Weather Risk, Information, and Agricultural Productivity: Evidence from Historical Data in the US

Vaibhav Anand and Honglin Li*

July, 2025

(Working Draft- Please do not cite or circulate)

(Click here for the latest version)

Abstract: In this paper, we examine the effect of access to information on local economic productivity in the context of agriculture and weather risk. The agricultural sector is particularly vulnerable to the impacts of weather and climate risk, and weather forecasts have been an important source of information for farmers' input decisions and adaptation to adverse weather shocks. We exploit the staggered establishment of weather stations across the US during 1870–1990 as a source of variation in access to weather information. Using a stacked difference-in-difference design, our results show that access to a weather station increases the agricultural productivity increases by approximately 15% after a county gains access to a weather station within a 25-mile radius. As the distance of a county to the nearest station increases, the agricultural productivity decreases.

*Anand: Tobin College of Business, St. John's University, anandv@stjohns.edu. Li: Wisconsin School of Business, University of Wisconsin-Madison, hli434@wisc.edu. We thank the participants at the Grainger Summer Research Workshop at the UW-Madison for helpful comments and suggestions.

1 Introduction

Weather risks pose significant challenges to economic productivity, particularly in sectors highly sensitive to climatic fluctuations, such as agriculture. This paper examines whether improved access to weather forecasts can help mitigate these risks and improve agricultural performance.

Agriculture is a vital sector for economic development and food security, yet it is among the most vulnerable to weather conditions. Farmers rely on information about future weather, both short term and long term, for their input decisions and adaptation to adverse weather shocks. Timely and accurate weather forecasts can help farmers make informed input decisions, reduce exposure to climate-related shocks, and improve overall productivity. Weather forecasts have been improving on account of significant public investments in meteorological operations and research (Bauer et al. 2015; Alley et al. 2019).Despite centuries of public investment in weather forecasting infrastructure and technology, there is limited empirical evidence on the productivity effects of access to weather information. This paper addresses that gap by examining the historical rollout of weather stations in the United States and their impact on local agricultural productivity.

We exploit the staggered establishment of weather stations across U.S. counties between 1870 and 1990 as a source of plausibly exogenous variation in access to weather forecasts. During this period, the number of weather stations increased from less than 10 to more than 10,000, driven by both civilian and military needs. These stations played a key role in producing and disseminating forecasts. In the early decades, proximity to a station was critical for accessing timely forecasts, as dissemination depended on local communication channels. Moreover, being closer to a station improves the accuracy of forecasts, since local observations improved the precision and relevance of weather predictions. Using newly digitized data from the National Centers for Environmental Information and agricultural census records, we estimate the effect of weather station proximity on county-level agricultural productivity over the twentieth century.

We begin with a generalized two-way fixed effects (TWFE) model that uses distance to the

nearest weather station as a continuous measure of treatment intensity. The results indicate that agricultural productivity decreases as distance from the nearest station increases: a 10-mile increase in distance is associated with a 2.4 percent decline in productivity. IIn a specification using distance bins, we find that counties with a weather station located within their boundaries experience a 23 percent increase in agricultural productivity relative to counties where the nearest station is over 100 miles away. The effects gradually diminish as the distance to the nearest station increases.

To address recent concerns about the validity of TWFE estimators under staggered treatment timing and treatment effect heterogeneity, we implement a stacked difference-in-differences (DID) design. The stacked event study shows that the effects of weather station access increase gradually over time, with productivity increasing by about 15 percent fifteen years after the station is established. The pre-treatment coefficients are close to zero, supporting the parallel trends assumption.

We First, it contributes to the literature on the value of weather forecasts and weather information (e.g., Rosenzweig and Udry, 2019; Shrader, 2020; Burlig et al., 2024;Anand, forthcoming). Using nearly a century of data, we provide novel evidence that improved access to more accurate weather forecasts enhances agricultural productivity. In this setting, weather forecasts function as a form of publicly provided information that shapes farmers' input choices, risk management strategies, and ultimately, their productivity. Second, the paper contributes to the literature on the long-run trajectory of agricultural productivity in the United States (Andersen et al., 2018). Kantor and Whalley (2019) discuss the role of technological innovation in shaping agricultural productivity. Our paper contributes to this literature by showing how access to weather forecasting infrastructure, as well as innovation in prediction technology, improved agricultural productivity over nearly a century.

The remainder of the paper is structured as follows. Section 2 provides historical background on weather forecasting and station, how farmers use weather forecasts, and why proximity to weather stations may matter. Section 3 describes the data. Section 4 outlines the empirical strategy

and presents results, and Section 5 concludes.

2 Weather Forecasts and Weather Stations

2.1 Developments of weather forecasts and weather stations.

Early Analog Forecasting and Weather Station Networks (Late 19th–Early 20th Century) Modern weather forecasting began in the late 1800s with the expansion of telegraph-linked weather stations across national networks. In the United States, Congress established a national weather service in 1870 under the Army Signal Corps, which set up dozens of observing stations to telegraph data to Washington, D.C.[1] By the 1880s, weather forecasting in the U.S. was based on synoptic methods: meteorologists manually plotted weather station observations on maps to identify pressure systems, fronts, and trends. For example, by 1878 the U.S. Signal Service operated 284 stations taking thrice-daily observations (temperature, pressure, winds, etc.) that were wired to Washington for daily forecasts. Data collection relied on a growing network of surface weather stations (often staffed by volunteers) transmitting observations via telegraph and telephone. In 1928, the teletype machine replaced the telegraph as the primary means of distributing weather information, speeding up data communication. Forecasters used these simultaneous reports to make short-term predictions, often drawing on pattern recognition and personal experience. Scientific advances from the Bergen School in Norway (e.g. the polar front theory of cyclones) began influencing U.S. practice, leading to the Weather Bureau's establishment of an Air Mass Analysis unit in 1934 – reflecting the adoption of modern frontal analysis to improve forecast accuracy.

These primitive forecasts – though often inaccurate – nonetheless aided day-to-day planning in farming, navigation, and commerce. Weather stations were thus the backbone of analog forecasting. Between 1900 and 1950, the U.S. Weather Bureau expanded the network of weather stations significantly, improving data coverage and forecast skill. The 1930s saw the addition of upper-air

[1] https://www.nps.gov/parkhistory/online_books/caha/shrs/index.htm.

observations, with the introduction of radiosonde systems in 1937, providing vertical profiles of temperature, humidity, and pressure.

