AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
http://ageconsearch.umn.edu
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Journal of Agricultural and Resource Economics 50(3):467–486 ISSN: 1068-5502 (Print); 2327-8285 (Online)
Copyright 2025 the authors
doi: 10.22004/ag.econ.347706

Misclassification Error in Remote Sensing Matters: The Effect of Ethanol Plants on Local Cropland Transitions

Nicholas J. Pates, Nathan P. Hendricks, and Tyler J. Lark

The quality of land cover data substantially affects estimates of treatment effects in land use change studies. We estimate the local impact of US ethanol plant expansions on cropland retention and conversion. Due to misclassification error, we show that using raw data from the Cropland Data Layer gives mixed results or results counter to economic expectations, while a cleaned version gives consistent and intuitive results. Our findings also highlight the importance of other methodological choices. Our preferred specification shows that plant expansions increase the probability of cropland conversion by 1.5 percentage points and cropland retention by 0.1 percentage points in plant neighborhoods.

Key words: cropland transition, misclassification error, remote sensing

Introduction

As economists increasingly use remote sensing data, there is a growing need to document how known problems with these data can impact econometric results. We demonstrate how spatiotemporally correlated misclassification errors can systematically bias difference-in-differences models in land use change (LUC) analyses. To show this, we consider the impact of ethanol plants on broad LUC and compare results using a raw and cleaned version of the Cropland Data Layer (CDL). Our results suggest that misclassed land covers in the raw CDL data can lead to conclusions that are statistically vague or at odds with economic theory.

Starting in 2008, the CDL has provided land cover data annually at a spatial resolution of 30 m^2 over the entire lower 48 US states. Produced by the US Department of Agriculture (USDA), CDL data facilitate rich, producer-level analyses (Hendricks, Smith, and Sumner, 2014; Hendricks et al., 2014; Donaldson and Storeygard, 2016; Pates and Hendricks, 2021). These data allow for local LUC effect estimates from policy using information on program participation or inception (Alix-Garcia, Shapiro, and Sims, 2012; Lark et al., 2022). However, working with these data presents challenges. The high resolution of spatial data increases its dimensionality. For all but very limited studies, this can lead to computational tractability and spatial correlation issues (Rashford, Albeke, and Lewis, 2013).

Although often ignored, remotely sensed data is often susceptible to misclassification error, a topic of interest in recent econometric LUC studies (Donaldson and Storeygard, 2016). Sandler and Rashford (2018) and Alix-Garcia and Millimet (2023) studied the impacts of classification errors in land use satellite data. They used econometric methods that directly model measurement

Nicholas J. Pates (corresponding author, nicholas.pates@uky.edu) is an assistant professor in the Department of Agribusiness and Applied Economics at North Dakota State University. Nathan P. Hendricks is a professor in the Department of Agricultural Economics at Kansas State University. Tyler J. Lark is a scientist at the Nelson Institute for Environmental Studies at the University of Wisconsin–Madison.

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

Review coordinated by Liang Lu.

errors in likelihood functions and relied on assumptions that data-generating processes are correctly specified for consistency (Hausman, Abrevaya, and Scott-Morton, 1998). Torchiana et al. (2023) developed a hidden Markov model that corrects for classification errors in land use transitions. In this related study, we show that misclassification can introduce substantial bias in econometric estimates. However, we document classification error that arises from differential errors across time and space. We show that these kinds of errors can bias difference-in-difference (DID) estimates in a systematic direction, not just through attenuation bias. Our approach to reduce this bias differs from prior studies and relies on remote sensing expertise to clean the data and reduce misclassification errors. While our approach is unlikely to eliminate all misclassifications, it allows us to address misclassification errors that are not easily modeled with a typical probability function.

In our application, we consider broad, field-level LUC following the 2007 Renewable Fuel Standard (RFS). Cropland conversion in the wake of the RFS has been a major research topic due to the negative environmental impacts associated with LUC (Donner and Kucharik, 2008; Searchinger et al., 2008; Roberts and Schlenker, 2013; Wright and Wimberly, 2013; Hendricks et al., 2014; Lark et al., 2022). The 2007 RFS substantially increased ethanol demand, crop demand, and crop prices, which contributed to a doubling of US ethanol production within a few years of its passage. The RFS could affect land use through general crop price increases (Carter, Rausser, and Smith, 2017). Alternatively—and the focus of our study—the construction of new ethanol plants could affect basis locally (McNew and Griffith, 2005).

The consensus in the LUC literature suggests a positive correlation between ethanol plant infrastructure and the local share of land in cropland (Miao, 2013; Motamed, McPhail, and Williams, 2016; Li, Miao, and Khanna, 2019; Wang et al., 2020). Towe and Tra (2012) found that land values near ethanol plants increased by 15%–30% following the construction of new facilities. This suggests that returns to farming rise after ethanol plants enter an area. Motamed, McPhail, and Williams (2016) found that a 1% increase in capacity within fixed procurement areas increased agricultural land shares by around 1.7%. Using a spatial competition model, Wang et al. (2020) found that the local ethanol competition index elasticity of corn land share was around 0.28. Li, Miao, and Khanna (2019) found that an ethanol capacity increase of 1 million gallons would increase a county's cropland by around 0.65%. At the county level, Miao (2013) found that ethanol plants increased Iowa county shares of corn acres by around 0.002%. Conversely, Ifft, Rajagopal, and Weldzuis (2018) found that ethanol plant capacity minimally affected county Conservation Reserve Program (CRP) enrollment or reenrollment.

We consider how method and study design choices might address misclassification bias in DID approaches. We compare results from two approaches: traditional two-way fixed effects (TWFE) and a method proposed by Callaway and Sant'Anna (2020) (CS). Recent literature shows that the TWFE method can introduce bias when treatment effects vary across time or groups (Goodman-Bacon, 2021). The CS method avoids this bias by restricting treated observation comparisons to a never-treated group. A popular variant of the CS approach also relaxes the parallel trend assumption with a doubly robust (DR) design. Simulated LUC studies suggest that the DR CS model might also address misclassification bias (Garcia and Heilmayr, 2024). We find that, while CS estimates were more economically consistent, the DR CS method did not entirely address misclassification bias.

DID models require groups of observations that receive a treatment and reference control groups. In settings with vague treatments, as in our application, researchers must define these groups. We use a variety of treated and control group definitions to consider how study design addresses misclassification bias. Like other studies, we assume that ethanol plants procure corn within a given distance from the plant. However, this distance likely depends on local rotational practices, productivity, and individual plant capacities. Our treated group definition allows fields to become treated as plants are constructed or expanded. We consider a range of rotation practices by adjusting these distances based on area shares of corn used for ethanol in regions surrounding plants. We also compare two alternative control group definitions. In one definition, we consider all fields outside of a procurement distance as controls. In a second definition, we restrict these control groups to fields within a given distance from the plant, creating doughnut-shaped control groups around plants. We find that using restricted control groups increases the credibility of the parallel trends assumption.

While prior research shows that ethanol plants likely induce LUC, there are distinct benefits to considering their effects at the field level rather than aggregating the data. Tying observations to the field level allows researchers to condition on variables at scales closer to decision makers. Tracking LUC at the field level also identifies specific transitions that compose the overall LUC. This can be useful for developing program priorities. An area's relative share of cropland can increase if cropland displaces noncropland (cropland conversion) or if cropland remains in cropland (cropland retention). If conversions are the primary LUC driver, then set-aside programs might prioritize keeping their existing participants enrolled to reduce conversion upon contract expiration. This could include lengthening contracts or encouraging reenrollment as contracts expire. Conversely, to address cropland retention, programs might prioritize attracting new participants. Tracking effects at the field level also allows researchers to better identify risk factors for cropland conversion and abandonment. This may help conservation programs reduce costs and increase program efficacy by targeting their efforts to plots at higher risk of conversion and abandonment.

We find that the group design and input data quality substantively affect our results. Our preferred model indicates that ethanol plants caused the probability of converting from noncropland to cropland to increase by about 2 percentage points and the probability of staying in cropland to increase by about 0.1 percentage points for fields in the neighborhood of an ethanol plant compared to those outside of the neighborhood. This result was contrary to those found using raw CDL data or when the model design did not restrict control groups to nearby areas.

Broad Land Use and Misclassification Error

Quantifying and correcting for measurement errors in the remotely sensed data is a major area of research in the geography literature (Johnson, 2013; Kline, Singh, and Dale, 2013; Lark, Salmon, and Gibbs, 2015; Lark et al., 2017). When remotely sensed data are used in econometric analysis, misclassification errors are a form of measurement error that could bias econometric estimates. Researchers commonly assess or address these errors by using multiple datasets to verify measurements and by using spatial and temporal correlation that likely exist in these data (Hansen et al., 2013; Lark et al., 2017; Alix-Garcia and Millimet, 2023). The latter approach requires researchers to consider the underlying context of the data. In this section, we justify the importance of addressing misclassification error in LUC studies and lay out our assumptions made in the cleaning process.

In this study, we focus on persistent changes in broad land use. These persistent changes have been controversial, a primary concern of land protection policies (e.g., the RFS and the crop insurance "Sodsaver" provision), and are especially significant in terms of carbon and biodiversity impacts. In our study, we define "persistent" land transitions as ones that endure for several years after a transition. Specifically, our cleaning procedure (described in a later section) defines persistent broad LUC transitions as those that occur only once during our study window of 9 years. This definition aligns with land enrolled in the Conservation Reserve Program (CRP), which requires participants to retire land for 10–15 years (US Department of Agriculture, 2022). While CRP represents a small portion of total land area, prior analyses show that changes in CRP represent a significant portion (typically around half) of the changes in cultivated cropland area in a given year (Hendricks, Er et al., 2018; Lark et al., 2022; US Department of Agriculture, 2020). The cost of transitioning from pasture to cropland varies significantly based on land conditions but can exceed \$100 per acre (Miao, Hennessy, and Feng, 2014). While transitions between cropland and noncropland can be more transitory, they are not the focus of our study. To account for these shorter-term transitions, we designate fields with more than one transition during our 9-year study window as "intermittently cultivated," and we exclude these fields from the analysis.

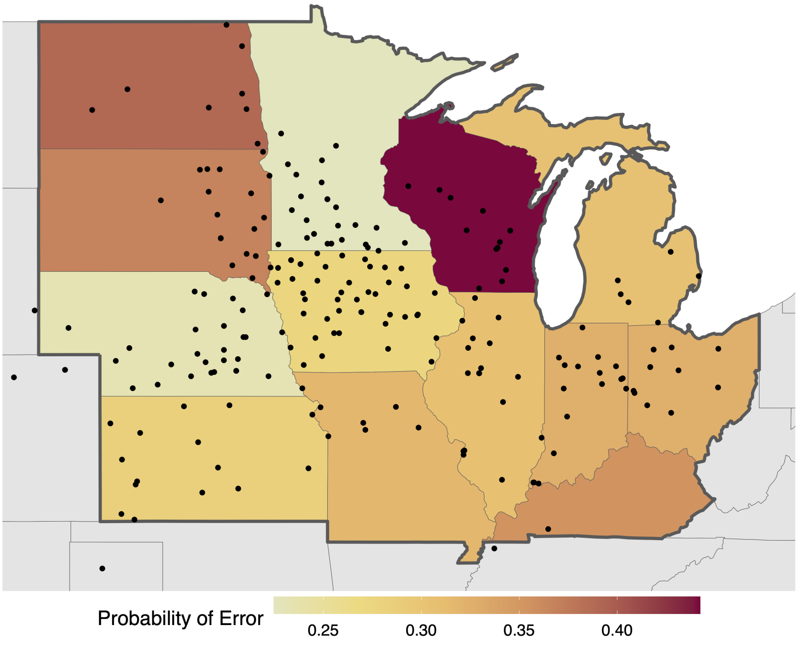

Figure 1. Probability of at Least One Crop Error (9-year sequences)

Notes: The map illustrates the average probability of having at least one misclassed broad land cover at the field level within a 9-year period for each state in our study region using state-specific metadata from the Cropland Data Layer. The thick boundary around the colored region defines our study area. The black dots denote the locations of ethanol plants within the Nebraska Department of Energy dataset.

