Cooling Externality of Large-Scale Irrigation

Thomas Braun* Wolfram Schlenker†

July, 2022

Abstract

We provide novel causal evidence that large-scale irrigation is partially responsible for explaining why the weather over agricultural areas in some parts of the US has seen less warming than what is projected by climate models. Relying on a triple-differences setup linking local temperature trends to changes in irrigation dynamics in neighboring areas, while differentiating between months during the growing season relative to months outside the growing season, we find that large-scale irrigation heterogeneously shifts the entire temperature distribution towards cooler temperatures. Using historic irrigation and weather data around the Ogallala aquifer, we find such cooling-by-irrigation effect to propagate mainly downwind and to be especially strong in limiting hotter temperatures, which are especially harmful for crops. Counties whose upwind neighbors intensively increased irrigation experienced a significant decrease in the 99% percentile of their distribution of summer temperatures, offsetting the projected warming in climate models to date. From an agricultural perspective, large-scale irrigation therefore not only directly benefits plant growth by controlling water supply and enhancing heat tolerance of crops but also indirectly by avoiding harmful degree-days over the growing season in downwind areas, a positive externality which we quantify in the form of avoided yield losses due to reduced exposure to extreme heat. While the former effect is on average larger, there is significant heterogeneity with some counties benefiting as much from the cooling effect of upwind irrigation as from the decrease in heat sensitivity due to their own irrigation.

Keywords: Irrigation, Temperature, Agriculture, Spatial first differences
JEL Codes: Q54, Q1, Q25

Acknowledgments: This work was supported by the U.S. Department of Energy, Office of Science, Biological and Environmental Research Program, Earth and Environmental Systems Modeling, MultiSector Dynamics, Contract No. DE-SC0022141.

*Columbia University, email: tpb2128@columbia.edu
†Columbia University and NBER, email: wolfram.schlenker@columbia.edu

1 Introduction

There is a strong consensus that human activity is causing the climate to change[1] and is imposing a stress on numerous natural systems. There have also been many studies documenting the reverse effect of climate change and natural resource depletion on human systems. A select list includes the effect on agriculture (Schlenker and Roberts 2009, Burke and Emerick 2016), energy demand (Auffhammer and Aroonruengsawat 2011), mortality (Barreca, Clay, Deschenes, Greenstone and Shapiro 2016), labor supply (Graff Zivin and Neidell 2014), GDP growth (Burke, Hsiang and Miguel 2015), amenity value of climate (Albouy, Graf, Kellogg and Wolff 2016) or freshwater supply (World Bank Group 2016).

We provide novel evidence of an overlooked additional channel, namely that adaptive behavior itself has a measurable effect on climate. In the context of agricultural production, adaptive behavior to better cope with hotter temperatures under the constraint of limited water resources is going to induce significant regional changes in future irrigation practices – both decreases as aquifers run dry and increases via irrigation expansion – which will affect local warming dynamics. We empirically investigate whether changes in large-scale irrigation over the Ogallala aquifer in the southern US have historically been responsible for a cooling externality affecting regional weather patterns by reducing extreme heat.

A common theme that has emerged from recent empirical studies involving temperature effects is the detrimental impact of the upper tail of the temperature distribution. While adaptive behavior may counter-balance some of the harmful effects of higher temperatures[2], such strategies come at their own cost (Schlenker, Roberts and Lobell 2013) and, in the case of irrigation intensification, remain subject to the availability of water suitable for sustainable irrigation.[3] As regional irrigation practices are bound to continue to evolve in the coming decades due to increasing constraints on freshwater resources or new irrigation technologies, it is important to understand the extent to which irrigation is *causing* a cooling externality on extreme temperatures in order to anticipate the implications of these changes on local weather patterns besides the primary effect on crops and with numerous indirect effects on populations, the economy as well as ecosystems.

Using historic weather records and irrigation data from a vast area surrounding the

[1]The American Association for the Advancement of Science put out a statement on December 9, 2006 that "the scientific evidence is clear: global climate change caused by human activities is occurring now."

[2]For instance, alternative crop varieties (Butler and Huybers 2013, Moscona and Sastry 2021) or the adaptation of irrigation (Haqiqi, Grogan, Hertel and Schlenker 2021) can reduce the sensitivity to extreme heat.

[3]While some world regions have been in the position to grow large-scale irrigation in the long run (e.g., along the Ganges or the Yangtze), others have dramatically affected freshwater resources with significant consequences on irrigation practices (e.g., around the Aral Sea).

Ogallala aquifer (see Figure 1), we empirically investigate the effect of large-scale irrigation on downwind summertime temperatures. In order to establish causality in this empirical research setting, we would ideally rely on a counterfactual world in which no irrigation were taking place in order to compare temperature outcomes. In practice, however, we have to rely on the best possible proxy for such counterfactual outcome. For this purpose, we exploit a tripe-differences strategy linking spatio-temporal changes in temperatures to spatio-temporal changes in irrigation during the growing season relatively to April – assuming that, at identical levels of irrigation trends, any pair of counties *would* have followed a parallel warming (or cooling) trajectory between April and any other chosen month of comparison. While this approach certainly helps to reduce the risk of omitted variable bias, we further enhance its performance and reliability by augmenting it in a spatial first differences setting (Druckenmiller and Hsiang 2018). Indeed, by solely relying on contrasts between *adjacent* counties, spatial first differences better exploit the data's spatial informational content and improve on the validity of the parallel trends assumption necessary for causal inference in the present setting, as neighboring counties are arguably the most comparable in both observed and unobserved features.

Our paper provides novel evidence that there is a substantial but relatively local (at the county-level scale) cooling effect of irrigation that heterogeneously affects the entire summertime temperature distribution and induces important spillover effects propagating mainly downwind. In the region of interest, our difference-in-differences specification in spatial first differences detects that this cooling-by-irrigation externality is strongest in July and August and with respect to the highest temperature percentiles. Quantitatively, we estimate that irrigating all of a county's surface causes the July temperature distributions in downwind counties to shift the hottest temperatures by -0.9°C and the coolest by -0.4°C in our preferred specification and even larger in alternate specifications. We find a symmetric relationship: while areas that expanded irrigation saw a lower increase in temperatures, areas that decreased the irrigated area saw a higher increase in temperatures, suggesting that continued irrigation temporarily limits temperature increases, but this cooling effect ceases as soon as irrigation ends.

By locally decreasing hot summertime temperatures, irrigation helps to safeguard yields not only in irrigated crops by increasing tolerance to hotter and drier weather, but also in agricultural-intensive areas located downwind by reducing exposure to harmful degree-days. We quantify this positive externality in the form of avoided yield losses for corn and find this indirect (externalized) effect *on average* to be about one order of magnitude smaller than the direct (internalized) effect of irrigation. Although it is small on average, it can be locally important as there is significant heterogeneity. Under certain circumstances, some counties can benefit as much from the cooling effect induced by irrigation in upwind neighbors as from their own irrigation efforts. The effect may protect up to 5% of the corn production in specific counties and during relatively hot years. Clearly, this cooling externality of large-scale irrigation on downwind temperatures may also benefit a full range of other crop varieties besides corn and help to locally dampen the detrimental effects of extreme heat on mortality, etc.

Our paper makes four major contributions to the literature. First, to the best of our knowledge, there has been no *causal* evidence that adaptation to climate-induced constraints – such as the need for irrigation in the western tail of the Corn Belt or its progressive decline in northwestern Texas – has itself a measurable effect on the climate. While there have been farming myths (e.g., "the rain will follow the plow"), we demonstrate that large-scale irrigation can cause sizable local temperature changes. Such cooling via irrigation manifests heterogeneously along the entire temperature distribution with largest effects on the hottest temperatures and subsists besides the reduction in maximal temperatures related to cropland intensification (Mueller, Butler, McKinnon, Rhines, Tingley, Holbrook and Huybers 2016). Our paper thus contributes to the understanding of why observed trends in extreme heat in agricultural areas in the High Plains have been lower than projected by climate models (Schlenker 2020).

Second, we quantify the positive externality that irrigation indirectly causes on downwind agricultural areas by reducing crop exposure to extreme heat, and compare it to the primary on-the-spot effect of irrigation that enhances heat tolerance of irrigated crops. On the example of corn, we find that certain counties substantially benefit from irrigation practiced by their upwind neighbors.

Third, we argue that the positive externality caused by large-scale irrigation has beneficial spillover effects on other sectors of the economy that also benefit from a decrease in extreme heat as soon as they are located downwind.

Fourth, building upon a set of falsification checks, we discuss the methodological advantage of combining simple difference-in-differences with spatial first differences in order to reduce the risk of omitted variable bias. We validate our approach by incorporating irrigation information from neighboring counties and find a cooling effect that propagates mainly downwind of irrigated areas.

Our paper is organized as follows. Section 2 describes the context of our study and motivates the research question. Section 3 introduces the variables of interest and the underlying data. Section 4 describes and motivates the empirical strategy relying on spatial first differences to establish causality. Section 5 presents and discusses the corresponding results while Section 6 examines the economic implications of the previous findings. Section 7 concludes and highlights perspectives for further investigation regarding the cooling-by-irrigation externality.

2 Study Setting and Motivation

Only a few regions benefit from an appropriate geology to sustain large-scale and intensive irrigation for agriculture. In the US, exceptional irrigation conditions are offered by one of the vastest aquifer system in the world, namely the High Plains Aquifer (see Figure 1). Also known as the "Ogallala Aquifer", this shallow groundwater reservoir covers a total surface of about 450,000 km^2, i.e. about the size of California or 1.8 times the area of the Great Lakes combined. Located in the Great Plains, it is shared by eight states in the central southern US. It is formed by three connected networks: the northern system covers most of Nebraska and only small portions of South Dakota, Wyoming, Colorado and Kansas; the central system is mostly located in Kansas and the northern tip of Texas but also covers the Oklahoma panhandle and small sections of Colorado and New Mexico; the southern system is shared between northwestern Texas and eastern New Mexico. The water is flowing, through porous soil and between impermeable layers, from the northern to the southern system at an average speed of only 45m per year under unperturbed conditions.[4]

Thanks to its unique positioning relative to the Ogallala aquifer in terms of overlapping area, underlying water volume[5] and upstream location on the northern High Plains system, Nebraska is the most vastly irrigated state in the US to this date, both in absolute irrigated area (33,600km^2 of irrigated land in 2012) and in relative proportion to its surface (17% in 2012).[6] Historically, however, and until the end of the 1950's, the irrigation hot-spot in eastern Nebraska was falling far behind several counties in northwestern Texas which had experienced the fastest growth in irrigation adoption after World War II (see Figure A4). In fact, large-scale irrigation in both Nebraska and Texas began in the second half of the 1940's, when the invention of center pivot irrigation and the widespread adoption of motorized groundwater pumps made it possible to fully exploit the Ogallala aquifer's unique irrigation potential (Hornbeck and Keskin 2014). While more than half of all Nebraskan counties have been in the position to consistently grow their irrigation capacities over the last 80 years with certain irrigating about 70% of their total land surface in 2012, some counties in

[4]See `http://www.hpwd.org/aquifers/`.

[5]About two thirds of all water in the Ogallala aquifer is physically stored in Nebraska.

[6]This compares to the 31,800km^2 of irrigated land in the Central Valley (8% of total surface in California) and the 19,400km^2 of irrigated land around the Mississippi river in Arkansas (14% of total surface in Arkansas). In terms of total water use for irrigation, Nebraska ranks only 7$^{\text{th}}$ (behind California, Idaho, Arkansas, Montana, Colorado and Wyoming). In the present paper, we exclusively focus on the measure of irrigation in terms of irrigated land area.

northwestern Texas reached even higher proportions already in 1959 but could not expand any further due to the geological constraints imposed by the limited water recharge upstream. Over the last 60 years, high-irrigating counties in this region subsequently experienced a progressive decline in their irrigated land proportion. The opposite irrigation trajectories followed by, on the one hand, Nebraska (increasing its irrigation coverage) and, on the other hand, northwestern Texas (decreasing irrigation coverage) – due to their locations at opposite ends of the Ogallala aquifer – illustrate the critical importance of a sustainable water use management in agreement with the constraints imposed by hydrology as well as geology, and also raise questions about the long-run viability of irrigation practices as currently practiced in the Great Plains in general (Harding and Snyder 2012).

From a climate perspective, Nebraska and South Dakota are of particular interest as these states belong to the few regions in the US (besides some areas in Iowa) which did not experience any warming in their summertime temperatures over the last few decades (Schlenker 2020). In fact, coupled general circulation models typically fail to replicate the so-called "warming hole" actually observed as of the 1950's in summertime temperature trends over the Central US (Kunkel, Liang, Zhu and Lin 2006), allegedly due to the phenomenon's relatively fine scale requiring regional circulation-precipitation models (Pan, Arritt, Takle, Gutowski, Anderson and Segal 2004).[7] While Pan et al. (2004) conceive that "[t]he observed cooling may be partly attributable to irrigation on local scales", they eventually largely dismiss the role of irrigation in explaining the overall warming hole phenomenon despite working with a regional model. Their main argument is the mismatch in spatial scales between the relatively small total surface of irrigated land in the Great Plains and the relatively large extent of the observed warming hole. In fact, climatologists typically relate the warming hole in the Central US to the interaction of decadal oscillations in sea surface temperature in the Pacific and Atlantic oceans (Kunkel et al. 2006, Wang, Schubert, Suarez, Chen, Hoerling, Kumar and Pegion 2009, Robinson, Ruedy and Hansen 2002) without considering the potentially additional local influence of irrigation. Robinson et al. (2002) find that warmer sea surface temperatures over the tropical Pacific decrease temperatures in the Central US by impeding insulation through increased cloud cover and precipitation. Weaver (2013) further examine the role of the Great Plains low-level jet to induce additional precipitation and subsequent cooling of surface temperatures in this region specifically in the summer, without consideration for potential fine-scale impacts via irrigation practices.