Numerical Weather Prediction Revolution (Mid-20th Century). The mid-20th century marked a paradigm shift in meteorology: the use of electronic computers to forecast the atmosphere's future state via mathematical equations, a method known as numerical weather prediction (NWP). The theoretical foundation for this approach was laid by Vilhelm Bjerknes in 1904, who proposed that accurate weather forecasts were possible if enough reliable observational data could be fed into the governing equations of atmospheric physics. This vision became feasible only after World War II with advances in computing. In 1950, Jule Charney and John von Neumann produced the first successful 24-hour numerical forecast using the ENIAC computer, marking the beginning of computer-based forecasting. By the mid-1950s, numerical weather prediction (NWP) had evolved from an experimental tool into an operational forecasting method.

With the rise of NWP, the 1950s–60s saw major advancements in atmospheric data collection in weather stations. The global upper-air network expanded with widespread radiosonde use, while radar technologies like the AN/CPS-9 and WSR-57 improved real-time storm tracking. These innovations enhanced forecasters' ability to collect the initial conditions required by numerical models and issue severe weather warnings. Station coverage grew—especially at airports and remote locations. The Weather Bureau partnered with military and academic institutions to improve observational infrastructure and forecast accuracy.

By the 1970s, improvements in computational power allowed forecast models to operate at higher spatial resolution—dropping from grid sizes of hundreds of kilometers in the 1960s to tens of kilometers in later decades. This refinement was enabled in part by denser and more accurate station data, which helped capture localized weather phenomena. Forecast lead times and accuracy improved markedly: whereas analog methods had limited reliability beyond a day or two, NWP models by the 1980s could produce three-day forecasts with the same skill as one-day forecasts in

the 1950s.

Satellite-Based Forecasting and Modern Observing Systems (1960s–Present) A major breakthrough in weather observation came with satellite-based forecasting. On April 1, 1960, the U.S. launched TIROS-1, the first weather satellite, which transmitted the first images of Earth's cloud cover from space. By the late 1960s, geostationary satellites like ATS-1 enabled continuous regional monitoring, while polar-orbiting satellites provided global high-resolution data twice daily—transforming hurricane tracking and large-scale weather surveillance. At the same time, weather radar evolved from rudimentary post-war systems to Doppler radar, which measures precipitation motion and storm rotation. The NEXRAD (Next Generation Radar) network, completed in the 1990s, significantly improved tornado and severe thunderstorm warning lead times. Together, satellite, radar, and surface station data feed into numerical models through data assimilation, allowing for accurate and timely forecasts.

2.2 Weather forecasts and agriculture

Agriculture—perhaps more than any other industry—illustrates the real-world value of weather stations, from the 19th-century farmer reading a posted "farmers' bulletin" at the local post office to today's precision farmer receiving targeted forecasts on a smartphone. Across time, the consistent thread is clear: better weather data empowers smarter decisions, directly enhancing farm productivity. In what follows, we discuss how farmers use forecast information in both the short term and the long term to guide their agricultural practices.

Short-Term Use (Daily/Weekly): Farmers heavily integrate short-term weather forecasts—ranging from daily to one-week outlooks—into operational decisions such as planting, spraying, harvesting, and irrigation timing to avoid adverse weather, and to take precautions ahead of expected adverse weather (Cabrini et al., 2023). Short-range forecasts also support pest and disease management and frost and heat mitigation. These immediate decisions help mitigate crop loss, optimize resource use, and avoid weather-related risks in the critical growth periods leading up to harvest.

For example, the Weather Bureau even established a special "Fruit-Frost" service during World War I to warn orchard growers of impending frosts.

Long-Term Planning (Seasonal/Annual): Beyond the day-to-day, farmers also consult longer-range forecasts – such as seasonal climate outlooks (monthly/seasonal trends, drought forecasts, El Niño/La Niña predictions) and even multi-year climate projections – to guide more strategic decisions like crop selection, acreage planning, irrigation provisioning, and input purchases. For example, Burlig et al. (2024) show that accurate long-range forecasts of the Indian summer monsoon's onset led farmers to substantially adapt their behavior along multiple margins, changing cultivated area, crop choice, farm input expenditure, and off-farm business.

2.3 Why Proximity Matters?

Weather stations performed primarily three tasks: collect data on local weather observations, generate weather forecasts, and disseminate forecasts and alerts to people and institutions. Not all weather stations were responsible for forecasting. However, forecasts, once generated, were shared with all stations for distribution to the public. Before the advent of radio around early 20th century, these forecasts were often displayed at the weather stations (using a notice or a signal flag) and public notice boards, and distributed through the local newspaper networks.[2] Weather stations established for a specific need, such as marine traffic or army, would also share these forecasts with the relevant institutions directly.

Counties that were closer to weather stations might have an advantage in accessing weather forecasts over counties that were farther for two reasons. First, counties closer to a station were more likely to have better access to forecasts that were communicated using physical notices or flags, and through local newspapers. In the late 1800s and early 1900s, the U.S. Weather Bureau used signal flags and bulletin boards to disseminate forecasts. Local offices or cooperative observers posted daily forecasts in post offices, railroad depots, and town centers. For example,

[2]The first forecast using radiotelephone was made in 1914 in North Dakota.

a farmer in a county near a signal station might see a white flag with a black square—signaling a cold wave warning—posted prominently by a Weather Bureau observer. Farmers farther away, especially those in remote counties without a local station or observer, would not have timely or reliable access to this visual signal. This physical proximity to flag-based alerts provided a tangible advantage for planning tasks like protecting crops from frost or delaying harvests..[3] Before the widespread adoption of radio and television around the 1940s, daily weather forecasts were often printed in local newspapers, which received forecast telegrams from the nearest Weather Bureau office. A county located near a first-order station (such as in Chicago or St. Louis) would typically receive more timely, detailed and location-specific forecast reports via the local press. In contrast, counties farther from a station might receive only generalized or delayed forecasts. However, advances in communication technology would decrease the cost and timeliness of distributing and accessing the forecasts. This would have likely reduced the proximity advantage.

