fire timing

may therefore be more directly informative than a post-fire scar. Active-fire detections also have limitations. Detection depends on overpass timing, cloud cover, fire size, fire temperature, sensor resolution, viewing geometry, and background surface conditions (Wooster et al., 2021). VIIRS improves the detection of smaller fires relative to earlier coarser products, but greater sensitivity does not eliminate omission error or the need for filtering and interpretation (Schroeder et al., 2014; Zhang et al., 2017). Higher-resolution sensors such as Landsat-8 provide finer spatial detail, but their lower revisit frequency makes them less suited for a consistent national annual panel of short-lived agricultural burning events (Schroeder et al., 2016). Validation reviews further show that active-fire product validation is difficult because reference data must be nearly coincident with the fire conditions observed by the satellite (Ramsey et al., 2024). These studies shape the interpretation of my measure. I do not claim to observe all crop-residue burning, and I do not treat VIIRS detections as a census of field-level burning events. The object I construct is therefore detected burning intensity, not true burning, burned acreage, or emissions.

Fourth, the paper uses crop-specific land-cover data to assign fire detections to agricultural systems. This step connects the paper to the crop-mapping literature and to work on the USDA Cropland Data Layer. Crop-specific land-cover products have become core infrastructure for agricultural remote sensing. A recent review describes the Cropland Data Layer as one of the most widely used operational crop-specific land-cover products, serving as a crop mask, training source, and benchmark dataset in remote-sensing applications (Zhang et al., 2025). This role is why the CDL is useful here. It provides annual crop classifications that allow generic fire detections to be linked to specific crop systems. The crop-mapping literature also makes clear that crop assignment should not be treated as error-free. Validation work shows that the CDL performs well for major U.S. crops and especially for broad cropland classification, but accuracy varies across crops, regions, and heterogeneous landscapes (Lark et al., 2021). Refinement studies show that some classification errors arise from isolated pixels, field boundaries, mixed pixels, roads, and small non-crop features embedded in agricultural landscapes (Lin et al., 2022). These issues matter when point-level fire detections are overlaid on a crop map. I therefore use the CDL as a crop-assignment layer rather than as field-level ground truth. The variable is defined by a reproducible overlay of fire detections, crop pixels, and harvested acreage. It is not an administrative record of burning. This interpretation also clarifies the role of validation. I test whether the resulting crop-specific fire measures line up with same-crop acreage, whether same-crop relationships are stronger than crop-placebo relationships, whether high-intensity systems appear in plausible agricultural regions, and whether seasonal patterns are consistent with crop-specific burning calendars. In this sense, the CDL is not merely a data input. It is part of the measurement

design that turns a physical fire signal into an agricultural variable.

Finally, recent economics and policy studies increasingly use satellite-derived fire or burn measures as outcomes in analyses of crop-residue-burning regulation, incentives, and externalities. Studies of India use satellite-based burning measures to examine bans, enforcement, payments for ecosystem services, and farmer responses (Jack et al., 2025; Khanna et al., 2025). Related work combines satellite, emissions, and survey information to correct fire-emissions inventories and quantify air-quality impacts from crop-residue burning (Liu et al., 2020; Lan et al., 2022). These papers demonstrate the policy and economic relevance of measuring agricultural burning. They also show why measurement choices matter: if the outcome variable changes, conclusions about compliance, enforcement, exposure, or damages may change as well. This paper contributes upstream of those applications. Existing policy and impact studies often take a satellite-derived fire or burn variable as an input to causal analysis. I focus instead on the construction of the variable itself. This is a narrower exercise, but it addresses a prior question: what should the agricultural burning variable be? Raw fire counts tend to identify places with many detected fires, which often means places with large agricultural systems. Crop assignment restricts the measure to detections on relevant crop pixels. Acreage normalization identifies crop systems where detected burning is high relative to harvested area. Seasonal decomposition separates concentrated crop-specific burning calendars from more diffuse fire activity. I then show that these choices change crop rankings, county-crop rankings, geographic classification, seasonal interpretation, and linkages to external environmental and labor-market data. The contribution of the paper is therefore not to show for the first time that crop-residue fires can be observed from space, nor to estimate emissions, pollution damages, or policy effects. The contribution is to define and evaluate an agricultural economic measurement object from satellite fire detections: detected crop-specific burning intensity per 1,000 harvested acres. By doing so, the paper connects the remote-sensing literature's ability to observe fire activity with the economics literature's need for comparable, interpretable, and reusable measures of agricultural practices.

3 Data

This section describes the datasets used to construct the crop-residue burning measure. The empirical object of interest is a county-crop-year measure of detected burning intensity for five major U.S. crops: corn, soybeans, wheat, cotton, and rice. The construction combines three types of information. VIIRS active-fire detections provide the physical fire signal. The USDA Cropland Data Layer provides annual crop-location information. USDA QuickStats harvested acreage provides the crop-specific denominator needed to compare fire activity

across agricultural systems of different scale. I also construct monthly versions of the fire panel to measure crop-specific seasonal concentration, and I link the resulting county-year measures to NOAA HMS smoke and QCEW labor-market data to illustrate how the measurement object can be connected to external environmental and economic datasets. The main sample covers 2012–2024. The data structure reflects the purpose of the paper. A pixel-level fire record is closest to the satellite observation, but it is not the unit at which most agricultural economic data are organized. A state-year measure would be simple to construct, but it would obscure local crop heterogeneity. I therefore use the county-crop-year as the main unit because it balances spatial detail, crop specificity, and compatibility with USDA agricultural statistics. This unit also allows the measure to be linked to county-level economic, environmental, and policy-relevant datasets without treating the satellite observation itself as the final economic variable.

3.1 VIIRS active fire detections

The fire data come from VIIRS active-fire products distributed through NASA's Fire Information for Resource Management System archive (NASA Fire Information for Resource Management System, 2026). FIRMS provides active-fire and thermal-anomaly detections from satellite sensors including MODIS, VIIRS, and Landsat. I use VIIRS detections because they provide repeated, spatially explicit observations of thermal anomalies over a long national panel. This makes them well suited for constructing a consistent measure of detected fire activity across crops, counties, and years.

The raw VIIRS observations are point-level detections. Each detection records the location and date of an observed thermal anomaly. I use detections from 2012 to 2024, matching the period of the constructed agricultural panel. These detections are the starting point for the measurement exercise, but they are not directly interpretable as crop-residue burning. Active-fire products identify fires burning at the time of satellite observation, and detection depends on the sensor's observation process, including overpass timing, fire size, fire temperature, cloud cover, background surface conditions, and fire duration (Wooster et al., 2021). This issue is especially relevant for agricultural burning because crop-residue fires are often small and short-lived. Some fires may be missed if they occur between satellite overpasses or are too small or obscured to be detected. Conversely, a detected thermal anomaly is not necessarily a residue-burning event. These limitations shape the interpretation of the measure. I do not treat VIIRS as a census of all burning events, and I do not use raw VIIRS fire counts as the main agricultural variable. Instead, I use VIIRS detections as a consistent, reproducible input into a broader measurement pipeline. The next steps assign detections

to crop systems and normalize the resulting crop-specific fire counts by harvested acreage. Throughout the paper, the constructed variables should therefore be interpreted as measures of detected fire activity on crop pixels, not as field-level records of all residue burning.

The choice of active-fire data also reflects a practical measurement tradeoff. Burned-area products can be useful for mapping the spatial extent of fire scars, but cropland burned-area mapping is difficult when fires are small, short-lived, and followed quickly by plowing or other field operations. Validation work has shown that coarse burned-area products can have high omission and commission errors in cropland settings (Hall et al., 2021). For the objective of this paper—a national county-crop-year panel of detected burning intensity—the timing and repeated coverage of VIIRS active-fire detections provide a more suitable starting point than attempting to infer burned acreage from coarse burn-scar products.

3.2 Cropland Data Layer

To identify the crop system associated with each fire detection, I use the USDA National Agricultural Statistics Service Cropland Data Layer (USDA National Agricultural Statistics Service, 2026a). The Cropland Data Layer is an annual, geo-referenced, crop-specific land-cover product for the United States. It is widely used in agricultural remote sensing as a crop mask, training source, and benchmark dataset (Zhang et al., 2025). In this paper, its role is more specific: I use it as a crop-assignment layer that links generic fire detections to economically meaningful crop systems.

For each VIIRS detection, I determine whether the detection falls on land classified as one of the five target crop groups: corn, soybeans, wheat, cotton, or rice. A detection that falls on a rice pixel is assigned to rice; a detection that falls on a corn pixel is assigned to corn. Detections outside these five crop groups are not used to construct the crop-specific burning measures. This step converts thermal detections into CDL-based crop-specific fire counts. Without it, a county-year fire count would mix crop-residue burning with other agricultural fires, grassland fires, forest fires, developed-land thermal anomalies, and fires unrelated to the five crop systems studied in this paper. The empirical importance of this crop-assignment step is large. Across 2012–2024, the balanced county-month fire file contains 6,054,814 VIIRS fire detections. Of these, 315,807 detections fall on pixels assigned to corn, soybeans, wheat, cotton, or rice. Target five-crop pixel detections therefore account for 5.2 percent of all VIIRS detections. Raw VIIRS fire counts are therefore much broader than the crop-residue burning object studied in this paper. Crop assignment narrows the satellite record to the part of detected fire activity that can be associated with the selected crop systems.

At the same time, crop assignment is only as accurate as the land-cover classification on which it relies. The CDL performs well for major U.S. crops, but validation studies document variation in accuracy across crops, regions, and landscapes (Lark et al., 2021). Refinement work further shows that pixel-level errors may arise near field boundaries, mixed pixels, roads, and other small non-crop features embedded in agricultural landscapes (Lin et al., 2022). These issues matter when point-level fire detections are overlaid on a raster crop map. I therefore interpret the resulting variables as CDL-based crop-assigned fire detections rather than field-level ground truth. This interpretation keeps the measure transparent: it is defined by a reproducible overlay of fire detections and annual crop pixels.

3.3 USDA QuickStats harvested acreage

The denominator for the detected burning-intensity measure comes from USDA NASS QuickStats (USDA National Agricultural Statistics Service, 2026b). QuickStats is the official NASS database for agricultural statistics and provides county-crop-year information by commodity, geography, and time period. I use harvested acreage for corn, soybeans, wheat, cotton, and rice. Harvested acreage measures the scale of each crop system in a county-year and provides the denominator needed to compare detected fire activity across crop systems of different size. This denominator is central to the economic interpretation of the measure. Raw fire counts combine at least two objects: the size of the crop system and the intensity of detected burning. A large corn-producing county may have many crop-assigned fire detections simply because it has many corn acres. A smaller rice-producing county may have fewer total detections but substantially more detections per acre. Dividing crop-assigned fire detections by harvested acreage separates detected burning intensity from crop scale. This normalization is what allows the paper to compare burning across crops, counties, and years.

The main analysis keeps county-crop-years with positive reported harvested acreage for the relevant crop. This restriction ensures that each denominator corresponds to an observed crop system in USDA statistics. It also aligns the satellite-based measure with the county-crop-year structure used in agricultural economics. The final positive-acreage panel contains 18,561 county-crop-year observations across the five crops. Because the measure is defined relative to harvested acreage, I also examine denominator robustness in the appendix. Those checks ask whether the crop-level intensity rankings are driven by very small acreage denominators rather than by systematic differences in detected burning intensity.

3.4 Constructed crop-residue burning panel

The core panel combines the three measurement inputs: VIIRS fire detections, CDL crop assignment, and QuickStats harvested acreage. For each crop c, county g, and year y, I count the number of VIIRS detections assigned to crop c and divide this count by harvested acreage for the same crop, county, and year. The baseline measure is fire detections per 1,000 harvested acres. The resulting variable is a county-crop-year index of detected burning intensity. This construction deliberately separates total detected fire activity from burning intensity relative to crop scale. Corn and soybeans account for large harvested acreage and many crop-assigned fire detections in levels. Rice appears in fewer counties and covers fewer acres, but it has much higher detected burning intensity per 1,000 harvested acres. The distinction between scale and intensity is one of the main reasons the paper treats measurement choice as substantive rather than mechanical. A raw count and an acreage-normalized measure do not simply express the same information in different units; they can identify different crops and different county-crop systems as high burning.

I also construct a county-crop-year-month version of the fire panel. Monthly crop-assigned detections are used to describe crop-specific seasonality and to define peak-season and off-peak burning measures. For each crop, I identify the three calendar months with the largest national shares of crop-specific fire detections. I then aggregate monthly detections into annual peak-season and off-peak components and link both components to the same county-crop-year harvested acreage denominator. This monthly structure allows the paper to ask whether detected fire activity is concentrated in crop-plausible burning windows or distributed more diffusely across the year. The constructed panel should not be interpreted as measuring emissions, burned acreage, or the number of true burning events. It is a satellite-based indicator of crop-specific detected burning intensity. Its value comes from being transparent, scalable, and aligned with the agricultural data system. The validation and measurement-consequences sections below ask whether this constructed variable contains crop-specific information and whether using less processed measures would lead researchers to identify different crops and counties as high burning.

3.5 Linkage to environmental and economic datasets

I use two additional datasets to illustrate how the constructed burning measure can be linked to external environmental and economic data. These linkages are descriptive. They are not used to estimate causal pollution, health, or labor-market effects.

First, I use NOAA Hazard Mapping System smoke data to describe whether county-years identified by different burning measures also differ in smoke exposure. The HMS product

identifies smoke plumes from satellite imagery and classifies smoke density. I aggregate the smoke information to county-year measures of smoke days and medium-or-heavy smoke days. These variables provide an environmental linkage for the burning measures. They should be interpreted cautiously because HMS smoke includes smoke from many sources, including non-agricultural fires and transported smoke from outside the county. The purpose is therefore not to attribute smoke to crop-residue burning. It is to show whether alternative burning measures identify different county-years when linked to an external environmental dataset.

Second, I use county-level employment and wage data from the Quarterly Census of Employment and Wages. I construct annual county-level measures of private agricultural employment, private agricultural establishments, agricultural wages, total private employment, and the agricultural share of private employment for 2012–2024. These variables provide an economic linkage for the burning measures. They allow me to ask whether acreage-normalized burning intensity identifies a different agricultural labor-market geography than raw fire counts or unnormalized crop-assigned fire counts.

The external datasets serve a limited but useful role. They show that the detected burning-intensity measure can be connected to standard environmental and economic data without changing the paper into a causal damages study. They also reinforce the central measurement point: different fire measures identify different sets of county-years, and those differences matter when the measure is brought into contact with other datasets.

3.6 Sample and crop coverage

The final analysis focuses on county-crop-years with positive harvested acreage for the relevant crop. Table 1 summarizes the main positive-acreage analysis sample and reports both crop scale and acreage-normalized detected burning intensity.