Acknowledging the above literature, Mueller et al. (2016) nevertheless envisage that agri-

[7]Without relying on any irrigation information, simulations by Pan et al. (2004) of maximum summertime temperature in the 2040's predict a cool spot over a region in southeastern Nebraska which roughly corresponds to the nowadays most densely irrigated area in the US.

cultural practices during the growing season may actually drive some degree of summertime cooling, especially with respect to hot temperature extremes in the Midwest.[8] While their preferred explanatory mechanism involves cropland intensification, they also find – specifically with respect to Nebraska, Arkansas and some parts in the western US – that greater irrigation trends are statistically significantly associated with some decline in *extreme* temperatures. In the present paper, which explores a similar question but in a causal framework and specifically in an irrigation-intensive region around the Ogallala, we detect a significant cooling-by-irrigation externality on the entire temperature distribution (and not only on the hot extremes) and mainly for downwind areas (and not only within-county effects). Therefore, while irrigation is unlikely to predominantly contribute to the long-run and large-scale cooling dynamics observed over the Central US as of the 1950's and which scientists attribute to hemispheric-scale natural forces shaping regional climate, we argue it may nevertheless leave some idiosyncratic, transient cooling signature in local weather patterns.

In fact, as illustrated in Figure 2, the observed 60-year-old evolution of average temperature changes[9] in the 99^{th}, 50^{th} and 5^{th} temperature percentiles in counties around the Ogallala already suggests the existence of a cooling phenomenon where irrigated land expanded in upwind counties (e.g., upstream of the Ogallala aquifer in eastern Nebraska) and, conversely, of a warming phenomenon where irrigated land declined in upwind counties (e.g., downstream of the Ogallala aquifer in northwestern Texas). Every bubble in Figure 2 corresponds to one of the 393 counties in the region of interest, with the size being proportional to the average irrigated land proportion as observed in the upwind county and the color[10] separating counties whose upwind neighbors have been intensifying (blue), maintaining (white) or curbing (red) irrigation.[11] Strikingly, counties with upwind irrigation in decline have been, on average, warming the most in their 99^{th} and 50^{th} temperature percentiles during summer months, while those with upwind irrigation in increase have seen the highest degree of cooling. In other words, there is a vertical stratification of the warming/cooling dynamics according to upwind irrigation practices that can only be seen during summer months, when

[8]Over their chosen time period, 1910-2014, Mueller et al. (2016) find evidence for some cooling of the 95^{th} percentile in the quantile regression in trends for daily temperature maxima. By contrast, they consider that "Midwest cooling is less evident" in the 50^{th} percentile and that there is some warming at the 5^{th} percentile of the quantile regression in trends for daily temperature maxima.

[9]Average changes in the 99^{th}, 50^{th} and 5^{st} temperature percentiles are calculated as the estimated trend in monthly county-level observations of the respective temperature percentile over the 1959-2019 period (in °C/y) multiplied by the duration of the time window (i.e., 61 years).

[10]The color code in Figure 2 is a binned version of the one used in Figure 1 with thresholds at -0.1%/y and +0.1%/y of additional county area being irrigated per year.

[11]Note that most (upwind) counties in the region of interest hardly irrigate at all, and only see a very limited change in their irrigated land proportion over time (small bubbles are typically white). Conversely, (upwind) counties that practice large-scale irrigation have been either on a steadily increasing or decreasing trend (large bubbles are either blue or red).

irrigation is actually taking place, but not during the remainder of the year. Such stratification appears to be mostly pronounced with respect to hotter temperature percentiles and almost absent from the evolution of the 5$^{\text{th}}$ temperature percentile. Figure 2 therefore suggests[12] the existence of a fine-scale mechanism – likely involving direct evaporation of irrigated water as well as stimulated transpiration by plants (Lobell, Bonfils, Kueppers and Snyder 2008, Mahmood, Foster, Keeling, Hubbard, Carlson and Leeper 2006, Harding and Snyder 2012), that is not only localized in space (visible at the county-level resolution) but also in time (as farmers only irrigate cropland during specific months of the year).

3 Data

3.1 Temperature Data

Temperature records are collected from an extended version of the fine-scaled dataset used by Schlenker and Roberts (2009), which combines PRISM data with information gathered by a fixed set of weather stations. The raw data correspond to daily minima and maxima in temperature over a 4km-by-4km grid covering the region of interest from January 1$^{\text{st}}$, 1959 to December 31$^{\text{st}}$, 2019. In the region of interest, the uptake of large-scale irrigation occurred as of the 1940's following the invention of center pivots and the introduction of motorized groundwater pumps. However, the full coverage of the region by meteorological stations is only complete by the end of the 1950's (see Figures A3, A2 and A7).[13]

For each grid cell and for each day within that period, we interpolate a (sinusoidal) temperature profile following Snyder (1985) to calculate the amount of time spent at each possible 1°C-wide temperature interval (e.g., amount of time in a given day and for a given grid cell for which the temperature falls between 19°C and 20°C, between 20°C and 21°C, etc.). These daily counts are then aggregated to a county-by-month resolution to construct specific distributions in temperature for each of the 393 counties around the Ogallala aquifer and for each of the 732 months between 1959 and 2019.[14] Specifically, the time spent at each

[12] Since (upwind) counties with the strongest decline/increase in irrigation are mostly in Texas/Nebraska (see Figure 1) and since Texas is likely to have warmed more than Nebraska irrespective of irrigation practices, this preliminary observation of a correlation in the raw data is only suggestive of the cooling-by-irrigation effect and needs to be complemented as explained in Section 4. We also note, however, that virtually none of the control counties (in white) have been heating/cooling in a similar fashion despite being distributed under all possible latitudes in the region of study – an important detail which motivates the existence of an irrigation-specific effect irrespective of the previous reservation.

[13] On top of this, 1959 is the first census year for which some irrigation information is available in every county of the region of interest (see below as well as Figures A4 and A5).

[14] While previous studies average the weather data by the cropland area in each PRISM grid, our main interest is the effect of irrigation on weather throughout a county and we hence do not weight grids.

1°C interval allows us to construct the cumulative density function (cdf) for each month. In order to detect potentially heterogeneous effects at different temperature levels, we extract from each of these individual county-month specific temperature distributions (cdfs) a set of temperature percentiles (namely, from warmest to coolest: the 99$^{\text{th}}$, 95$^{\text{th}}$, 90$^{\text{th}}$, 75$^{\text{th}}$, 50$^{\text{th}}$, 25$^{\text{th}}$, 10$^{\text{th}}$, 5$^{\text{th}}$ and 1$^{\text{st}}$ temperature percentiles). Eventually, for each county and month of the year, we estimate the long-run warming or cooling behavior by computing time trends[15] in each of our chosen temperature percentiles p over the period of interest. Formally, for any given county i (e.g., Lincoln) and month m (e.g., April), we estimate the slope coefficients $\hat{\theta}^p_{im}$ for each of the $i \times m$ independent regressions:

$$T^p_{imy} = \theta^p_{imo} + \theta^p_{im} y + e^p_{imy} \tag{1}$$

where T^p_{imy} is the temperature level at some chosen percentile p, in county i during month m of year y. θ^p_{imo} is the county-month-percentile specific intercept, θ^p_{im} measures the county-month-percentile specific warming (or cooling) trend and e^p_{imy} is the county-month-percentile specific error term. Figure 2 illustrates the distributions of $\hat{\theta}^{99}_{\cdot m}$, $\hat{\theta}^{50}_{\cdot m}$ and $\hat{\theta}^{5}_{\cdot m}$ (up to a factor 2019-1959+1=61) for each month m of the year and Figure A1 shows maps of $\hat{\theta}^{99}_{\cdot April}$, $\hat{\theta}^{99}_{\cdot July}$, $\hat{\theta}^{50}_{\cdot April}$ and $\hat{\theta}^{50}_{\cdot July}$.

3.2 Irrigation Data

We extract county-level information about irrigated acres from each of the 19 USDA censuses which occurred between 1900 and 2012[16], before normalizing by county area. The widespread adoption of large-scale irrigation has only been possible after World War II and the introduction of both pivot irrigation and motorized groundwater pumps (Hornbeck and Keskin 2014) – with only a few documented exceptions in Colorado and in Scotts Bluff county, Western Nebraska.[17] Figures A4 and A5 show, for each of the eight states in the

[15]Even in the long-run, the pairwise comparison of temperature percentiles for any given month and county remains highly sensitive to short-run weather shocks in the arbitrarily chosen start- and end-year because, in practice, the standard deviation of temperature shocks from one year to another remains substantial even in front of long-run differences. For our purpose, long-run comparisons are thus inappropriate to quantify the amount of warming or cooling as these may artificially compensate or exacerbate the true long-run evolution. An alternative to our proposed approach in trends could be the long-run comparison in temperature measures averaged over a short set of several start- and end-years. However, such approach remains subject to the choice of the bandwidth of years to average over.

[16]Agricultural censuses took place in 1900, 1910, 1920, 1930, 1940, 1950, 1954, 1959, 1964, 1969, 1974, 1978, 1982, 1987, 1992, 1997, 2002, 2007 and 2012, years which are marked by vertical lines in Figures A4 and A5

[17]Early irrigation experiments have been conducted for instance in Scotts Bluff as early as 1890 by using an atypical network of canals and aeromotor windmills for the irrigation and production of sugar beets. See https://www.nps.gov/nr/travel/scotts_bluff/essay_agriculture.html

region of study, the county-by-county evolution of irrigated land proportion as well as the number of observed counties for each available census year. 1959 is the first year in the data for which (at least) some irrigation information is available in every county of the region of interest.

As the opportunity to irrigate cropland primarily depends on the proximity to a river or on the availability of water in an underlying aquifer, large-scale irrigation remains spatially restricted to specific counties, with record levels only found in Nebraska or Texas – although at different points in time as explained above. In fact, the irrigated land proportion has been steadily increasing in most of Nebraskan counties since 1959: by 2012, half of all counties in Nebraska irrigate more than 15% of their land, with some counties in eastern Nebraska reaching US-wide records of approximately 70%. However, such evolution is in stark contrast with the situation in (northwestern) Texas, where multiple counties achieved their peak irrigation levels (up to 75%) in 1959 and have observed a steadily declining trajectory since then. Several counties in Kansas have been irrigating more and more land since 1959 but in moderate proportions (in any event below 25%) when compared to Nebraska, and the few counties in Colorado that were already irrigating in 1959 have been maintaining their moderate irrigation levels throughout the period. Most counties in Colorado, however, never adopted large-scale irrigation, similarly to counties in New Mexico, South Dakota and Wyoming. In light of the above, it appears that the evolution of irrigated proportions displays significant serial as well as spatial correlation.

In line with our earlier approach adopted towards the pre-processing of temperature data, we estimate long-run time trends in irrigation as the slope coefficients $\hat{\iota}_i$ for the i independent regressions:

$$I_{iy} = \iota_{io} + \iota_i y + e_{iy} \tag{2}$$

where I_{iy} measures the proportion of irrigated area in county i in census year y after 1959. The quantity ι_{io} is the county-specific intercept, while the coefficient ι_i measures the county-specific trend in irrigation and e_{iy} is the county-specific error term. Figure 1 shows the resulting map of $\hat{\iota}_.$ in the region of interest.

3.3 Controls

Several factors may be correlated with local trends in both temperature and irrigation levels. For instance, more precipitation is associated with higher temperature and lower need for irrigation, elevation is associated with cooler temperatures and more difficult irrigation conditions, thickness of the sedimentary deposit correlates with groundwater potential and may influence temperature via the influence of soil moisture. If they are omitted from the analysis, such variables could bias the naively estimated effect of irrigation on temperature. Therefore, all of our specifications control for observed county-level measures of precipitation trends, elevation, soil depth, water capacity, clay content, permeability, soil erodibility, proportion of high-quality top soil, proportion of cropland area and vegetation index in April.

Precipitation data come from the same source as our temperature variables (see Section 3.1). The raw data consist in daily observations of total precipitation levels on a 4km-by-4km grid over the region of interest, which we adjust to monthly totals at the county level. Similarly to the pre-processing of temperature and irrigation data, our specifications rely on long-run changes in precipitation as measured by the slope coefficients ρ_{im} in the $i \times m$ independent regressions:

$$R_{imy} = \rho_{imo} + \rho_{im} y + e_{imy} \tag{3}$$

where R_{imy} is the total rainfall (precipitation) in county i during month m in year y. The quantity ρ_{imo} measures the county-month specific intercept, while the coefficient ρ_{im} is the county-month specific trend in rainfall and e_{imy} is the county-month specific error term. Figure A6 illustrates the resulting maps for $\hat{\rho}_{\cdot April}$ and $\hat{\rho}_{\cdot July}$.