Second, counties closer to stations were also likely to have access to better forecasts, since weather observations collected as well as forecasts generated were likely to be more accurate for geographies closer to the stations. In the early days of the U.S. Signal Service (1870s–1890s), forecasts were produced in Washington, D.C., using telegraphed observations from a sparse network of stations. Areas near these stations, such as Chicago or St. Louis, received forecasts directly informed by local measurements, while distant rural counties had to rely on extrapolated and less reliable projections. From the 1900s through the 1920s, forecasters drew daily synoptic weather maps by hand, but their ability to predict storm movements or pressure changes was only as good as the density of local station data. The introduction of air mass analysis and frontal theory by Norwegian scientists in 1917, along with the Norwegian Cyclone Model for interpreting surface pressure maps and frontal boundaries, provided a scientific foundation for storm pre-

[3]. In fact, around 1901 it was noted that the Post Office Department was delivering slips of paper with the daily forecast, along with frost and cold-wave warnings, to people's doors with their morning mail. The one disadvantage to the system was the mail carriers started their routes about 7:00 a.m. and that day's forecast was not issued until 10:00 a.m., so the previous night's forecasts were used. See details in https://www.weather.gov/timeline.

diction that U.S. Weather Bureau forecasters quickly adopted. These methods greatly improved forecast accuracy in counties near stations, where detailed observations were available. Around the same time, stations began launching upper-air observations—initially with kites and later with balloons—enabling vertical atmospheric profiling that enhanced predictions of cold fronts, thunderstorms, and storm intensity. Later innovations, such as the deployment of radar in the 1950s and Doppler radar in the 1990s, further improved local severe weather forecasting. Beginning in the 1950s, the rise of numerical weather prediction (NWP) made station data even more critical, as models required dense, high-quality surface and upper-air observations to initialize forecasts. In recent decades, automated surface stations and satellite-linked networks have expanded real-time coverage, allowing for highly localized and timely predictions. Throughout these developments, counties closer to stations consistently benefitted from more accurate and geographically relevant forecasts due to better observational input.

3 Data and Descriptive Statistics

Weather station data. Historical weather station data are obtained from the Master Station History Report (MSHR) database of the National Center for Environmental Information's (NCEI) Historical Observing Metadata Repository (HOMR). The MSHR data provide a variety of information for more than 30,000 land-based weather stations in the United States during the years 1738–2021. This includes data on location, ownership, nature of operations, and operating dates throughout the lifespan of a weather station. We also augment the MSHR data with information collected from the textual historical account of all the existing and defunct weather stations operated by the National Weather Service (NWS).[4]

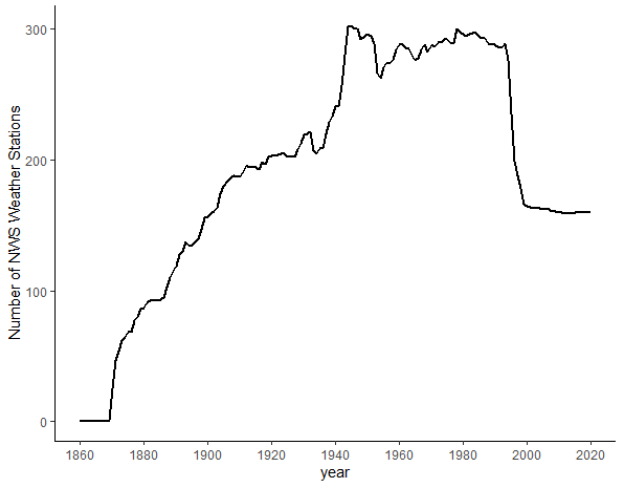
Figure 1 shows the number of weather stations established by the Weather Bureau (WB) and, later, by the National Weather Service (NWS). The first weather station was set up by the Army Signal Services in July 1869 in San Luis Obispo, California.[5] The number of weather stations

[4]Textual histories are obtained from https://www.weather.gov/ilx/nws-wb-history.

[5]The meteorological responsibilities of the Signal Service were transferred to U.S. Weather Bureau, then created

gradually increased until 1930s. During the World War II the number of new stations increased sharply over a short period of time. While the overall number of weather stations remains somewhat stable after the 1950s until 1990s, some stations were closed and the new ones were opened during these years. Figure 2 shows the location of weather stations in 1910 and 1990. The figure also highlights the counties within 25 miles (shaded dark red) and 50 miles (shaded light red) of the nearest weather station. Consistent with the overall increase in the number of stations, we observe a broader geographic spread of stations over time.

Figure 1: Number of weather stations set up the NWS/WB



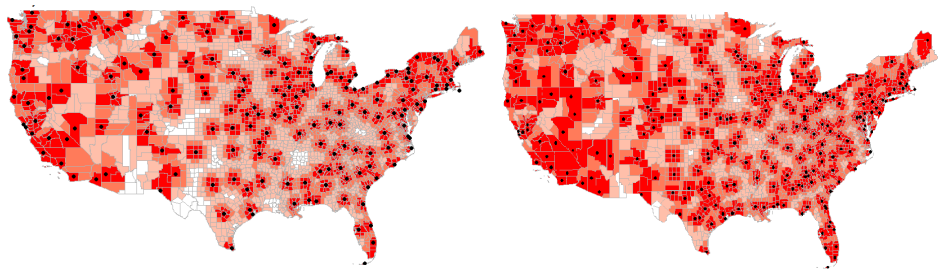
Notes: The figure plots the number of active weather stations setup and operated by the Weather Bureau or NWS in a year.

After 1990s, a major restructuring at the NWS resulted in consolidation of severe weather prediction and weather forecasting activity among a reduced number of stations. We limit our analysis to the new weather stations opened between 1909–1992.[6] Our analysis exploits variation in the distance to the nearest weather station across counties and over time. Figure 3 displays the distribution of distances from county centroids to the nearest weather station. The median distance in our sample is approximately 28 miles.

as a unit under the Department of Agriculture, in 1890. The bureau became the National Weather Service in 1970.

[6]The census data on crop revenue is available from the year 1909 onward.