Table 1: Crop-assigned fire detections and acreage-normalized intensity by crop, 2012–2024

Crop	Positive county-years	Counties	Harvested acres (million)	Crop-assigned fire detections	Mean fire detections per county-year	Aggregate intensity fires/1,000 acres	Mean intensity fires/1,000 acres	Median intensity fires/1,000 acres	P95 intensity fires/1,000 acres	Max. intensity fires/1,000 acres
Rice	301	85	12.3	13,680	45.45	1.111	1.171	0.784	3.227	6.8
Cotton	1,434	403	40.1	5,161	3.60	0.129	0.175	0.077	0.685	5.9
Wheat	4,405	1,571	123.3	15,103	3.43	0.122	0.146	0.000	0.662	21.7
Soybeans	6,567	1,685	382.0	34,259	5.22	0.090	0.098	0.032	0.375	19.7
Corn	5,854	1,941	262.4	20,395	3.48	0.078	0.146	0.027	0.716	10.0

Notes: The table summarizes VIIRS active fire detections assigned to crop pixels and USDA QuickStats harvested acreage for county-crop-years with positive harvested acreage from 2012 to 2024. Crop-assigned fire detections are VIIRS active fire detections falling on Cropland Data Layer pixels classified as the listed crop. Acreage-normalized intensity is crop-assigned fire detections divided by harvested acreage and multiplied by 1,000. Aggregate intensity is computed using total crop-assigned fire detections and total harvested acreage by crop. Mean, median, P95, and maximum intensity summarize the county-crop-year distribution of fire detections per 1,000 harvested acres.

The sample contains 18,561 positive-acreage county-crop-years across the five target crops. Corn and soybeans have the broadest coverage, with 5,854 corn county-crop-years and 6,567 soybean county-crop-years. Wheat accounts for 4,405 county-crop-years, cotton for 1,434, and rice for 301. The smaller rice sample reflects the geographic concentration of rice production, while the larger corn and soybean samples reflect their broader national coverage. The table illustrates the central measurement distinction in the paper. Corn and soybeans account for much larger harvested acreage than rice and cotton, so they can appear important in raw crop-assigned fire counts. Soybeans have 34,259 crop-assigned detections and corn has 20,395 detections in the positive-acreage sample. Rice has fewer total detections, 13,680, because it is observed in fewer counties and covers much less acreage. After normalizing by harvested acreage, however, rice has by far the highest aggregate detected burning intensity, at 1.111 detections per 1,000 harvested acres. Cotton ranks second at 0.129, followed by wheat at 0.122, soybeans at 0.090, and corn at 0.078. This difference shows why acreage normalization is a substantive measurement correction rather than a cosmetic rescaling. A measure based only on raw fire counts would emphasize the scale of agricultural systems. The acreage-normalized measure instead asks where detected burning is high relative to the crop system being measured. This distinction motivates the formal measurement framework in the next section.

Figure 1 provides a geographic overview of the acreage-normalized burning-intensity measure. The map includes 2,285 counties with positive harvested acreage for at least one of the five target crops in the measurement panel. Of these counties, 1,920 have at least one crop-assigned VIIRS detection during 2012–2024. The spatial pattern reinforces the distinction between agricultural scale and detected burning intensity: high-intensity counties are not simply the largest agricultural counties, but specific crop-producing areas where detected burning is high relative to harvested acreage. This geography motivates the formal measurement framework in the next section.

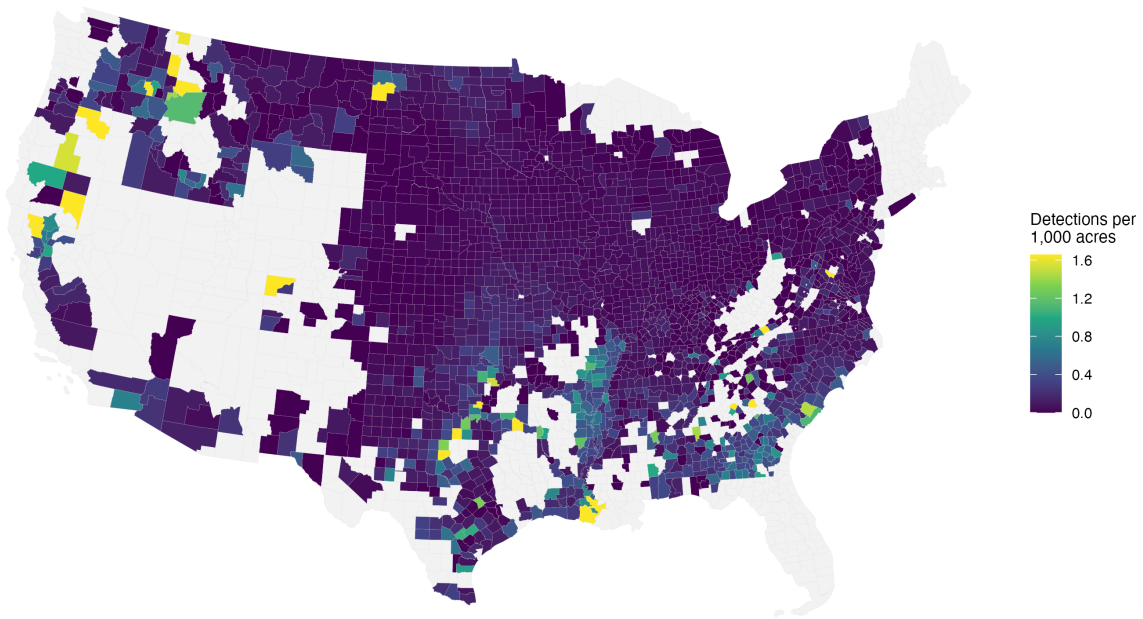


Figure 1: Geography of detected crop-residue burning intensity, 2012–2024

Notes: The figure maps county-level detected crop-residue burning intensity from 2012 to 2024. The measure is crop-assigned VIIRS active fire detections per 1,000 harvested acres, aggregated across corn, soybeans, wheat, cotton, and rice. Counties without positive harvested acreage for any of the five target crops in the measurement panel are shown in light grey. Values are capped at the 99th percentile for display.

4 Measurement Framework: From Fire Detections to County-Crop-Year Burning Intensity

This section defines the measurement object used in the paper. The starting point is a set of VIIRS active-fire detections. The measurement object is different: a crop-specific, acreage-normalized, and seasonally decomposed measure of detected residue-burning intensity at the county-crop-year level. This distinction is important. Active-fire products record thermal anomalies observed by a satellite sensor at particular places and times. They do not directly identify crop systems, burning practices, burned acreage, emissions, or economic exposure. As emphasized in the remote-sensing literature, active-fire detections are best understood as observations generated by a sensor-specific detection process, with omissions and classification challenges shaped by overpass timing, fire size, cloud cover, background conditions, and algorithm thresholds (Wooster et al., 2021). For crop-residue burning, these issues are especially relevant because fires can be small, short-lived, and followed quickly by field operations.

The framework below treats VIIRS detections as an input into economic measurement rather than as the final variable. I do not attempt to estimate emissions, burned area, or the number of true burning events. Those objects require additional assumptions about burned area, fuel loads, combustion completeness, emission factors, or field-level validation, as in emissions-inventory and burned-area studies of agricultural fires (McCarty, 2011; Hall et al., 2021). Instead, I construct a reproducible agricultural measurement variable that can be linked to standard county-crop-year agricultural statistics. The framework addresses three measurement problems. First, fire detections must be assigned to crop systems. Second, crop-specific fire counts must be separated from the scale of crop production. Third, fire activity must be interpreted in relation to crop-specific seasonal timing. Together, these steps transform raw satellite detections into a measure of detected burning intensity that is comparable across crops, counties, and years.

The county-crop-year is the main unit because it is the level at which many U.S. agricultural statistics are reported and the level at which the measure can be linked to crop acreage, production, local labor markets, smoke exposure, weather, and other agricultural and environmental variables. A pixel-level measure would remain closer to the satellite observation, but it would be less aligned with the economic data system. A state-year measure would be simpler to construct, but it would obscure local crop heterogeneity. The county-crop-year therefore balances spatial detail, crop specificity, and compatibility with agricultural economic datasets. Let c index crops, g counties, y years, and m months. The five crop groups are corn, soybeans, wheat, cotton, and rice. Let F_{cgy} denote the number of VIIRS active-fire

detections assigned to crop c in county g and year y, and let A_{cgy} denote USDA QuickStats harvested acreage for the same crop, county, and year. The framework begins with F_{cgy}, but the economic object is not the raw count. It is the number of crop-assigned detections relative to the size of the crop system.

4.1 Crop assignment

The first measurement problem is crop attribution. A satellite fire detection is a physical observation of a thermal anomaly, not a direct observation of an agricultural practice. By itself, the detection does not identify whether the fire occurred on cropland, whether it was associated with crop-residue management, or which crop system was present at the detection location. This is the first point at which economic measurement requires an additional layer beyond the satellite product.

I assign fire detections to crop systems using the USDA Cropland Data Layer. For each detection, I identify whether the detection falls on pixels classified as one of the five crop groups studied in the paper. A detection assigned to rice contributes to $F_{rice,gy}$, while a detection assigned to corn contributes to $F_{corn,gy}$. Detections outside these five crop groups are not used to construct the crop-specific burning measures. This step converts generic thermal detections into crop-specific fire counts. The crop-assignment step follows the logic of the crop-mapping literature, which treats products such as the Cropland Data Layer as operational land-cover information that can be used as crop masks, training data, and benchmark layers (Zhang et al., 2025). At the same time, I do *not* treat the crop map as error-free. Validation work shows that the CDL performs well for major U.S. crops, while also documenting crop- and region-specific classification error (Lark et al., 2021). Refinement studies further show that some errors arise near field boundaries, mixed pixels, roads, and small non-crop features (Lin et al., 2022). These issues are relevant because point-level fire detections may fall near these locations. For this reason, F_{cgy} should be interpreted as a CDL-based count of VIIRS detections assigned to crop c in county g and year y. It is not a field-level administrative record of burning, and it is not a complete census of residue-burning events. The advantage of the assignment rule is that it is transparent, reproducible, and scalable across the national panel. It creates the crop-specific satellite input needed for the burning-intensity index, while leaving room for empirical validation of whether the resulting crop-specific measures behave as expected.

4.2 Acreage normalization

The second measurement problem is scale. Raw crop-specific fire counts combine two objects: the extent of the crop system and the intensity of detected burning. A county with many corn acres may have many corn-assigned fire detections even if detected burning is uncommon relative to corn area. A smaller rice-producing county may have fewer detections in levels but more detections per acre. For economic measurement, this distinction is central. Raw counts answer where many crop-pixel fires are detected. They do not answer where detected burning is intensive relative to the crop system being measured.

To separate detected burning intensity from crop scale, I normalize crop-specific fire detections by USDA QuickStats harvested acreage. The baseline burning-intensity measure is:

$$BI_{cgy} = \frac{F_{cgy}}{A_{cgy}} \times 1000.$$

Here, BI_{cgy} is detected burning intensity for crop c in county g and year y. The factor of 1,000 expresses the measure as fire detections per 1,000 harvested acres. This normalization changes the interpretation of the satellite measure. Raw fire counts measure the level of detected fire activity. BI_{cgy} measures detected fire activity relative to the size of the crop system.

I define the baseline measure only for county-crop-years with positive reported harvested acreage. This restriction ensures that the denominator corresponds to an observed crop system in USDA statistics. It also avoids assigning burning intensity to crop-county-year cells where the crop is not reported as harvested. County-crop-years with no reported harvested acreage are therefore not assigned a burning-intensity value for that crop. The normalization is not a cosmetic rescaling. It changes the empirical object. Large crop systems can dominate raw fire-count rankings because they occupy more land and create more opportunities for detections. Acreage-normalized intensity instead identifies crop systems where detected burning is high relative to harvested area. This is the object most closely aligned with the paper's purpose: measuring crop-residue burning as an agricultural practice rather than measuring total fire activity.

4.3 Seasonal decomposition

The third measurement problem is seasonal interpretation. Crop-residue burning is not only crop-specific; it is also tied to the agricultural calendar. Fires occurring during post-harvest or field-preparation windows have a different measurement interpretation from detections scattered throughout the year. Earlier remote-sensing studies of agricultural burning show

that cropland fires are strongly seasonal and vary by crop and region (McCarty et al., 2007, 2009). Policy and emissions studies of crop-residue burning also emphasize that timing matters for interpretation, exposure, and environmental accounting (Liu et al., 2020; Lan et al., 2022). I therefore use seasonal timing as an additional diagnostic layer rather than treating all crop-assigned detections as equivalent.

Let F_{cgym} denote the number of fire detections assigned to crop c in county g, year y, and month m. For each crop, I identify the three months with the largest national shares of crop-specific fire detections. These months define the crop's peak burning season. Let $\mathcal{P}_c$ denote the set of peak months for crop c.

Peak-season fire detections are:

$$F^P_{cgy} = \sum_{m \in \mathcal{P}_c} F_{cgym}.$$

Off-peak fire detections are:

$$F^O_{cgy} = \sum_{m \notin \mathcal{P}_c} F_{cgym}.$$

Using these quantities, I define peak-season burning intensity as:

$$PSBI_{cgy} = \frac{F^P_{cgy}}{A_{cgy}} \times 1000,$$

and off-peak burning intensity as:

$$OPBI_{cgy} = \frac{F^O_{cgy}}{A_{cgy}} \times 1000.$$

The comparison between $PSBI_{cgy}$ and $OPBI_{cgy}$ separates temporally concentrated fire activity from more diffuse fire activity. Peak-season burning intensity is not intended to replace annual burning intensity in all settings. Annual burning intensity captures the full crop-pixel fire signal per harvested acre. Peak-season burning intensity captures the portion of that signal occurring in the crop's high-burning months. Off-peak burning intensity captures the remaining crop-pixel fire activity and helps assess whether the fire signal is concentrated or diffuse.

Finally, I define the peak fire share:

$$PeakShare_{cgy} = \frac{F^P_{cgy}}{F_{cgy}}.$$

This share measures the fraction of annual crop-specific fire detections that occurs during

peak months. It is defined for county-crop-years with positive annual crop-assigned fire detections. A high value indicates that detected fire activity is concentrated in a crop-specific seasonal window. A lower value indicates that the crop-pixel fire signal is more diffuse across the year. This distinction is useful because seasonal concentration is likely to be crop-specific. A narrow peak-season measure may be informative for crops with a well-defined burning calendar, but less informative for crops with multiple harvest periods, regional variation in timing, or more heterogeneous field-management practices.