Each of the remaining observed control variables are assumed to be time-invariant spatial characteristics and do not require any processing beyond adjustment to the county-level spatial resolution. Elevation raster data comes from the USGS[18] from which we extract median elevation levels for each county. Average soil depth measures the thickness of the sedimentary deposit and originates from a gridded dataset by Pelletier, Broxton, Hazenberg, Zeng, Troch, Niu, Williams, Brunke and Gochis (2016). As explained in Taylor (2021), regions with deeper sedimentary deposit demonstrate higher groundwater potential for irrigation. Average water capacity, percent clay content, minimum permeability, erodibility factor, proportion of high-quality top soil correspond to the soil characteristics used in Schlenker, Hanemann and Fisher (2006). We work with the same cropland area data as in Schlenker and Roberts (2009). Eventually, Enhanced Vegetation Index (EVI) data for the months of April are extracted over Nebraska from NASA's MODIS product MOD13Q1v006 over the 2000-2021 period. We average the 44 bi-monthly measures to obtain county-level mean April EVI in order to control for county-specific characteristics in vegetation landcover.[19] Individual maps for each of the above proposed time-invariant control variables are provided in Figure A8.

[18] https://www.usgs.gov/u.s.-board-on-geographic-names/download-gnis-data

[19] We do not allow the EVI control variable to change by month because, as of May, such time-varying vegetation index becomes largely collinear with irrigation features and would constitute an inappropriate (post-treatment) control.

3.4 Wind data

In order to explore the extent to which irrigation induces (cooling) spillover effects on neighboring counties, we exploit additional information about dominant wind directions. Since the dependent variable measures drybulb air temperature, any potential cooling effect of irrigation should only propagate downwind. Verifying the absence of any externality on counties located crosswind or upwind may thus serve as a falsification check to validate our specifications.

For this purpose, we use hourly measures of latitudinal and longitudinal wind speeds between 1979 and 2019 on a 0.125° grid from the North American Land Data Assimilation System (NLDAS). From the orthogonal wind speed components averaged at county level, we deduce the hourly dominant direction of wind flow and, for each county in the region of interest and month of the year, we eventually select among all adjacent counties – even if located *outside* the region of interest – those which are most often located up-, cross- and downwind. Importantly, we do not construct the average wind direction but rather choose the county that is *most frequently* upwind as "upwind" county, etc. Wind direction often reverses between daytime and nighttime, so an *average* wind direction would be less meaningful than selecting the county that is most often associated with a chosen direction. We assume that wind speed observations made between 1979 and 2019 are representative of overall wind flow conditions for the entire period between 1959 to 2019, an assumption we can only partially validate by restricting the data to smaller subsets and verifying this does not materially influence the classification into up-, cross- and downwind counties.[20]

Figure A9 illustrates the monthly distributions of hourly wind directions aggregated across all counties in the area of interest and over the entire 1979-2019 period. We note that wind is mostly blowing from South to North during summer months. This observation, which holds true at an aggregated level, is also individually valid for most counties in the region of interest – a feature which facilitates our identification strategy (see Section 4).

4 Empirical Strategy

We estimate the effect of irrigation on local temperature in the region of interest around the Ogallala by combining a "continuous" difference-in-differences strategy[21] with spatial first differences (Druckenmiller and Hsiang 2018) in the cross-sectional data in trends. While the

[20]For instance, our classification is hardly affected when considering a much smaller proportion of the wind dataset (e.g., 2007-2019, representing the most recent third).

[21]While one of our time dimensions is binary (involving the comparisons between any individual month and April), the dimension measuring irrigation is *continuous*.

stand-alone difference-in-differences approach already allows to control for some unobservable features, we argue it may still suffer from omitted variable bias due to its flexibility in the comparison of counties irrespective of their relative positioning. In other words, the application of a difference-in-differences strategy to assess county-level temperature outcomes is perfectible because it fails to leverage any information about the spatial disposition of counties, thereby requiring some rather strong reliance on the parallel trends assumption. By contrast, the source of identification in spatial first differences exclusively originates from contrasts between adjacent counties. As a consequence, we should expect spatial first differences to more finely control for unobservable features which vary relatively smoothly in space over distances as long as typical county dimensions in the region of interest. Nevertheless, in order to facilitate the comparison, we also provide results in Appendix for the stand-alone difference-in-differences strategy and, in Section 5, compare its performance to our preferred specification augmented in spatial first differences. Moreover, we show in Appendix that, for our data, the stand-alone difference-in-differences strategy based on *cross-sectional data in trends* provides highly similar results to those obtained via a fixed effect triple-difference model based on *panel data in levels* where we linearly interpolate the irrigation quantity between Census years. Since most of the variation is in the trend over time, we prefer the trend specification.

In a continuous difference-in-differences setting, we compute the effect of irrigation on temperature for each month of the year by interacting month-level dummies with irrigation trends observed in current, up-, cross- and downwind counties. We fix April as a reference month for comparison, since cropland in the region of interest is not yet irrigated during this early season month. The specification with controls for precipitation and time-invariant soil characteristics is given by[22]

$$\theta^p_{im} = \alpha^p + \sum_{\substack{n=1 \\ n\neq 4}}^{12} \boldsymbol{\beta}^p_n \boldsymbol{\iota}_{in} \mathbb{1}_{\{n=m\}} + \sum_{\substack{n=1 \\ n\neq 4}}^{12} \gamma^p_n \mathbb{1}_{\{n=m\}} + \boldsymbol{\delta}^p \boldsymbol{\iota}_{im} + \zeta^p \rho_{im} + \boldsymbol{\eta}^p \boldsymbol{S}_i + \epsilon^p_{im} \tag{4}$$

where θ^p_{im} represents the yearly trend over 1959-2019 in any chosen temperature percentile p during month m as observed in county i, the vector $\boldsymbol{\iota}_{im} = (\iota_i, \iota_{i_{up}(m)}, \iota_{i_{cross}(m)}, \iota_{i_{down}(m)})'$ measures the yearly trends[23] over the same period of time in the proportions of irrigated area in county i as well as in its (month-specific) up-, cross- and downwind neighbors ($i_{up}(m)$,

[22] Note that, without the inclusion of controls ρ_{im} and $\boldsymbol{S}_i$, the coefficients $\hat{\boldsymbol{\beta}}^p_n$ can be equivalently estimated via the fixed effects specification $\theta^p_{im} = \sum_{\substack{n=1 \\ n\neq 4}}^{12} \boldsymbol{\beta}^p_n \boldsymbol{\iota}_{in} \mathbb{1}_{\{n=m\}} + \boldsymbol{\delta}^p \boldsymbol{\iota}_{im} + \lambda^p_m + \mu^p_i + \varepsilon^p_{im}$, in which the month fixed effects λ^p_m absorb any potential month-specific but spatially invariant confounders and the county fixed effects μ^p_i control for any potential county-specific but time-invariant confounders such as any of our soil characteristics.

[23] An alternative to the proposed model adapted to the *cross-sectional data in trends*, would be to interpolate irrigation levels between census years and exploit a fixed effect model for the *panel data in levels* such as

$i_{cross}(m)$ and $i_{down}(m)$ respectively). $\mathbb{1}_{\{n=m\}}$ is a dummy variable equal to 1 in month m and zero otherwise, ρ_{im} is the yearly trend in precipitation in county i and in month m, $\boldsymbol{S}_i$ is a vector of time-invariant soil characteristics for county i, namely water capacity, percent clay content, minimum permeability, erodibility factor, proportion of high-quality top soil, mean soil depth, median elevation above sea level, proportion of cropland surface and average April EVI level, and ϵ^p_{im} is the error term.

The coefficients of interest (i.e., individual elements of the row vectors $\boldsymbol{\beta}^p_n$) measure the difference, between April and month n, in the differences, between counties differing by a unit increase in irrigation trends, in some chosen temperature percentile trend while controlling for precipitation trend and the chosen time-invariant soil characteristics. By construction, and similarly to a distributed lag model in time series analysis, such difference-in-differences coefficients are estimated *jointly* for current, up-, cross- and downwind counties so that spillover effects in various directions are simultaneously estimated, net from each other and net from the same-county effect. Given that both temperature percentiles and irrigation are expressed in *yearly* trends, we may as well interpret the elements of $\boldsymbol{\beta}^p_n$ as the difference, between April and month n, in the differences, between non- and fully-irrigated counties, in temperature percentile p. Importantly, identification originates from the joint comparison along two orthogonal dimensions in time (across months) and space (across counties with contrasting trends in irrigation) but without appreciation of geography. Any simple comparison of temperature percentiles between April and month n in a given irrigated county would be inappropriate as observed differences may in fact result from changes unrelated to irrigation and that occur in *all* counties within the region of interest. Similarly, any simple comparison between irrigated and non-irrigated counties in any given month n would fail to account for the fact that irrigated counties are typically already hotter in the spring and before irrigation actually takes place. The underlying identification assumption is that, in the absence of irrigation, an actually fully irrigated county *would* have experienced a parallel warming (or cooling) between April and month n when compared to a non-irrigated county – irrespective of their proximity within the region of interest.

In practice, we are only interested in estimating effects in May, June, July and August, as intensive irrigation can be reasonably expected to take place only during these few months of the year. By contrast, estimated coefficients for the remainder of the year, i.e., when farmers do not irrigate, should ideally be estimated close to zero, or at least be statistically

---

$T^p_{imy} = \sum_{\substack{n=1 \\ n\neq 4}}^{12} \boldsymbol{\beta}^p_n \boldsymbol{I}_{iny} \mathbb{1}_{\{n=m\}} + \pi^p_{im} + \sigma^p_{my} + \tau^p_{iy} + \varepsilon^p_{imy}$, where $\boldsymbol{I}_{imy} = (I_{iy}, I_{i_{up}(m)y}, I_{i_{cross}(m)y}, I_{i_{down}(m)y})'$ and π^p_{im}, σ^p_{my} and τ^p_{iy} respectively denote county-month, month-year and county-year fixed effects. As a matter of fact and for the purpose of modelling our data, this panel model specification in levels yields highly similar estimations for the coefficients of interest (compare Figures A10 and A11).

indistinguishable from zero. Similarly, we are only interested in measuring the cooling effect in some irrigating (current) county as well as the cooling externality coming from its irrigating neighbor(s) located upwind. In fact, since our dependent variable relates to drybulb air temperature, we should expect the spillover effects coming from cross- and downwind counties to be close to zero for each month of the year. This set of intuitive predictions with respect to winter months and to the expected direction of spillover effects from irrigating neighbors located cross- and downwind allow to empirically appreciate and assess the performance of the estimation strategy in continuous difference-in-differences in contrast to our preferred strategy in spatial first differences.

In order to account for the limitation of continuous difference-in-differences in accounting for the spatial disposition of counties, we augment the previous specification by exploiting spatial first differences as introduced by Druckenmiller and Hsiang (2018). In this alternative setting, identification of the empirical effect of irrigation adoption on long-run trends in temperature percentiles for each month of the year derives from contrasts between adjacent counties. As it relies on the confrontation of differenced county-level features between each possible pair of neighboring counties, spatial first differences allow to control for all – both observed and unobserved – spatially-invariant features that vary sufficiently smoothly across adjacent pairs of counties. Theoretically, this should reduce the risk of omitted variable bias since adjacent counties can be expected to be most similar to one another in terms of both observed and unobserved characteristics. Formally, let Δ_{ij} denote the spatial first difference operator between adjacent counties i and j such that, for any observed, spatially-dependent variable of interest $y_{\cdot} \in \{\theta^p_{\cdot m}, \iota_{\cdot}, \rho_{\cdot m}\}$, $\Delta_{ij} y$ measures the difference $y_j - y_i$ and, for any observed vector $\boldsymbol{v}_{\cdot} \in \{\boldsymbol{\iota}_{\cdot m}, \boldsymbol{S}_{\cdot}\}$, $\Delta_{ij}\boldsymbol{v}$ refers to the vector of element-wise application of the spatial first differences operator. After application of the spatial first difference operator, specification (4) translates into:

$$\Delta_{ij}\theta^p_m = \sum_{\substack{n=1 \\ n\neq 4}}^{12} \boldsymbol{\beta}^p_n \Delta_{ij}\boldsymbol{\iota}_n \mathbb{1}_{\{n=m\}} + \boldsymbol{\delta}^p \Delta_{ij}\boldsymbol{\iota}_m + \zeta^p \Delta_{ij}\rho_m + \boldsymbol{\eta}^p \Delta_{ij}\boldsymbol{S} + \Delta_{ij}\epsilon^p_m \tag{5}$$

because $\Delta_{ij}\alpha^p = 0$ and, for any pair of months m and n, $\Delta_{ij}\mathbb{1}_{\{n=m\}} = 0$. Following Druckenmiller and Hsiang (2018) and Taylor and Schlenker (2021), we fit a more flexible[24] model including an intercept κ^p_0, on top of which we further include a set of month-level dummies κ^p_n to obtain:

$$\Delta_{ij}\theta^p_m = \kappa^p_0 + \sum_{\substack{n=1 \\ n\neq 4}}^{12} \boldsymbol{\beta}^p_n \Delta_{ij}\boldsymbol{\iota}_n \mathbb{1}_{\{n=m\}} + \sum_{\substack{n=1 \\ n\neq 4}}^{12} \kappa^p_n \mathbb{1}_{\{n=m\}} + \boldsymbol{\delta}^p \Delta_{ij}\boldsymbol{\iota}_m + \zeta^p \Delta_{ij}\rho_m + \boldsymbol{\eta}^p \Delta_{ij}\boldsymbol{S} + \epsilon^p_m \tag{6}$$

[24]In practice, this choice is without consequences as the two models (5) and (6) yield virtually indistinguishable coefficients of interest $\hat{\beta}^p_n$ and *de minimis* values for the auxiliary coefficients $\hat{\kappa}^p_n$.