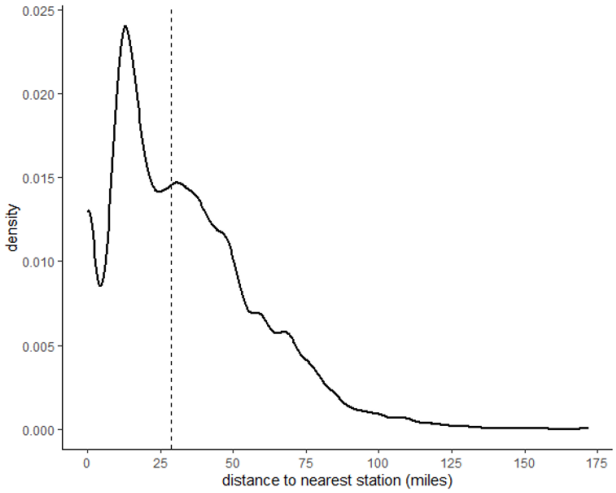
Figure 2: Weather Stations during 1910 and 1990



(a) Stations in 1910 **(b)** Stations in 1990

Notes: The figures plots the location of NWS weather stations (black dots) in year 1910 (Panel A) and 1990 (Panel B). Counties shaded dark red (light red) are within 25 miles (50 miles) of the nearest station. Counties with no weather station within 50 miles of their centroid are shaded white.

Figure 3: Distribution of distance to the nearest station from county centroid



Notes: The figure plots the distribution of distance to the nearest weather station from county centroid in our sample.

Agricultural data. Data on crop revenue and farmland area are obtained from the Agricultural Census for the years 1909 to 2012 (Haines et al., 2018).[7] We obtain county-level data on farm size and crop revenue and construct our main output variable, farm productivity, as the ratio of farm

[7]The agricultural productivity data are available for the years 1909, 1919, 1924, 1939, 1944, 1949, 1954, 1959, 1964, 1969, 1974, 1978, and between 1982–2012 every five years.

revenue to farmland area, followed by (Kantor and Whalley, 2019).

Final sample. Our primary sample includes 45,288 county-year observations from 49 states in the US. For each county-year, we observe the distance to the nearest station, measures of farm productivity, and other county-level characteristics.

4 Empirical Design and Results

4.1 Effects of Proximity to Weather Stations: Generalized Two-Way Fixed Effects Approach

To estimate the effect of proximity to weather stations on agricultural output and productivity, we begin with a generalized two-way fixed effects (TWFE) approach that leverages the staggered roll-out of weather stations across counties over time. Specifically, we exploit within-county variation in the distance to the nearest weather station. Unlike standard TWFE models that typically rely on binary treatment indicators, our specification treats distance as a continuous measure reflecting the intensity of exposure to the treatment.

The generalized difference-in-differences specification is

$$y_{ct} = \beta Distance_{ct} + \alpha_c + \delta_{st} + \varepsilon_{ct}, \tag{1}$$

where y_{ct} is the log of agricultural output per acre of county c in year t, $Distance_{ct}$ is the distance in miles of the county c centroid from the nearest weather station in year t. δ_{st} and α_c are state–year and county fixed effects that control for time-series trends in agricultural output and productivity within a state, and the time-invariant county-specific characteristics that may affect agricultural outcomes, respectively. The coefficient β measures the percentage change in agricultural productivity due to an additional mile in distance from the nearest weather station on the agricultural proximity to a weather station. We weight the regression by log of farmland area in a county.

Errors are clustered at the county level.

The previous estimates assume that the effect of distance to the nearest weather station is constant across all distances. However, the impact of distance may be nonlinear due to several factors, such as diminishing informational value with distance. To explore these potential nonlinear effects, we categorize distance into bins and estimate the following specification:

$$y_{ct} = \beta_1 \mathbf{1}_{(distance_{it}=0)} + \beta_2 \mathbf{1}_{(0<distance_{it}<25)} + \beta_3 \mathbf{1}_{(25<distance_{it}<50)} + \beta_4 \mathbf{1}_{(50<distance_{it}<100)} + \alpha_c + \delta_{st} + \varepsilon_{ct}, \tag{2}$$

where each coefficient β_i represents the average effect on agricultural productivity of having a weather station within the specified distance range, relative to the omitted category of areas located more than 100 miles away.

Column (1) of Table 1 presents the result from estimating equation 1. It suggests that agricultural productivity declines as the distance to the nearest weather station increases. Specifically, a 10-mile increase in distance is associated with an estimated 2.4% decrease in productivity. Column (2) of Table 1 presents the results from estimating equation 2, and Figure 4 visualizes the corresponding estimated effects. The coefficient on $\mathbf{1}_{(distance=0)}$ indicates that counties with a weather station located within their boundaries experienced, on average, a 23.1% increase in agricultural productivity relative to counties situated more than 100 miles away. The effect gradually declines with distance: counties within 0–25 miles see a 13.8% increase, those 25–50 miles away experience a 13.5% increase, and those 50–100 miles away still benefit with a 9.7% increase. All estimates are statistically significant. These findings point to strong gains from close proximity to weather stations.

4.2 Effects of Proximity to Weather Stations: Stacked DID Approach

Recent methodological critiques have highlighted important limitations of the standard two-way fixed effects (TWFE) model in DID designs with staggered treatment timing. Goodman-Bacon

Table 1: Estimated effect of distance on agricultural productivity

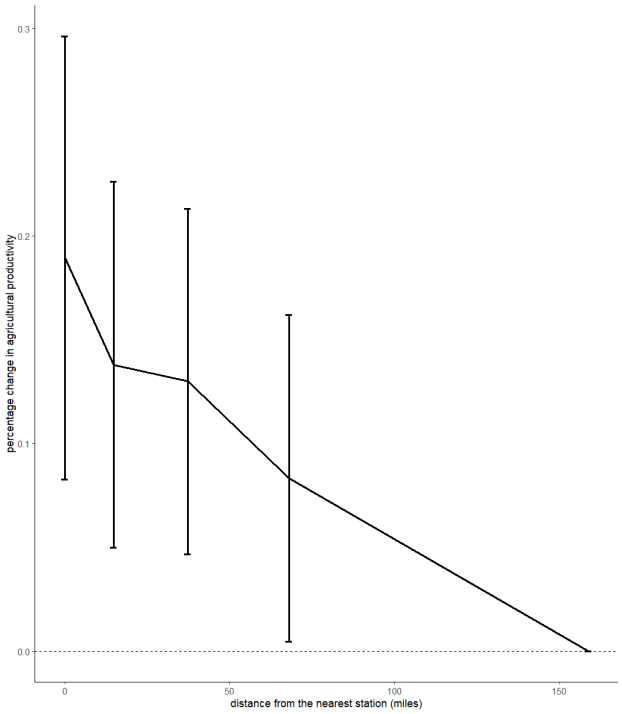
Dependent Variable:	log of agricultural output per acre of farmland	
Model:	(1)	(2)
distance	-0.0024***	
	(0.0006)	
$\mathbf{1}_{(distance=0)}$		0.2313***
		(0.0617)
$\mathbf{1}_{(0<distance<25)}$		0.1376***
		(0.0483)
$\mathbf{1}_{(25<distance<50)}$		0.1353***
		(0.0455)
$\mathbf{1}_{(50<distance<100)}$		0.0969**
		(0.0429)
county fixed-effects	Yes	Yes
state-year fixed effects	Yes	Yes
Observations	45,288	45,288