4.4 Measurement comparisons

The framework produces three nested measures of fire activity. The first is the unfiltered measure: all VIIRS fire detections in a county-year. This measure captures total detected fire activity but does not distinguish agricultural fires from other fires. The second is the crop-assigned measure: VIIRS fire detections that fall on pixels assigned to the five crop groups studied in the paper. This measure restricts attention to detected fire activity associated with the selected crop systems. The third is the acreage-normalized measure: crop-assigned fire detections per 1,000 harvested acres. This measure adjusts for the scale of the crop system. I use the comparison among these measures as part of the measurement exercise. If raw fire counts and acreage-normalized intensity identified the same crops and counties, the transformations would mainly change units. If they identify different crops, county-crop units, and high-burning locations, then measurement choice changes the empirical object. This comparison turns the paper from a data-construction exercise into a measurement-consequences exercise. It asks not only whether an acreage-normalized measure can be built, but also what would be learned from alternative measures.

The three measures answer different questions. Raw fire counts identify high-fire places. Crop-assigned counts identify detected fire activity on the five crop systems studied in the paper. Acreage-normalized intensity identifies crop systems where detected burning is high relative to harvested acreage. The last object is the most appropriate measure for comparing detected burning practices across crops and counties of different scale. The results below therefore compare raw VIIRS fire counts, crop-assigned fire counts, and acreage-normalized detected burning intensity to show how each measurement step changes the interpretation of agricultural burning.

4.5 Validation logic

The framework produces three related crop-specific objects: annual burning intensity, peak-season burning intensity, and off-peak burning intensity. I use several validation exercises

to evaluate whether these variables behave like crop-specific agricultural practice measures. These exercises are not causal tests, and they are not event-level audits of individual fire detections. They are construct-validity checks conducted at the level at which the measure is designed to be used: the county-crop-year.

First, I examine the consequences of the measurement choices. I compare raw fire-count rankings with acreage-normalized intensity rankings across crops and county-crop units. This exercise tests whether normalization merely rescales raw detections or instead changes the empirical object being measured. I also examine county-crop comparisons, geographic plausibility, and denominator robustness to assess whether high-intensity crop systems appear in plausible locations and whether the crop-level ranking is sensitive to very small harvested-acreage denominators.

Second, I test same-crop construct validity. If fire detections assigned to crop c capture activity associated with that crop system, then counties and years with more harvested acreage of crop c should also tend to have more crop-c fire detections, conditional on county and year fixed effects. The baseline validation regression is:

$$\log(A_{cgy}) = \beta_c \log(1 + F_{cgy}) + \alpha_g + \lambda_y + \varepsilon_{cgy}.$$

This regression is estimated separately by crop. The coefficient β_c is not interpreted causally. It is a construct-validity statistic. It asks whether crop-specific fire detections line up with crop-specific harvested acreage within counties over time. I write the regression with harvested acreage as the dependent variable because the diagnostic asks whether the satellite-based crop assignment recovers variation in the corresponding USDA crop system, not because fire detections are assumed to determine acreage. The direction of the regression is therefore diagnostic rather than behavioral: it evaluates whether the satellite-assigned fire signal contains same-crop information at the county-crop-year level. The test is intentionally modest. It does not prove that every fire is residue burning; it checks whether the aggregate crop-assigned measure contains same-crop information.

Third, I estimate crop-by-crop placebo tests. In these tests, harvested acreage for crop c is regressed on fire detections assigned to crop j. The same-crop relationship corresponds to $c = j$, while different-crop relationships correspond to $c \neq j$. A useful crop-specific measure should show stronger and more consistently positive same-crop relationships than different-crop relationships. The purpose is not to require all off-crop relationships to be zero. Crop systems are spatially correlated within counties, and agricultural counties may have multiple crops, shared land-management practices, and common fire conditions. The question is whether the same-crop signal is more systematic than the off-crop signal.

Fourth, I compare peak-season and off-peak fire counts. For each crop, I estimate:

$$\log(A_{cgy}) = \beta_{1c} \log(1 + F^{P}_{cgy}) + \beta_{2c} \log(1 + F^{O}_{cgy}) + \alpha_g + \lambda_y + \varepsilon_{cgy}.$$

This joint specification tests whether peak-season fire detections contain crop-specific information after accounting for off-peak detections. The purpose is not to show that the peak-season measure dominates for every crop. A one-size-fits-all seasonal correction would be inappropriate if crops differ in harvest timing, regional calendars, or field-management practices. Instead, the seasonal decomposition reveals whether the crop-specific fire signal is concentrated in crop-plausible burning windows or spread more diffusely across the year.

Together, these steps define the paper's economic measurement object: crop-specific, acreage-normalized, and seasonally decomposed detected residue-burning intensity at the county-crop-year level. The next section evaluates whether this object behaves like a crop-specific agricultural practice measure. The following section then asks how much the empirical picture changes when researchers use alternative fire measures instead.

5 Construct Validity of the Burning-Intensity Measure

This section evaluates whether the constructed measures behave like crop-specific agricultural practice variables at the county-crop-year level. The validation task differs from the task in remote-sensing algorithm studies. I do *not* attempt to verify each individual thermal anomaly as a residue-burning event, *nor* do I estimate the true burned area of each field. Those exercises require near-coincident reference data, field-level labels, or high-resolution burn-scar mapping (Wooster et al., 2021; Hall et al., 2021). The object developed here is more aggregated and economic: a crop-specific measure of detected burning intensity that can be linked to USDA county-crop-year data. The validation question is therefore whether this constructed variable contains crop-specific and agriculturally meaningful information. This distinction matters because much of the existing literature validates different objects. Remote-sensing studies show that agricultural fires can be detected from space and that crop-residue burning is geographically and seasonally structured (McCarty et al., 2007, 2009). Emissions-inventory work uses satellite-derived fire or burned-area information to estimate pollution-relevant quantities (McCarty, 2011). Policy studies often take satellite fire or burn indicators as outcomes in evaluations of bans, incentives, or payments for ecosystem services (Jack et al., 2025; Khanna et al., 2025). These papers demonstrate the usefulness of satellite-based burning measures, but they do not directly answer the measurement question posed here: whether fire detections, once assigned to crop pixels and normalized by har-

vested acreage, behave like a crop-specific agricultural measure. I focus on that intermediate object.

The exercises below are construct-validity checks. They ask whether the measure is internally consistent with crop assignment, crop acreage, crop specificity, and the distribution of burning across crop systems. I focus on four dimensions. First, I document the scale of the crop-assignment restriction, which defines the crop-specific fire-count input used throughout the paper. Second, I test whether same-crop fire detections move with same-crop harvested acreage within counties over time. Third, I compare these same-crop relationships with crop-by-crop placebo relationships. Fourth, I examine whether the measures capture meaningful differences in the prevalence and concentration of detected burning across crops. Taken together, these checks do not establish field-level ground truth. They evaluate whether the aggregate measure has the structure expected of a crop-specific agricultural practice variable.

5.1 Crop assignment benchmark

The first check documents the crop-assignment step. Raw VIIRS fire detections capture all detected thermal activity, not only fires on cropland and not only fires associated with the five crop systems studied in this paper. Without land-cover assignment, a county-year fire count would combine target crop-pixel detections with fires on other agricultural land, forests, grasslands, developed land, and other non-target land covers. Such a measure may be useful for describing total fire activity, but it is too broad to serve as a crop-specific residue-burning measure for corn, soybeans, wheat, cotton, or rice. Appendix Table 14 reports the benchmark. Across 2012–2024, the balanced county-month file contains 6,054,814 VIIRS fire detections. Of these, 315,807 detections fall on pixels assigned to the five target crop groups. Target crop-pixel detections therefore account for 5.2 percent of all VIIRS detections, with annual shares ranging from 3.6 percent to 7.7 percent.

This benchmark clarifies the object being validated in the rest of the section. The validation exercises do not ask whether all VIIRS detections are agricultural, or whether the national fire record can be interpreted as crop-residue burning. They ask whether the subset of detections assigned to target crop pixels behaves like a crop-specific agricultural measure once combined with USDA acreage data. Among target crop-pixel detections, soybeans and corn account for the largest numbers of detections, followed by wheat, rice, and cotton. These raw crop-pixel counts describe the distribution of crop-assigned detections, but they do not yet show whether the resulting crop-specific variables line up with crop systems in the agricultural data. The next checks address that question directly.

5.2 Same-crop construct validity

The second check asks whether crop-specific fire detections move with same-crop harvested acreage within counties over time. If detections assigned to crop c contain information about activity associated with that crop system, then counties and years with more harvested acreage of crop c should tend to have more detections assigned to crop c, after accounting for time-invariant county differences and common year shocks. This is a deliberately modest test. It does not ask whether acreage causes fire detections, or whether fire detections cause acreage. It asks whether the crop-assignment procedure produces a measure that lines up with the agricultural system to which it is assigned.

Table 2 reports the same-crop validation regressions. The dependent variable is log harvested acreage for each crop. The explanatory variable is log one plus fire detections assigned to the same crop. The regressions include county and year fixed effects, and standard errors are clustered by county.

Table 2: Crop-specific fire detections and same-crop harvested acreage

Crop	Log harvested acreage
Corn	0.046***
	(0.009)
Soybeans	0.033***
	(0.006)
Wheat	0.103***
	(0.017)
Cotton	0.085***
	(0.023)
Rice	0.107***
	(0.035)
County fixed effects	Yes
Year fixed effects	Yes
Standard errors	Clustered by county

Notes: Each row reports a separate county-year fixed-effects regression for the listed crop. The dependent variable is log harvested acreage for that crop. The explanatory variable is log one plus CDL-filtered VIIRS agricultural fire detections assigned to the same crop. The sample covers county-years from 2012 to 2024. Standard errors are clustered by county. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The same-crop coefficients are positive and statistically significant for all five crops. This pattern indicates that crop-assigned detections contain crop-specific information rather than

only generic fire-prone conditions. The relationship is strongest for wheat, rice, and cotton, and remains positive for corn and soybeans. These differences across crops are informative. They suggest that the strength of the satellite-acreage relationship is not uniform, which is consistent with the broader remote-sensing literature showing that agricultural burning differs across crops, regions, and management calendars (McCarty et al., 2007, 2009).

The same-crop exercise also has an important limitation. Harvested acreage is mechanically related to the opportunity for crop-assigned fire detections: more acres create more locations where fires could be observed. A positive same-crop coefficient is therefore necessary but not sufficient for validation. It shows that the constructed measure is not disconnected from the crop system, but it does not by itself prove crop specificity. For this reason, I treat the same-crop regressions as a first construct-validity check and then move to a more demanding crop-by-crop placebo exercise. That exercise asks whether fire detections assigned to a given crop are more systematically related to that crop than to other crops grown in the same agricultural landscape.

5.3 Crop-by-crop placebo evidence

The third check compares same-crop relationships with different-crop relationships. This exercise addresses a limitation that is common in satellite-based agricultural measurement. Agricultural systems are spatially correlated. A county with many soybean acres may also have corn acres, wheat acres, agricultural fires, and other land-management activity. If a fire measure simply captures general agricultural scale or regional fire conditions, it should be similarly related to several crops. If the crop-assignment procedure is meaningful, the same-crop relationship should be more systematic than the off-crop relationships.

Table 3 summarizes the crop-by-crop placebo exercise. Same-crop coefficients are diagonal entries from regressions that pair a crop's harvested acreage with fire detections assigned to the same crop. Different-crop coefficients are off-diagonal entries from regressions that pair harvested acreage for one crop with fire detections assigned to another crop.

Table 3: Same-crop and different-crop placebo diagnostics

Coefficient group	Coefficients	Mean	Median	Min	Max	Positive	Negative	Sig. 10%	Sig. 5%	Sig. 1%
Same crop	5	0.075	0.085	0.033	0.107	5	0	5	5	5
Different crop	20	0.001	-0.002	-0.101	0.099	9	11	9	7	0

Notes: The table summarizes the crop-by-crop placebo matrix. Same-crop coefficients are diagonal entries where the harvested-acreage crop and fire-detection crop are the same. Different-crop coefficients are off-diagonal entries. Significance counts report the number of coefficients statistically significant at each threshold.

The same-crop coefficients are uniformly positive and statistically significant. By con-

trast, the different-crop coefficients are mixed in sign and centered closer to zero. This pattern supports the interpretation that the constructed variables capture crop-specific burning signals rather than only general fire activity or broad agricultural scale. The placebo evidence is useful because it evaluates crop assignment at the level where classification matters for economic analysis: not whether a single pixel is labeled correctly, but whether the aggregate crop-specific measure behaves differently from off-crop alternatives.

The placebo results also clarify the limits of the measure. Some off-crop relationships remain statistically significant. This is not surprising. Counties often specialize in related crop systems, crop rotations create correlations across crops, and regional fire conditions can affect multiple land uses. A strict requirement that all off-crop coefficients be zero would be unrealistic. The relevant test is comparative: whether same-crop relationships are more consistently positive and stronger on average than different-crop relationships. This is the pattern observed in the table. The result supports the use of the measure as a crop-specific aggregate indicator, while preserving the cautious interpretation emphasized in the active-fire and CDL literatures (Wooster et al., 2021; Lark et al., 2021).

5.4 Extensive and intensive margins

The final check examines whether the measure captures meaningful differences in the prevalence and concentration of detected burning across crops. A crop can have high aggregate burning intensity for different reasons. Detected burning may be widespread, with many county-crop-years recording at least one fire. Alternatively, total fire activity may be concentrated in a small number of high-burning county-crop-years. These two patterns have different measurement implications. A widespread extensive margin suggests that detected burning is a common feature of the crop system. A concentrated intensive margin suggests that the national crop signal is driven by a smaller set of local agricultural systems.

Figure 2 reports these two margins by crop. The first measure is the share of county-crop-years with any detected fire. The second is the share of total crop-specific fires accounted for by the top 10 percent of county-crop-years. Rice has a high extensive margin, meaning that detected burning is common across rice county-crop-years. Wheat has a more concentrated pattern, with a large share of detections occurring in a relatively small number of high-burning county-crop-years.