The coefficients of interest (i.e., individual elements of the row vector $\boldsymbol{\beta}_n^p$) in this alternative specification can be directly compared to the coefficients estimated earlier under the stand-alone difference-in-differences approach, although its source of identification is now restricted to the pairwise comparisons between adjacent counties. As mentioned above, we argue that these revised coefficients are less prone to omitted variable bias, a claim which we motivate – at least for our data – in the following section which describes our results.

5 Results

For each of the 99$^{\text{th}}$, 95$^{\text{th}}$, 90$^{\text{th}}$, 75$^{\text{th}}$, 50$^{\text{th}}$, 25$^{\text{th}}$, 10$^{\text{th}}$, 5$^{\text{th}}$ and 1$^{\text{st}}$ temperature percentiles, our preferred difference-in-differences specification in spatial first differences (6) consistently suggests the existence – exclusively in the summer (i.e., when farmers actually irrigate) and with a peak effect centered around July – of some highly statistically significant *cooling* externality offered by irrigation practiced in upwind neighbors. Results are graphically summarized for the 99$^{\text{th}}$, 90$^{\text{th}}$, 75$^{\text{th}}$, 50$^{\text{th}}$, 25$^{\text{th}}$ and 5$^{\text{th}}$ temperature percentiles in Figure 3. Although it is empirically detected for the different percentiles throughout the temperature distribution, such cooling externality appears to be most pronounced with respect to warmer temperatures and to progressively decrease in magnitude with cooler temperatures. Quantitatively, the fact to fully irrigate an upwind county is estimated to translate, in July and relatively to April, into a -0.86°C drop in the highest temperature percentile and only into a -.37°C decrease in the lowest, with intermediate values for the other percentiles (e.g., -.52°C for median temperatures). Surprisingly perhaps, same-county irrigation is not detected to cause any cooling in the upper half of the temperature distribution, but only in the lower half of it. In fact, as the influence of upwind irrigation decreases with lower temperature percentiles, the impact of same-county irrigation becomes progressively more sizeable and more significant. Given that most of the economic implications involving temperature effects in the summer relate to extreme heat, the significant externality coming from upwind irrigation is of greatest interest and will be discussed in more details in section 6 with respect to the example of corn production.

In order to account for spatial correlation in the data, we compute Conley standard errors with uniform kernel and 0.5° bandwidth for both models. The fact that no statistically significant effect is detected outside of the irrigation season with respect to irrigation within the same or upwind county constitutes strong validation of our first falsification test for the model specification as explained in Section 4. A second falsification check consists in the absence of statistically significant effects detected during summer months with respect to irrigation practiced in cross- and downwind counties. This second prediction is also largely validated in overall for the model specification (6) in spatial first differences.

Although the continuous difference-in-differences specification (4) (as well as the fixed effects specification described in footnote 23, for which results are highly similar) also detects a similar relationship with largest cooling effects for upwind irrigation levels on the higher temperature percentiles (see Figures A10 and A11 respectively), it does not resist to the proposed falsification checks as satisfactorily as specification (6). For instance, a large and statistically significant cooling effect is typically detected for October temperatures with respect to upwind (summer) irrigation status – a characteristic which does not result from the previous specification in spatial first differences and which does not *a priori* have any meaningful interpretation. Worse, specification (4) predicts a statistically significant *warming* effect due to same-county irrigation in the highest temperature percentiles as well as several statistically significant coefficients with respect to cross- and downwind irrigation levels including during summer months. Eventually, we observe that the estimated cooling externality via continuous difference-in-differences is *much* larger than previously estimated via spatial first differences. In other words, our preferred specification gives conservative estimates of the cooling effect, which would be even larger under alternative specifications. However, given that specification (4) does not incorporate information about the counties' relative positioning, it allows for comparisons between counties that are potentially far apart from each other and for which the parallel trends assumption is less likely to hold. As a consequence, this specification runs more risk to suffer from omitted variable bias since unobserved features may be sub-optimally controlled for under these conditions.

Despite generally validating the falsification tests with respect to cross- and downwind irrigation levels during summer months, our preferred specification in spatial first differences seems to spuriously detect some "warming" effect for a few temperature percentiles during winter months. In fact, these apparent "warming" effects are largely due to the *ex ante* choice of April as a reference month. In order to verify this hypothesis, we re-run the same specification but after binning individual month-dummies into wider seasons, namely (i) a control season (running from September to April), (ii) an early-irrigation season (May and June) and (iii) a peak irrigation season (July and August). Numerical results are shown in Table 1 for each of the 99^{th}, 90^{th}, 75^{th}, 50^{th}, 25^{th}, 10^{th} and 1^{st} temperature percentiles.

Under this alternative approach by season (and no longer by individual months), for which the control season is more likely to be representative of weather conditions during the non-irrigation period, we observe that each statistically significant coefficient involving irrigation reflects a *cooling* effect as intuitively expected. The previous observations pertaining to same-county and upwind county irrigation effects remain largely valid. In particular, upwind irrigation is consistently responsible for a significant cooling of summer temperatures, throughout the temperature distribution, with largest effects in July and August and especially with respect to hotter temperatures. Same-county irrigation, by contrast, remains smaller in magnitude – except for cooler temperatures as previously observed. Importantly, however, the seasonal specification also suggests the existence of a cooling externality originating from cross- and downwind irrigation, exclusively during the peak irrigation season and until the 75$^{\text{th}}$ temperature percentile. We observe, however, that these coefficients are much smaller in magnitude and much less precisely estimated than the main effect detected for upwind irrigation. Also, we argue that these apparent spillover effects propagating orthogonally and against the dominant wind flow may reflect the limitation of our classification into up-, cross- and downwind counties. In fact, such classification is only most frequently true and does not account for the fact that wind flow may largely diverge from its modal direction over the duration of a month. For instance, our classification does not capture bi-modal distributions of wind directions, although 180° shifts are quite common in practice. As such, the cooling externality coming from upwind counties is – for some fraction of time – erroneously classified as coming from cross- or (especially) downwind, a lack of precision which manifests itself the most when the main effect is large. In the remainder, we will therefore make the simplifying assumption that the cooling-by-irrigation externality is restricted to the upwind direction. Also, we argue that this approach is conservative as some of the cooling effect is actually lost towards other (misclassified) wind directions.

In light of the above results, it is legitimate to ask how far the cooling externality actually travels with the wind. In order to explore this question, we augment the previous seasonal model by adding irrigation information from one further layer upwind. Formally, we replace $\boldsymbol{\iota}_{im} = (\iota_i, \iota_{i_{up}(m)}, \iota_{i_{cross}(m)}, \iota_{i_{down}(m)})'$ by $\tilde{\boldsymbol{\iota}}_{im} = (\iota_i, \iota_{i_{up}(m)}, \iota_{i_{2\times up}(m)}, \iota_{i_{cross}(m)}, \iota_{i_{down}(m)})'$ in specification (6). Numerical results are shown in Table A1 and suggest that the cooling-by-irrigation externality does not travel beyond the first layer of upwind neighbors. Indeed, in this revised specification, the cooling externality from upwind irrigation remains significant and sizeable throughout the temperature distribution, while the effect of the 2-step upwind irrigation is not typically found to be statistically significant – except on the 99$^{\text{th}}$ temperature percentile for which it is highly significant although much smaller in size. In the remainder, we will therefore make the simplifying assumption that the cooling-by-irrigation externality is restricted to the direct upwind neighbor.

In order to further test the robustness of our model findings, we run the previous (seasonal) specification but on separate subsets of the initial data. Indeed, as shown on Figure 1, the northern and southern halves of the dataset have seen fundamentally different irrigation dynamics, and we propose to verify whether each of these two observed dynamics yield an identification of about the same cooling-by-irrigation externality. Results for the northern part (including Nebraska, Colorado, Wyoming, South Dakota and Kansas) and the southern part (including Texas, Oklahoma and New Mexico) are displayed in Tables A2 and A3 respectively and confirm the overall conclusions made previously – in particular with respect to the cooling-by-irrigation externality originating predominantly in upwind irrigation areas. Importantly, we note that the magnitudes of the upwind irrigation externality are in rough agreement between the two segments of the data – despite the fact that the northern part consists of counties with a strong increase in irrigation, while the southern part consists of counties with a strong decrease in irrigation. In other words, this confirms the preliminary observation already made on Figure 2 regarding the fact that counties which historically saw a decline in irrigation induced some warming in neighboring areas.

6 Discussion

Large-scale irrigation has historically allowed farmers to grow crops in parts of the Great Plains that would otherwise be locally unsuitable for intensive agriculture. For instance, corn requires a substantial amount of water for healthy growth and is very vulnerable to droughts due to its shallow root system. In the US, large-scale corn production is primarily found in the Corn Belt – where corn is predominantly rainfed with the notable exception of its western tail located in Nebraska and which is heavily irrigated – as well as in specific regions with access to abundant water supply for irrigation such as above the Ogallala aquifer or in the Arkansas delta. In fact, the very disposition of the Ogallala aquifer can be easily recognized from a map of average corn area grown in the US[25].

Besides allowing farmers to adjust water supply to precipitation conditions, irrigation serves as an adaptation strategy against heat (Taylor 2021) as it increases the vegetation's tolerance to high temperatures. While crop yields only marginally benefit from moderate degree-days over the growing season, they severely suffer from particularly hot degree-days (Schlenker and Roberts 2009). A primary consequence of irrigation is the alteration of this non-linear response function linking crop yields to heat exposure so that plants become more resilient to extreme temperatures (the "Direct Effect").[26] On top of this beneficial effect which occurs on the spot, we have shown above that irrigation is also causally responsible for cooling downwind temperatures. This secondary and unintended effect on temperature patterns in turn indirectly benefits plant growth in counties located downwind by cooling

[25] See, for instance, Map 1 in Schlenker (2020)

[26] For instance, when fitting the piecewise linear specification proposed by Schlenker and Roberts (2009) on historic log yields for corn between 1940 and 2019 in the region of interest, we find (i) a steeper slope coefficient for the detrimental effect of heat and (ii) an earlier threshold for such effects to manifest in non-irrigating counties when compared to irrigating counties.

the entire temperature distribution and thereby protecting crops from extreme heat exposure (the "Indirect Effect"). In other words, farmers who practice large-scale irrigation do not only benefit from such investment in the form of higher yields in their own fields but also favor production levels in neighboring fields located downwind. In the remainder, we quantify this positive externality at the county level in the form of avoided corn production losses and compare the magnitudes of the Direct and the Indirect Effects in the region of interest. In order to appreciate each of these effects in isolation, we simulate counterfactual corn production outcomes respectively (i) under actual heat exposure *but for alternative response functions to heat* (i.e., with and without the Direct Effect of irrigation) and (ii) for actual response functions to heat *but under alternative heat exposure* (i.e., with and without the Indirect Effect of irrigation).

In order to quantify the Direct Effect of irrigation on corn production, we compute the individual heat exposure of every county in the region of interest, over the entire growing season for corn (i.e., from March to August) and for each year between 1940 and 2019. Then, we confront these distributions to an empirically determined irrigation-dependent heat response function under actual irrigation levels as well as under a counterfactual scenario without irrigation.[27] In order to model the sensitivity of the heat response function to irrigation, we adapt the piece-wise linear specification in (log) yield proposed by Schlenker and Roberts (2009) by allowing the effect of harmful degree days to vary linearly in a county's irrigation level. Using the same methodology as in Schlenker and Roberts (2009) and the same break-point temperature separating beneficial from harmful degree-days (i.e., 29°C for corn), we find a positive slope of 0.00012 (p-value below 1%) for degree-days below 29°C and a linearly-varying slope of $-0.00452 + 0.00852 \times I_{iy}$ (with respective p-values below 1%) for harmful degree-days. The latter slope suggests that the substitution of a full day at 33°C with a full day at 34°C during the growing season is associated, on average, with an approximate decline of 0.45% in yields for non-irrigated counties, but a decline of only 0.28% for counties irrigated at 20% of their surface. Left panels in Figure 4 illustrate the spatial distribution of avoided corn production losses in absolute levels as well as relatively to total corn production. On aggregate and according to the proposed methodology, the Direct Effect is responsible for the production of 6.3 billion bushels of corn, representing about 7% of total corn production over the 1959-2019 period. Virtually all of these avoided production losses occur within the boundary of the Ogallala aquifer, in agreement with areas where large-scale irrigation is possible. 65% of the total corn production in the region of interest comes from Nebraska, and a similar proportion (67%) of the beneficial effect of irrigation on

[27] In order to simulate counterfactual heat response functions for *each* year since 1940, we linearly interpolate county-level irrigation information between consecutive census years.

crop yields is also found in this state. Nevertheless, the Direct Effect also achieves sizeable proportions (up to 25%) of corn production in counties located in northwestern Texas and western Kansas, where temperatures are higher and total corn production is more limited on average than in eastern Nebraska.