Notes: The table shows the results from estimating the regression models in Equation 1. The dependent variable is log of agricultural output per acre of farmland. All specifications include county and state-year fixed effects. All regressions are weighted by log of total farmland in county. Standard-errors clustered at county level are in parentheses. Significance Level: ***: 0.01, **: 0.05, *: 0.1

(2021) shows that when treatment is adopted at different times across units and the treatment effects are heterogeneous over time, TWFE models can produce biased estimates due to inappropriate comparison groups and problematic weighting. These concerns also extend to the generalized TWFE specification used in Equation 1, which relies on variation in distance to the nearest weather station over time (Callaway et al., 2025). In our context, the impact of weather stations on agricultural productivity may not be constant but may evolve with technological changes or improvements in forecast accuracy. For example, the introduction of weather radar in the 1950s likely enhanced the precision of forecasts, potentially amplifying the productivity effects of nearby stations. As such, the staggered establishment of weather stations across counties poses a potential source of bias that motivates a more robust empirical approach.

To address the limitations of the generalized TWFE model in the presence of staggered treatment timing and potentially heterogeneous treatment effects, we adopt a stacked DID design (Cengiz et al., 2019; Deshpande and Li, 2019; Wing et al., 2024). This approach re-structures the data into a series of cohort-specific panels, each centered on the "event" year—defined as the year a county first receives a weather station. For each treated cohort, we construct an event-time win-

Figure 4: Estimated effect of distance on agricultural productivity



Notes: The figure plots the estimated effect of distance from nearest station on agricultural productivity based on estimating the Equation 1. The dependent variable is log of agricultural output per acre of farmland. The markers with whiskers plot the estimated effect of distance (in miles on x-axis) on agricultural productivity (y-axis) along with the associated 95% confidence interval. The x-axis coordinate of markers correspond to the average distance in the corresponding bin. Standard Errors are clustered at county level.

dow spanning certain years before and after treatment, and compare outcomes in treated counties to a control group of counties. By stacking these cohort-specific event studies and estimating them jointly, we avoid contamination from inappropriate control groups and the problematic implicit weighting schemes that often arise in TWFE models with staggered treatment timing.

To be specific, we contract stacks as follows. Using the historical record of weather station establishments, we classify counties in each year into two categories: those with an active weather station located within a critical radius of 25 miles ("treated counties") and those without any station within that radius ("control counties").[8] For each new weather station established since 1909, we generate a separate stack of county-year observations. A county is considered treated in a given

[8]For our main analysis, we assume a county has access to a weather station if the nearest weather station to the county is within a radius of 25 miles. We vary the value of critical radius to test the sensitivity of our estimates.

stack if it falls within 25 miles of the newly established station and has not had any station within that radius in the prior 20 years. Control counties, in contrast, are those that had no station within 25 miles during the 20 years before and after the station opening in that stack. For each stack h, we two sets of variables: (1) an indicator variable $stack_h$ which is 1 for the stack h and 0 otherwsie; (2) a set of indicator variables for event period, D^{τ}_{th}, which is 1 if the year t in stack h is τ years after the new station is opened in that stack, and 0 otherwise. In total, we construct 114 such stacks to form our final analysis panel. Appendix Figure 6 displays the number of weather stations used to construct the stacked event study sample.

Following the approach of Cengiz et al. (2019) and Deshpande and Li (2019), we first estimate the following specification using the stacked panel to assess the dynamic effects of treatment:

$$y_{cth} = \sum_{\tau} \beta_{\tau}(Treated_{ch} \times D^{\tau}_{th}) + \alpha_{ch} + \delta_{th} + \varepsilon_{cth} \tag{3}$$

where $Treated_{ch}$ is 1 if county c is a treated county in stack h, and 0 otherwise. D^{τ}_{th} is 1 if year t is τ years since the event in stack h, else it is 0. δ_{th} and α_{ch} are year-stack and county-stack fixed effects, respectively. The coefficients of interest are β_{τ} that measure the effect of access to a weather station on agricultural outcome τ years after (or before if τ is negative) the county c gets access to a weather station.

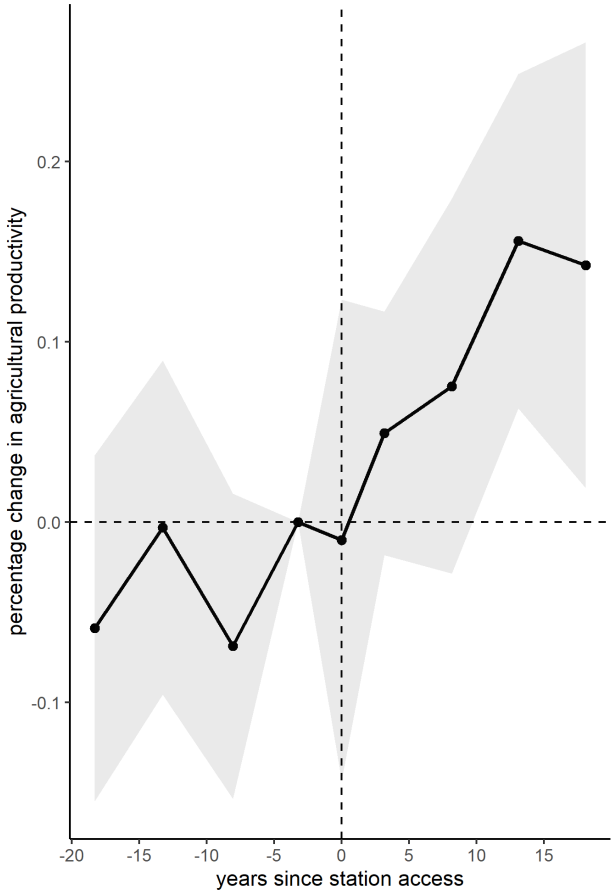
Identification assumption. Our identification assumption is that, in the absence of treatment, treated and control counties would have followed parallel trends in agricultural productivity. A potential concern is that the placement of weather stations may depend on factors unobserved to the econometrician. For instance, if counties with more intensive agricultural activity were more likely to receive a station, this could introduce selection bias. However, this form of selection—based on levels rather than trends—does not necessarily undermine our approach. The stacked DID design remains valid in the presence of selection on levels, as long as the trends in outcomes would have been similar in the absence of treatment. That is, if weather stations were placed in areas