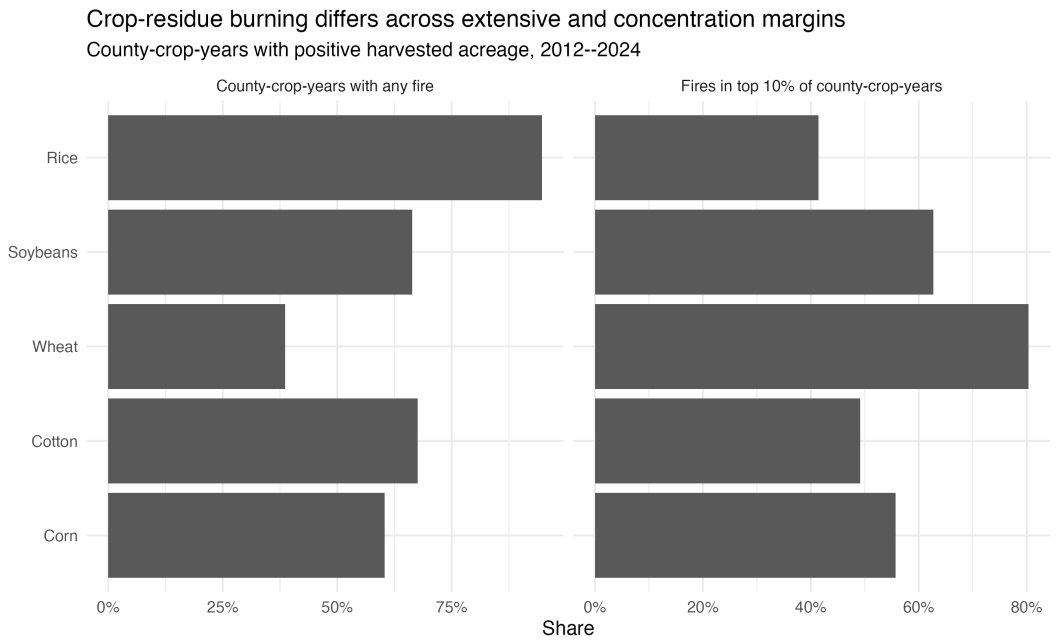


Figure 2: Extensive margin and spatial concentration of crop-specific fire detections

Table 4 reports the corresponding statistics. The results reinforce a point that is often visible in local remote-sensing studies but less explicit in aggregate agricultural data: crop-residue burning is not a single uniform practice across crop systems. Rice burning is relatively widespread within rice-growing county-years, while wheat burning is more spatially concentrated. Corn, soybeans, and cotton occupy intermediate positions. These differences would be difficult to detect from crop-specific totals alone because totals combine crop scale, geography, and burning concentration.

Table 4: Extensive and intensive margins of crop-residue burning

Crop	County-years	Harvested acres (million)	Fire detections	Any fire share	Aggregate intensity fires/1,000 acres	Conditional intensity fires/1,000 acres	Top 1% fire share	Top 5% fire share	Top 10% fire share
Rice	301	12.3	13,680	0.947	1.111	1.237	0.074	0.254	0.414
Cotton	1434	40.1	5,161	0.676	0.129	0.259	0.116	0.333	0.492
Wheat	4405	123.3	15,103	0.386	0.122	0.379	0.380	0.666	0.803
Soybeans	6567	382.0	34,259	0.664	0.090	0.148	0.214	0.489	0.627
Corn	5854	262.4	20,395	0.604	0.078	0.242	0.159	0.400	0.557

Notes: The table summarizes extensive and intensive margins of crop-specific fire detections for county-crop-years with positive harvested acreage. Any fire is the share of county-crop-years with at least one CDL-filtered VIIRS fire detection assigned to the crop. Aggregate intensity is total fire detections divided by total harvested acreage and multiplied by 1,000. Conditional intensity is the mean county-crop-year intensity among county-crop-years with at least one fire detection. Top 1%, 5%, and 10% columns report the share of crop-specific fire detections accounted for by the highest-fire county-crop-years within each crop.

The extensive and intensive margins also help connect the measurement exercise to the literature on agricultural burning. Earlier studies showed that cropland fires vary across

regions and seasons (McCarty et al., 2007, 2009). Emissions and policy studies then used fire or burn measures to analyze environmental consequences or behavioral responses (McCarty, 2011; Jack et al., 2025; Khanna et al., 2025). The margin decomposition here fills a narrower but important gap. It shows how detected burning is distributed within crop systems once detections are assigned to crops and expressed in county-crop-year form. This distribution matters for economic measurement because a national crop total may conceal whether burning is a broad practice or a localized one.

Taken together, the validation results support the measurement framework while also clarifying its scope. The crop-assignment benchmark defines the subset of the VIIRS fire record used to construct the crop-specific measures. Same-crop regressions show that crop-assigned detections move with the corresponding crop systems within counties over time. Crop-by-crop placebo tests show that the same-crop signal is more systematic than off-crop alternatives. Extensive and intensive margin checks show that detected burning differs across crops not only in levels, but also in prevalence and concentration. These results do not convert VIIRS detections into a census of residue-burning events. They show that the constructed county-crop-year variables contain crop-specific and agriculturally meaningful information. The next section asks what changes when researchers use less processed fire measures instead of the acreage-normalized intensity measure.

6 Consequences of Measurement Choice

The validation results show that the crop-assigned measures contain crop-specific information. This section asks a different question: how much does the empirical picture change when researchers use alternative fire measures? This question is central to the paper's measurement contribution because satellite fire products are often aggregated directly to administrative units. Such aggregation is transparent, scalable, and easy to link to other datasets, but it can combine several distinct objects: general fire activity, agricultural land cover, crop scale, and burning intensity. I compare three measurement objects. The first is the raw VIIRS fire record, which captures all detected fire activity regardless of land cover or crop system. The second is the crop-assigned fire count, which restricts detections to pixels classified as corn, soybeans, wheat, cotton, or rice in the Cropland Data Layer. The third is the acreage-normalized measure, which divides crop-assigned fire detections by USDA harvested acreage and expresses the result as detections per 1,000 harvested acres. These measures answer different questions. Raw fire counts measure total detected fire activity. Crop-assigned counts measure detected fire activity on the five crop systems studied in the paper. Acreage-normalized intensity measures detected burning relative to the size of the

crop system.

The comparison is a measurement-consequences exercise rather than another validation exercise. The previous section asked whether the constructed crop-specific variables behave like agricultural practice measures. This section asks whether alternative aggregations would lead researchers to identify the same high-burning crops and places. If raw counts, crop-assigned counts, and acreage-normalized intensity ranked crops and counties similarly, the transformations would mainly change units. If they rank them differently, measurement choice changes the substantive interpretation of agricultural burning. Previous studies have shown that cropland fires can be detected, mapped, and used in emissions, air-quality, and policy applications (McCarty et al., 2007, 2009; McCarty, 2011; Zamanialaei et al., 2023; Jack et al., 2025; Khanna et al., 2025). The exercise here is prior to those applications. It asks whether different ways of constructing the fire measure lead researchers to identify the same agricultural systems. I therefore treat measurement choice as a comparison problem. A measure that identifies large crop systems with many total detections is not the same as a measure that identifies crop systems where detected burning is intensive relative to acreage.

6.1 Crop rankings and county-crop rankings

Table 5 summarizes the consequences of these measurement choices. Panel A reports the crop-assignment restriction documented in the previous section and uses it as the starting point for the comparison. Only 5.2 percent of all VIIRS detections fall on pixels assigned to the five crop groups studied in the paper, so the first measurement choice already separates broad fire activity from detected fire activity on the selected crop systems. Panel B shows that acreage normalization changes crop rankings. In raw crop-assigned counts, soybeans rank first and corn ranks second, reflecting the large scale of these crop systems. After normalization by harvested acreage, rice ranks first and cotton ranks second. Rice moves from fourth in raw counts to first in intensity, while cotton moves from fifth to second. By contrast, soybeans move from first to fourth, and corn moves from second to fifth. Wheat remains third under both rankings. The transformation therefore does not simply rescale fire detections. It changes the interpretation from where many crop-pixel fires are detected to where detected burning is most intensive relative to harvested acreage. Panel C shows that the same distinction matters within crops across counties. Raw-count and intensity rankings are positively correlated, as expected if both measures contain information about crop-pixel fire activity. However, they identify different top county-crop units. For corn and soybeans, none of the top 10 raw-count county-crop units are also top 10 intensity units. For cotton, rice, and wheat, the corresponding overlap is only 10 percent. Top-25 overlap is

also limited, ranging from 7.4 percent for wheat to 40.0 percent for rice. Median rank gaps are especially large for corn, soybeans, and wheat, indicating that many county-crop units move substantially when the measure changes from counts to intensity.

Table 5: Consequences of measurement choice, 2012–2024

Panel A. Crop assignment narrows the raw fire record

Measure	Value
All VIIRS fire detections	6,054,814
Crop-assigned fire detections, five target crops	315,807
Crop-assigned share of all detections	5.2%

Panel B. Acreage normalization changes crop rankings

Crop	Raw-count rank	Intensity rank	Rank change	Acreage-normalized intensity fires/1,000 acres
Rice	4	1	+3	1.111
Cotton	5	2	+3	0.129
Wheat	3	3	0	0.122
Soybeans	1	4	-3	0.090
Corn	2	5	-3	0.078

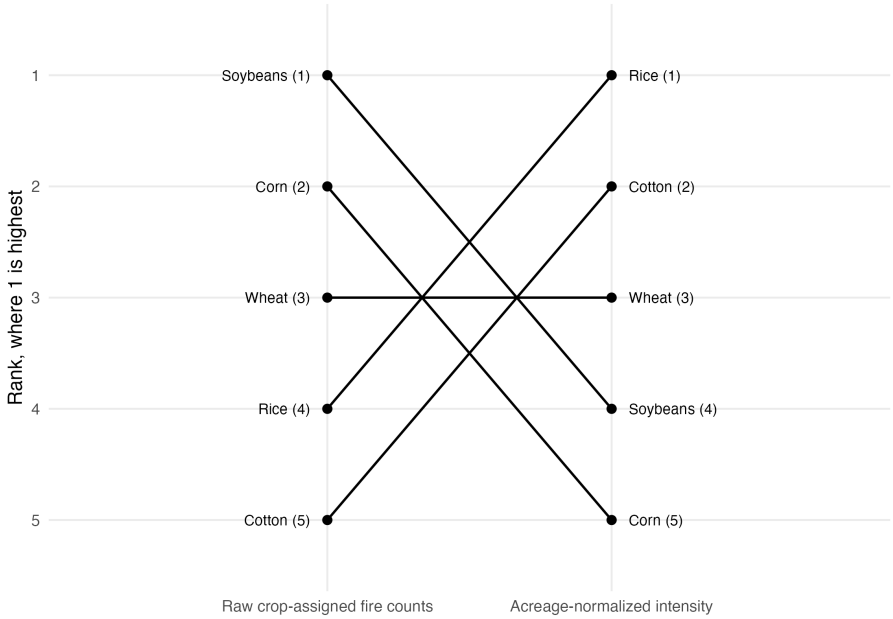
Panel C. Raw-count and intensity rankings identify different county-crop units

Crop	Spearman correlation	Top-10 overlap	Top-25 overlap	Median rank gap
Corn	0.707	0.0%	11.5%	819
Cotton	0.729	10.0%	28.0%	141
Rice	0.541	10.0%	40.0%	20
Soybeans	0.640	0.0%	8.0%	690
Wheat	0.912	10.0%	7.4%	615

Notes: Panel A compares all VIIRS active fire detections with crop-assigned detections on Cropland Data Layer pixels classified as corn, soybeans, wheat, cotton, or rice. Panel B compares crop rankings based on raw crop-assigned fire detections and acreage-normalized detected burning intensity. Acreage-normalized intensity is crop-assigned fire detections divided by harvested acreage and multiplied by 1,000. Panel C compares county-crop rankings based on raw crop-assigned fire counts and acreage-normalized intensity within each crop. Top-10 and top-25 overlap report the share of highest-ranked county-crop units under the raw-count ranking that also appear among the highest-ranked units under the intensity ranking.

Figure 3 visualizes the crop-level ranking shift. The figure shows that the high-count crops are *not* necessarily the high-intensity crops. Soybeans and corn dominate raw crop-assigned detections, while rice and cotton dominate the acreage-normalized measure. This is the central empirical result of the section. Raw counts do not simply measure the same object with noise. They often emphasize a different object: crop scale rather than burning intensity.

Figure 3: Measurement choice changes crop rankings



Notes: The figure compares crop rankings based on raw crop-assigned fire detections and acreage-normalized detected burning intensity. Rank 1 denotes the highest value under each measure. The acreage-normalized measure is fire detections per 1,000 harvested acres.

Figure 4 extends the ranking comparison to the county level. The first panel identifies counties in the top decile of raw crop-assigned fire detections, while the second identifies counties in the top decile of acreage-normalized detected burning intensity. The overlap between the two groups is limited: only 72 counties appear in both top-decile groups. This means that most counties identified as high burning under raw crop-assigned counts are not identified as high burning under the acreage-normalized intensity measure. The figure therefore provides a geographic version of the paper's main measurement result. Raw crop-assigned fire counts identify places with many total crop-pixel detections, while acreage-normalized intensity identifies places where detected burning is high relative to harvested acreage. Appendix Figure 7 complements this comparison by showing that the geography of detected burning intensity also differs sharply across crop systems.

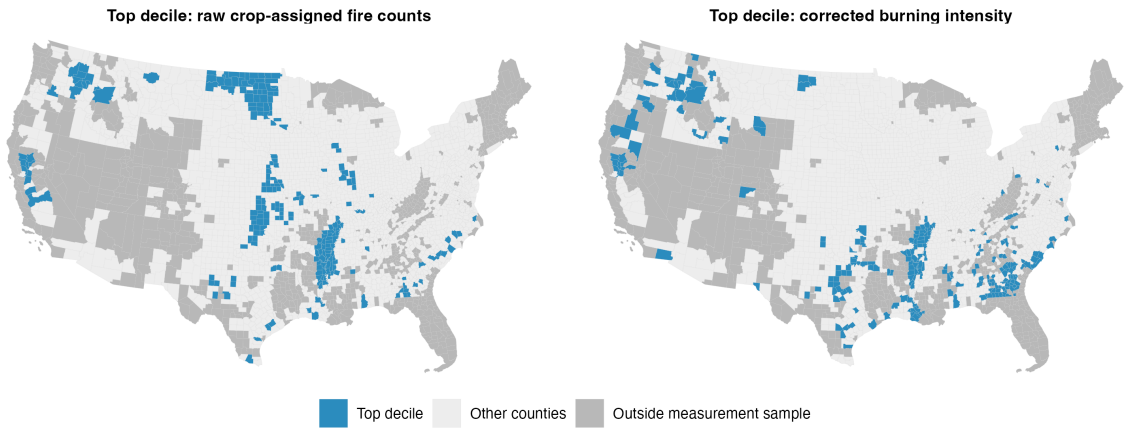


Figure 4: Measurement choice changes county classification

Notes: The figure compares counties in the top decile of raw crop-assigned fire detections with counties in the top decile of acreage-normalized detected burning intensity. Acreage-normalized intensity is crop-assigned VIIRS active fire detections per 1,000 harvested acres, aggregated across corn, soybeans, wheat, cotton, and rice from 2012 to 2024. The comparison sample includes counties with positive harvested acreage for at least one of the five crop groups in the measurement panel. Counties outside the measurement sample are shown separately.