Our focus on persistent transitions emphasizes the importance of data accuracy and consistency. Tracking persistent transitions requires accurate land use observations over many periods. Consider a sequence of land uses with two states—cropland (C) and noncropland (N)—for parcel i at time t, $\{y_{i,t}\}_{t=1}^{T} : y_{i,t} \in \{N, C\}$. If a persistent transition occurs at time $t^{\star}$, then $y_{i,t^{\star}} = y_{i,t \geq t^{\star}}$ and would have identical land states in $y_{i,t<t^{\star}}$. Based on our definition of persistent land use, producers can make only one land use transition within our sample, limiting the relevant observable sequences to a small subset of potential sequences.

The probability that a parcel's measured broad land state is correctly classified as state s is referred to as producer accuracy in the remote sensing literature and is written as $1 - \lambda_i = Pr[\tilde{y}_{it} = s \mid y_{it} = s]$, where $\tilde{y}_{it}$ denotes the measured broad land state. The 2020 CDL metadata have cropland (noncropland) producer accuracy of 98.3% (94.3%). While these percentages are respectable, this means that, on average and if errors were not spatially correlated, we might expect $1 - (\frac{0.983+0.943}{2})^9 = 29\%$ of the broad land sequences to have at least one misclassification over a 9-year period.

The CDL's classification errors are also not randomly distributed. Lark et al. (2017) showed that misclassification errors are temporally correlated because the CDL's broad land use accuracy improved over time. Further troubling, these errors also appear to be spatially correlated, meaning that they could be correlated with our variables of interest (Lark, Schelly, and Gibbs, 2021). Using confusion tables from the CDL's metadata, we construct the producer accuracy for broad cropland classes by state.[1] We then compute the likelihood of at least one cropland status in a 9-year series being misclassed (Figure 1). These values range between 20% to 45% and, importantly, states on the periphery of ethanol networks are more prone to classification errors.

[1] We include an illustrative explanation of confusion tables in the online supplement. See www.jareonline.org.

Data

Our study covers a 13-state region, (see Figure 1), containing a majority of the nation's ethanol plants and a large share of its corn acreage. Historically, most of the transitions to cropland in this region were from pasture and hay production (Lubowski et al., 2006).

Cropland Transitions

The CDL observations come from classification models using satellite imagery and ground-truthed data from the Farm Service Agency (FSA) for training and accuracy assessment. To illustrate the effect of the CDL's misclassification error, we compare our results using two datasets. We use a version of the CDL where Lark et al. (2020) cleaned inconsistent cropland trajectories and removed trajectories that were at greater risk of being misclassed (i.e., "cleaned CDL"). We then create a second "raw" dataset using the CDL. In the raw data, we group land covers into cropland and non-cropland classes and make minor adjustments to fallowed covers to be consistent with Lark et al.

As described in Lark et al. (2020), the cleaned CDL tracks crop, noncrop, and fallow land covers between 2008 and 2016 and has five types of trajectories: (i) continuous cropland, (ii) continuous noncropland, (iii) land transitioning once from noncropland to cropland, (iv) land transitioning once from cropland to noncropland, and (v) intermittent cropland.[2] Transitioning pixels note the year of the transition. Lark et al. also applied spatial filters and used the National Land Cover Database (NLCD) to remove misclassifications. Lark et al. corrected "noisy" trajectories that had a single transitory change in cropland status. They considered these transitions as likely classification errors and maintain the broad CDL observation of the predominant land use classification.[3] Lark et al. classified trajectories as intermittent cropland if they transitioned to or from cropland and remained consistent for at least 2 years and then transitioned back to the original land use.[4] In our analysis, we remove fields with intermittent trajectories to avoid classification accuracy issues and to focus on persistent LUC.

In the raw CDL, fields can transition to or from cropland more than once. We do not alter noisy trajectories or remove intermittent cropland sequences in the raw CDL data. Although these data are available more recently than 2016, we consider LUC between 2008 and 2016 with the raw CDL for consistency with the cleaned CDL.

We construct field-level data by sampling each dataset at field centroids. We define fields using the 2008 Common Land Unit (CLU) boundaries from the FSA (Woodard, 2016). Each CLU has a common producer, is contiguous, and is used for a common purpose. In areas where CLU boundaries are unavailable, we use remotely sensed field boundaries from Yan and Roy (2016). We measure field-level land use and characteristics using the same set of field centroids. Since smaller fields are more prone to misclassification error, we remove fields of less than 15 acres (Lark et al., 2020).

Land use transitions are our dependent variables. Let $y_{i,t} = 1$ ($y_{i,t} = 0$) denote that field i was cropland (noncropland) at time t. For the raw CDL data, we construct indicators for cropland retention and conversion equation (1). For the cleaned data, we condition on the land use in the initial period (see equation 2) since land, by construction, can only transition once over the sample period. Table 1 reports summary statistics for the transitions in the raw and cleaned CDL data. The probability of cropland abandonment is nearly 9 times larger and the probability of cropland conversion is about 2 times larger in the raw data. Like Lark et al. (2017), we find that the raw CDL classifies more fields as noncropland and fewer fields as cropland than the cleaned sample does.

[2] We list crop, noncrop, and fallow land cover categories in Table S4 in the online supplement.
[3] For example, the noisy trajectory $\{C,C,C,C,N,C,C,C,C\}$ would be corrected to $\{C,C,C,C,C,C,C,C,C\}$.
[4] For example, $\{N,N,C,C,C,C,N,N,N\}$ is an intermittent trajectory: The crop cover was consistent for more than 1 year.

Table 1. Transition Summaries for Raw and Cleaned CDL Data

		Current Status		
Data	**Transition**	**Crop**	**Noncrop**	***N***
Raw CDL lagged status	Crop	0.9577	0.0423	23,317,981
	Noncrop	0.0628	0.9372	21,173,231
Cleaned CDL trajectory type	Always crop or abandon	0.9952	0.0048	23,398,152
	Always noncrop or convert	0.0316	0.9684	20,665,008

Notes: The first two rows of the table shows the share of observed land covers transitioning from a prior broad land cover (crop, and noncrop) to another and the total number of observations of each broad land cover in the raw Cropland Data Layer data. The bottom two rows of the table show the share of observations transitioning to a broad land cover (crop, and noncrop) given an initial land cover and the count of observations by initial land cover in the cleaned Cropland Data Layer dataset.

$$(1) \qquad y^{CC}_{i,t,raw} = \begin{cases} 1 \mid y_{i,t} = 1 \;\&\; y_{i,t-1} = 1 \\ 0 \mid y_{i,t} = 0 \;\&\; y_{i,t-1} = 1 \end{cases} ; \; y^{NC}_{i,t,raw} = \begin{cases} 1 \mid y_{i,t} = 1 \;\&\; y_{i,t-1} = 0 \\ 0 \mid y_{i,t} = 0 \;\&\; y_{i,t-1} = 0 \end{cases}$$

$$(2) \qquad y^{CC}_{i,t,clean} = \begin{cases} 1 \mid y_{i,t} = 1 \;\&\; y_{i,1} = 1 \\ 0 \mid y_{i,t} = 0 \;\&\; y_{i,1} = 1 \end{cases} ; \; y^{NC}_{i,t,clean} = \begin{cases} 1 \mid y_{i,t} = 1 \;\&\; y_{i,1} = 0 \\ 0 \mid y_{i,t} = 0 \;\&\; y_{i,1} = 0 \end{cases}$$

Ethanol Plant Location and Capacity

Data on ethanol plant location, capacity, and construction dates come from the “Ethanol Production Capacity by Plant” dataset from the (Nebraska Department of Energy, 2019). We use these data to track nameplate capacities of over 166 ethanol plants between 2005 and 2018.[5] Observing ethanol plant capacity changes over a longer time period than cropland transitions provides more information to estimate effects before and after the plant became operational in TWFE models.

There were isolated instances of an incoming plant locating extremely close to an existing plant, usually owned by the same company. For these plants, we assume that producers would respond to pockets of new constructions as a change in capacity as opposed to a separate plant. We represent these plant locations using cluster centroids. We construct these clusters using a complete-linkage hierarchical algorithm with a cutoff range of 10 miles.[6] Most of these clusters consist of a single plant or a small group of plants managed by a single company.

We posit that ethanol plants influence LUC in nearby fields by increasing local basis. Persistent and positive basis effects from ethanol capacity changes would increase returns to crop production and may in turn lead to persistent shifts toward cropland. While the literature estimating the effect on local basis is relatively limited, McNew and Griffith (2005) found that ethanol plants strengthened local basis. Ethanol plants generally acquire grain locally via truck (Chang et al., 2019). Conceptually, the required area needed to fulfill a plant’s annual nameplate capacity should drive its geographic influence on local basis. The average plant cluster had an annual nameplate capacity of just over 90 million gallons, and annual capacities ranged from 50 million to 585 million gallons. To account for variation in plant size, we allow our treatment neighborhoods to vary by plant capacity.

We create plausible treated areas using plant nameplate capacities, corn throughput, county-level corn yields, and assumptions about the percentage of land surrounding plants that produce corn for

[5] We only consider ethanol plants that use conventional crops (i.e., corn or sorghum) as their primary input. We do not consider plants that use cellulosic biomass, straw, or waste products.

[6] This algorithm creates clusters using nearest neighbors and then expands clusters so that no plant within a cluster is over 10 miles from another plant within the same cluster.

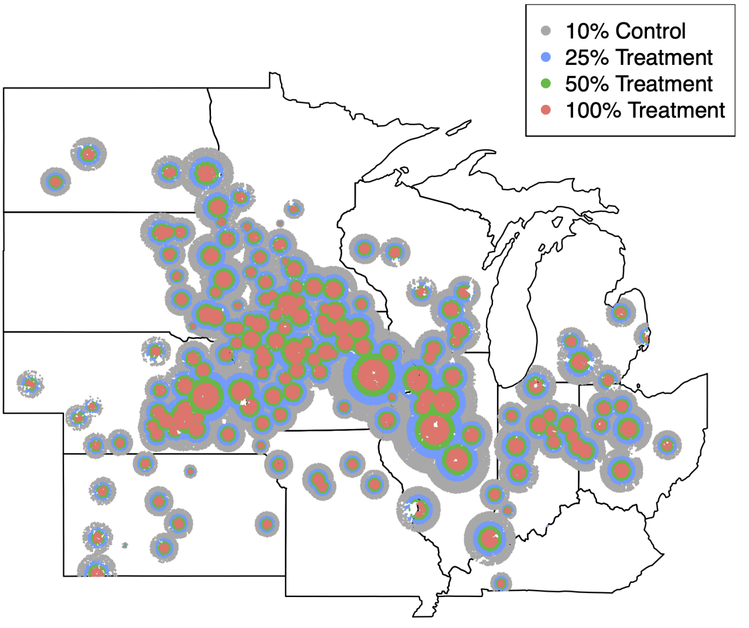

Figure 2. Ethanol Plant Treatment Neighborhoods in 2016

Notes: The map illustrates the treatment and control areas under various procurement assumptions using plant construction as of 2016. Ted regions indicate the treated areas around each plant, assuming that 100% of land is used to completely fill nameplate capacity under average yields; areas in green show the additional area needed if only 50% of land surrounding the plant is used to fulfill nameplate capacity, and areas in blue show the additional area of land if only 25% of land surrounding the plant is used to fulfill nameplate capacity. The gray area illustrates the 10% control group, area that is beyond blue the 25% procurement area but within the 10% procurement area.

ethanol.[7] We assume that 1 bushel of corn produces 2.8 gallons of ethanol (Jackson, 2018) and use the local corn yield to determine the area of corn required to fulfill the nameplate capacity of the plant.[8] To construct the radii of treatment neighborhoods, we use equation (3)—where cap_{it} is the cluster's total capacity in gallons per year, $yield_i$ is the average yield measured at the county level between 2009 to 2016, and *prop* is the proportion of area used to produce corn for ethanol—and consider different values of *prop*. We allow the neighborhoods to vary over time only through nameplate capacity changes.[9] The year of treatment for a neighborhood is the year in which the ethanol plant began operating.

$$r_{i,t}^{prop} = \left(\frac{cap_{i,t}}{2.8}\frac{yield_i}{640}\frac{1}{prop}\frac{1}{\pi}\right)^{\frac{1}{2}}. \tag{3}$$

Figure 2 shows the treatment neighborhoods in 2016. Observations in red, green, and blue are treated under the assumptions that 100%, 50%, and 25% of the total land area is planted to corn for ethanol each year.[10] Observations in gray are outside 25% areas but inside the 10% area that we use as a "doughnut" control group. We use observations in the white area as controls in the full control

[7] While we can determine the amount of corn planted in the area surrounding the plant, we would not know if it went to ethanol production. Further, tying observed corn plantings to treated areas may flaw the analysis since a parcel's cropland status is a necessary condition for corn production.