with high—but stable—agricultural productivity, our approach remains valid, and the relative null pre-trends support this.

The concern becomes more serious if station placement is driven by trends - for example, if stations were preferentially established in counties experiencing a boom in agricultural productivity. This would violate the parallel trends assumption and bias our estimates. Most weather stations in our sample period were established to support aviation operations, which provides some degree of exogeneity with respect to agricultural activity. The placement of stations was often driven by the needs of civil and military aviation, rather than local farming conditions. This institutional feature helps mitigate concerns that station locations were systematically chosen based on agricultural productivity. However, some caution is still warranted. Airports—where many weather stations were co-located—may themselves have been established in areas experiencing broader economic growth, including agricultural booms. If this is the case, then the location of weather stations could still be correlated with unobserved upward trends in agricultural productivity, which would violate the parallel trends assumption and bias our estimates. We directly address these concerns through a set of robustness checks and falsification tests, discussed in detail in Section 4.3.

Results. Figure 5 presents the results from estimating Equation 3. The figure shows that agricultural productivity in treated counties increases following the establishment of a weather station, with treatment effects growing larger over time. This pattern suggests that the effects of weather stations may take time to materialize—potentially due to gradual public investment in agricultural risk management, delayed adaptation in crop choice and farming practices, or a learning process among farmers. It may also reflect improvements in forecast accuracy over time driven by technological advancements. Approximately 15 years after a station's establishment, agricultural productivity has increased by about 15%. Moreover, the figure provides support for the parallel trends assumption: in the periods prior to treatment, the differences in productivity between treated and control counties are statistically indistinguishable from zero.

Figure 5: Estimated effect of weather station access on agricultural productivity



Notes: The figure shows the estimated effect of weather station establishment on agricultural productivity in affected counties (counties within 25 miles of the new station) in the years before and after the station establishment. The figure plots the estimated coefficients β_τ from equation 3. The dependent variable is log of agricultural output per acre of farmland. The markers plot the estimated coefficients and the shaded region is the 95% confidence interval. The regression is weighted by the log of farmland area in county. The regression includes stack-county and stack-year fixed effects. Standard Errors are clustered at stack-county level.

To estimate the average effect of weather station access on the treated counties, we estimate the following specification:

$$y_{cth} = \beta(Treated_{ch} \times Post_{th}) + \alpha_{ch} + \delta_{th} + \varepsilon_{cth} \tag{4}$$

where $Post_{th}$ is an indicator equal to 1 if year t is after the weather station establishment year in stack h. Table Table 2 reports the results from estimating Equation 4. The estimates indicate that,

on average, agricultural productivity increases by approximately 15% after a county gains access to a weather station within a 25-mile radius.

Table 2: Estimated effect of distance on agricultural productivity

Dependent Variable:	log of agriculatural output per acre of farmland
Treated × *Post*	0.1487***
	(0.0431)
county fixed-effects	Yes
state-year fixed effects	Yes
Observations	1,029,248

Notes: The table shows the results from estimating the regression models in Equation 4. The dependent variable is log of agricultural output per acre of farmland. All specifications include stack-county and stack-year fixed effects. All regressions are weighted by log of total farmland in county. Standard-errors clustered at stack-county level are in parentheses. Significance Level: ***: 0.01, **: 0.05, *: 0.1

4.3 Robustness Tests

Alternative definitions of treated and control groups by different radii. Our baseline specification defines treated counties as those located within a 25-mile radius of a weather station. To assess the robustness of our findings, we repeat the analysis using alternative distance thresholds. Appendix Figure 7 presents the results based on varying the treatment radius. We find that the positive effects of weather station access on agricultural productivity remain consistent across different definitions of treatment distance, although the estimated magnitude tends to be slightly larger when using a more restrictive (e.g., 15-mile) radius and somewhat attenuated when the radius is expanded (e.g., to 50 miles). This suggests that the proximity to weather stations plays an important role in driving the observed productivity gains.

Potential confounding from concurrent local investments. Another potential concern is that the establishment of weather stations may coincide with other local investments—such as infrastructure upgrades, agricultural extension services, or rural electrification—that could independently affect agricultural productivity. To address this possibility, we conduct a robustness check in which we exclude counties where the weather station was directly established. By doing so, we

rule out the influence of investments that are strictly localized to the counties receiving the station. Appendix Figure 8 presents the results, which remain consistent with our baseline estimates. However, it is important to note that this test only eliminates concerns about local confounding. If the establishment of a weather station is part of a broader regional investment program that generates spillover effects—such as improved transportation networks or market access—this strategy may not fully account for those externalities.

US Army Observations - TBA. As discussed before, a potential threat to our identification is that, if station placement is driven by trends—for example, if stations were preferentially established in counties experiencing a boom in agricultural productivity. Though most weather stations in our sample period were established to support aviation operations, which provides some degree of exogeneity with respect to agricultural activity, some caution is still warranted. Many weather stations were co-located with airports, and airport siting may itself have been influenced by broader regional growth, including agricultural expansion. If so, the location of weather stations could still be correlated with unobserved upward trends in productivity.

To further address this concern, we plan to exploit variation in the historical purpose of station placement. For example, some of the earliest weather stations were established by the U.S. Army and Navy in locations chosen for strategic or meteorological coverage purposes, rather than local economic activity. We use this institutional heterogeneity to isolate plausibly exogenous variation in access to weather stations. Additionally, we conduct robustness checks using alternative control groups and pre-treatment trends to validate our identification strategy. Specifically, we will estimate the regression Equation 1 using a sub-sample of county observation where the nearest weather station's purpose is to cater to the US Army.