These results clarify the economic role of the measurement framework. The paper does not claim that one satellite aggregation is universally appropriate for all questions. A researcher interested in total detected fire activity may want raw fire counts; a researcher interested in fires on specific crop systems needs crop assignment; and a researcher interested in residue burning as a crop management practice relative to crop scale needs acreage normalization. The last object is the focus of this paper because it is the measure most closely aligned with comparing detected burning practices across crops and counties of different agricultural scale.

6.2 County-crop classification

The county-level comparison in Figure 4 shows that measurement choice changes the broad geography of high-burning counties. I next examine the same issue within crop systems by asking whether the highest-ranked county-crop units are the same under raw counts and acreage-normalized intensity. Table 6 compares the highest-ranked county-crop unit within each crop under the two measurement choices. To avoid selecting very small denominator cases, the comparison is restricted to county-crop units observed for at least three years, with at least 10,000 average annual harvested acres and at least 10 crop-assigned fire detections

over the sample period. Even with these safeguards, the top county-crop unit changes for four of the five crops. Corn is the only crop for which the same county is selected under both raw counts and intensity. For cotton, rice, soybeans, and wheat, the raw-count and intensity measures identify different top county-crop systems.

Table 6: County-crop targeting consequences of measurement choice, 2012–2024

Crop	Top county by raw counts	Fires	Intensity	Top county by intensity	Fires	Intensity	Same
Corn	Mississippi, MO	439	1.60	Mississippi, MO	439	1.60	Yes
Cotton	Reagan, TX	121	0.77	Decatur, GA	99	0.79	No
Rice	Colusa, CA	1,262	1.67	Washington, MS	175	3.11	No
Soybeans	Poinsett, AR	1,067	0.72	Lawrence, AR	374	1.03	No
Wheat	Williams, ND	722	0.60	Columbia, WA	658	1.83	No

Notes: The table compares the highest-ranked county-crop unit within each crop under two measurement choices. The raw-count target is the county-crop unit with the largest number of crop-assigned VIIRS fire detections over 2012–2024. The intensity target is the county-crop unit with the largest number of crop-assigned fire detections per 1,000 harvested acres over 2012–2024. The targeting sample is restricted to county-crop units observed for at least three years, with at least 10,000 average annual harvested acres and at least 10 crop-assigned fire detections over the sample period. Intensity is computed as aggregate crop-assigned fire detections divided by aggregate harvested acres and multiplied by 1,000. The table illustrates how measurement choice can change which locations appear to be the highest-burning county-crop systems.

The table illustrates why the measurement distinction is substantively important. Raw counts tend to identify large crop systems with many total detections. The intensity measure can identify county-crop systems with fewer total detections but substantially higher detections per 1,000 harvested acres. For example, Colusa County, California has the largest number of rice fire detections, while Washington County, Mississippi has the highest rice burning intensity in the restricted comparison sample. Similarly, Poinsett County, Arkansas leads soybeans by raw counts, while Lawrence County, Arkansas leads by intensity. These differences show that acreage normalization changes not only crop rankings but also which local agricultural systems appear most relevant for measurement, monitoring, and subsequent empirical analysis. This local comparison is useful because many economic applications begin by defining high-exposure or high-activity units. In subsequent empirical work, environmental accounting, or descriptive agricultural monitoring, the selected set of "high-burning" places may become the basis for later comparisons. If raw counts emphasize high-scale systems while intensity emphasizes high-burning-practice systems, then the initial measurement choice can shape the empirical exercise before any regression is estimated.

6.3 Seasonal timing as a measurement consequence

Measurement choice also affects how seasonal burning patterns are interpreted. Raw crop-assigned fire counts identify where many crop-pixel detections occur. Acreage-normalized intensity identifies where detected burning is high relative to harvested acreage. Seasonal timing adds a further dimension: some crop systems have concentrated peak-season fire activity, while others have more diffuse fire activity across the year.

Table 7 links the ranking comparison to crop-specific seasonal timing. Rice is the clearest case. It moves from fourth in raw crop-assigned fire counts to first in acreage-normalized intensity, and 80.0 percent of rice fire detections occur in the crop's peak months. In the joint validation regression, the rice peak-season estimate is positive and precisely estimated, while the off-peak estimate is close to zero. This pattern suggests that the acreage-normalized measure identifies a crop system with both high detected burning intensity and a concentrated seasonal signal.

Table 7: Seasonal consequences of measurement choice, 2012–2024

Crop	Raw rank	Intensity rank	Peak months	Peak share	Peak/off-peak estimates
Rice	4	1	Oct, Sep, Nov	80.0%	0.108 (0.027) / 0.001 (0.020)
Cotton	5	2	Mar, Oct, Sep	33.9%	0.077 (0.022) / 0.049 (0.023)
Wheat	3	3	Sep, Aug, Jul	40.2%	0.072 (0.016) / 0.077 (0.019)
Soybeans	1	4	Oct, Apr, Sep	46.4%	0.016 (0.004) / 0.033 (0.006)
Corn	2	5	Apr, Sep, Mar	47.0%	0.052 (0.008) / 0.018 (0.010)

Notes: The table links the measurement-choice comparison to crop-specific seasonal timing. Raw rank is based on crop-assigned fire counts, while intensity rank is based on crop-assigned fire detections per 1,000 harvested acres. Peak months are the three months with the largest national shares of crop-specific fire detections. Peak share is the aggregate share of crop-specific fire detections occurring in the crop's peak months. Peak/off-peak estimates report coefficients from the joint validation regression using peak-season and off-peak fire detections; standard errors are in parentheses. The table is intended as a compact bridge between the measurement-choice results and the seasonal concentration analysis.

The other crops show why seasonal concentration should be treated as a diagnostic refinement rather than a universal correction. Cotton also moves up after acreage normalization, from fifth in raw counts to second in intensity, and its peak-season estimate is larger than its off-peak estimate. However, off-peak cotton fire detections remain informative. Corn moves down after acreage normalization, but its peak-season signal remains stronger than its off-peak signal. By contrast, soybeans and wheat have more diffuse seasonal patterns: their off-peak estimates are at least as large as their peak-season estimates in the joint validation exercise. These differences show that measurement choice changes not only crop and county rankings, but also the interpretation of seasonal fire timing. A single national

fire calendar would miss this heterogeneity. The seasonal results therefore preview the more detailed seasonal decomposition in Section 7, where timing is evaluated as a crop-specific diagnostic rather than as a mechanical filter.

6.4 Illustrative linkage to HMS smoke

The acreage-normalized burning-intensity measure is designed to be linked to external agricultural, environmental, and economic datasets. As an illustrative environmental linkage, I merge the county-year burning measures with NOAA Hazard Mapping System smoke data. This exercise is descriptive and does not estimate the causal effect of crop-residue burning on smoke exposure. Its purpose is to show how alternative burning measures identify different county-years when brought into contact with an external environmental dataset.

Table 8 compares county-years with at least 10,000 observed acres in the five crop groups to three top-decile groups: county-years with high all-fire counts, high crop-assigned fire counts, and high acreage-normalized burning intensity. The differences across groups reinforce the measurement argument. County-years in the top decile of all VIIRS fires have many total detections, but these detections are not necessarily concentrated on pixels in the five crop groups. County-years in the top decile of crop-assigned fire counts have the largest number of crop fires and the largest crop acreage. By contrast, county-years in the top decile of acreage-normalized burning intensity have smaller crop acreage but the highest detections per 1,000 crop acres.

Table 8: Illustrative linkage between burning measures and HMS smoke, 2012–2024

Group	County-years	Counties	Crop acres	All fires	Crop fires	Intensity	Smoke days	Medium/heavy
All sufficient-area county-years	12,622	1,783	95.3	99.1	16.5	0.229	78.6	24.7
Top decile: all VIIRS fires	1,267	325	84.6	611.5	65.4	0.711	85.6	25.2
Top decile: crop-assigned fires	1,267	295	154.5	274.8	100.9	1.020	87.5	27.3
Top decile: corrected intensity	1,263	441	61.9	313.2	82.9	1.285	82.7	20.6

Notes: The table reports descriptive means for county-years with at least 10,000 observed target-crop acres. Crop acres are observed target-crop acres in thousands. All fires are all VIIRS fire detections in the county-year. Crop fires are detections assigned to the five target crop groups. Intensity is crop-assigned fire detections per 1,000 observed target-crop acres, winsorized at the 99th percentile among county-years with sufficient target-crop area. Smoke days are annual counts of days with any NOAA HMS smoke polygon over the county. Medium/heavy reports the sum of days with medium or heavy HMS smoke. Top-decile groups are defined separately using all VIIRS fire counts, crop-assigned fire counts, and corrected burning intensity. The table is descriptive and does not estimate a causal effect of burning on smoke.

The smoke patterns should be interpreted cautiously. All sufficient-area county-years have an average of 78.6 HMS smoke days. The top-decile groups have somewhat higher smoke-day counts, ranging from 82.7 days for acreage-normalized-intensity county-years to 87.5 days for crop-assigned fire-count county-years. Medium and heavy smoke days are highest for the crop-assigned fire-count group. These differences are descriptive and may reflect many sources of smoke, including non-agricultural fires and transported smoke. I therefore do not interpret the table as evidence that crop-residue burning causes smoke exposure. The

main point is instead that the measurement framework creates a crop-specific object that can be linked transparently to environmental datasets while preserving the distinction between total fire activity, crop-assigned fire activity, and acreage-normalized detected burning intensity.

This linkage is useful because many environmental applications begin by choosing the fire measure before estimating exposure or pollution relationships. Emissions and air-quality studies often move from burning measures to pollution consequences (McCarty, 2011; Lan et al., 2022; Zamanialaei et al., 2023). The exercise here stops one step earlier. It asks whether the acreage-normalized agricultural measure identifies a different set of county-years before any causal pollution model is estimated. That distinction is important because a pollution or exposure analysis built on raw fire counts may begin from a different geography than one built on acreage-normalized detected burning intensity.

6.5 Illustrative linkage to agricultural labor markets

The same measurement object can also be linked to economic datasets. As a second descriptive linkage, I merge the county-year burning measures with county-year employment and wage measures from the Quarterly Census of Employment and Wages. This exercise does not estimate employment effects. Instead, it asks whether the acreage-normalized burning-intensity measure identifies county-years that differ in agricultural labor-market structure from those identified by raw fire counts or crop-assigned fire counts.

Table 9 compares county-years with at least 10,000 observed acres in the five crop groups and matched QCEW data to three top-decile groups. The top decile by all VIIRS fire counts has the largest average agricultural employment and agricultural wages, but this group reflects all fire activity and is not restricted to crop-pixel detections. The top decile by crop-assigned fire counts has the largest crop acreage and the largest number of crop fires. By contrast, the top decile by acreage-normalized burning intensity has smaller crop acreage than the crop-assigned fire-count group but higher detections per 1,000 crop acres.

Table 9: Illustrative linkage between burning measures and agricultural labor markets, 2012–2024

Group	County-years	Counties	Crop acres	Crop fires	Intensity	Ag. emp.	Ag. emp. share	Ag. estab.	Ag. wages
All sufficient-area county-years	12,565	1,780	95.3	16.6	0.229	288.4	2.8%	37.8	9.8
Top decile: all VIIRS fires	1,260	324	84.8	65.6	0.712	1404.3	3.2%	92.1	44.5
Top decile: crop-assigned fires	1,258	295	154.9	101.3	1.023	236.8	3.7%	54.3	8.7
Top decile: corrected intensity	1,257	440	61.9	83.1	1.287	474.1	4.1%	56.8	16.3

Notes: The table reports descriptive means for county-years with at least 10,000 observed target-crop acres and matched QCEW labor-market data. Crop acres are observed target-crop acres in thousands. Crop fires are VIIRS detections assigned to the five target crop groups. Intensity is crop-assigned fire detections per 1,000 observed target-crop acres, winsorized at the 99th percentile among county-years with sufficient target-crop area. Ag. emp. is average annual private agricultural employment. Ag. emp. share is private agricultural employment divided by total private employment. Ag. estab. is average annual private agricultural establishments. Ag. wages are annual private agricultural wages in millions of dollars. Top-decile groups are defined separately using all VIIRS fire counts, crop-assigned fire counts, and corrected burning intensity. The table is descriptive and does not estimate employment effects.

The labor-market patterns reinforce the measurement interpretation. County-years in the acreage-normalized-intensity top decile have a higher agricultural employment share than the full sufficient-area sample and the top-decile all-fire group. They also include more counties than the crop-assigned fire-count top decile. Thus, the acreage-normalized measure identifies a broader set of agricultural counties where detected burning is intensive relative to crop scale, not only counties with the largest number of crop-pixel fires. This is the sense in which the measure is economically useful: it can be joined to standard county-year data while retaining a clear interpretation as a crop-specific burning-intensity measure.

Taken together, the results in this section show that measurement choice changes the empirical picture of crop-residue burning. Raw fire counts emphasize total fire activity and large agricultural systems. Crop assignment restricts the object to detections on pixels in the five crop groups studied in the paper. Acreage normalization identifies crop systems and counties where detected burning is high relative to crop scale. Seasonal decomposition shows that the timing of the fire signal is crop-specific rather than uniform. The HMS and QCEW linkages then show that these measurement choices carry through when the burning measure is connected to external environmental and economic datasets. The main lesson is not that one measure is universally best for every application. It is that researchers must match the fire measure to the economic object being studied. For crop-residue burning as an agricultural practice, the acreage-normalized crop-specific measure provides a substantially different interpretation of crops, counties, and linked datasets than raw fire counts alone.

7 Seasonal Concentration

The previous section showed that measurement choice changes crop rankings, county-crop classification, seasonal interpretation, and external linkages. This section develops the seasonal component of the measure in more detail. The motivation is straightforward: crop-residue burning is not only crop-specific, but also calendar-specific. A fire detection that occurs during a crop's post-harvest or field-preparation window has a different measurement interpretation from a detection that occurs in a month with little connection to the crop calendar. This does *not* mean that all peak-season detections are residue burning, or that all off-peak detections are irrelevant. It means that timing provides an additional diagnostic for assessing whether crop-assigned satellite detections behave like an agricultural practice measure.