[8] For example, an ethanol plant with a capacity of 300 million gallons per year in a county with a yield of 160 bushels per acre requires $\frac{\text{300 million gal}}{\text{2.8 gallons/bushel}} = 107.14$ million bushels of corn. This would require $\frac{\text{107.14 million bushel}}{\text{160 bushels/acre}} = 669,625$ acres, or $\frac{\text{669,625 acres}}{\text{640 acres per sq. mile}} = 1,046.29$ square miles of corn.

[9] While yield increases over time reduce the area needed to fulfill nameplate capacity, using an average yield serves as a long-run expectation of input procurement areas.

[10] We consider land in cropland and noncropland since our goal is to estimate broad transitions. We also use static measures of average yields over our study window to isolate the impact of changes in local capacity. Higher annual corn planting probabilities mean that ethanol plants can source more corn from fields closer to the facility.

group model. On average, the sizes of the 100%, 50%, and 25% radii were 9.5 miles, 13 miles, and 19 miles, respectively. These average distances are in line with fixed distances considered in the literature (Towe and Tra, 2012; Wang et al., 2020).

A relatively small share of fields is within the neighborhood of multiple plants. For example, given a medium-sized treatment radius (50%), 13% of all fields were in the neighborhood of a single plant and only 1% were in the neighborhood of two or more plants. The DID approaches compare outcomes of observations that change their treatment status to those that do not. Considering larger treatment radii (smaller planting probabilities) means that more of the overall sample will be treated over our analysis. However, considering smaller radii increases the size of untreated areas between ethanol plants and more of the sample will become newly treated later in the study as new plants expand into these gaps. The sizes of the newly treated groups span a wide range from less than 1% to over 10% of all fields in a given year.[11]

Controls for Weighting and Regression Adjustment

Our preferred specification (CS) uses Callaway and Sant'Anna's (2020) doubly robust variant of the DID method, which requires a set of controls for both propensity score weighting and regression adjustment. As our focus is on persistent LUC decisions, our set of controls consists of either static or pretreatment field characteristics. We use all control variables for both weighting and regression adjustment.

We consider four categories of control variables: infrastructure, soils, climate, and prices. Infrastructural controls consist of distances to elevators, railroads, and population centers of at least 25,000 people. These are controls used in prior literature (Motamed, McPhail, and Williams, 2016) and address the influence from market access of both crop producers and ethanol plants as well as infrastructural factors for plant development. An area's soil and climate can also jointly influence plant siting and land use. To account for crop productivity, we merge Soil Survey Geographic Database (SSURGO) data and include field slope, soil texture, and the National Commodity Crop Productivity Index (NCCPI). We also include field-level hydric status, and water and wind erodibility indices since these are eligibility criteria for set-aside programs. To control for climate, we use data from the 800 m rasterized Parameter-elevation Regressions on Independent Slopes Model (PRISM) and include prior 20-year historical means and variances of growing-season growing degree days, extreme degree days, and cumulative precipitation. To control for pre-RFS storage and crop prices, we use 2004 expected corn price estimates, which we construct using elevator-level data from over 1,000 locations in the region using Data Transfer Network (DTN) and Cash Grain Bids (CGB) data. Table 2 shows the summary statistics for our control variables.[12]

Methods

We estimate the local effect of ethanol plants on LUC using a set of DID models. These models isolate the influence of ethanol plants from other contemporaneous shocks to LUC that affect all fields. We consider the dynamic effect of ethanol plants using construction and expansion times to construct event study estimates. Employing DIDs in our application is challenging. Plant construction can take several years, and producers may preemptively transition land before plants begin operating. Conversely, it may take several years for producers to make transitions after a plant enters. Land under set-aside contracts cannot freely transition until contracts expire. Producers may also take several years to determine whether the basis change from an ethanol plant justifies a transition. We will discuss how recent advances in the DID methods can address these issues.

---

[11] We include a breakdown of the number of observations and sample shares of each newly treated group in Table S5 in the online supplement.

[12] We provide a more detailed discussion of our data in the online supplement).

Table 2. Summary Statistics of Control Variables (*N* = 5,565,926 fields)

Control	Mean	Std. Dev.	Min.	Max.
Clay %	25.73	9.92	0.00	100.00
Sand %	27.98	22.57	0.00	98.52
Silt %	46.25	16.75	0.00	92.31
Hydric rating	0.18	0.34	0.00	1.00
Slope	5.80	8.00	0.00	90.00
Wind erodibility index	63.18	28.69	0.00	310.00
NCCPI	0.58	0.23	0.00	0.99
Distance to rail (miles)	5.89	6.34	0.00	63.88
Distance to population (miles)	43.85	30.67	0.01	192.57
Mean precipitation (inches)	532.50	103.66	235.67	950.19
Std. dev. precipitation (inches)	124.15	28.85	43.96	288.81
Mean 10°C growing degree days	1,639.85	311.21	628.95	2,418.37
Std. dev. 10°C growing degree days	135.51	9.90	89.47	173.56
Mean 30°C extreme degree days	25.78	21.27	0.02	137.65
Std. dev. 30°C extreme degree days	14.87	6.50	0.04	42.51
2004 corn basis ($)	2.55	0.12	2.30	2.89

Notes: The table shows the summary statistics for control variables at the field level. Soil-related control variables come from SSURGO data, distances to infrastructure are estimated using TIGER data, distance to population centers are estimated using US census data, data on climate are estimated using 20-year histories of PRISM data, 2004 corn basis is estimated using spatially interpolated elevator-level measures provided by DTN and CGB. All controls are measured at the field centroid.

DID approaches require treatment and control groups. Control groups should provide relevant counterfactuals to benchmark against treated outcomes. Defining both groups is challenging. Ethanol marketing areas (our treated groups) are vague and likely vary with the size of the plant and the productivity of the surrounding area. Following the literature, we consider several treated area assumptions. Our chosen treated areas span plausible marketing areas based on common crop rotations coupled with ethanol demand based on plant-specific nameplate capacities.

The entry of ethanol plants should not affect the returns to cropland in control groups, and control groups should also behave similarly to treated groups prior to treatment. To assess this, we consider two control group regimes. The first uses the entire untreated sample outside of the treated area as a control group (i.e., the “full control group”). However, as Figure 2 illustrates, this includes land far outside of ethanol procurement areas (e.g., the western Dakotas and eastern Kentucky). These areas have different economic, agronomic, and climatic profiles relative to ethanol-producing regions such that parallel trends may not hold. The second strategy uses fields that are outside of the 25% procurement areas but within the 10% procurement area (i.e., the gray area in Figure 2). We refer to this as the doughnut control groups since they form doughnut-shaped rings around ethanol plants. Fields in doughnut control groups are far enough from ethanol plants to be uninfluenced by expansions yet close enough that parallel trends with treated fields are likely to hold. To the extent that local basis was affected in our control group by the entry of an ethanol plant, we underestimate the impact of ethanol plants on land use transitions.

We assume that once a field has been exposed to an ethanol plant’s marketing area, it does not lose its treated status. In the DID literature, this is known as “staggered adoption.” Plant construction is expensive and time-consuming. While plant closures led to local ethanol plant capacity reductions in some areas, we assume that producers respond to local capacity changes as permanent intentions of increased local demand for corn. With these closures, our DID designs should provide conservative estimates of the average treatment effects on the treated (ATTs).

Two-Way Fixed Effects

The recent DID literature shows that TWFE models may be biased when treatment effects are heterogeneous across time or groups. However, TWFE models have been a common method in prior similar studies (Towe and Tra, 2012; Arora et al., 2016; Ifft, Rajagopal, and Weldzuis, 2018; Wang et al., 2020). Recent LUC simulations also suggest that the CS method is robust to misclassification error (Garcia and Heilmayr, 2024). We compare TWFE and CS model estimates to assess the degree to which this new method resolves bias from misclassification. For comparison, we construct a set of ATTs using a set of TWFE event-study regressions:

$$y_{i,t}^{trans} = \alpha_i + \gamma_t + \sum_{\ell} \mu_\ell \mathbf{1}(t - E_i = \ell) + \varepsilon_{i,t}. \tag{4}$$

Here, $y_{i,t}^{trans} \in \{y_{i,t}^{CC}, y_{i,t}^{NC}\}$ denotes transitions from cropland to cropland or noncropland to cropland as defined in equations (1) and (2) for the raw and cleaned CDL datasets, respectively. Here, α_i are field fixed effects, γ_t are year fixed effects, and E_i is the initial time that field i was treated. The term $\mathbf{1}(t - E_i = \ell)$ indicates that field i was treated within ℓ periods at time t. The μ_ℓ terms are the parameters of interest which measure the treatment effect ℓ years relative to the initial treatment time.[13]

Recent decompositions of TWFE estimators in settings with multiple treatment times lead to serious criticism of the approach (De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021). Goodman-Bacon (2021) found that TWFE estimates embed convoluted comparisons of treated, never-treated, soon-to-be-treated, and already-treated outcomes. He showed that if there is treatment effect heterogeneity across groups or time, TWFE estimates will be biased.[14] Heterogeneity bias in TWFE models arises from outcome comparisons of already-treated groups to newly treated groups made in the model.

Several recent papers derived alternative DID estimators to resolve this bias (De Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2020; Sun and Abraham, 2021; Athey and Imbens, 2022). Given assumptions on parallel trends, sampling, treatment effect homogeneity, and anticipation effects, these alternative methods allow for ATE heterogeneity across treated groups, time periods, or both. Many of the newer methods' assumptions are similar standard TWFE assumptions but allow for heterogeneous treatment effects by restricting the groups "compared" in the models.

The Callaway and Sant'Anna (2020) (CS) Method

To illustrate its potential bias, we compare the results of the TWFE and Callaway and Sant'Anna (2020) (CS) models. The CS method allows for both group and time heterogeneity and can be doubly robust (DR) by including propensity score weighting and regression adjustments. The DR property means that the estimator is consistent even if either the propensity score weights or the regression adjustment are misspecified (Sant'Anna and Zhao, 2020).

The CS method's DR property differentiates it from competing models (e.g., De Chaisemartin and d'Haultfoeuille, 2020; Sun and Abraham, 2021; Athey and Imbens, 2022). This is a desirable property because ethanol plant locations are likely endogenous. Ethanol plants are expensive to construct, and developers consider a variety of geographic and economic factors common to the crop industry (Shapouri and Gallagher, 2005). A productive local crop industry is an important criterion for plant location choices (McAloon et al., 2000). The CS method was also shown to be robust to misclassification errors in simulated land cover studies (Garcia and Heilmayr, 2024). The CS model

[13] The ℓ terms are called "event time" in the literature.

[14] Here "groups" are defined as cohorts that are treated in the same time period.

corrects for heterogeneous treatment effects by restricting the DID comparison groups to never-treated or last-treated cohorts. Since our data have a sizable never-treated cohort and construction slowed over time, we construct our ATT estimates using never-treated cohort comparisons.