In order to estimate avoided production losses due to the Indirect Effect, we empirically determine the heat response function of corn using the piece-wise linear specification in (log) yield proposed by Schlenker and Roberts (2009), i.e., without consideration for irrigation information: the effect of beneficial degree-days becomes 0.00016 (p-value below 1%) and the effect of harmful degree-days is fixed at -0.00443 (p-value below 1%). The latter value suggests that the substitution of a full day at 33°C with a full day at 34°C is associated, on average, with an approximate decline of 0.44% in yields – which is consistent with the previous model under no-irrigation conditions. This (unique) response function is then separately applied to the actual heat distribution and to a simulated counterfactual scenario from which the cooling-by-irrigation effect from upwind irrigation has been subtracted.[28] More precisely, month-specific counterfactual temperature percentiles are predicted based on the effects obtained from the difference-in-differences specification in spatial first differences (6) and using county-level irrigation levels.[29] A cumulative distribution function of heat exposure is linearly interpolated between the 0.001^{st}, 1^{st}, 5^{th}, 10^{th}, 25^{th}, 50^{th}, 75^{th}, 90^{th}, 95^{th}, 99^{th} and 99.999^{th} temperature percentiles for each individual month of the growing season in order to construct the cumulative distribution function of heat exposure over the entire growing season. Since irrigation is responsible for some cooling, the counterfactual distributions (without irrigation) display a slightly heavier exposure to warmer temperatures, and especially so with respect to the warmest temperature percentiles.

Right panels in Figure 4 illustrate the spatial distribution of avoided corn production losses due to the Indirect Effect in absolute levels as well as relatively to total corn production. On aggregate and according to the proposed methodology, the Indirect Effect is responsible for 440 million bushels of corn, representing about 0.5% of total corn production over the 1959-2019 period. Again, the majority of avoided losses due to the Indirect Effect can be allocated to Nebraska (74%) and almost all of these avoided production losses occur

[28]For simplicity and in order to produce conservative order of magnitude estimates, we neglect to take into account the cooling effect detected for irrigation in current, cross- and downwind counties (see Table 1). First, the cooling effect originating from same-county irrigation is only statistically significant with respect to cooler temperatures, for which there is only *de minimis* impact on corn yields. Second, the cooling effect measured with respect to cross- and downwind irrigation levels is only statistically significant in the July-August time bracket (and not in May-June) and with much smaller coefficients (in absolute value) than for upwind irrigation levels.

[29]In order to simulate counterfactual temperature percentiles for *each* year since 1940, we linearly interpolate county-level irrigation information between consecutive census years.

within the boundary of the Ogallala, with large year-to-year and county-level variability depending on extreme weather conditions. For instance, during extreme heat episodes, avoided losses may constitute up to 5% of total corn production in some selected counties. Figures A12 and A13 show the yearly evolution of avoided production losses in absolute terms and in proportions to total production for the overall region of interest and for Nebraska in particular. Avoided losses have been rising until the 1970s before stabilizing at around .9% on average. While this long-run evolution of avoided losses roughly follows the overall trend in total production, short-run shocks (e.g., during the heat wave of 2012) evolve in opposite directions. Indeed, particularly hot years are associated both with a marked decline in total production and a surge in the amount of avoided crop losses due to the strong non-linearity in the yield response to extreme temperatures. By contrast, years deprived from extreme heat episodes (e.g., in years following the eruption of Mount Pinatubo in 1991) see relatively high yields irrespective of the cooling-by-irrigation effect. Importantly, most of the avoided production losses occur in more recent decades: over the second half of the observation window, avoided losses due to the cooling-by-irrigation effect amount to 10 million bushels per year on average against only 1 million bushels per year for the first half. In 2012, avoided losses culminated at 15 million bushels, a record level due to an extreme heat episode. Although exceptional in the historic data, such situations may in fact become more an more frequent in the context of an overall warming climate.

In light of the above, the Indirect Effect appears to be, on average, one order of magnitude *smaller* than the Direct Effect. However, in contrast to the location of the Direct Effect, the Indirect Effect consists in an *externality* and is typically shifted by one county northwards – in agreement with the fact that wind predominantly blows from South to North in the region of interest during the growing season (see Figure A9). This raises the question as to whether the Indirect Effect occasionally benefits individual counties *more* than the Direct Effect. Figure 5 answers this point by illustrating the ratio of the Indirect Effect to the Direct Effect for Nebraskan counties, which concentrate the bulk of the corn production in the region of interest.[30] In several counties, the externality caused by irrigation in upwind neighbors and measured in the form of avoided corn production losses is as important as (or even more important than) the direct effect of same-county irrigation. Such situations are typically observed in counties located immediately downwind from significant irrigation zones (for instance, north of the Platte River) and that do not significantly irrigate themselves. While these areas do not encompass major corn-producing counties in Nebraska, it is important to stress that the cooling-by-irrigation effect is further responsible for additional positive externalities on the economy via the spillover it generates on local weather (for instance

[30] The similar map but for the entire region of interest is provided in Figure A14.

with respect to the production of other crops, the reduction in energy demand, the protection of populations and/or ecosystems during heat waves, the preservation of labor productivity, etc.), an effect which is not quantified in the previous (conservative) estimation of the positive cooling-by-irrigation externality and which is restricted to corn yields.

7 Conclusions

By exploiting county-level data in temperature and irrigation trends over the last 60 years in a vast region around the Ogallala aquifer, we empirically demonstrate that large-scale irrigation is responsible – mainly downwind – for a significant cooling externality which affects the entire distribution of summer temperatures and is most effective at cooling the higher end of the temperature distribution. Using a continuous difference-in-differences strategy augmented in spatial first differences, we argue that such effect is causal in nature and may participate in explaining the larger-scale warming hole observed in the Central US as of the 1950's. Conceptually, the unintended feedback mechanism which we identify with respect to large-scale irrigation is remarkable because it suggests that adaptive behavior to regional climate (and by extension also to climate change or rising constraints in water availability) can have measurable and sizeable effects on the climate – a channel which, to our knowledge, has not yet been considered in the adaptation literature but certainly deserves more attention.

Pragmatically, from an agricultural perspective, we show that large-scale irrigation does not only benefit plant growth directly by controlling water supply and enhancing heat tolerance of crops but also indirectly by avoiding harmful degree-days in downwind areas. Although being comparatively small on average, this positive externality appears to be occasionally as important as the direct effect of irrigation on corn production with respect to its adaptation to heat. Moreover, since extreme heat is also detrimental to other spheres of the economy besides agriculture, such as labor productivity or local population and ecosystem health in general, we argue that the benefits of the cooling-by-irrigation externality may in fact be much vaster than expected from a purely agricultural perspective, a perspective which opens new areas of research with respect to areas where populations live relatively close to irrigated land.

Despite having improved crop yields in the Great Plains over the last 60 years, large-scale irrigation as historically allowed by the Ogallala aquifer is unlikely to be replicable in other agricultural settings within the US given fundamental geological constraints in freshwater availability. On the contrary, the progressive depletion of the Ogallala aquifer may result not only in profound changes in local agricultural production and productivity, but also in a localized warming of summertime temperatures, a situation which is already empirically observed downstream of the Ogallala aquifer.

References

Albouy, David, Walter Graf, Ryan Kellogg, and Hendrik Wolff, "Climate Amenities, Climate Change, and American Quality of Life," *Journal of the Association of Environmental and Resource Economists*, March 2016, *3* (1), 205–246. Publisher: The University of Chicago Press.

Auffhammer, Maximilian and Anin Aroonruengsawat, "Simulating the impacts of climate change, prices and population on California's residential electricity consumption," *Climatic Change*, December 2011, *109* (1), 191–210.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro, "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century," *Journal of Political Economy*, February 2016, *124* (1), 105–159. Publisher: The University of Chicago Press.

Burke, Marshall and Kyle Emerick, "Adaptation to Climate Change: Evidence from US Agriculture," *American Economic Journal: Applied Economics*, August 2016, *8* (3), 106–140.

____ **, Solomon M. Hsiang, and Edward Miguel**, "Global non-linear effect of temperature on economic production," *Nature*, November 2015, *527* (7577), 235–239. Number: 7577 Publisher: Nature Publishing Group.

Butler, Ethan E. and Peter Huybers, "Adaptation of US maize to temperature variations," *Nature Climate Change*, January 2013, *3* (1), 68–72.

Druckenmiller, Hannah and Solomon Hsiang, "Accounting for Unobservable Heterogeneity in Cross Section Using Spatial First Differences," Technical Report w25177, National Bureau of Economic Research, Cambridge, MA October 2018.

Haqiqi, Iman, Danielle S. Grogan, Thomas W. Hertel, and Wolfram Schlenker, "Quantifying the impacts of compound extremes on agriculture," *Hydrology and Earth Systmen Sciences*, February 2021, *25* (2), 551–564.

Harding, Keith J. and Peter K. Snyder, "Modeling the Atmospheric Response to Irrigation in the Great Plains.: Part I: General Impacts on Precipitation and the Energy Budget," *Journal of Hydrometeorology*, 2012, *13* (6), 1667–1686. Publisher: American Meteorological Society.

Hornbeck, Richard and Pinar Keskin, "The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought," *American Economic Journal: Applied Economics*, January 2014, *6* (1), 190–219.

Kunkel, Kenneth E., Xin-Zhong Liang, Jinhong Zhu, and Yiruo Lin, "Can CGCMs Simulate the Twentieth-Century "Warming Hole" in the Central United States?," *Journal of Climate*, September 2006, *19* (17), 4137–4153.

Lobell, David B., Celine J. Bonfils, Lara M. Kueppers, and Mark A. Snyder, "Irrigation cooling effect on temperature and heat index extremes," *Geophysical Research Letters*, 2008, *35* (9). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2008GL034145.

Mahmood, R, S Foster, T Keeling, K Hubbard, C Carlson, and R Leeper, "Impacts of irrigation on 20th century temperature in the northern Great Plains," *Global and Planetary Change*, November 2006, *54* (1-2), 1–18.

Moscona, Jacob and Karthik A. Sastry, "Does Directed Innovation Mitigate Climate Damage? Evidence from US Agriculture," 2021.

Mueller, Nathaniel D., Ethan E. Butler, Karen A. McKinnon, Andrew Rhines, Martin Tingley, N. Michele Holbrook, and Peter Huybers, "Cooling of US Midwest summer temperature extremes from cropland intensification," *Nature Climate Change*, March 2016, *6* (3), 317–322.

Pan, Zaitao, Raymond W. Arritt, Eugene S. Takle, William J. Gutowski, Christopher J. Anderson, and Moti Segal, "Altered hydrologic feedback in a warming climate introduces a "warming hole": ALTERED FEEDBACK IN A WARMING CLIMATE," *Geophysical Research Letters*, September 2004, *31* (17), n/a–n/a.

Pelletier, Jon D., Patrick D. Broxton, Pieter Hazenberg, Xubin Zeng, Peter A. Troch, Guo-Yue Niu, Zachary Williams, Michael A. Brunke, and David Gochis, "A gridded global data set of soil, intact regolith, and sedimentary deposit thicknesses for regional and global land surface modeling," *Journal of Advances in Modeling Earth Systems*, 2016, *8* (1), 41–65. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/2015MS000526.

Robinson, Walter A., Reto Ruedy, and James E. Hansen, "General circulation model simulations of recent cooling in the east-central United States," *Journal of Geophysical Research: Atmospheres*, 2002, *107* (D24), ACL 4–1–ACL 4–14. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2001JD001577.

Schlenker, Wolfram, "Environmental Drivers of Agricultural Productivity Growth and Socioeconomic Spillovers," 2020, p. 24.

___ and Michael J. Roberts, "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change," *Proceedings of the National Academy of Sciences*, September 2009, *106* (37), 15594–15598.

___, ___, and David B. Lobell, "US maize adaptability," *Nature Climate Change*, August 2013, *3* (8), 690–691.

___, W Michael Hanemann, and Anthony C Fisher, "THE IMPACT OF GLOBAL WARMING ON U.S. AGRICULTURE: AN ECONOMETRIC ANALYSIS OF OPTIMAL GROWING CONDITIONS," *THE REVIEW OF ECONOMICS AND STATISTICS*, 2006, p. 13.

Snyder, R, "Hand calculating degree days," *Agricultural and Forest Meteorology*, October 1985, *35* (1-4), 353–358.

Taylor, Charles A, "Irrigation and Climate Change: Long-run Adaptation and its Externalities," 2021, p. 49.

Taylor, Charles and Wolfram Schlenker, "Environmental Drivers of Agricultural Productivity Growth: CO2 Fertilization of US Field Crops," Technical Report w29320, National Bureau of Economic Research, Cambridge, MA October 2021.

Wang, Hailan, Siegfried Schubert, Max Suarez, Junye Chen, Martin Hoerling, Arun Kumar, and Philip Pegion, "Attribution of the Seasonality and Regionality in Climate Trends over the United States during 1950–2000," *Journal of Climate*, 2009, *22* (10), 2571–2590. Publisher: American Meteorological Society.

Weaver, Scott J., "Factors Associated with Decadal Variability in Great Plains Summertime Surface Temperatures," *Journal of Climate*, January 2013, *26* (1), 343–350.

World Bank Group, "High and Dry: Climate Change, Water, and the Economy," Working Paper, World Bank, Washington, DC May 2016. Accepted: 2016-01-14T22:38:01Z.