Seasonality has been central to previous remote-sensing studies of agricultural burning. Work on agricultural fires in the southeastern United States and the contiguous United States shows that cropland fire activity is not randomly distributed across the year, but

varies with crop systems, harvest timing, and regional land-management practices (McCarty et al., 2007, 2009). Emissions and air-quality studies also treat timing as important because burning in different months can have different implications for emissions inventories, atmospheric transport, and exposure (McCarty, 2011; Lan et al., 2022). The question in this paper is narrower. I do not use seasonality to estimate emissions or exposure. I use it to evaluate whether the constructed county-crop-year fire measure contains temporally structured information consistent with crop-specific burning calendars. I construct crop-specific peak burning seasons using the national monthly distribution of fire detections assigned to each crop. For each crop, I define the peak season as the three calendar months with the largest shares of crop-specific fire detections. These peak months are then used to decompose annual fire detections into peak-season and off-peak components. This choice is intentionally simple. It avoids imposing a single national burning calendar across crops, and it avoids relying on crop-specific agronomic assumptions that may vary across regions. The resulting seasonal measures should therefore be interpreted as diagnostic components of detected burning intensity, not as definitive classifications of residue-burning events.

7.1 Crop-specific fire seasonality

Figure 5 shows the monthly distribution of crop-specific fire detections. The figure reveals substantial differences across crops. Rice has the clearest seasonal pattern, with detections concentrated in September, October, and November. Corn and soybeans also show visible seasonal peaks, although their detections are less concentrated than rice. Wheat and cotton have more diffuse seasonal patterns. These differences are informative because they show that the crop-assigned fire signal is not simply a uniform annual background process. It has crop-specific temporal structure. The seasonal patterns also show why a single seasonal correction would be inappropriate. Rice detections are sharply concentrated, while soybean, wheat, and cotton detections are more spread across the year. A narrow peak-season measure is therefore likely to be more informative for some crops than for others. This point is consistent with the broader literature on crop-residue burning, where the timing of burning depends on crop calendars, local management practices, and regional production systems (McCarty et al., 2007; Liu et al., 2020). I therefore treat seasonality as a crop-specific measurement dimension rather than as a universal filter.

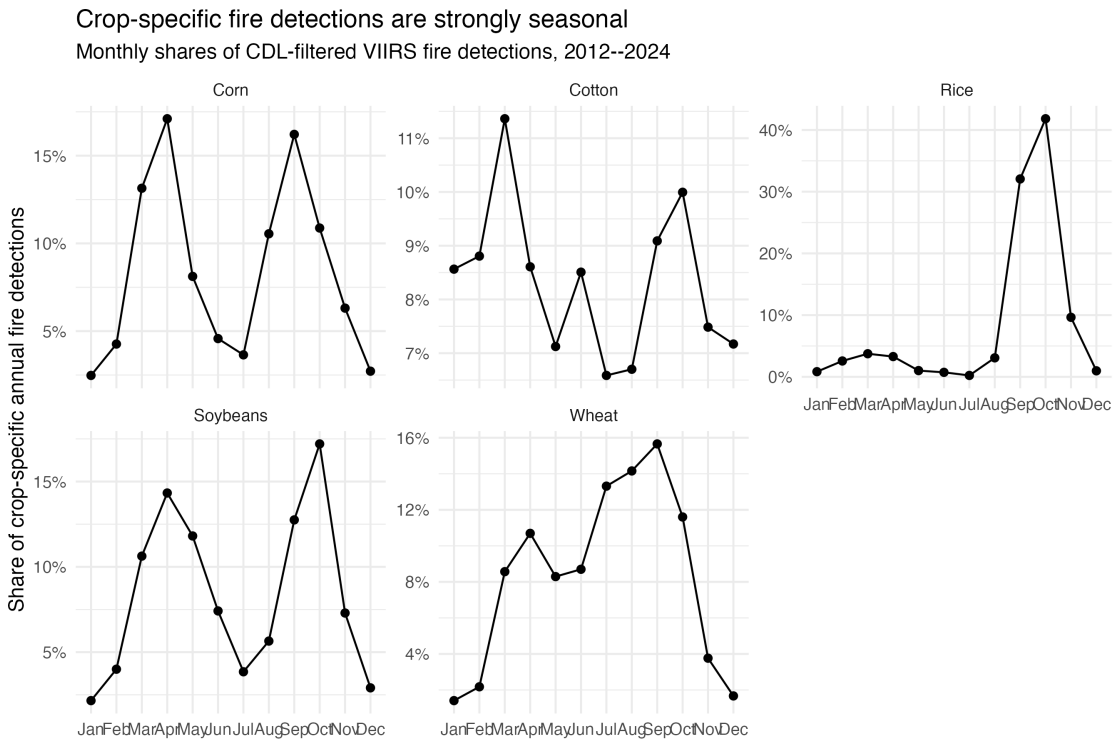


Figure 5: Crop-specific fire seasonality, 2012–2024

The peak fire shares reported in Table 7 reinforce this interpretation. Rice has by far the most concentrated seasonal signal, with 80.0 percent of crop-specific detections occurring in its three peak months. Corn and soybeans have peak shares below one half, while wheat and cotton have more diffuse seasonal distributions. Seasonal concentration is therefore not a mechanical feature of all crop-assigned fire detections. It varies across crop systems and provides a useful diagnostic for distinguishing temporally concentrated burning signals from more diffuse crop-pixel fire activity.

7.2 Peak-season and off-peak burning measures

Using the crop-specific peak months, I decompose annual fire detections into peak-season and off-peak components. For crop c in county g and year y, peak-season fire detections are the sum of monthly detections in the three peak months for that crop. Off-peak fire detections are the sum of monthly detections in the remaining months. Appendix Table 18 reports the crop-specific peak months used in this decomposition.

The construct-validity regression compares these two components directly:

$$\log(A_{cgy}) = \beta_{1c} \log(1 + F^{P}_{cgy}) + \beta_{2c} \log(1 + F^{O}_{cgy}) + \alpha_g + \lambda_y + \varepsilon_{cgy}.$$

Here, A_{cgy} is harvested acreage, F^{P}_{cgy} is peak-season fire detections, and F^{O}_{cgy} is off-peak fire detections. The regression is estimated separately by crop and includes county and year fixed effects. The coefficients are not interpreted causally. They are measurement diagnostics that ask whether peak-season and off-peak detections contain different amounts of same-crop information.

Table 10 reports the joint construct-validity results. Peak-season fire detections are positive and statistically significant for all five crops after controlling for off-peak detections. This pattern indicates that the seasonal component contains independent crop-specific information. The comparison with off-peak detections, however, varies substantially across crops. For rice, the peak-season coefficient is large and statistically significant, while the off-peak coefficient is close to zero. For corn, the peak-season coefficient is also substantially larger than the off-peak coefficient. For cotton, peak-season detections are more strongly related to acreage than off-peak detections, although off-peak detections remain informative. For soybeans and wheat, both components contain crop-specific information, indicating a more diffuse seasonal fire signal.

Table 10: Seasonal validation of crop-specific burning measures

Crop	Peak-season fires	SE	Off-peak fires	SE	Observations
Corn	0.052***	(0.008)	0.018*	(0.010)	5854
Soybeans	0.016***	(0.004)	0.033***	(0.006)	6567
Wheat	0.072***	(0.016)	0.077***	(0.019)	4405
Cotton	0.077***	(0.022)	0.049**	(0.023)	1434
Rice	0.108***	(0.027)	0.001	(0.020)	301
County fixed effects	Yes		Yes		
Year fixed effects	Yes		Yes		

Notes: Each row reports a separate county-year fixed-effects regression for the listed crop. The dependent variable is log harvested acreage. The explanatory variables are log one plus peak-season fire detections and log one plus off-peak fire detections assigned to the same crop. Peak-season months are the three months with the largest national shares of crop-specific fire detections. Standard errors are clustered by county. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

These results support a careful interpretation of the seasonal index. Peak-season burning intensity is useful, but it should not be treated as a universal replacement for annual burning intensity. For crops with sharply concentrated burning calendars, especially rice and to a lesser extent corn, peak-season detections provide a more concentrated crop-specific signal. For crops with more diffuse fire timing, annual burning intensity remains informative because relevant fire activity is not confined to three months. This distinction matters because a seasonal filter that improves measurement for one crop could discard meaningful information for another. The results also clarify what is missing from a purely annual measure. Annual

burning intensity summarizes total crop-pixel fire activity per harvested acre, but it does not reveal whether that activity is temporally concentrated in plausible management windows. Peak-season and off-peak measures add this diagnostic layer. They allow the paper to distinguish a high-intensity crop system with a concentrated burning calendar from one where fire detections are spread across the year.

7.3 Seasonal measurement performance by crop

Table 11 summarizes the crop-level performance of the seasonal decomposition. The table reports the aggregate peak fire share, the share of county-crop-years with at least one peak-season fire, and the coefficients from the joint peak-season and off-peak construct-validity regressions.

Table 11: Seasonal measurement performance by crop

Crop	Peak fire share	Peak-fire county-years	Peak coef.	SE	Off-peak coef.	SE
Rice	80.0%	88.4%	0.108	(0.027)	0.001	(0.020)
Corn	47.0%	41.5%	0.052	(0.008)	0.018	(0.010)
Soybeans	46.4%	44.0%	0.016	(0.004)	0.033	(0.006)
Wheat	40.2%	22.2%	0.072	(0.016)	0.077	(0.019)
Cotton	33.9%	40.4%	0.077	(0.022)	0.049	(0.023)

Notes: Peak fire share is the share of crop-specific fire detections occurring in the crop's three peak burning months. Peak-fire county-years is the share of county-crop-years with at least one fire detection in a peak month. Peak and off-peak coefficients come from crop-specific county-year fixed-effects regressions of log harvested acreage on log one plus peak-season fire detections and log one plus off-peak fire detections, estimated jointly. Standard errors are clustered by county.

The crop-level patterns are informative. Rice has the strongest seasonal concentration: 80.0 percent of fire detections in the acreage-normalized analysis sample occur in peak months, and 88.4 percent of rice county-crop-years with positive harvested acreage have at least one peak-season fire. The construct-validity results show that the rice fire signal is almost entirely a peak-season signal. This pattern is consistent with a relatively well-defined seasonal rice burning signal. It also helps explain why rice becomes the highest-intensity crop after acreage normalization: the rice signal is not only high relative to acreage, but also temporally concentrated. Corn also shows a stronger peak-season than off-peak relationship, although the seasonal concentration is less extreme than for rice. Cotton has a stronger peak-season coefficient as well, but off-peak detections remain significant, suggesting that cotton-related fire activity is not fully captured by a narrow peak window. Soybeans and wheat show more diffuse seasonal patterns. For soybeans, the off-peak coefficient is larger than the peak-season coefficient. For wheat, the two coefficients are similar. These patterns

suggest that the relevant burning calendar for soybeans and wheat is not fully captured by a national three-month peak-season definition. This may reflect regional differences in crop calendars, multiple management windows, or fire detections on crop pixels that are less tightly linked to a single residue-burning season.

The seasonal results refine the measurement framework in two ways. First, they show that temporal concentration is itself an informative feature of the crop-specific fire signal. Second, they show that crop-residue burning should not be measured with a one-size-fits-all seasonal adjustment. Annual burning intensity, peak-season burning intensity, and off-peak burning intensity are complementary measures. Their relative usefulness depends on how sharply the fire signal is concentrated within each crop system. Overall, the seasonal decomposition strengthens the interpretation of the satellite-based measure. It shows that the constructed fire variables are not only crop-specific and acreage-normalized, but also temporally structured. This temporal structure is strongest for rice and corn, more moderate for cotton, and more diffuse for soybeans and wheat. The broader implication is that agricultural measurement from satellite fire data should respect crop calendars, but should not impose the same calendar rule on all crops. Seasonality improves the measure most when it is treated as crop-specific evidence rather than as a mechanical filter.

8 Conclusions

This paper develops a framework for measuring crop-residue burning as an agricultural practice in the United States. The starting point is a measurement gap in standard agricultural statistics. County-level data provide rich information on acreage, production, yields, prices, and farm income, but they provide limited information on field-level residue-management practices. Satellite active-fire detections can help fill this gap, but raw fire counts are not themselves agricultural economic variables. They must be linked to crop systems, separated from the scale of production, and interpreted in relation to crop-specific seasonal timing. I address these measurement problems by combining VIIRS active fire detections, the USDA Cropland Data Layer, and USDA QuickStats harvested acreage. The resulting panel measures crop-specific detected burning intensity at the county-crop-year level for corn, soybeans, wheat, cotton, and rice from 2012 to 2024. The baseline measure is fire detections per 1,000 harvested acres. It is designed to measure detected burning intensity relative to the size of the crop system, rather than total detected fire activity.

The main result is that measurement choice changes the empirical picture of crop-residue burning. Raw fire counts emphasize large crop systems, especially soybeans and corn. After crop assignment and acreage normalization, rice has the highest detected burning intensity,

followed by cotton and wheat. The same distinction appears within crops across counties. County-crop units with the largest number of detections are often not the same units with the highest detections per 1,000 harvested acres. Acreage normalization therefore does more than change the units of the variable. It changes the object from total detected fire activity to detected burning intensity relative to crop scale. The construct-validity exercises support this interpretation while also clarifying the limits of the measure. Same-crop fire detections are positively related to same-crop harvested acreage within counties over time. Crop-by-crop placebo tests show that same-crop relationships are more systematic than off-crop relationships. Extensive and intensive margin checks show that detected burning differs across crop systems in both prevalence and concentration. These exercises do not turn VIIRS detections into field-level audit records of residue burning. Instead, they evaluate the measure at the level at which it is designed to be used: the county-crop-year. At that level, the constructed variables contain crop-specific and agriculturally meaningful information.

The measurement-consequences results show why these choices matter. Raw VIIRS fire counts, crop-assigned fire counts, and acreage-normalized detected burning intensity identify different crops, county-crop units, and high-burning locations. Local rankings change for most crops when the measure changes from raw counts to intensity. Seasonal timing also differs across crops. Rice has a highly concentrated peak-season signal, while soybeans and wheat display more diffuse fire activity across the calendar year. These patterns show that crop-residue burning should not be measured with a one-size-fits-all seasonal adjustment. Annual burning intensity, peak-season burning intensity, and off-peak burning intensity are complementary measures whose usefulness differs across crop systems. Descriptive linkages to NOAA HMS smoke and QCEW labor-market data further illustrate the value of the measurement object. These linkages are not causal analyses of smoke exposure or employment. They show that the acreage-normalized burning-intensity measure can be connected to external environmental and economic datasets while preserving a clear distinction between total fire activity, crop-assigned fire activity, and detected burning intensity relative to crop scale. This distinction matters because many empirical applications begin by identifying high-activity or high-exposure places. If raw fire counts and acreage-normalized burning intensity identify different places, measurement choice can shape the empirical exercise before any outcome model is estimated.