Callaway and Sant'Anna (2020) provided a detailed description of variants of their model. We include specifics of our chosen CS variant in the online supplement. In brief, our model uses fields that were always outside of input procurement areas as our comparison group (the never-treated group) and used inverse propensity score weighting and regression adjustment to refine our group-time ATT estimates. The CS method also allows us to consider pretreatment effects. That is, it allows for anticipatory LUC effects prior to plant construction. Since we found that local LUC did not exhibit pretreatment effects, we use the year of construction to delineate treatment. To make our CS estimates comparable with the event-study TWFE model estimates, we aggregate group-time effects to event-time ATT effects using the procedures outlined by Callaway and Sant'Anna (2020).

Assumptions and Justifications of the CS Estimator

The CS method shares most of the identification assumptions of Abadie's (2005) semiparametric TWFE (SPTWFE) model. Both models relax the parallel trend assumption by propensity score weighting and require the weaker conditional parallel trends to hold. Since both models use propensity scores, they both require propensity score overlap for each of the treated groups and the control group across all time periods. The CS method also uses regression adjustment to correct for time-invariant control imbalances, making the conditional parallel trends more plausible. The models also share the standard technical assumptions that cross-sectional observations of outcomes, time-invariant control variables, and treatment status are drawn from an independent and identically distributed sample, allowing researchers to regard potential outcomes as random.

While the CS method is well-suited for addressing endogenous treatment, it requires assumptions about the timing and persistence of treatments. Like many TWFE alternatives, the CS method is restricted to "staggered" adoption designs. This assumes that, once treated, treated observations never lose their treatment status. Since ethanol plants closed over our study, this would result in more conservative effect estimates. Unlike many competing models, the CS model can consider anticipatory effects, when soon-to-be treated groups adjust their behavior in anticipation of treatment. While De Chaisemartin and d'Haultfoeuille (2020) provide a test for anticipation effects, the CS method allows researchers to model them directly by varying year prior to treatment. However, our results using the year prior to construction as the pretreatment year did not exhibit strong pretreatment effects. This suggests that producers do not respond to capacity changes prior to construction.

The CS method is more flexible than recent DID alternatives. Under Sun and Abraham's (2021) and Athey and Imbens's (2022) methods, effects can differ across groups but must be linear over time. Linear effects are unlikely in our case due to concurrent land set-aside program changes over ethanol's construction boom. Changes to the CRP after 2007 led to a greater emphasis on continuous CRP contracts and temporary reenrollment extension (REX) contracts (Ifft, Rajagopal, and Weldzuis, 2018). In select counties, this made reenrolling easier for producers and made more farmers eligible for continuous CRP contracts. While this likely affected broad LUC, these program changes were county-specific; reenrollment efforts were especially strong in ethanol-producing counties.[15] De Chaisemartin and d'Haultfoeuille's (2020) method allows for either ATT heterogeneity across groups or across time, but the CS method allows for both. Further, the DR property of the CS ATTs adds robustness to plant location endogeneity.

[15] While county-level CRP controls could be used, plant marketing areas can envelope entire county boundaries. A circular area with a 13-mile radius would envelope the typical Iowa county.

Treatment and Control Groups By Model

As the TWFE and the CS methods differ, it is important to understand the groups that each method uses to produce their respective event time estimates. With our data, the TWFE method can provide estimates for event times from 9 years pretreatment to 11 years posttreatment. Because the CS method uses the last outcome before treatment to construct long differences, it can not estimate treatment effects for observations that were never treated or always treated. Since we estimate the effect on land transitions, every estimable treated group must have land use observations at least 2 years prior to treatment. Always-treated observations lack pretreatment counterfactual untreated values and are dropped in CS models. This restricts our treated groups to those initially treated in 2010, 2011, 2013, 2014, and 2015.[16] Unlike the TWFE method, the CS method does not use a left-out reference period and compares the first difference prior to treatment to the difference post treatment. Therefore, the CS method loses the initial first difference in the process and produces event time estimates from 5 periods prior to treatment to 6 years posttreatment.

Impacts of Misclassification Error

Before examining our treatment results, we consider how misclassification error and its subsequent improvement can impact our estimates. DID methods compare outcome changes in treated and control groups before and after treatment. Differences in misclassification error and error improvement over time and across space can bias these models.

In classification problems, errors of omission and errors of commission are assessed using producer and user accuracy, respectively. Let z indicate a transition to some land cover given an observed prior land cover.[17] Let z^{obs} indicate the observed transition and z^{true} indicate the underlying true transition. Producer accuracy (*PA*), where $PA = Pr[z^{obs} \mid z^{true}]$, measures the probability that the classifier indicates the transition given the transition truly occurred. User accuracy (*UA*), where $UA = Pr[z^{true} \mid z^{obs}]$, measures the probability that the transition truly occurred given the classifier indicates the transition. Using Bayes Theorem, one can relate UAs and PAs as conditional relationships. With this, we can construct a useful term, the accuracy ratio (*AR*), as $AR = \frac{PA}{UA} = \frac{Pr[z^{obs}]}{Pr[z^{true}]}$.[18] Intuitively, AR is a measure of gross over- or undercounting of true transitions. When AR is below (above) 1, the classifier undercounts (overcounts) true observations.

We can then construct the standard DID estimator as the change in the probability of observing a given transition before and after treatment between the treated and control groups in equation (5). Here, *trt* (*con*) denotes treated (control) observations and *pre* (*post*) denotes observations before (after) treatment. This equation shows how changes in UA and PA between the treated and control groups can systematically bias the model.

$$(5) \qquad DID = \left(Pr\left[z^{true}_{trt,post}\right]AR_{trt,post} - Pr\left[z^{true}_{trt,pre}\right]AR_{trt,pre}\right) - \left(Pr\left[z^{true}_{con,post}\right]AR_{con,post} - Pr\left[z^{true}_{con,pre}\right]AR_{con,pre}\right).$$

We consider the direction of misclassification error bias when transition probabilities are constant across time and groups (i.e., $Pr[z^{true}_{..}]$). In this case, the true treatment effect is 0, but the estimated effect with observed data is $Pr[z^{true}_{..}] \times \widetilde{DID}_{AR}$, where $\widetilde{DID}_{AR}$ is defined in equation (6).[19] If $\widetilde{DID}_{AR}$ is above 0, then the raw CDL tended to overstate transitions after treatment in the treated group relative to the control group, biasing the treatment effect upwards. While the cleaned

[16] Table S6 in the online supplement shows the treated groups and land use transitions that correspond with each event time estimate.

[17] For example, z can be a cropland retention ($y_{i,t} = 1 \mid y^{obs}_{i,t-1} = 1$) transition. We condition the probabilities on observed lagged values throughout since they determine the model each observation enters.

[18] See the online supplement.

[19] The tilde indicates that the ratio was derived using the cleaned CDL as a reference dataset.

transitions are not necessarily true, they are validated using other observations in land use trajectories and the context of broad LUC. We therefore consider the impact of misclassification error by calculating equation (6) using the cleaned data as the "true" reference data:

$$\widetilde{DID}_{AR} = \left(\widetilde{AR}_{trt,post} - \widetilde{AR}_{trt,pre}\right) - \left(\widetilde{AR}_{con,post} - \widetilde{AR}_{con,pre}\right). \tag{6}$$

We consider the magnitude and direction of misclassification bias using a full control group and 50% treated areas for each treated cohort.[20] The raw CDL consistently understated retention in all treated cohorts after treatment. This indicates that misclassification error will bias our cropland retention effect estimates downward. Generally, the raw CDL also understated cropland conversion in the treated cohorts. However, it was more prone to overstate cropland conversion in the 2013 treated cohort. With this, we expect misclassification error will primarily bias our conversion effects downward but that cohort-specific biases should offset one another around the middle of our study window.[21]

Results

We now present the results of our analysis. We compared our local LUC estimates across four methodological dimensions: (i) data quality (raw and cleaned CDL), (ii) technique (TWFE and CS), (iii) control group (full and 10% doughnut controls), and (iv) treated group (100%, 50%, and 25% procurement assumptions).

Two-Way-Fixed Effects

Figure 3 shows the event-time ATTs from the TWFE models. Within each row, all plots share a common scale and all TWFE estimates are relative to five periods prior to treatment.

We first discuss our cropland conversion estimates (top four plots). The plots in the first row of Figure 3, which show our estimated conversion effects using raw CDL data, exhibit a saw-toothed pattern. As our CS results will show, this is a recurring pattern when we used raw CDL data but is not as pronounced when we use cleaned data, suggesting that this pattern is a symptom of misclassification error in the CDL. Posttreatment, all significant effects from the raw CDL were negative or insignificant. This suggests that, counter to expectations, ethanol plants reduce cropland conversion. The second row of Figure 3 suggests that misclassification error is likely the dominant source of this downward bias. While the saw-toothed pattern in the first row is present in the second, it is less pronounced when we used cleaned CDL data.

Over our study period, the raw CDL initially underpredicted aggregate cropland relative to survey-based estimates but became better aligned with these estimates over time. To reach parity with these other sources, the raw CDL likely overstated cropland conversion. Raw CDL data suggested that US cropland expanded by around 20 million acres between 2008 and 2010. The National Resources Inventory (NRI) suggested an expansion of less than a 5 million acres over this same period. Reaching near parity with NRI estimates, the raw CDL suggested that cropland expanded by 10 million acres between 2010 and 2012. In our models, error corrections towards cropland are incorrectly considered cropland conversions. Between 2008 and 2012, aggregate cropland in the cleaned CDL trended similarly to the NRI.[22] Our accuracy analysis in the prior section suggests that errors and subsequent corrections in the raw CDL were more prevalent in the control group. This likely explains why the raw conversion effects oscillate between positive and negative. We also observe significant and negative pretreatment effects. However, the pattern suggests that these are related to our choice of our left-out time period and overall

[20] These are groups treated in 2010, 2011, 2013, 2014, and 2015.

[21] See Table S7 in the online supplement for the $\widetilde{DID}_{AR}$ results.

[22] Lark et al. (2017) provided an insightful comparison of aggregate cropland trends across various of datasets.