Zivin, Joshua Graff and Matthew Neidell, "Temperature and the Allocation of Time: Implications for Climate Change," *Journal of Labor Economics*, January 2014, *32* (1), 1–26. Publisher: The University of Chicago Press.

Tables

Table 1: Estimated effect of irrigation on temperature percentiles ($\hat{\beta}_n^p$) in early (May-June) and peak (July-August) irrigation seasons

p	n	Irrigation information from				R^2	N
		current	upwind	crosswind	downwind		
99th	May-June	0.15	-0.35**	-0.12	-0.02		
		(0.12)	(0.15)	(0.11)	(0.10)		
	July-August	0.02	-0.63***	-0.31**	-0.26**		
		(0.17)	(0.16)	(0.12)	(0.13)	8%	13500
90th	May-June	0.06	-0.33***	-0.11	0.02		
		(0.09)	(0.08)	(0.07)	(0.07)		
	July-August	-0.06	-0.48***	-0.17**	-0.22**		
		(0.13)	(0.12)	(0.07)	(0.09)	8%	13500
75th	May-June	0.01	-0.22***	-0.05	0.06		
		(0.07)	(0.07)	(0.06)	(0.07)		
	July-August	-0.14	-0.47***	-0.15**	-0.20**		
		(0.10)	(0.10)	(0.06)	(0.09)	7%	13500
50th	May-June	-0.10*	-0.25***	-0.04	0.03		
		(0.06)	(0.06)	(0.05)	(0.07)		
	July-August	-0.21**	-0.41***	-0.08	-0.17*		
		(0.09)	(0.09)	(0.05)	(0.09)	6%	13500
25th	May-June	-0.12*	-0.25***	-0.02	-0.01		
		(0.07)	(0.06)	(0.05)	(0.08)		
	July-August	-0.30***	-0.35***	-0.04	-0.10		
		(0.11)	(0.08)	(0.06)	(0.11)	5%	13500
10th	May-June	-0.13**	-0.20***	-0.04	-0.07		
		(0.06)	(0.07)	(0.06)	(0.09)		
	July-August	-0.28***	-0.27***	-0.01	-0.04		
		(0.10)	(0.09)	(0.07)	(0.11)	4%	13500
1st	May-June	-0.11	-0.09	0.06	0.01		
		(0.09)	(0.11)	(0.09)	(0.12)		
	July-August	-0.30***	-0.24**	-0.06	-0.16		
		(0.11)	(0.12)	(0.11)	(0.14)	3%	13500

Notes: coefficients are estimated separately for each of the 99th, 90th, 75th, 50th, 25th, 10th and 1st temperature percentiles by continuous difference-in-differences in spatial first differences with irrigation information from current, upwind, crosswind and downwind counties. Reference season runs from September to April. All specifications include an intercept, a linear term in irrigation trend and controls for precipitation trends and soil characteristics. Units are in °C/(% of irrigated county area). Conley standard errors (given in parentheses) are computed using a uniform kernel and a bandwidth of .5 degrees. ***, ** and * denote statistical significance respectively at the 1%, 5% and 10% levels.

Figures

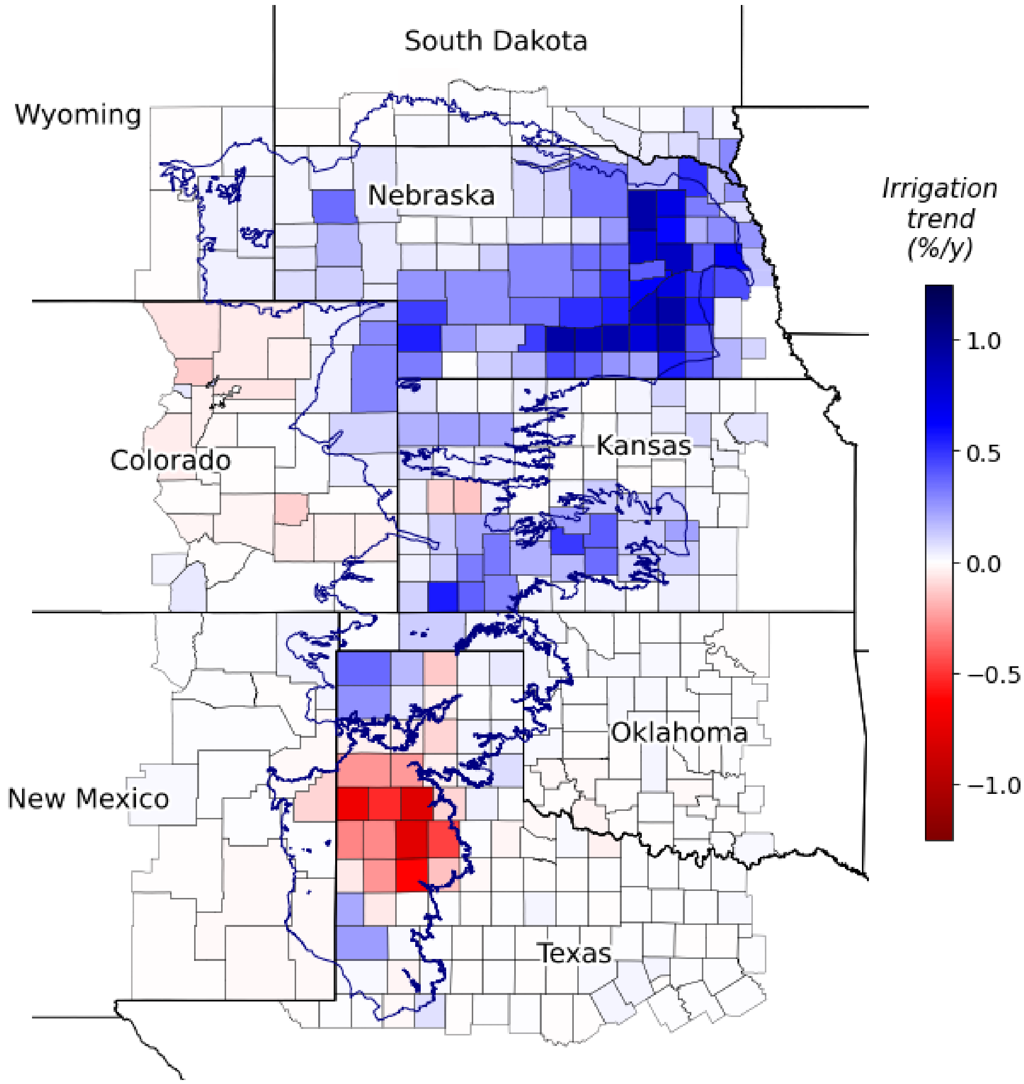

Figure 1: *Extent of the High Plains ("Ogallala") aquifer (blue contour line) covering parts of South Dakota, Wyoming, Nebraska, Colorado, Kansas, Oklahoma, New Mexico and Texas.*

Notes: The 393 counties (i) with centroid coordinates falling between the minimum/maximum latitude and longitude of the Ogallala (or crossing its boundary) constitute the spatial area of interest and are colored according to the irrigation trend ($\hat{\iota}_i$) as observed over the 1959-2019 period: counties which saw an increase/decrease in irrigation are in blue/red respectively.

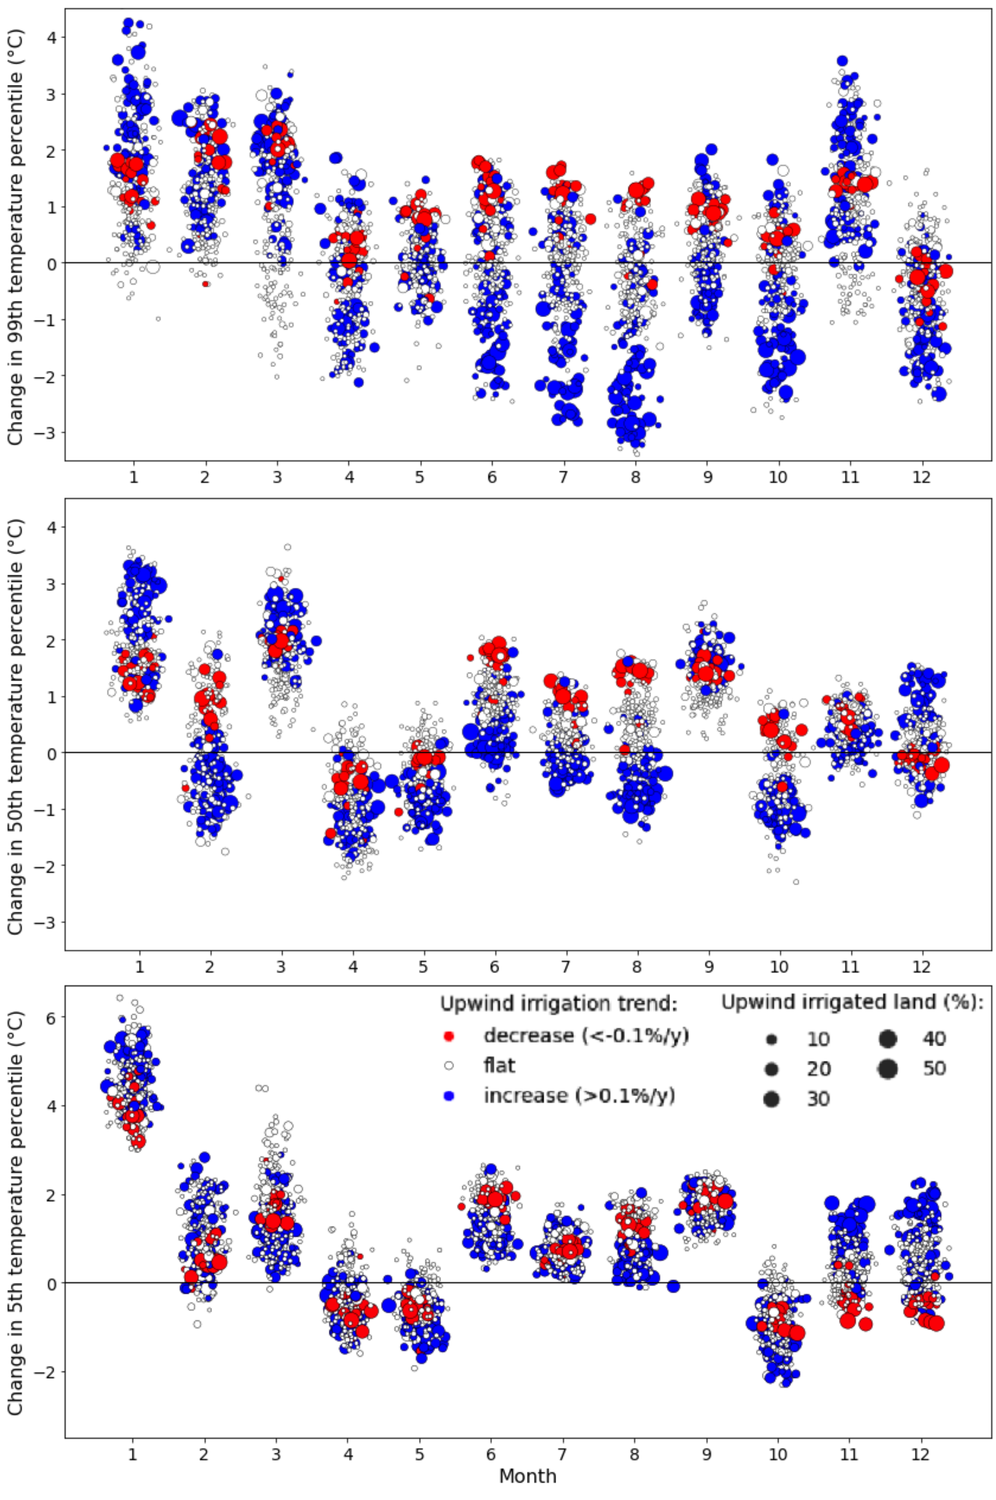

Figure 2: *Average change in the 99^{th} (**top**), 50^{th} (**middle**) and 5^{th} (**bottom**) temperature percentiles over the 1959-2019 period as a function of the month (horizontally jittered) for each of the 393 counties in the region of interest.*

Notes: Each county is represented by a bubble, the size and color of which respectively indicate the average proportion of irrigated land and the irrigation trend observed in the county's upwind neighbor.