The contribution of the paper is therefore to translate satellite fire detections into an agricultural economic measurement object: crop-specific, acreage-normalized, and seasonally decomposed detected residue-burning intensity. The measure is not intended to replace burned-area products, emissions inventories, or ground-based records where they exist. It *complements* them by creating a transparent and scalable panel variable aligned with the

county-crop-year structure used in agricultural economics. Because the measure is built from public satellite, crop-map, and USDA acreage data, the framework is transparent, reproducible, and extendable to later years or to other crop systems where comparable land-cover and acreage data are available. More broadly, the paper illustrates how remote-sensing data can be converted into economic measurement tools for agricultural practices that are otherwise difficult to observe. The main lesson is that satellite detections are physical observations, while agricultural economists need economic measurement objects. For crop-residue burning, constructing that object changes what researchers identify as the relevant crops, counties, seasons, and agricultural systems.

References

Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–198.

Hall, J. V., Argueta, F., and Giglio, L. (2021). Validation of MCD64A1 and FireCCI51 cropland burned area mapping in ukraine. *International Journal of Applied Earth Observation and Geoinformation*, 102:102443.

Jack, B. K., Jayachandran, S., Kala, N., and Pande, R. (2025). Money (not) to burn: Payments for ecosystem services to reduce crop residue burning. *American Economic Review: Insights*, 7(1):39–55.

Khanna, S., Mahajan, A., and Ramanujam, S. (2025). Are crop residue burning bans effective? evidence from india. Technical Report Working Paper 428, Grantham Research Institute on Climate Change and the Environment, London School of Economics.

Kubitza, C., Krishna, V. V., Schulthess, U., and Jain, M. (2020). Estimating adoption and impacts of agricultural management practices in developing countries using satellite data: A scoping review. *Agronomy for Sustainable Development*, 40(16).

Lan, R., Eastham, S. D., Liu, T., Norford, L. K., and Barrett, S. R. H. (2022). Air quality impacts of crop residue burning in india and mitigation alternatives. *Nature Communications*, 13(1):6537.

Lark, T. J., Schelly, I. H., and Gibbs, H. K. (2021). Accuracy, bias, and improvements in mapping crops and cropland across the united states using the USDA cropland data layer. *Remote Sensing*, 13(5):968.

Lin, C., Zhong, L., Song, X., Dong, J., Lobell, D. B., and Jin, Z. (2022). Validation and refinement of cropland data layer using a spatial-temporal decision tree algorithm. *Scientific Data*, 9:63.

Liu, T., Mickley, L. J., Singh, S., Jain, M., DeFries, R. S., and Marlier, M. E. (2020). Crop residue burning practices across north india inferred from household survey data: Bridging gaps in satellite observations. *Atmospheric Environment: X*, 8:100091.

McCarty, J. L. (2011). Remote sensing-based estimates of annual and seasonal emissions from crop residue burning in the contiguous united states. *Journal of the Air & Waste Management Association*, 61(1):22–34.

McCarty, J. L., Justice, C. O., and Korontzi, S. (2007). Agricultural burning in the southeastern united states detected by MODIS. *Remote Sensing of Environment*, 108(2):151–162.

McCarty, J. L., Korontzi, S., Justice, C. O., and Loboda, T. (2009). The spatial and temporal distribution of crop residue burning in the contiguous united states. *Science of the Total Environment*, 407(21):5701–5712.

NASA Fire Information for Resource Management System (2026). Fire Information for Resource Management System (FIRMS). `https://firms.modaps.eosdis.nasa.gov/`. Accessed July 2026.

Ramsey, R. et al. (2024). Review of approaches and challenges for the validation of satellite-based active fire products in savanna ecosystems. *International Journal of Wildland Fire*.

Schroeder, W., Oliva, P., Giglio, L., and Csiszar, I. (2014). The new VIIRS 375 m active fire detection data product: Algorithm description and initial assessment. *Remote Sensing of Environment*, 143:85–96.

Schroeder, W., Oliva, P., Giglio, L., Quayle, B., Lorenz, E., and Morelli, F. (2016). Active fire detection using Landsat-8/OLI data. *Remote Sensing of Environment*, 185:210–220.

Travis, K. R. et al. (2023). Emission factors for crop residue and prescribed fires in the eastern US during FIREX-AQ. *Journal of Geophysical Research: Atmospheres*, 128.

USDA National Agricultural Statistics Service (2026a). Cropland Data Layer. `https://www.nass.usda.gov/Research_and_Science/Cropland/`. Accessed July 2026.

USDA National Agricultural Statistics Service (2026b). Quick Stats. `https://www.nass.usda.gov/Quick_Stats/`. Accessed July 2026.

Wooster, M. J., Roberts, G. J., Giglio, L., Roy, D. P., Freeborn, P. H., Justice, C. O., et al. (2021). Satellite remote sensing of active fires: History and current status, applications and future requirements. *Remote Sensing of Environment*, 267:112694.

Wuepper, D., Oluoch, S., and Hadi, A. (2025). Satellite data in agricultural and environmental economics: Theory and practice. *Agricultural Economics*.

Zamanialaei, M., Shew, A. M., Fain, J. J., Borkowski, A., and McCarty, J. L. (2023). Crop residue burning from high-resolution satellite imagery and $PM_{2.5}$ dispersion: A case study of mississippi county, arkansas, usa. *Sustainable Environment*, 9(1).

Zhang, C., Kerner, H., Wang, S., Hao, P., Li, Z., Hunt, K. A., Abernethy, J., Zhao, H., Gao, F., Di, L., et al. (2025). Remote sensing for crop mapping: A perspective on current and future crop-specific land cover data products. *Remote Sensing of Environment*, 330:114995.

Zhang, T., Wooster, M. J., and Xu, W. (2017). Approaches for synergistically exploiting VIIRS I- and M-band data in regional active fire detection and FRP assessment: A demonstration with respect to agricultural residue burning in eastern china. *Remote Sensing of Environment*, 198:407–424.

A Additional Tables and Figures

This appendix reports supplementary figures and tables that support the measurement exercises in the main text. The appendix has four purposes. First, it documents annual variation in the constructed crop-specific burning-intensity measure. Second, it reports the full crop-by-crop placebo matrix underlying the compact construct-validity evidence in the main text. Third, it provides additional checks on the measurement choices, including crop assignment, geographic plausibility, and denominator robustness. Fourth, it reports the crop-specific peak-month definitions used to construct the seasonal decomposition.

These appendix materials do not introduce a separate empirical argument. They provide supporting evidence for the paper's main claim that crop assignment, acreage normalization, and crop-specific timing change the interpretation of satellite fire detections.

A.1 Annual crop-residue burning intensity trends

Appendix Figure 6 plots annual crop-level trends in acreage-normalized detected burning intensity. The figure reports crop-specific fire detections per 1,000 harvested acres by year. These trends describe how the constructed measure varies across crops and over time. I use

the figure to characterize the time-series behavior of the measure, not to explain the causes of annual changes.

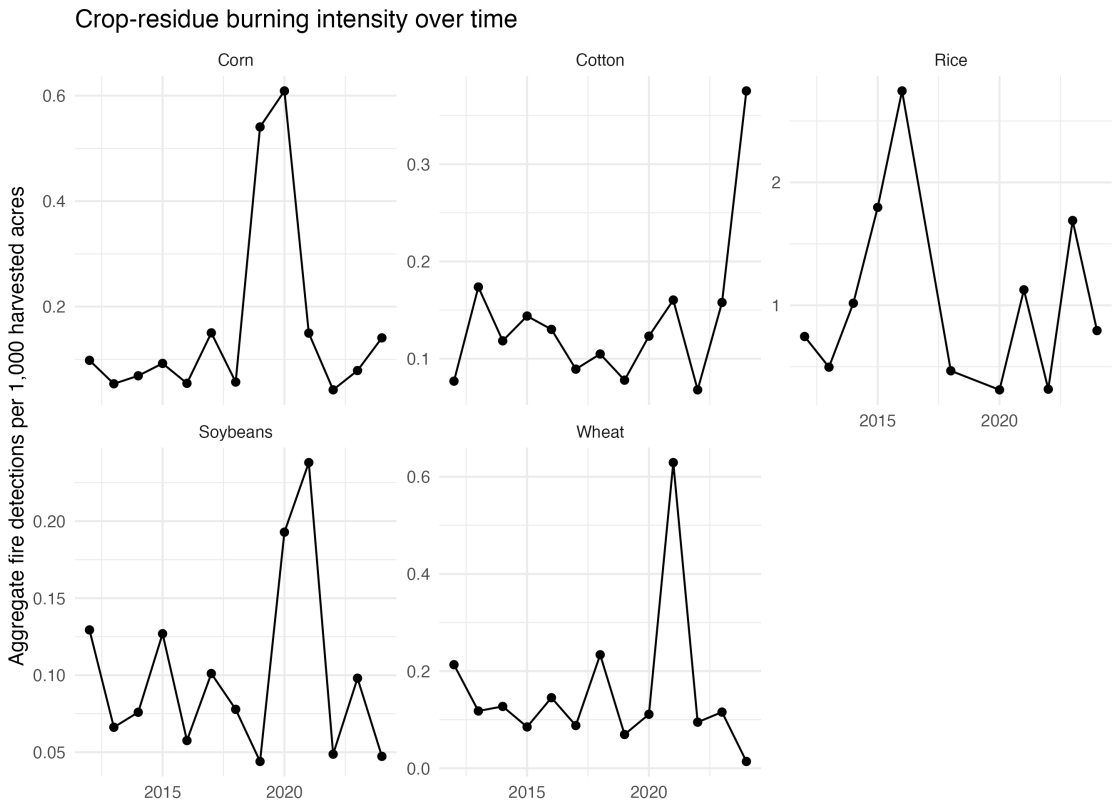


Figure 6: Annual crop-residue burning intensity trends, 2012–2024

Appendix Table 12 reports the corresponding annual trend statistics. The table summarizes aggregate detected burning intensity by crop and the distribution of county-level burning intensity within crop-years. It documents that the measure contains both cross-crop and within-crop variation. The table should be interpreted descriptively. Its purpose is to characterize the constructed panel, not to estimate the drivers of year-to-year changes in detected burning intensity.

Table 12: Annual crop-residue burning intensity trends, 2012–2024

Crop	Years	Harvested acres (million)	Fire detections	Mean aggregate fires/1,000 acres	Min aggregate fires/1,000 acres	Max aggregate fires/1,000 acres	Mean county fires/1,000 acres	Median county fires/1,000 acres	P90 county fires/1,000 acres
Rice	11	12.3	13,680	1.046	0.312	2.744	1.135	0.891	2.101
Corn	13	262.4	20,395	0.164	0.042	0.609	0.223	0.097	0.526
Wheat	13	123.3	15,103	0.157	0.014	0.629	0.217	0.024	0.573
Cotton	13	40.1	5,161	0.138	0.068	0.375	0.169	0.088	0.414
Soybeans	13	382.0	34,259	0.100	0.044	0.238	0.099	0.032	0.232

Notes: The table summarizes annual crop-level burning intensity for county-crop-years with positive harvested acreage. Aggregate intensity is calculated as total crop-specific fire detections divided by total harvested acreage and multiplied by 1,000. County moments summarize the distribution of county-level fire detections per 1,000 harvested acres across county-crop-years within each crop-year.

A.2 Full crop-by-crop placebo matrix

Appendix Table 13 reports the full crop-by-crop placebo matrix. Each cell comes from a separate county-year fixed-effects regression of log harvested acreage for the row crop on log one plus fire detections assigned to the column crop. Diagonal entries correspond to same-crop relationships. Off-diagonal entries correspond to different-crop placebo relationships.

Table 13: Crop-by-crop placebo matrix: harvested acreage and crop-specific fire detections

	Fire-detection crop				
Harvested-acreage crop	Corn	Soybeans	Wheat	Cotton	Rice
Corn	0.046***	-0.019**	-0.033**	-0.003	-0.055**
	(0.009)	(0.008)	(0.013)	(0.018)	(0.021)
Soybeans	0.000	0.033***	-0.016*	-0.001	-0.017
	(0.006)	(0.006)	(0.010)	(0.021)	(0.013)
Wheat	0.008	0.003	0.103***	0.043*	-0.101**
	(0.015)	(0.015)	(0.017)	(0.026)	(0.041)
Cotton	0.006	-0.014	-0.016	0.085***	0.035
	(0.020)	(0.024)	(0.028)	(0.023)	(0.037)
Rice	0.040**	0.099**	-0.005	0.058**	0.107***
	(0.019)	(0.043)	(0.024)	(0.023)	(0.035)

Notes: Each cell reports a separate county-year fixed-effects regression. The dependent variable is log harvested acreage for the row crop. The explanatory variable is log one plus CDL-filtered VIIRS agricultural fire detections assigned to the column crop. Diagonal entries therefore compare same-crop harvested acreage and same-crop fire detections, while off-diagonal entries are placebo comparisons across different crops. The sample covers county-years from 2012 to 2024. Standard errors are clustered by county. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The full matrix provides the detailed results underlying the compact same-crop versus different-crop diagnostic reported in the main text. The diagonal coefficients are uniformly positive and statistically significant. Off-diagonal coefficients are mixed in sign and centered closer to zero. This pattern is consistent with the interpretation that the constructed measures contain crop-specific information, while also reflecting the fact that crop systems are spatially correlated within counties. Counties often grow multiple crops, share regional fire conditions, and operate within common agricultural management environments. The placebo exercise should therefore not be read as requiring all off-diagonal coefficients to be zero. It is a construct-validity check that asks whether same-crop relationships are more systematic than off-crop alternatives.

A.3 Crop assignment benchmark

Appendix Table 14 reports the crop-assignment benchmark before the analysis sample is restricted to county-crop-years with positive USDA QuickStats harvested acreage. This table evaluates the first measurement step: moving from all VIIRS fire detections to detections assigned to pixels in the five crop groups studied in the paper. Across 2012–2024, the balanced county-month fire file contains 6,054,814 VIIRS active fire detections. Of these, 315,807 detections fall on pixels assigned to corn, soybeans, wheat, cotton, or rice. Target five-crop pixel detections therefore account for 5.2 percent of all VIIRS detections.

Table 14: Crop assignment benchmark, 2012–2024

Panel A. Assignment of VIIRS fire detections	
All VIIRS fire detections	6,054,814
Target five-crop pixel fire detections	315,807
Non-target or other detections	5,739,007
Target five-crop share of all detections	5.2%
Panel B. Composition of target five-crop pixel detections	
Crop	Fire detections
Soybeans	94,058
Corn	82,638
Wheat	63,464
Rice	43,489
Cotton	32,158

Notes: Panel A compares all VIIRS active fire detections in the balanced county-month fire file with detections assigned to pixels in the five target crop groups. Panel B reports the composition of these five-crop detections. The benchmark is computed before restricting the analysis to county-crop-years with positive USDA QuickStats harvested acreage.