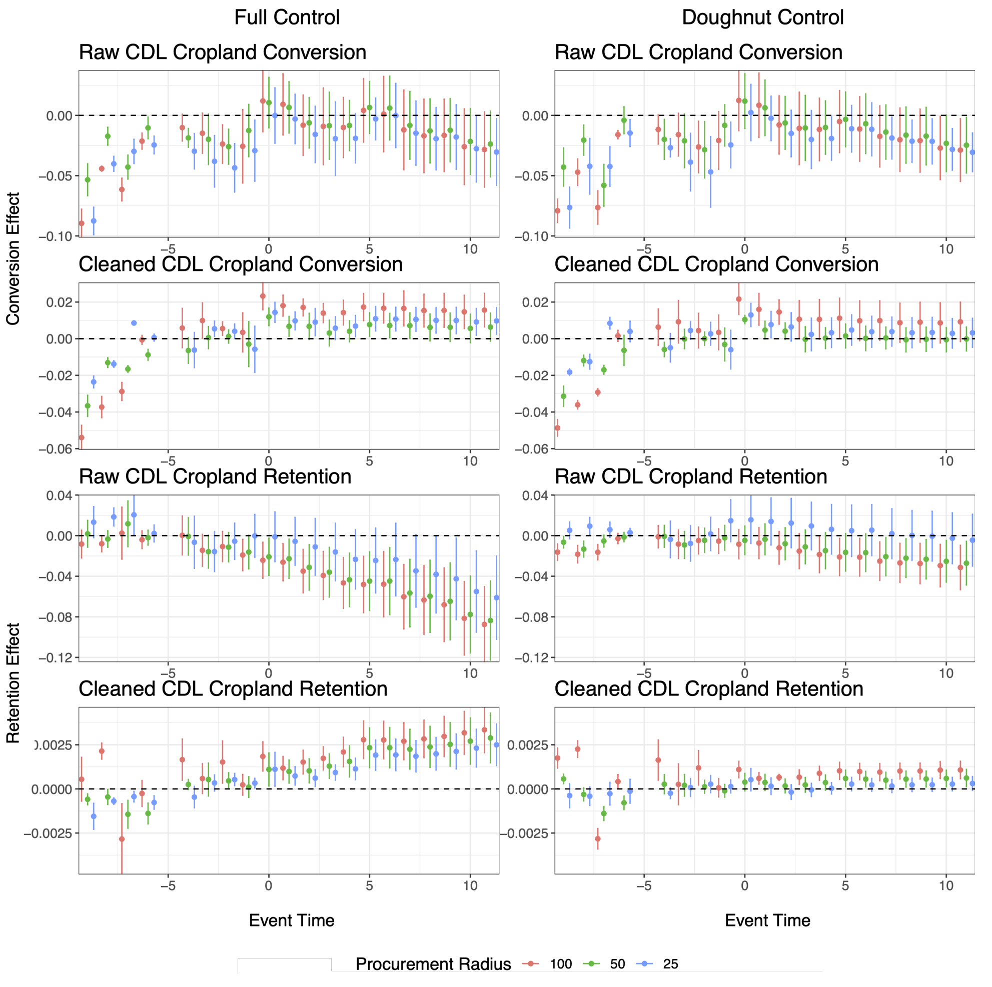

Figure 3. Two-Way-Fixed Effect (TWFE) Estimates

Notes: Point-and-whisker plots show the mean ATT of ethanol plants at each time period since treatment and their 95% confidence interval using the whiskers. Plots in red, green, and blue illustrate the treatment effect of ethanol plants in the 100%, 50%, and 25% procurement areas, respectively. Plots in the left column show the effects using the entire nontreated sample as a control group. Plots in the right column use just the 10% doughnut control group. Plots in the first two rows show cropland conversion effects. Plots in the bottom two rows show cropland retention effects. Plots in odd-numbered rows show estimates using the raw CDL data. Plots in even-numbered rows show estimates using the cleaned CDL data.

misclassification error. The pattern of negative pretreatment effects follows the same saw-toothed pattern posttreatment.

As expected, posttreatment conversion effects were positive and significant when we used the cleaned CDL (second row of Figure 3). This suggests that, shortly after construction, ethanol plants caused a short-lived, 2-percentage-point increase in cropland conversion. These positive conversion effects waned over time and were not as pronounced on fields further from the plant. Using doughnut control groups as a reference (second column), these conversion effects become insignificant after 2 years postconstruction. Although there are concerns with using pretreatment effects from TWFE to verify the parallel trends assumption (see Sun and Abraham, 2021), we find negative and significant pretreatment effects across all TWFE conversion models, which may indicate parallel trend violations. Small pretreatment samples may explain this pattern. Pretreatment effects 7 or more years prior to treatment come from a small cohort of fields that were treated by a single plant in 2018.[23] Notably, these pretreatment effects were less significant in our cleaned data models and when using the doughnut control group, particularly 6 years prior to treatment. Because the 2018 treated group lacks posttreated values, we drop this group from the sample under the CS method.[24]

The last two rows of Figure 3 show our TWFE retention effects. Our results using the raw CDL (third row) suggest that, counter to expectations, ethanol plants decreased cropland retention. However, misclassification error may contribute to a downward bias in posttreatment effects. Our $\widetilde{DID}_{AR}$ results showed that errors and error corrections were more prevalent in the control group.[25] This relatively overstated cropland conversion and understated cropland retention in the control group. Subsequent accuracy improvements would then bias the retention effect estimates downward. While posttreatment retention effects were negative throughout, the full-control results suggest plants decrease retention by 8 percentage points, while our doughnut-control results show that plants reduced retention by only 2 percentage points. Since observations in the doughnut control group are closer to treated observations, this may indicate spatial correlation in CDL accuracy improvement.

The fourth row of Figure 3 shows the TWFE retention estimates using the cleaned CDL data. Both plots suggest that, postconstruction, ethanol plants increased cropland retention, with larger effects in the full-control model. Retention effects were much smaller than the conversion effects in the cleaned data, suggesting that ethanol plants caused an increase in cropland retention of around 0.1 percentage points. Like our conversion estimates, we observe statistically significant pretreatment retention effects more than 6 years prior to treatment. This again could be a result of the small size of the cohort that was initially treated in 2018. Like our cleaned conversion effects, cleaned retention effects were larger for treated groups closer to the plant. Using the doughnut control group, 25% treated area effects were insignificant. This suggests that with the doughnut control group, parallel trends are more likely and effect estimates are more conservative.

Our retention results between the third and fourth rows of Figure 3 differ considerably. These differences demonstrate the importance of cleaning data when effect sizes are small. Cropland retirement was rare in the cleaned version of the CDL. Only 0.5% of initial cropland ended up retiring by the end of our study. Conversely, the raw CDL data suggested that over 4% of cropland in a given year would retire the following year. This suggests that the raw data poorly identify changes in persistent land use in the sample. While accuracy improvements in are typically desirable, such improvements can—in cases with potentially small potential effects— provide strong, seemingly statistically significant, and counterintuitive results. This would explain our strong, negative posttreatment retention effects in the raw CDL model and small, positive, but statistically significant effects in the cleaned CDL model.

Callaway and Sant'Anna Results

Figure 4 shows the results of the CS DID models. The CS method can account for heterogeneous treatment effects and relaxes the parallel trends assumption. Comparing these plots with our TWFE plots in Figure 3 illustrates the degree to which these improvements remediate misclassification bias.

We now consider whether the CS method alone added robustness to misclassification error. The first and third rows of Figure 4 show our results with the CS method and raw CDL data, respectively. The plots in the first row show our conversion effect estimates. Like the TWFE estimates, the conversion effects under the CS method exhibit a similar saw-toothed pattern with the raw CDL data and confusingly indicate that posttreatment effects were positive, negative, and insignificant. As expected, the pretreatment effects were considerably smaller using the CS method and doughnut control groups. A similar saw-toothed pattern is present in the raw retention effects in the third row of plots. Like our TWFE retention estimates, our raw CS estimates also suggest that,

[23] This treated group consists of around 1,500 observations for the 100% radius and 25% radius treated areas and around 2,500 observations in the 50% radius treated area.

[24] Table S6 provides a description of treated cohorts used to construct event time effect estimates.

[25] See Table S7 in the online supplement.

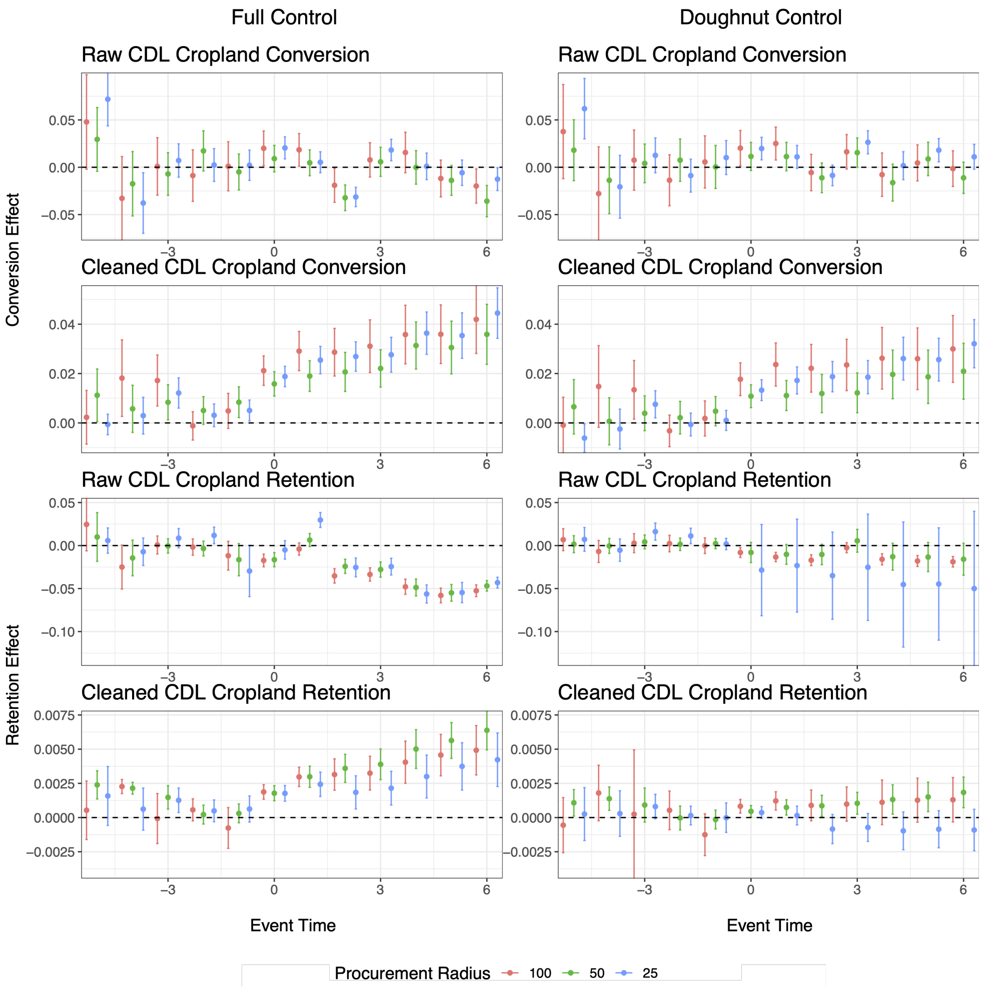

Figure 4. Callaway-Sant'Anna (CS) DID Estimates

Notes: Point-and-whisker plots show the mean ATT of ethanol plants at each time period since treatment and their 95% confidence interval using the whiskers. Plots in red, green, and blue illustrate the treatment effect of ethanol plants in the 100%, 50%, and 25% procurement areas, respectively. Plots in the left column show the effects using the entire nontreated sample as a control group. Plots in the right column use just the 10% doughnut control groups. Plots in the first two rows show cropland conversion effects. Plots in the bottom two rows show cropland retention effects. Plots in odd-numbered rows show estimates using the raw CDL data. Plots in even-numbered rows show estimates using the cleaned CDL data.

counter to expectations, ethanol plants reduced cropland retention. The smaller and less significant pretreatment effects across the models suggest that the conditional parallel trend assumption of the CS method is more plausible and that the CS method provides more valid DID estimates. However, the prevalence of ambiguous or economically inconsistent posttreatment effects suggests that using the CS method alone did not correct for the differential accuracy changes in the raw CDL data.

The plots in the second and fourth rows of Figure 4 show the results of preferred specification, using the CS method with the cleaned CDL. Like our TWFE estimates, we did find saw-toothed effect patterns when we modeled with the cleaned CDL data. Consistent with expectations, these models show that, postconstruction, ethanol plants increased local cropland conversion and retention. Corn–soybean rotations are popular in our study area, making the 50% and 25% treated groups the most plausible procurement areas for ethanol plants. Our results using these treated groups suggest that ethanol plants increased cropland conversion between 1 and 3 percentage points and increased cropland retention by 0.1 percentage points. The TWFE estimates are generally about half as large as CS estimates. This downward bias is expected because the treatment effect increases over time and TWFE erroneously uses the early treated as a comparison for the late treated. Further, pretreatment effects were less significant when we used the doughnut control group. While there were sporadic pretreatment significance issues in the doughnut models, these transitory effects were small in magnitude and significance. Pretreatment effects were smaller and less significant under the CS method, implying that this method improved the validity of the DID design. Both pre- and posttreatment effects were also smaller and less significant when we used the doughnut control group. This suggests that applied researchers should take steps to correct measurement error in the data, use a modeling strategy that allows for heterogeneity in treatment effects, and consider control groups near the treated samples of interest.