5 Conclusion and Next Steps

Our preliminary results show that proximity of a county to a weather station increases the agricultural productivity relative to counties that do not have access to a station. We also propose to examine how these proximity effects change over time. For example, it is likely that during the first half of the 20th century when forecasting technology improved with the advent of better computing systems, counties closer to the weather station may have access to more accurate weather forecasts. However, this proximity effect may have declined over time due to the advances in better communication technologies. We also plan to examine how these effects vary for different types of crops. For example, crops which are more sensitive to weather events (such as frost) may benefit more from access to weather information relative to other crops. Overall, our results will provide evidence on to what extent the access to weather information have meaningful effect on agricultural productivity, and if these effects depend on the proximity to the source of forecast information.

References

Richard B. Alley, Kerry A. Emanuel, and Fuqing Zhang. Advances in weather prediction. *Science*, 363(6425):342–344, January 2019. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aav7274. URL https://science.sciencemag.org/content/363/6425/342.

Vaibhav Anand. Does getting forecasts earlier matter? evidence from winter advisories and vehicle crashes. *American Economic Journal: Economic Policy*, forthcoming.

Matthew A. Andersen, Julian M. Alston, Philip G. Pardey, and Aaron Smith. A century of u.s. farm productivity growth: A surge then a slowdown. *American Journal of Agricultural Economics*, 100(4):1092–1111, 2018.

Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction. *Nature*, 525(7567):47–55, September 2015. ISSN 1476-4687. doi: 10.1038/nature14956. URL https://www.nature.com/articles/nature14956.

Fiona Burlig, Amir Jina, Erin M Kelley, Gregory V Lane, and Harshil Sahai. The value of forecasts: Experimental evidence from india. Working Paper 32173, National Bureau of Economic Research, February 2024. URL http://www.nber.org/papers/w32173.

Silvina Cabrini, Gary Schnitkey, and Scott Irwin. The use of climate information in midwest agriculture: Results from a farmer survey — part i. farmdoc daily (13):12, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, January 2023. URL https://origin.farmdocdaily.illinois.edu/2023/01/the-use-of-climate-information-in-midwest-agriculture-results-from-a-farmer-survey-par
html.

Brantly Callaway, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Difference-in-differences with a continuous treatment, 2025. URL https://arxiv.org/abs/2107.02637.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The Effect of Minimum Wages on Low-Wage Jobs*. *The Quarterly Journal of Economics*, 134(3):1405–1454, 05 2019. ISSN 0033-5533. doi: 10.1093/qje/qjz014. URL https://doi.org/10.1093/qje/qjz014.

Manasi Deshpande and Yue Li. Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy*, 11(4):213–48, November 2019. doi: 10.1257/pol.20180076. URL https://www.aeaweb.org/articles?id=10.1257/pol.20180076.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021. ISSN 0304-4076. doi: https://doi.org/10.1016/j.jeconom.2021.03.014. URL https://www.sciencedirect.com/science/article/pii/S0304407621001445. Themed Issue: Treatment Effect 1.

Michael Haines, Price Fishback, and Paul Rhode. United states agriculture data, 1840 - 2012. *Inter-university Consortium for Political and Social Research [distributor]*, August 2018. URL https://doi.org/10.3886/ICPSR35206.v4.

Shawn Kantor and Alexander Whalley. Research proximity and productivity: Long-term evidence from agriculture. *Journal of Political Economy*, 127(2):819–854, 2019. doi: 10.1086/701035. URL https://doi.org/10.1086/701035.

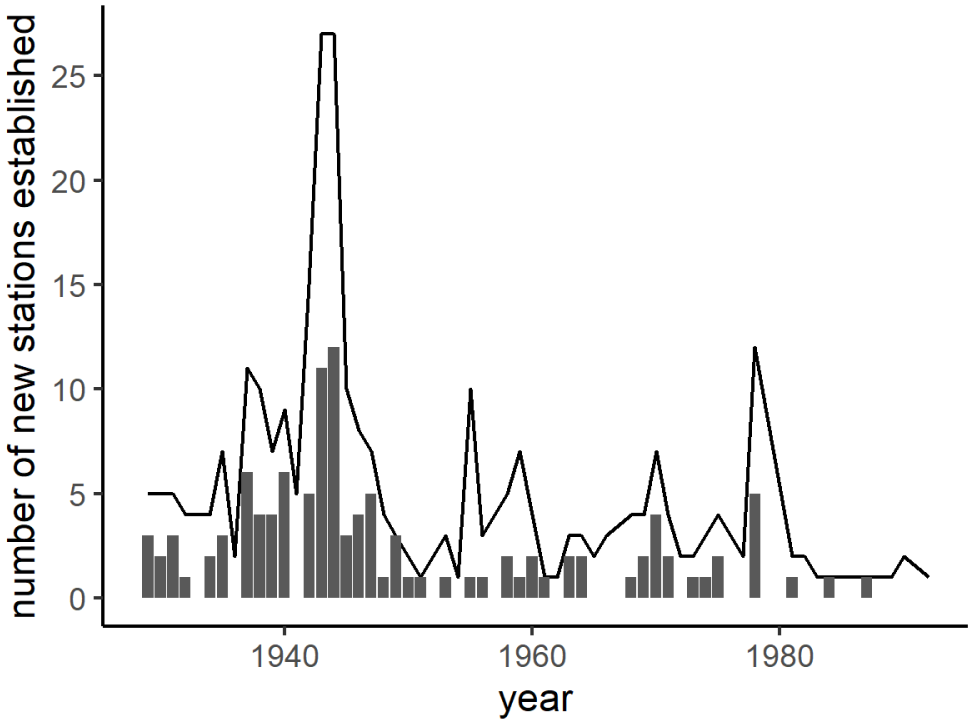
Mark R Rosenzweig and Christopher R Udry. Assessing the benefits of long-run weather forecasting for the rural poor: Farmer investments and worker migration in a dynamic equilibrium model. Working Paper 25894, National Bureau of Economic Research, May 2019. URL http://www.nber.org/papers/w25894.