This benchmark shows why crop assignment is a necessary first step. Raw VIIRS fire counts are much broader than the crop-residue burning object studied in this paper. They include detected fire activity from many land covers and fire types. Crop assignment narrows the satellite record to detections occurring on pixels in the five crop groups used in the paper. Among these target five-crop detections, soybeans account for the largest number of assigned detections, followed by corn, wheat, rice, and cotton. The benchmark is not the final burning-intensity measure. It is computed before merging to the positive-acreage county-crop-year analysis sample and before normalizing by USDA harvested acreage. For this reason, the crop-assignment totals in Appendix Table 14 differ from the positive-acreage analysis sample summarized in Table 1. The benchmark documents the crop-assignment step, while the main burning-intensity measure asks where detected burning is high relative to the size of the crop

system being measured.

A.4 Geographic plausibility of high-intensity crop systems

Appendix Tables 15 and 16 provide geographic plausibility checks for the acreage-normalized burning-intensity measure. These checks ask whether high-intensity crop systems appear in locations that are plausible given crop geography and regional residue-management patterns. They do not validate individual fire detections as residue-burning events. Instead, they examine whether the aggregate measure identifies credible crop-state and crop-county systems at the spatial scale used in the paper. Appendix Figure 7 complements these checks by mapping detected burning intensity separately by crop.

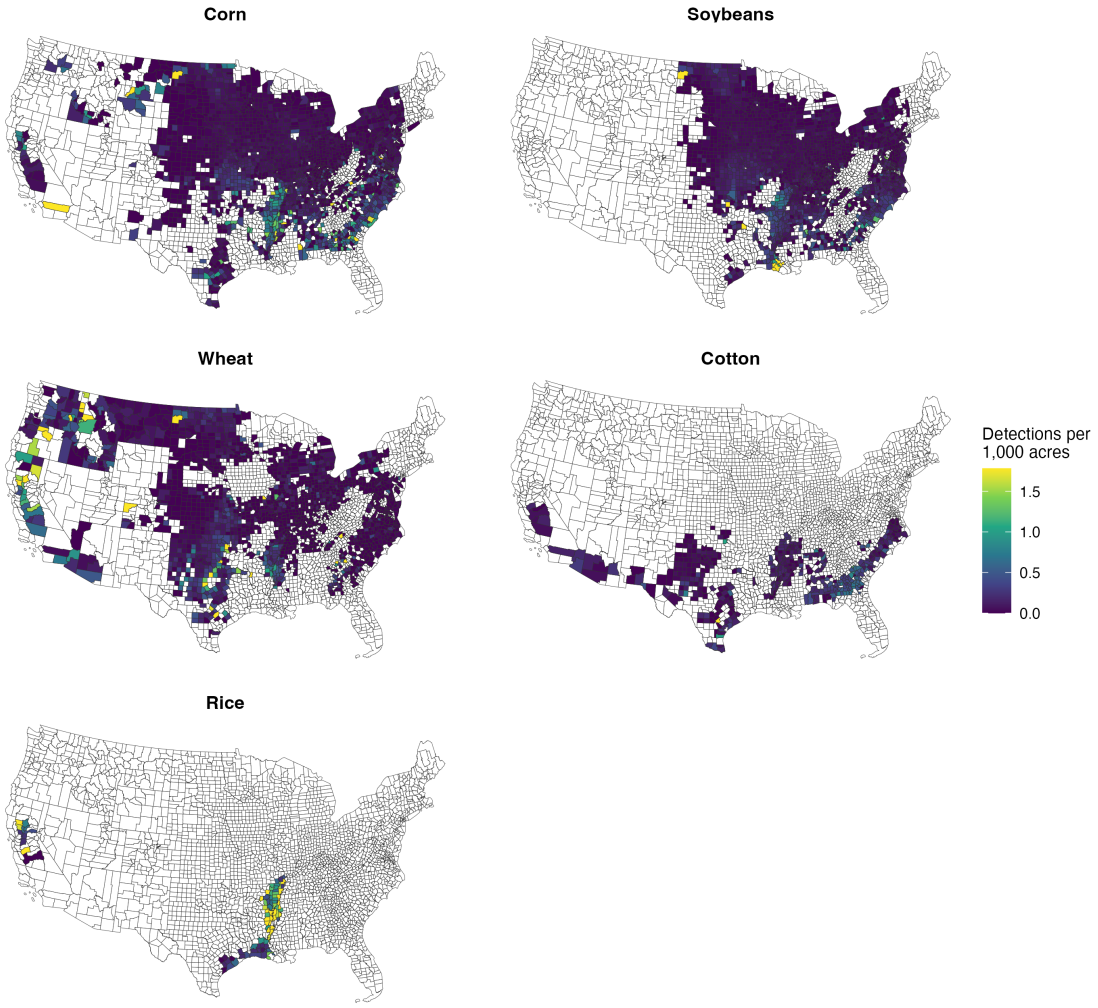


Figure 7: Crop-specific geography of detected burning intensity, 2012–2024

Notes: The figure maps county-level detected burning intensity separately for corn, soybeans, wheat, cotton, and rice. Intensity is crop-assigned VIIRS active fire detections per 1,000 harvested acres, aggregated over 2012–2024. Counties without positive harvested acreage for the crop in the measurement panel are shown in white but remain outlined. Values are capped at the 99th percentile of the pooled crop-specific intensity distribution for display.

Appendix Figure 7 shows that the geography of detected burning differs substantially across crop systems. Corn and soybeans cover broad parts of the Midwest and South, but most counties have low acreage-normalized intensity. Rice is observed in far fewer counties, but its detected burning intensity is higher and more spatially concentrated. Cotton intensity is concentrated mainly in southern cotton-growing counties, while wheat shows a broader but more uneven geography. These patterns support the interpretation that high-intensity burning is crop-specific and geographically concentrated rather than a simple reflection of national crop acreage.

Appendix Table 15 reports the five highest-intensity states within each crop. High-intensity rice systems appear in Mississippi, California, Arkansas, Missouri, and Louisiana. Cotton, corn, and soybean intensity is concentrated largely in southern and Mid-South states. Wheat intensity is more geographically mixed, with high-intensity states in the West, Pacific Northwest, and South. These patterns are consistent with the interpretation that the measure captures crop-specific burning systems rather than only national crop scale.

Table 15: Top states by crop-specific burning intensity

Crop	Rank	State	Counties	Harvested acres (million)	Fire detections	Intensity fires/1,000 acres	Any-fire share
Corn	1	MS	37	0.92	722	0.785	68.7%
Corn	2	AR	37	5.08	3,542	0.697	91.5%
Corn	3	GA	72	0.53	336	0.635	64.0%
Corn	4	LA	18	0.82	484	0.591	84.4%
Corn	5	WA	4	0.22	89	0.400	100.0%
Soybeans	1	AR	40	23.32	10,140	0.435	91.0%
Soybeans	2	LA	34	2.31	738	0.319	94.7%
Soybeans	3	SC	35	2.05	438	0.214	67.1%
Soybeans	4	MS	50	9.67	2,010	0.208	79.6%
Soybeans	5	GA	54	0.21	41	0.191	44.4%
Wheat	1	CA	21	0.70	460	0.654	90.7%
Wheat	2	ID	34	1.20	594	0.495	70.9%
Wheat	3	MS	33	0.62	209	0.336	43.6%
Wheat	4	WA	21	8.16	2,663	0.326	73.3%
Wheat	5	AR	34	1.09	316	0.289	61.0%
Cotton	1	GA	61	4.61	1,986	0.431	95.8%
Cotton	2	FL	7	0.17	59	0.349	60.0%
Cotton	3	MO	3	0.15	34	0.225	100.0%
Cotton	4	AL	36	1.95	376	0.192	72.7%
Cotton	5	SC	21	0.70	130	0.186	72.2%
Rice	1	MS	9	0.74	1,586	2.136	100.0%
Rice	2	CA	11	2.46	2,999	1.217	90.2%
Rice	3	AR	27	7.02	8,237	1.173	99.3%
Rice	4	MO	6	0.28	239	0.865	90.0%
Rice	5	LA	21	1.44	528	0.365	94.1%

Notes: The table reports the five highest-intensity states within each crop. Intensity is total crop-specific VIIRS fire detections divided by total harvested acreage and multiplied by 1,000. Any-fire share is the share of county-crop-years in the state with at least one crop-specific fire detection.

Appendix Table 16 reports the five highest-intensity counties within each crop after applying a denominator screen. County-crop units are included only if they have at least five observed years and total harvested acreage above the crop-specific 10th percentile of county-level total harvested acreage. This screen reduces the risk that the highest-intensity county rankings are driven by very small denominators.

Table 16: Top counties by crop-specific burning intensity

Crop	Rank	County FIPS	State	Years	Harvested acres (thousand)	Fire detections	Intensity fires/1,000 acres	Any-fire share
Corn	1	29133	MO	5	274.4	439	1.600	100.0%
Corn	2	48123	TX	5	24.1	25	1.038	100.0%
Corn	3	06067	CA	8	103.4	106	1.025	100.0%
Corn	4	05075	AR	5	38.4	39	1.016	100.0%
Corn	5	05079	AR	9	169.0	170	1.006	100.0%
Soybeans	1	05075	AR	6	364.7	374	1.026	100.0%
Soybeans	2	45049	SC	6	17.1	16	0.933	100.0%
Soybeans	3	05055	AR	7	482.9	393	0.814	100.0%
Soybeans	4	05121	AR	5	179.4	139	0.775	100.0%
Soybeans	5	05111	AR	8	1488.9	1,067	0.717	100.0%
Wheat	1	53013	WA	5	360.3	658	1.826	100.0%
Wheat	2	06011	CA	6	35.9	64	1.781	100.0%
Wheat	3	06101	CA	6	28.1	37	1.314	100.0%
Wheat	4	06077	CA	6	35.1	34	0.969	100.0%
Wheat	5	53023	WA	5	288.7	277	0.959	100.0%
Cotton	1	48383	TX	6	157.6	121	0.768	100.0%
Cotton	2	01005	AL	5	31.4	14	0.447	100.0%
Cotton	3	01003	AL	6	70.4	29	0.412	100.0%
Cotton	4	01097	AL	5	38.3	12	0.314	80.0%
Cotton	5	01067	AL	6	124.4	39	0.314	100.0%
Rice	1	28151	MS	5	56.3	175	3.110	100.0%
Rice	2	28143	MS	5	146.2	389	2.661	100.0%
Rice	3	05041	AR	6	130.2	309	2.373	100.0%
Rice	4	28011	MS	6	238.1	556	2.335	100.0%
Rice	5	05003	AR	6	50.9	107	2.101	100.0%

Notes: The table reports the five highest-intensity counties within each crop after requiring at least five observed county-crop-years and excluding counties below the crop-specific 10th percentile of total harvested acreage. Intensity is total crop-specific VIIRS fire detections divided by total harvested acreage and multiplied by 1,000. Any-fire share is the share of observed years in which the county-crop unit has at least one crop-specific fire detection. Counties are reported by five-digit county FIPS code.

The county-level results complement the state-level patterns. High-intensity rice counties are concentrated in Mississippi and Arkansas, with additional high-intensity rice activity in California. High-intensity soybean county-crop units are concentrated mainly in Arkansas. Cotton high-intensity counties appear in Texas and Alabama, while wheat high-intensity counties appear in Washington and California. These patterns show that high detected burning intensity is concentrated in specific crop-county systems rather than evenly distributed within crops. This geographic concentration supports the interpretation of the measure as a

crop-specific agricultural practice indicator, while also showing that the relevant geography differs across crops.

A.5 Denominator robustness

Appendix Table 17 tests whether the crop-level intensity ranking is sensitive to small harvested-acreage denominators. This check is important because the main measure divides crop-assigned fire detections by harvested acreage. Very small denominators could mechanically create extreme intensity values and distort crop rankings. The table recomputes aggregate crop-level detected burning intensity after excluding county-crop-years below the crop-specific 5th, 10th, and 25th percentiles of harvested acreage. The crop ranking remains unchanged across these restrictions: rice has the highest detected burning intensity, followed by cotton, wheat, soybeans, and corn. This robustness check supports the interpretation that the crop-level intensity ranking is not an artifact of very small acreage denominators. Instead, the ranking reflects systematic differences in crop-assigned detected burning relative to crop scale.

Table 17: Robustness of crop rankings to small acreage denominators

Crop	Full intensity	Full rank	Excl. 5% intensity	Excl. 5% rank	Excl. 10% intensity	Excl. 10% rank	Excl. 25% intensity	Excl. 25% rank
Corn	0.078	5	0.078	5	0.077	5	0.076	5
Soybeans	0.090	4	0.090	4	0.090	4	0.089	4
Wheat	0.122	3	0.122	3	0.122	3	0.122	3
Cotton	0.129	2	0.128	2	0.128	2	0.126	2
Rice	1.111	1	1.109	1	1.109	1	1.110	1

Notes: The table tests whether crop-level burning-intensity rankings are driven by small harvested-acreage denominators. The full sample includes all county-crop-years with positive harvested acreage. The remaining columns recompute aggregate crop-level intensity after excluding county-crop-years below the crop-specific 5th, 10th, or 25th percentile of harvested acreage.

A.6 Crop-specific peak burning months

Appendix Table 18 reports the crop-specific peak burning months used to construct the seasonal concentration measures. For each crop, the peak months are defined as the three calendar months with the largest national shares of crop-specific fire detections.

Table 18: Peak months of crop-specific fire detections

Crop	Fire detections	Top-three month share	Top three months
Rice	43,489	83.5%	Oct (41.8%); Sep (32.1%); Nov (9.6%)
Corn	82,638	46.5%	Apr (17.1%); Sep (16.2%); Mar (13.1%)
Soybeans	94,058	44.3%	Oct (17.2%); Apr (14.3%); Sep (12.8%)
Wheat	63,464	43.1%	Sep (15.7%); Aug (14.2%); Jul (13.3%)
Cotton	32,158	30.4%	Mar (11.4%); Oct (10.0%); Sep (9.1%)

Notes: The table reports the three calendar months with the largest shares of crop-specific CDL-filtered VIIRS fire detections from 2012 to 2024. Month shares are calculated within crop, so they sum to one across the twelve months for each crop.

These peak-month definitions are used to split annual crop-specific fire detections into peak-season and off-peak components. Rice has the most concentrated seasonal pattern, with most detected fires occurring in September, October, and November. Other crops have more diffuse seasonal distributions. This is why the main text treats the seasonal index as a diagnostic refinement rather than a universal replacement for annual burning intensity. The peak-month table documents the rule used in the seasonal decomposition and makes clear that the timing adjustment is crop-specific rather than imposed uniformly across all crops.