Conclusions

While remotely sensed data allow researchers to observe land use at broad scopes, establishing causal estimates of broad LUC remains challenging. Broad and persistent LUC is rare, which increases the influence of even minute misclassification errors on model results. Recent advances in the DID literature provide useful tools that can properly estimate causal effects and ensure that identification conditions are met. However, researchers must still be mindful of their added assumptions and ensure that they carefully define treatment and reference groups. The effects under alternative specifications, particularly those using raw CDL data, were noisier and suggested counterintuitive and sometimes contradictory effect patterns.

Our preferred estimates using the CS DID approach and the cleaned CDL dataset show that ethanol plants increased cropland conversion and weakly increased cropland retention within plausible input procurement areas. While nominally small, our estimates of an approximately 1.5-percentage-point increase in conversion and a 0.1-percentage-point increase in retention are sizable relative to sample averages. Our cleaned sample showed that only 3.16% of lands starting in noncropland in 2008 converted to cropland by 2016 and only 0.48% of cropland was abandoned by 2016. This suggests that lands within ethanol markets had around a 47% greater risk of converting to cropland $\left(\frac{0.015+0.0316}{0.0316}\right)$ and around a 21% greater risk $\left(\frac{0.001+0.0048}{0.0048}\right)$ of retaining cropland than the overall sample average.

Our results support the general findings of the literature that ethanol production leads to more intensive agricultural land use (Miao, 2013; Motamed, McPhail, and Williams, 2016; Li, Miao, and Khanna, 2019; Wang et al., 2020). However, directly comparing the magnitude of our results with previous literature is challenging due to differences in outcome variables, units of analysis, definitions of treatment, and statistics of interest. Importantly, our results using the raw CDL also follow field-level studies LUC studies that also found both positive and negative corn planting effects of ethanol capacity at the parcel level (Arora et al., 2016). We find that when classification errors are not addressed, cropland conversion and retention effects may, against economic intuition, falsely appear to be negative and significant, offering a suggestion as to how misclassification errors influence DID models.

Modeling at the field level presents challenges but adds rigor to LUC analyses. One limitation of our study is that we remove fields classified as intermittent cropland from the analysis because we are less certain about the accuracy of this classification. Differentiating conversion from retention is useful for conservation programs to develop more impactful policy objectives and is only possible in field-level studies. Modeling at the field level also strengthens the microfoundations of LUC analyses and allows researchers to control for confounders at the level at which decisions are made. Combining these data, the best practices of geography, and the latest advances in DID methodology provides a template for researchers to explore causal relationships in LUC.

[First submitted January 2024; accepted for publication January 2025.]

References

Abadie, A. 2005. “Semiparametric Difference-in-Differences Estimators.” *Review of Economic Studies* 72(1):1–19. doi: 10.1111/0034-6527.00321.

Alix-Garcia, J., and D.L. Millimet. 2023. “Remotely Incorrect? Accounting for Nonclassical Measurement Error in Satellite Data on Deforestation.” *Journal of the Association of Environmental and Resource Economists* 10(5):1335–1367. doi: 10.1086/723723.

Alix-Garcia, J.M., E.N. Shapiro, and K.R.E. Sims. 2012. “Forest Conservation and Slippage: Evidence from Mexico’s National Payments for Ecosystem Services Program.” *Land Economics* 88(4):613–638. doi: 10.1353/lde.2012.0052.

Arora, G., P.T. Wolter, H. Feng, and D.A. Hennessy. 2016. *Role of Ethanol Plants in Dakotas Land Use Change: Incorporating Flexible Trends in the Difference-in-Difference Framework with Remotely-Sensed Data*. Working Paper 584. Center for Agriculture and Rural Development.

Athey, S., and G.W. Imbens. 2022. “Design-Based Analysis in Difference-in-Differences Settings with Staggered Adoption.” *Journal of Econometrics* 226(1):62–79. doi: 10.1016/j.jeconom.2020.10.012.

Callaway, B., and P.H. Sant’Anna. 2020. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics* 225(2):200–230. doi: 10.1016/j.jeconom.2020.12.001.

Carter, C.A., G.C. Rausser, and A. Smith. 2017. “Commodity Storage and the Market Effects of Biofuel Policies.” *American Journal of Agricultural Economics* 99(4):1027–1055. doi: 10.1093/ajae/aaw010.

Chang, K.L., P. Caffarelli, J. Gastelle, and A. Sparger. 2019. *Transportation of US Grains: A Modal Share Analysis, 1978-2016 Update*. USDA Agricultural Marketing Service. doi: 10.22004/ ag.econ.313485.

De Chaisemartin, C., and X. d’Haultfoeuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review* 110(9):2964–96. doi: 10.1257/aer.20181169.

Donaldson, D., and A. Storeygard. 2016. “The View from Above: Applications of Satellite Data in Economics.” *Journal of Economic Perspectives* 30(4):171–98. doi: 10.1257/jep.30.4.171.

Donner, S.D., and C.J. Kucharik. 2008. “Corn-Based Ethanol Production Compromises Goal of Reducing Nitrogen Export by the Mississippi River.” *Proceedings of the National Academy of Sciences* 105(11):4513–4518. doi: 10.1073/pnas.0708300105.

Garcia, A., and R. Heilmayr. 2024. “Impact Evaluation with Nonrepeatable Outcomes: The Case of Forest Conservation.” *Journal of Environmental Economics and Management* 125:102971. doi: 10.1016/j.jeem.2024.102971.

Goodman-Bacon, A. 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics* 225(2):254–277. doi: 10.1016/j.jeconom.2021.03.014.

Hansen, M.C., P.V. Potapov, R. Moore, M. Hancher, S.A. Turubanova, A. Tyukavina, D. Thau, S.V. Stehman, S.J. Goetz, T.R. Loveland et al. 2013. “High-Resolution Global Maps of 21st-Century Forest Cover Change.” *Science* 342(6160):850–853. doi: 10.1126/science.1244693.

Hausman, J.A., J. Abrevaya, and F.M. Scott-Morton. 1998. “Misclassification of the Dependent Variable in a Discrete-Response Setting.” *Journal of Econometrics* 87(2):239–269. doi: 10.1016/S0304-4076(98)00015-3.

Hendricks, N.P., E. Er et al. 2018. “Changes in Cropland Area in the United States and the Role of CRP.” *Food Policy* 75(C):15–23. doi: 10.1016/j.foodpol.2018.02.001.

Hendricks, N.P., S. Sinnathamby, K. Douglas-Mankin, A. Smith, D.A. Sumner, and D.H. Earnhart. 2014. “The Environmental Effects of Crop Price Increases: Nitrogen Losses in the US Corn Belt.” *Journal of Environmental Economics and Management* 68(3):507–526. doi: 10.1016/j.jeem.2014.09.002.

Hendricks, N.P., A. Smith, and D.A. Sumner. 2014. "Crop Supply Dynamics and the Illusion of Partial Adjustment." *American Journal of Agricultural Economics* 96(5):1469–1491. doi: 10.1093/ajae/aau024.

Ifft, J., D. Rajagopal, and R. Weldzuis. 2018. "Ethanol Plant Location and Land Use: A Case Study of CRP and the Ethanol Mandate." *Applied Economic Perspectives and Policy* 41(1):37–55. doi: 10.1093/aepp/ppy007.

Jackson, S. 2018. *SP700-B-Ethanol: A Primer*. Bulletin 5-2008. University of Tennessee Ag Research Extension. Available online at https://trace.tennessee.edu/utk_agexbiof/4.

Johnson, D.M. 2013. "A 2010 Map Estimate of Annually Tilled Cropland Within the Conterminous United States." *Agricultural Systems* 114:95–105. doi: 10.1016/j.agsy.2012.08.004.

Kline, K.L., N. Singh, and V.H. Dale. 2013. "Cultivated Hay and Fallow/Idle Cropland Confound Analysis of Grassland Conversion in the Western Corn Belt." *Proceedings of the National Academy of Sciences* 110(31):E2863–E2863. doi: 10.1073/pnas.1306646110.

Lark, T.J., N.P. Hendricks, A. Smith, N. Pates, S.A. Spawn-Lee, M. Bougie, E.G. Booth, C.J. Kucharik, and H.K. Gibbs. 2022. "Environmental Outcomes of the US Renewable Fuel Standard." *Proceedings of the National Academy of Sciences* 119(9):e2101084119. doi: 10.1073/pnas.2101084119.

Lark, T.J., R.M. Mueller, D.M. Johnson, and H.K. Gibbs. 2017. "Measuring Land-Use and Land-Cover Change Using the US Department of Agriculture's Cropland Data Layer: Cautions and Recommendations." *International Journal of Applied Earth Observation and Geoinformation* 62:224–235. doi: 10.1016/j.jag.2017.06.007.

Lark, T.J., J.M. Salmon, and H.K. Gibbs. 2015. "Cropland Expansion Outpaces Agricultural and Biofuel Policies in the United States." *Environmental Research Letters* 10(4):044003. doi: 10.1088/1748-9326/10/4/044003.

Lark, T.J., I.H. Schelly, and H.K. Gibbs. 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. doi: 10.3390/rs13050968.

Lark, T.J., S.A. Spawn, M. Bougie, and H.K. Gibbs. 2020. "Cropland Expansion in the United States Produces Marginal Yields at High Costs to Wildlife." *Nature Communications* 11(1): 4295. doi: 10.1038/s41467-020-18045-z.

Li, Y., R. Miao, and M. Khanna. 2019. "Effects of Ethanol Plant Proximity and Crop Prices on Land-Use Change in the United States." *American Journal of Agricultural Economics* 101(2): 467–491. doi: 10.1093/ajae/aay080.

Lubowski, R.N., S. Bucholtz, R. Claassen, M.J. Roberts, J.C. Cooper, A. Gueorguieva, and R. Johansson. 2006. *Environmental Effects of Agricultural Land-Use Change*. Economic Research Report ERR-25. USDA Economic Research Service.

McAloon, A., F. Taylor, W. Yee, K. Ibsen, and R. Wooley. 2000. *Determining the Cost of Producing Ethanol from Corn Starch and Lignocellulosic Feedstocks*. National Renewable Energy Lab. Available online at https://www.nrel.gov/docs/fy01osti/28893.pdf.

McNew, K., and D. Griffith. 2005. "Measuring the Impact of Ethanol Plants on Local Grain Prices." *Review of Agricultural Economics* 27(2):164–180.

Miao, R. 2013. "Impact of Ethanol Plants on Local Land Use Change." *Agricultural and Resource Economics Review* 42(2):291–309. doi: 10.1017/S106828050000438X.

Miao, R., D.A. Hennessy, and H. Feng. 2014. "Sodbusting, Crop Insurance, and Sunk Conversion Costs." *Land Economics* 90(4):601–622.

Motamed, M., L. McPhail, and R. Williams. 2016. "Corn Area Response to Local Ethanol Markets in the United States: A Grid Cell Level Analysis." *American Journal of Agricultural Economics* 98(3):726–743. doi: 10.1093/ajae/aav095.

Nebraska Department of Energy. 2019. *State and Plant Capacity*. NDOE. Available online at https://catenergy.ne.gov/programs/stats/info/122.htm [Accessed December 2019].

Pates, N.J., and N.P. Hendricks. 2021. “Fields from Afar: Evidence of Heterogeneity in United States Corn Rotational Response from Remote Sensing Data.” *American Journal of Agricultural Economics* 103(5):1759–1782. doi: 10.1111/ajae.12208.

Rashford, B.S., S.E. Albeke, and D.J. Lewis. 2013. “Modeling Grassland Conversion: Challenges of Using Satellite Imagery Data.” *American Journal of Agricultural Economics* 95(2):404–411. doi: 10.1093/ajae/aas110.

Roberts, M.J., and W. Schlenker. 2013. “Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate.” *American Economic Review* 103(6):2265–2295. doi: 10.1257/aer.103.6.2265.

Sandler, A.M., and B.S. Rashford. 2018. “Misclassification Error in Satellite Imagery Data: Implications for Empirical Land-Use Models.” *Land Use Policy* 75:530–537. doi: 10.1016/j.landusepol.2018.04.008.