Jeffrey Shrader. Improving Climate Damage Estimates by Accounting for Adaptation. SSRN Scholarly Paper ID 3212073, Social Science Research Network, Rochester, NY, June 2020. URL https://papers.ssrn.com/abstract=3212073.

Coady Wing, Seth M Freedman, and Alex Hollingsworth. Stacked difference-in-differences. Working Paper 32054, National Bureau of Economic Research, January 2024. URL http://www.nber.org/papers/w32054.

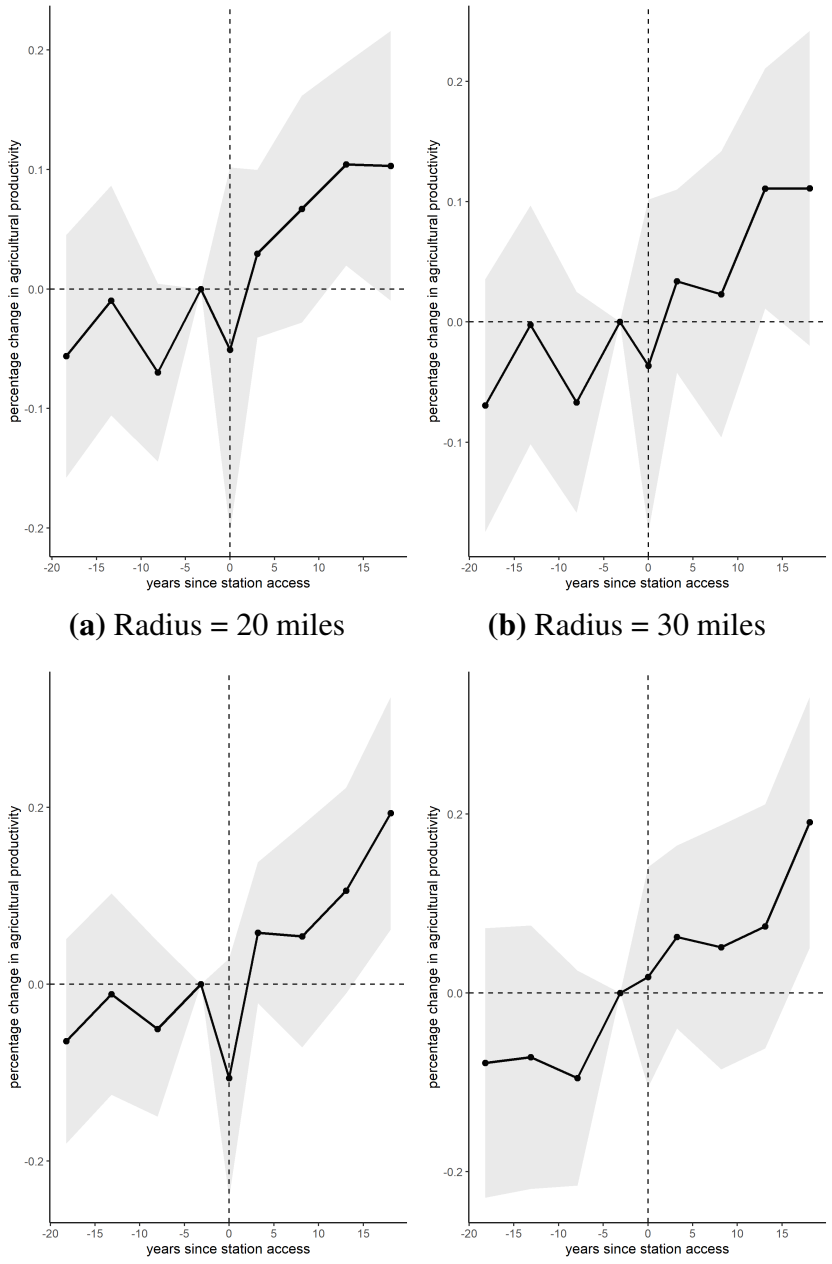
A Additional Tables and Figures

Figure 6: Distribution of Stacks Over Time



Note: The figure displays the total number of new weather stations established each year and the subset used to construct the stacks for the stacked DID study.

Figure 7: Estimated effect of weather station access on agricultural productivity by different distance radii

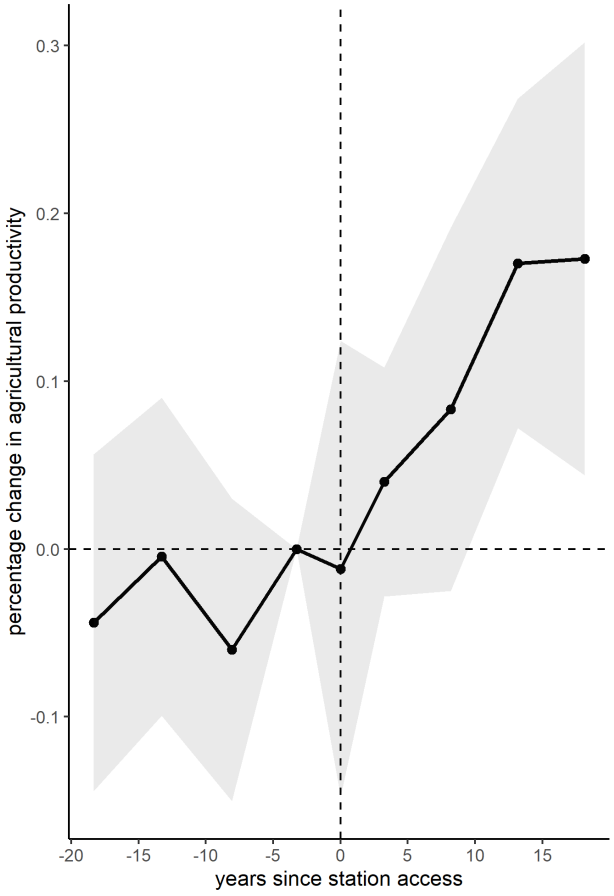


(a) Radius = 20 miles **(b)** Radius = 30 miles

(c) Radius = 40 miles **(d)** Radius = 50 miles

Note: The figures display the estimated stacked event study coefficients using alternative definitions of treatment radius. Each panel corresponds to a different distance threshold used to define access to a weather station.

Figure 8: Estimated effect of weather station access on agricultural productivity by excluding counties where the station is located



Note: The figures display the estimated stacked event study coefficients using treatment counties by excluding counties where the station is located.