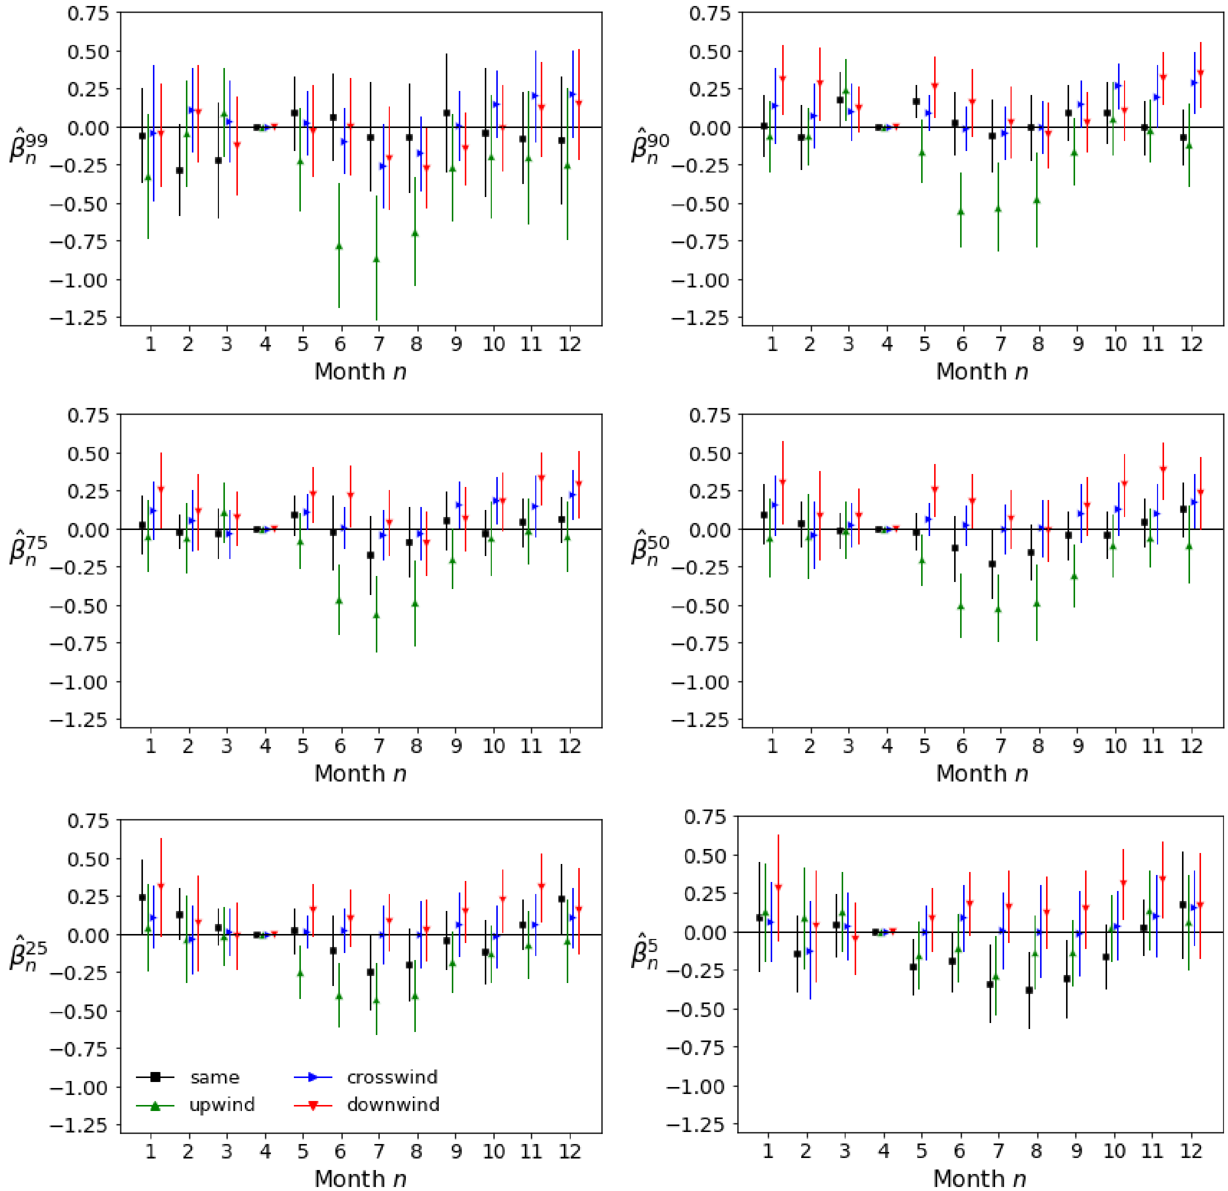

Figure 3: *Effects of irrigation from current (black), upwind (green), crosswind (blue) and downwind (red) counties on the 99th, 90th, 75th, 50th, 25th and 5th temperature percentiles estimated jointly via spatial first differences for each month of the year and with 95% confidence intervals.*

Notes: See specification (6). Units are in °C/(% of irrigated county area). Compare with Figure A10 (different scale).

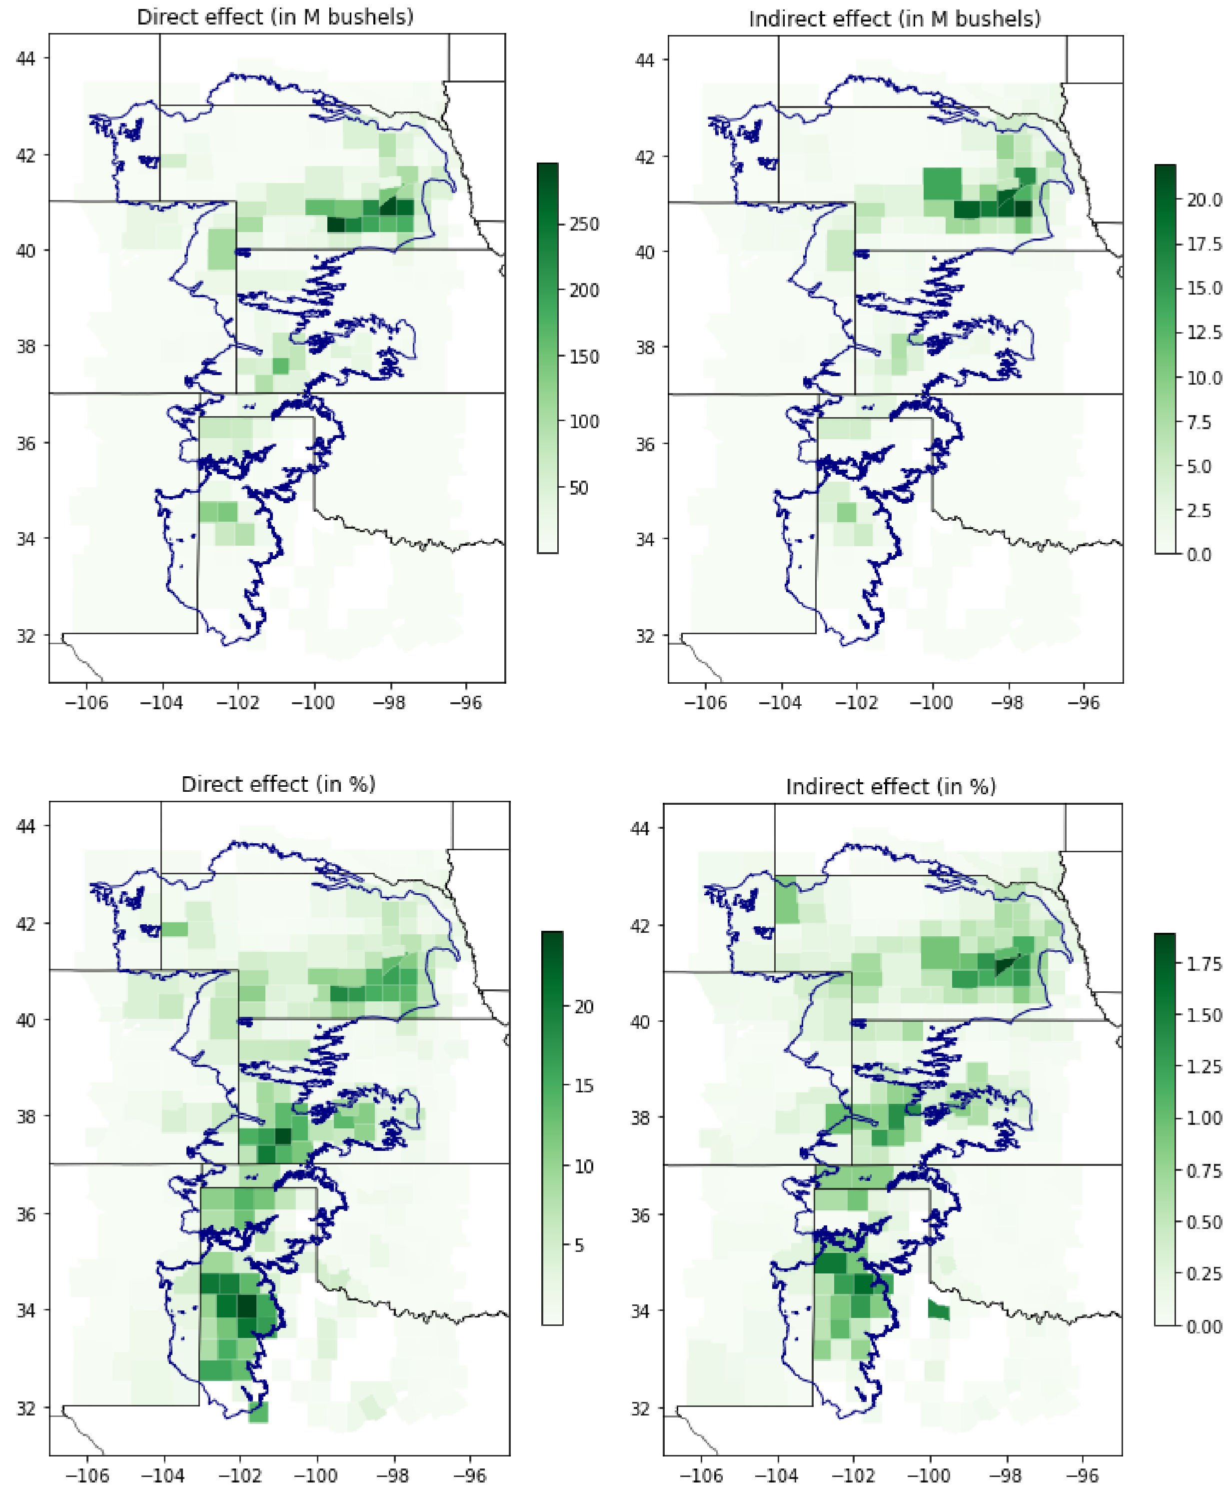

Figure 4: *Avoided corn production losses via the primary effect of within-county irrigation ("direct effect", **left column**) and via the cooling-by-irrigation effect induced by irrigation in upwind counties "indirect effect", **right column**), estimated in million bushels (**top row**) and in proportion to total corn production (**bottom row**) over the 1959-2019 period.*

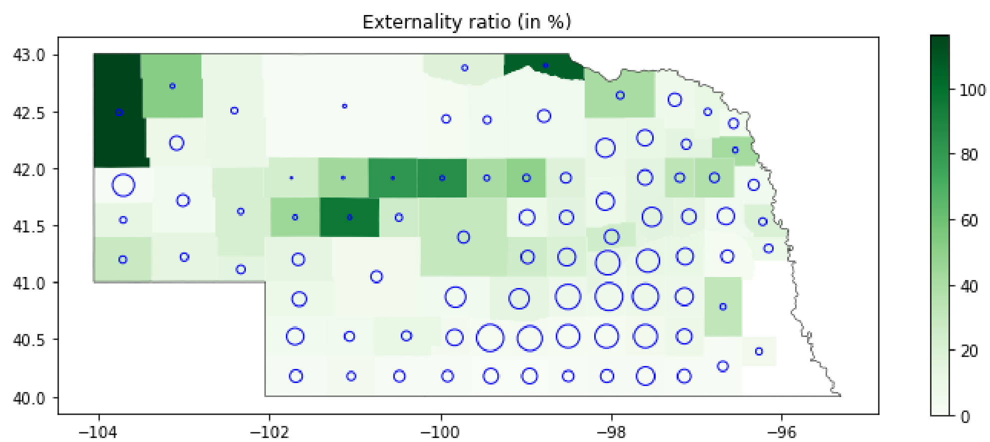

Figure 5: *Ratio of avoided corn production losses due to unintended cooling induced by irrigation in upwind neighbors to avoided corn production losses due to the primary effect of within-county irrigation in Nebraska.*

Notes: Size of blue bubbles is proportional to average irrigation intensity over the 1959-2019 period.