Sant’Anna, P.H., and J. Zhao. 2020. “Doubly Robust Difference-in-Differences Estimators.” *Journal of Econometrics* 219(1):101–122. doi: 10.1016/j.jeconom.2020.06.003.

Searchinger, T., R. Heimlich, R.A. Houghton, F. Dong, A. Elobeid, J. Fabiosa, S. Tokgoz, D. Hayes, and T.H. Yu. 2008. “Use of US Croplands for Biofuels Increases Greenhouse Gases Through Emissions from Land-Use Change.” *Science* 319(5867):1238–1240. doi: 10.1126/science. 1151861.

Shapouri, H., and P. Gallagher. 2005. *USDA’s 2002 Ethanol Cost-of-Production Survey*. Agricultural Economic Report AER-841. US Department of Agriculture Office of the Chief Economist.

Sun, L., and S. Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics* 225(2):175–199. doi: 10.1016/j.jeconom.2020.09.006.

Torchiana, A.L., T. Rosenbaum, P.T. Scott, and E. Souza-Rodrigues. 2023. “Improving Estimates of Transitions from Satellite Data: A Hidden Markov Model Approach.” *Review of Economics and Statistics* :1–45doi: 10.1162/rest_a_01301.

Towe, C., and C.I. Tra. 2012. “Vegetable Spirits and Energy Policy.” *American Journal of Agricultural Economics* 95(1):1–16.

US Department of Agriculture. 2020. *Summary Report: 2017 National Resources Inventory*. Natural Resources Conservation Service and Iowa State University Center for Survey Statistics and Methodology. Available online at https://www.nrcs.usda.gov/sites/default/files/2022-10/2017NRISummary_Final.pdf.

———. 2022. “Conservation Reserve Program.” Fact Sheet. USDA Farm Service Agency. Available online at https://www.fsa.usda.gov/tools/informational/fact-sheets/conservation-reserve-program-crp.

Wang, Y., M.S. Delgado, J. Sesmero, and B.M. Gramig. 2020. “Market Structure and the Local Effects of Ethanol Expansion on Land Allocation: A Spatially Explicit Analysis.” *American Journal of Agricultural Economics* 102(5):1598–1622. doi: 10.1111/ajae.12119.

Woodard, J.D. 2016. “Data Science and Management for Large Scale Empirical Applications in Agricultural and Applied Economics Research.” *Applied Economic Perspectives and Policy* 38(3):373–388. doi: 10.1093/aepp/ppw009.

Wright, C.K., and M.C. Wimberly. 2013. “Recent Land Use Change in the Western Corn Belt Threatens Grasslands and Wetlands.” *Proceedings of the National Academy of Sciences* 110(10): 4134–4139. doi: 10.1073/pnas.1215404110.

Yan, L., and D. Roy. 2016. “Conterminous United States Crop Field Size Quantification from Multi-Temporal Landsat Data.” *Remote Sensing of Environment* 172:67–86. doi: 10.1016/j.rse.2015.10.034.

Journal of Agricultural and Resource Economics 25(3):S1–S7
ISSN: 1068-5502 (Print); 2327-8285 (Online)
doi: 10.22004/ag.econ.347706
Copyright 2025 the authors

Online Supplement: Misclassification Error in Remote Sensing Matters: The Effect of Ethanol Plants on Local Cropland Transitions

Nicholas J. Pates, Nathan P. Hendricks, and Tyler J. Lark

An Illustrative Example of Confusion Tables, Accuracies, and Table Aggregations

Each CDL dataset provides metadata to assess its accuracy. NASS trains and validates the CDL using ground-truthed measures and land use reports. Confusion tables are a common way to assess "agreement" between two sets of data. For each land cover, they show the congruency between CDL land use estimates and validated datasets. Consider the three-class example below where land use consists of corn, soybeans, and wheat across 2000 pixels with ground-truthed data.

Table S1. Illustrative Example of a Three Class Confusion Table

	CDL Estimate			
Ground Truthed Observations	**Corn**	**Soybean**	**Wheat**	**Total Observed**
Corn	992	6	2	1,000
Soybean	30	400	20	450
Wheat	20	30	500	550
Total Estimated	1,042	436	522	**2,000**

In this example, the CDL classed 1042 pixels as "corn", 436 as "soybean" and 522 as "wheat". In reality, there were 1000 pixels of corn, 450 of soybeans, and 550 of wheat. The first row shows the CDL confused 6 corn pixels for soybeans and 2 corn pixels for wheat. The first column shows the CDL incorrectly stated 30 soybean pixels and 20 wheat pixels were corn. User accuracies indicate the reliability of CDL estimates to users and are measured as the share of total classed pixels that were "correct" for a given crop. Producer accuracies track how well the CDL measures true land covers and are the shares of ground-truthed pixels that were correctly estimated by the CDL for each land cover. In this example, the CDL over-predicted corn ($1042 > 1000$) and under-predicted soybeans ($436 < 450$) and wheat ($522 < 550$). Later in the paper, we construct our own measures of over- and under -counting using the cleaned data which we call "accuracy ratios" $\left(AR = \frac{\text{Producer Accuracy}}{\text{User Accuracy}}\right)$. Table S2 shows the accuracy values and ratios from the example above.

Table S2. Accuracy Statistics from Illustrative Example

Crop	**User Accuracy**	**Producer Accuracy**	**Accuracy Ratio**
Corn	$0.952\ \left(\frac{992}{1042}\right)$	$0.992\ \left(\frac{992}{1000}\right)$	$1.042\ \left(\frac{1042}{1000}\right)$
Soybean	$0.917\ \left(\frac{400}{436}\right)$	$0.889\ \left(\frac{400}{450}\right)$	$0.969\ \left(\frac{436}{450}\right)$
Wheat	$0.958\ \left(\frac{500}{522}\right)$	$0.909\ \left(\frac{500}{550}\right)$	$0.949\ \left(\frac{522}{550}\right)$

Aggregating Confusion Tables

To assess broad land use accuracy, we aggregated categories of the confusion table. Suppose we were interested in classifying between either corn or soybean and wheat pixels in the above example. Following the table, we would sum the values accordingly:

Table S3. Illustrative Example of Confusion Table Aggregation

	CDL Estimate		
Ground Truthed Observations	**Corn/Soybean**	**Wheat**	**Total Observed**
Corn/Soybean	1,428	22	1,450
Wheat	50	500	550
Total Estimated	1,478	522	**2,000**

In our application, we performed this aggregation over land covers designated by the red and blue crops in table S4. After aggregating to our broad land use categories, we constructed our accuracy ratios using the same procedures shown in the three-crop setting discussed in the prior section.

Table S5 shows the percent of the sample that is newly treated in each year under different radius assumptions. This illustrates how our treated group choices impact the size and composition of the treatment cohorts which are important in DID event studies. The value for 2005 shows the fields that were in the neighborhood of an ethanol plant that began operations in 2005 or earlier. The subsequent years show the number of fields that were newly treated in that year (i.e., in the neighborhood of a newly operational plant). Newly treated observation percentages fluctuate according to the incoming ethanol capacity.

Discussion of Control Variables

Infrastructure

Railroad infrastructure is an important consideration for ethanol plants since plants primarily transport ethanol to blenders via rail car. Rail networks may also influence basis prices through elevators, which are also commonly on or near rail networks. Since the general footprint of rail lines was well-established prior to the rise of the ethanol industry, we use it as a control for propensity score weighting and regression adjustment (Motamed, McPhail, and Williams, 2016). We measure the distance to rail using the Euclidean distance between each field and rail line maps.

Ethanol production requires water and wastewater infrastructure for heating and cooling. Plants require between 2 and 6 gallons of water to produce each gallon of ethanol. Plants do not require access to potable sources but can make use of wastewater from municipal sources. Locating closer to municipal wastewater facilities can reduce pumping costs for the plant while offering wastewater facilities an alternative to discharging it on their own (Keeney and Muller, 2006). The spatial correlation between these facilities and population centers may also influence producer decisions. Population centers offer both off-farm employment opportunities and market access. We proxy for wastewater infrastructure and market access using Euclidean distances to population centers with at least 25,000 people in the 2000 US Census.

Soil

Soil quality is important for both ethanol plant locations and producers' land use decisions. To control for soil features, we merged the Soil Survey Geographic Database (SSURGO) to our field locations. Soil characteristics serve as both an indication of cropland productivity and as a criteria for enrollment in land set-aside programs. We include slope, soil texture, and National Commodity

Table S4. Cropland Data Layer Cover Designations

Crop (ID)	Crop (ID)	Crop (ID)	Crop (ID)
Corn (1)	Sugarcane (45)	Prunes (210)	Celery (245)
Cotton (2)	Sweet Potatoes (46)	Olives (211)	Radishes (246)
Rice (3)	Misc Vegs and Fruits (47)	Oranges (212)	Turnips (247)
Sorghum (4)	Watermelons (48)	Honeydew Melons (213)	Eggplants (248)
Soybeans (5)	Onions (49)	Broccoli (214)	Gourds (249)
Sunflower (6)	Cucumbers (50)	Peppers (216)	Cranberries (250)
Peanuts (10)	Chick Peas (51)	Pomegranates (217)	Dbl Crop Barley/Soybeans (254)
Tobacco (11)	Lentils (52)	Nectarines (218)	Other Hay/Non Alfalfa (37)
Sweet Corn (12)	Peas (53)	Greens (219)	Forest (63)
Pop or Orn Corn (13)	Tomatoes (54)	Plums (220)	Shrubland (64)
Mint (14)	Caneberries (55)	Strawberries (221)	Barren (65)
Barley (21)	Hops (56)	Squash (222)	Clouds/No Data (81)
Durum Wheat (22)	Herbs (57)	Apricots (223)	Developed (82)
Spring Wheat (23)	Clover/Wildflowers (58)	Vetch (224)	Water (83)
Winter Wheat (24)	Sod/Grass Seed (59)	Dbl Crop WinWht/Corn (225)	Wetlands (87)
Other Small Grains (25)	Switchgrass (60)	Dbl Crop Oats/Corn (226)	Nonag/Undefined (88)
Dbl Crop WinWht/Soybeans (26)	Cherries (66)	Lettuce (227)	Aquaculture (92)
Rye (27)	Peaches (67)	Pumpkins (229)	Open Water (111)
Oats (28)	Apples (68)	Dbl Crop Lettuce/Durum Wht (230)	Perennial Ice/Snow (112)
Millet (29)	Grapes (69)	Dbl Crop Lettuce/Cantaloupe (231)	Developed/Open Space (121)
Speltz (30)	Christmas Trees (70)	Dbl Crop Lettuce/Cotton (232)	Developed/Low Intensity (122)
Canola (31)	Other Tree Crops (71)	Dbl Crop Lettuce/Barley (233)	Developed/Med Intensity (123)
Flaxseed (32)	Citrus (72)	Dbl Crop Durum Wht/Sorghum (234)	Developed/High Intensity (124)
Safflower (33)	Pecans (74)	Dbl Crop Barley/Sorghum (235)	Barren (131)
Rape Seed (34)	Almonds (75)	Dbl Crop WinWht/Sorghum (236)	Deciduous Forest (141)
Mustard (35)	Walnuts (76)	Dbl Crop Barley/Corn (237)	Evergreen Forest (142)
Alfalfa (36)	Pears (77)	Dbl Crop WinWht/Cotton (238)	Mixed Forest (143)
Camelina (38)	Pistachios (204)	Dbl Crop Soybeans/Cotton (239)	Shrubland (152)
Buckwheat (39)	Triticale (205)	Dbl Crop Soybeans/Oats (240)	Grass/Pasture (176)
Sugarbeets (41)	Carrots (206)	Dbl Crop Corn/Soybeans (241)	Woody Wetlands (190)
Dry Beans (42)	Asparagus (207)	Blueberries (242)	Herbaceous Wetlands (195)
Potatoes (43)	Garlic (208)	Cabbage (243)	Fallowed/Idled (61)
Other Crops (44)	Cantaloupes (209)	Cauliflower (244)	

Notes: Blue cells indicate cropland, red noncropland, and yellow replaced/withheld.

Table S5. Newly Treated Sample Proportions by Year

	100% Radius		50% Radius		25% Radius		10% Radius	
Year	**Obs.**	**Share**	**Obs.**	**Share**	**Obs.**	**Share**	**Obs.**	**Share**
Untreated	5,196,407	92.33%	4,805,594	85.39%	4,218,233	74.95%	3,197,559	56.82%
2005	213,878	3.80%	427,252	7.59%	808,485	14.37%	1,587,104	28.20%
2006	102,459	1.82%	191,826	3.41%	305,184	5.42%	409,525	7.28%
2007	57,971	1.03%	106,423	1.89%	162,596	2.89%	252,273	4.48%
2008	32,853	0.58%	54,490	0.97%	76,060	1.35%	113,247	2.01%
2009	4,808	0.09%	9,996	0.18%	19,448	0.35%	33,188	0.59%
2010	10,859	0.19%	17,361	0.31%	21,132	0.38%	15,097	0.27%
2011	729	0.01%	1,532	0.03%	2,954	0.05%	6,913	0.12%
2013	2,600	0.05%	4,938	0.09%	7,822	0.14%	13,006	0.23%
2014	727	0.01%	1,034	0.02%	4	0.00%	–	–
2015	3,171	0.06%	5,012	0.09%	4,414	0.08%	–	–
2018	1,450	0.03%	2,454	0.04%	1,580	0.03%	–	–

Notes: The table shows the number of observations and shares of the total sample that were newly treated by year and input procurement radius assumption (share of land surrounding plants that is used to fulfill nameplate capacity). Dashed lines indicate zero newly-treated observations for the given year.

Crop Productivity Index (NCCPI) to control for cropland productivity and suitability. Other soil features directly affect eligibility in land set-aside programs. To account for this, we also include the field's hydric status, water erosion index, and wind erosion index. Hydric soils are "sufficiently wet in the upper part to develop anaerobic conditions during the growing season". These soils are likely to retain excessive moisture in the topsoil and hamper plant development. We include indicators of water and wind erosion since erosion potential is a criterion for acceptance into CRP.

Climate

Climate can affect land use by restricting the suitable set of crops and is a major determinant of crop productivity. To control for climate, we use daily measures of the 800-meter Parameter-elevation Regressions on Independent Slopes Model (PRISM) evaluated at field centroids between 1984 and 2004. We then construct mean and variance estimates of growing-season extreme degree days, growing degree days, and cumulative precipitation in inches using historic 20 year histories.[1] Extreme degree days are a measure of exposure to extreme temperatures at or above 30°C and growing degree days are a measure of exposure to suitable heat units during a season between temperatures of 10°C and 30°C. We then condition means and variances across seasons derived from annual seasonal values.

Price

Local corn markets can influence both crop producers and ethanol plants. Basis patterns are a function of distance to major markets and local storage supply and demand. While ethanol plants likely increase local basis, historic basis can proxy for storage supply. Ample storage in an area can help ensure plants will receive a consistent corn supply throughout the year. To control for local pre-ethanol prices, we use the 2004 field-level expected corn prices. We construct the expected price as the sum of the nearby planting-harvest corn futures contract spread plus local corn prices averaged over January and February in 2004. We collected daily spot corn prices from over 1,000 locations in the region from Data Transfer Network (DTN) and Cash Grain Bids (CGB). After computing the average corn price in January and February, we interpolated them to the centroid of each field using

[1] We assume a growing season between April 1st and September 30th.

ordinary kriging. As this was prior to the RFS polices of 2005 and 2007, this should approximate local storage capacity prior to the uptick in ethanol construction.

Discussion of the Callaway and Sant'Anna Method

Equation S1 shows the DR version of the CS ATT estimate in group-time form. The field subscript i is suppressed here for brevity. G_g is an indicator variable equal to one if the individual was first treated in period g.[2] This ATT term gives distinct estimates of the ATT by both the period of initial treatment (g) and a given period of the analysis (t). For this reason, Callaway and Sant'Anna refer to this as a "group-time" ATT estimate.

$$\text{(S1)} \qquad ATT_{dr}^{nev}(g,t;\delta) = \mathbb{E}\left[\left(\frac{G_g}{\mathbb{E}\left[G_g\right]} - \frac{\frac{p_g(X)C}{1-p_g(X)}}{\mathbb{E}\left[\frac{p_g(X)C}{1-p_g(X)}\right]}\right)\left(y_t^{trans} - y_{g-\delta-1}^{trans} - m_{g,t,\delta}^{nev}(X)\right)\right]$$

This design only compares treated observations with never-treated observations. These are represented by the indicator variable C which is equal to one if the individual is never treated throughout the duration of the study. The first term in the equation illustrates the weighing portion of the DR method. All of the regression adjusted first-differenced outcome variables from treated observations are included in the ATT estimation and are normalized by the expected value of being initially treated in period g $\left(\mathbb{E}\left[G_g\right]\right)$ where $g \in \mathcal{G}$. Never-treated observations are weighted based on the odds ratio using a propensity score of being initially treated $\left(p_g(X)\right)$ in period g. For instance, if, based on the controls (X), a never-treated observation were twice as likely to be initially treated in period g than not, $\left(\frac{p_g(X)}{1-p_g(x)}\right) = 2$, it would receive roughly twice the weight of an observation that was initially treated in period g. As in observations first treated in g, this weight is normalized by the expected propensity odds ratio.

The second term shows the regression adjustment component of the DR estimate. Here, the observed outcome in period t is compared with the outcome in period $g - \delta - 1$, the treatment reference period (g) minus the number of periods of anticipation (δ) minus one. The regression adjustment arises through $m_{g,t,\delta}^{nev}$. This is the predicted value from an OLS regression using the never-treated group as the training sample. The dependent variable in this regression is the "long difference" of the outcome variable $\left(y_t^{trans} - y_{g-\delta-1}^{trans}\right)$ and the regressors are same the set of controls (X) used to compute the propensity score. The purpose of the $m_{g,t,\delta}^{nev}$ term is to provide a conditional expectation of the change in untreated potential outcomes. Unlike the propensity weighting component, this term applies to all observations, not just the control group.[3] For treated observations, the regression adjustment term will partial out the predicted change in outcomes that would be expected if the observation had been in the control group. For the control observations, the final terms in braces will be its error term from the regression used to construct $m_{g,t,\delta}^{nev}$ (Callaway and Sant'Anna, 2020). The CS method estimates group-time ATTs by comparing the "long difference" between observed pre-treatment outcome $y_{g-\delta-1}$ and the observed outcome from calendar time t. This means the CS method cannot produce estimates for observations with missing pre-treatment values (e.g., always-treated observations).

The technique produces summary ATT estimates using aggregations of group-time ATT estimates. Our focus in this paper is on event-time aggregations, shown in equation S2. Here, (e) denotes the time elapsed since treated and is negative for pre-treated observations and positive for post-treated observations. The first term in the summation requires that the calendar time considered

[2] Here g indexes the initial treatment time.
[3] Note that it is not multiplied by the "C" variable.

Table S6. Calendar Year Transitions Event-Time Correspondence by Group

	Event Time																				
Group (g)	**−9**	**−8**	**−7**	**−6**	**−5**	**−4**	**−3**	**−2**	**−1**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
2005	—	—	—	—	—	—	—	—	—	—	—	—	—	08-09	09-10	10-11	11-12	12-13	13-14	14-15	15-16
2006	—	—	—	—	—	—	—	—	—	—	—	—	08-09	09-10	10-11	11-12	12-13	13-14	14-15	15-16	—
2007	—	—	—	—	—	—	—	—	—	—	—	08-09	09-10	10-11	11-12	12-13	13-14	14-15	15-16	—	—
2008	—	—	—	—	—	—	—	—	—	—	08-09	09-10	10-11	11-12	12-13	13-14	14-15	15-16	—	—	—
2009	—	—	—	—	—	—	—	—	—	08-09	09-10	10-11	11-12	12-13	13-14	14-15	15-16	—	—	—	—
2010	—	—	—	—	—	—	—	—	**08-09**	**09-10**	**10-11**	**11-12**	**12-13**	**13-14**	**14-15**	**15-16**	—	—	—	—	—
2011	—	—	—	—	—	—	—	**08-09**	**09-10**	**10-11**	**11-12**	**12-13**	**13-14**	**14-15**	**15-16**	—	—	—	—	—	—
2013	—	—	—	—	—	**08-09**	**09-10**	**10-11**	**11-12**	**12-13**	**13-14**	**14-15**	**15-16**	—	—	—	—	—	—	—	—
2014	—	—	—	—	**08-09**	**09-10**	**10-11**	**11-12**	**12-13**	**13-14**	**14-15**	**15-16**	—	—	—	—	—	—	—	—	—
2015	—	—	—	**08-09**	**09-10**	**10-11**	**11-12**	**12-13**	**13-14**	**14-15**	**15-16**	—	—	—	—	—	—	—	—	—	—
2018	08-09	09-10	10-11	11-12	12-13	13-14	14-15	15-16	—	—	—	—	—	—	—	—	—	—	—	—	—

Notes: The ‘Event Time’ row shows the possible event-time estimates. The Group (g) column shows the treated groups by year of initial treatment. Transitions in red show the calendar time transitions used to estimate the group-time treatment effects under the Callaway-Sant’Anna method. The two-way-fixed effects model uses both the red and black transitions to construct treatment effects.

$g + e$ is not outside of the sample. The second term weights the group-time ATT by the size of the cohort initially treated at time g. The last term considers only the group-time ATT from group g at calendar time $g + e$.

$$\theta_e = \sum_{g \in G} 1\,(g + e \leq T)\, P\,(G = g \mid G + e \leq T)\, ATT\,(g, g + e) \tag{S2}$$

Using Bayes Theorem to Derive Accuracy Ratios

In the paper, we describe the producer accuracy (PA) and user accuracy (UA) measures where:

$$PA = Pr\left[z^{obs} \mid z^{true}\right]; \quad UA = Pr\left[z^{true} \mid z^{obs}\right] \tag{S3}$$

Using Bayes Theorem one can arrange these conditional relationships as follows.

$$Pr\left[z^{true} \mid z^{obs}\right] = \frac{Pr\left[z^{obs} \mid z^{true}\right] Pr\left[z^{true}\right]}{Pr\left[z^{obs}\right]} \Rightarrow Pr\left[z^{obs}\right] = Pr\left[z^{true}\right] AR \tag{S4}$$

$$\Rightarrow AR = \frac{Pr\left[z^{obs}\right]}{Pr\left[z^{true}\right]} \tag{S5}$$

Table S7. Difference-in-Differences Accuracy Ratio Measures by Treated Cohort

	$\widetilde{DID}_{AR}$	
Cohort↓/Model→	**Retention**	**Conversion**
2010	-0.0212	-0.117
2011	-0.0271	-0.0279
2013	-0.00688	0.110
2014	-0.0288	-0.588
2015	-0.0290	-0.316

Notes: The table shows the yearly $\widetilde{DID}_{AR}$ estimates using only variation in the accuracy ratios between the raw and cleaned versions of the Cropland Data Layer and a 50% input procurement radius, and the full-control group.

References

Callaway, B., and P.H. Sant'Anna. 2020. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225(2):200–230. doi: 10.1016/j.jeconom.2020.12.001.

Keeney, D., and M. Muller. 2006. *Water use by ethanol plants: Potential challenges*. Minneapolis, MN: Institute for Agriculture and Trade Policy.

Motamed, M., L. McPhail, and R. Williams. 2016. "Corn Area Response to Local Ethanol Markets in the United States: A Grid Cell Level Analysis." *American Journal of Agricultural Economics* 98(3):726–743. doi: 10.1093/ajae/aav095.