A Appendix

Table A1: Estimated effect of irrigation on temperature percentiles ($\hat{\beta}_n^p$) in early (May-June) and peak (July-August) irrigation seasons

p	n	Irrigation information from					R^2	N
		current	upwind	upwind (2×)	crosswind	downwind		
99th	May-June	0.15	-0.33**	-0.06	-0.12	-0.03		
		(0.12)	(0.15)	(0.14)	(0.11)	(0.11)		
	July-August	0.01	-0.53***	-0.29***	-0.31**	-0.27**		
		(0.17)	(0.16)	(0.10)	(0.12)	(0.12)	8%	13500
90th	May-June	0.06	-0.35***	0.09	-0.11	0.02		
		(0.09)	(0.09)	(0.10)	(0.07)	(0.08)		
	July-August	-0.06	-0.44***	-0.13	-0.17***	-0.22**		
		(0.13)	(0.12)	(0.10)	(0.07)	(0.09)	8%	13500
75th	May-June	0.01	-0.24***	0.07	-0.05	0.06		
		(0.07)	(0.07)	(0.10)	(0.06)	(0.07)		
	July-August	-0.15	-0.44***	-0.12	-0.15**	-0.20**		
		(0.10)	(0.10)	(0.10)	(0.06)	(0.09)	7%	13500
50th	May-June	-0.10*	-0.25***	0.02	-0.04	0.03		
		(0.06)	(0.07)	(0.08)	(0.05)	(0.07)		
	July-August	-0.21**	-0.37***	-0.11	-0.09*	-0.17*		
		(0.09)	(0.09)	(0.08)	(0.05)	(0.09)	6%	13500
25th	May-June	-0.11	-0.25***	0.01	-0.02	-0.01		
		(0.06)	(0.07)	(0.07)	(0.05)	(0.08)		
	July-August	-0.29***	-0.33***	-0.05	-0.04	-0.10		
		(0.11)	(0.09)	(0.07)	(0.06)	(0.11)	5%	13500
10th	May-June	-0.12*	-0.19**	-0.02	-0.04	-0.06		
		(0.06)	(0.08)	(0.08)	(0.06)	(0.10)		
	July-August	-0.26**	-0.27***	-0.01	-0.02	-0.04		
		(0.10)	(0.10)	(0.08)	(0.07)	(0.11)	4%	13500
1st	May-June	-0.08	-0.06	-0.06	0.04	0.02		
		(0.09)	(0.12)	(0.13)	(0.09)	(0.12)		
	July-August	-0.26**	-0.27**	0.06	-0.07	-0.15		
		(0.11)	(0.13)	(0.12)	(0.11)	(0.14)	3%	13500

Notes: coefficients are estimated separately for each of the 99th, 90th, 75th, 50th, 25th, 10th and 1st temperature percentiles by continuous difference-in-differences in spatial first differences with irrigation information from current, upwind, 2-step upwind, crosswind and downwind counties. Reference season runs from September to April. All specifications include an intercept, a linear term in irrigation trend and controls for precipitation trends and soil characteristics. Units are in °C/(% of irrigated county area). Conley standard errors (given in parentheses) are computed using a uniform kernel and a bandwidth of .5 degrees. ***, ** and * denote statistical significance respectively at the 1%, 5% and 10% levels.

Table A2: Estimated effect of irrigation on temperature percentiles ($\hat{\beta}_n^p$) in early (May-June) and peak (July-August) irrigation seasons for counties in **Nebraska**, **Colorado**, **Wyoming**, **South Dakota** and **Kansas**

p	n	Irrigation information from				R^2	N
		current	upwind	crosswind	downwind		
99^{th}	May-June	0.09	-0.35*	-0.16	0.03		
		(0.15)	(0.19)	(0.15)	(0.15)		
	July-August	-0.10	-0.70***	-0.31*	-0.44***		
		(0.20)	(0.22)	(0.17)	(0.17)	13%	7344
90^{th}	May-June	-0.04	-0.38***	-0.20**	0.03		
		(0.09)	(0.11)	(0.09)	(0.11)		
	July-August	-0.23*	-0.51***	-0.15*	-0.42***		
		(0.14)	(0.16)	(0.09)	(0.11)	10%	7344
75^{th}	May-June	-0.08	-0.28***	-0.11	0.09		
		(0.07)	(0.09)	(0.08)	(0.09)		
	July-August	-0.30***	-0.45***	-0.09	-0.35***		
		(0.10)	(0.14)	(0.08)	(0.10)	8%	7344
50^{th}	May-June	-0.15**	-0.30***	-0.09	0.06		
		(0.07)	(0.09)	(0.06)	(0.09)		
	July-August	-0.33***	-0.36***	-0.03	-0.24**		
		(0.10)	(0.12)	(0.06)	(0.10)	7%	7344
25^{th}	May-June	-0.16**	-0.31***	-0.03	0.03		
		(0.08)	(0.08)	(0.06)	(0.09)		
	July-August	-0.40***	-0.32***	0.03	-0.08		
		(0.13)	(0.10)	(0.07)	(0.12)	6%	7344
10^{th}	May-June	-0.20***	-0.25***	-0.00	-0.03		
		(0.07)	(0.10)	(0.07)	(0.11)		
	July-August	-0.39***	-0.22***	0.08	-0.01		
		(0.12)	(0.11)	(0.08)	(0.13)	5%	7344
1^{st}	May-June	-0.09	-0.01	0.15	0.03		
		(0.11)	(0.13)	(0.10)	(0.14)		
	July-August	-0.38***	-0.14	0.05	-0.24		
		(0.13)	(0.15)	(0.10)	(0.17)	5%	7344

Notes: coefficients are estimated separately for each of the 99^{th}, 90^{th}, 75^{th}, 50^{th}, 25^{th}, 10^{th} and 1^{st} temperature percentiles by continuous difference-in-differences in spatial first differences with irrigation information from current, upwind, crosswind and downwind counties. Reference season runs from September to April. All specifications include an intercept, a linear term in irrigation trend, dummies for the May-June and July-August periods and controls for precipitation trends and soil characteristics. Units are in °C/(% of irrigated county area). Conley standard errors (given in parentheses) are computed using a uniform kernel and a bandwidth of .5 degrees. ***, ** and * denote statistical significance respectively at the 1%, 5% and 10% levels.

Table A3: Estimated effect of irrigation on temperature percentiles ($\hat{\beta}_n^p$) in early (May-June) and peak (July-August) irrigation seasons for counties in **Texas**, **Oklahoma** and **New Mexico**

p	n	Irrigation information from				R^2	N
		current	upwind	crosswind	downwind		
99th	May-June	0.44***	-0.45**	-0.06	-0.04		
		(0.16)	(0.19)	(0.10)	(0.13)		
	July-August	0.25	-0.52***	-0.09	0.16		
		(0.21)	(0.19)	(0.18)	(0.12)	8%	5868
90th	May-June	0.43***	-0.25*	0.06	0.05		
		(0.10)	(0.15)	(0.07)	(0.10)		
	July-August	0.27*	-0.43***	-0.07	0.26***		
		(0.16)	(0.14)	(0.13)	(0.10)	14%	5868
75th	May-June	0.32***	-0.13	0.05	0.08		
		(0.06)	(0.12)	(0.06)	(0.12)		
	July-August	0.14	-0.53***	-0.21*	0.19		
		(0.17)	(0.14)	(0.12)	(0.15)	16%	5868
50th	May-June	0.10	-0.15	0.04	0.04		
		(0.07)	(0.11)	(0.07)	(0.13)		
	July-August	-0.01	-0.50***	-0.18	0.05		
		(0.14)	(0.15)	(0.11)	(0.16)	14%	5868
25th	May-June	-0.11	-0.16	-0.01	-0.07		
		(0.10)	(0.12)	(0.08)	(0.15)		
	July-August	-0.02	-0.44***	-0.15	-0.10		
		(0.14)	(0.15)	(0.13)	(0.17)	11%	5868
10th	May-June	0.15	-0.17	-0.15	-0.17		
		(0.14)	(0.11)	(0.10)	(0.17)		
	July-August	0.00	-0.44***	-0.23	-0.10		
		(0.16)	(0.16)	(0.15)	(0.16)	9%	5868
1st	May-June	-0.10	-0.30*	-0.18	-0.02		
		(0.20)	(0.16)	(0.19)	(0.21)		
	July-August	-0.14	-0.53***	-0.34	0.04		
		(0.21)	(0.15)	(0.26)	(0.20)	7%	5868

Notes: coefficients are estimated separately for each of the 99th, 90th, 75th, 50th, 25th, 10th and 1st temperature percentiles by continuous difference-in-differences in spatial first differences with irrigation information from current, upwind, crosswind and downwind counties. Reference season runs from September to April. All specifications include an intercept, a linear term in irrigation trend, dummies for the May-June and July-August periods and controls for precipitation trends and soil characteristics. Units are in °C/(% of irrigated county area). Conley standard errors (given in parentheses) are computed using a uniform kernel and a bandwidth of .5 degrees. ***, ** and * denote statistical significance respectively at the 1%, 5% and 10% levels.

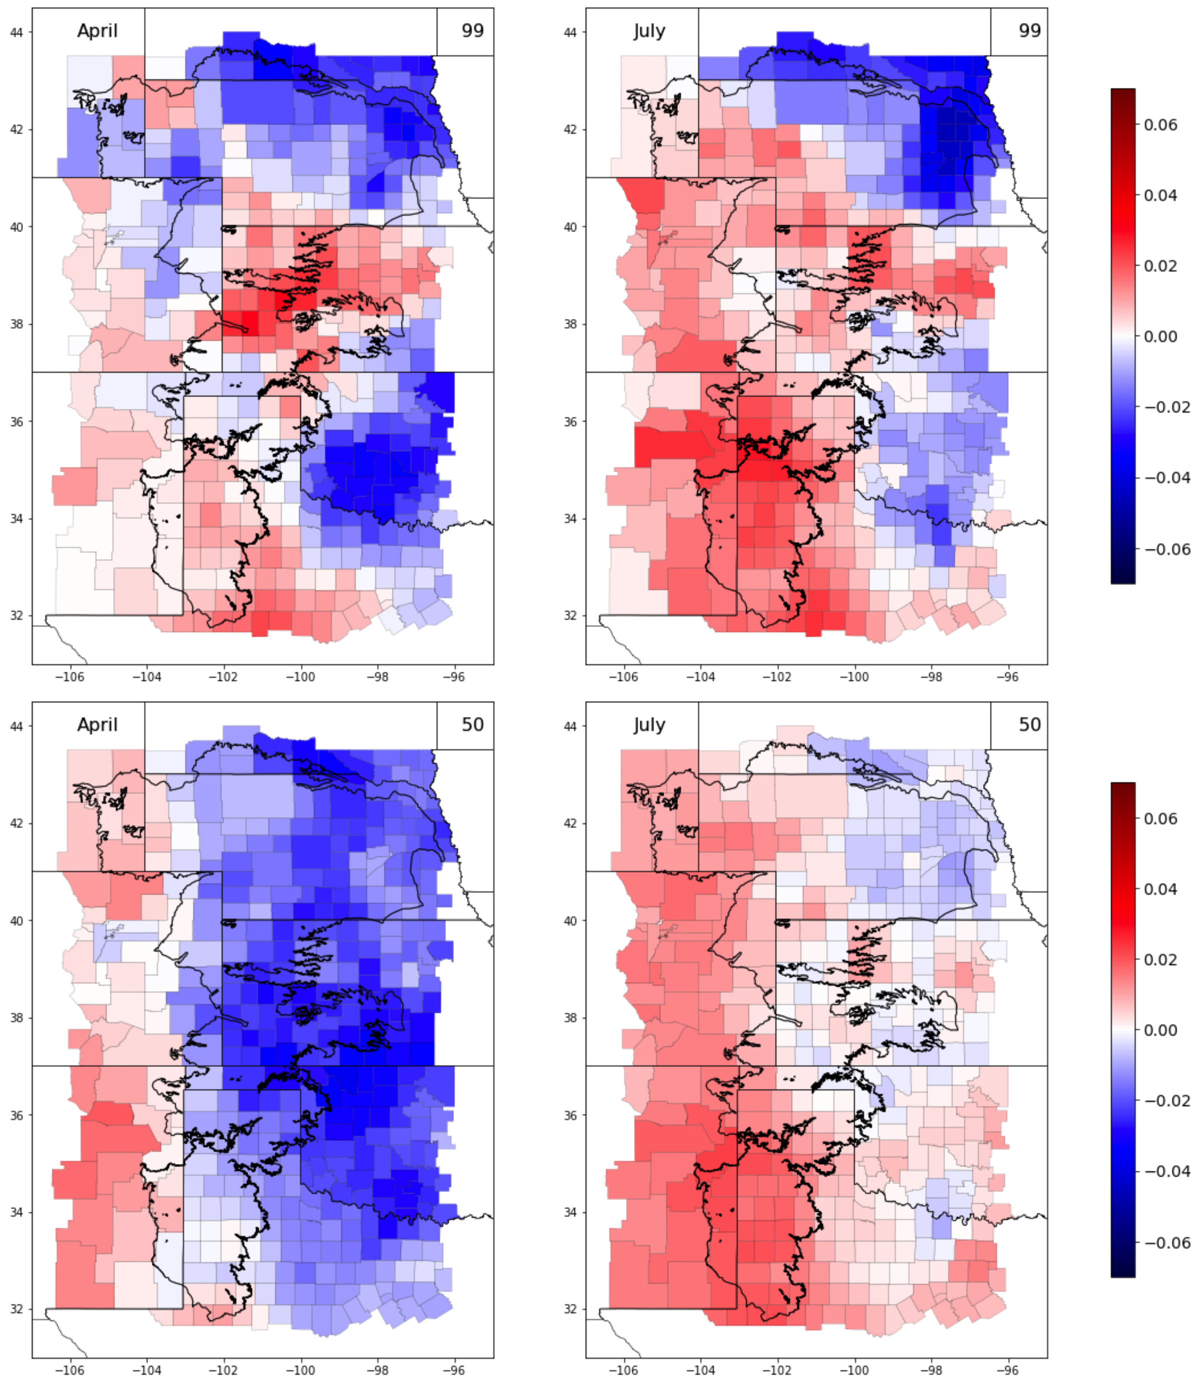

Figure A1: *Observed trends* ($\hat{\theta}^p_{im}$) *in the 99*th *(**top row**) and 50*th *(**bottom row**) temperature percentiles* (p) *in the 393 counties* (i) *of the area of interest for months* (m) *of April (**left column**) and July (**right column**).*

Notes: Counties having experienced an average cooling or warming over the 1959-2019 period are colored in blue or red respectively. The color bar is on a common scale and given in units of $^{\circ}C/y$.

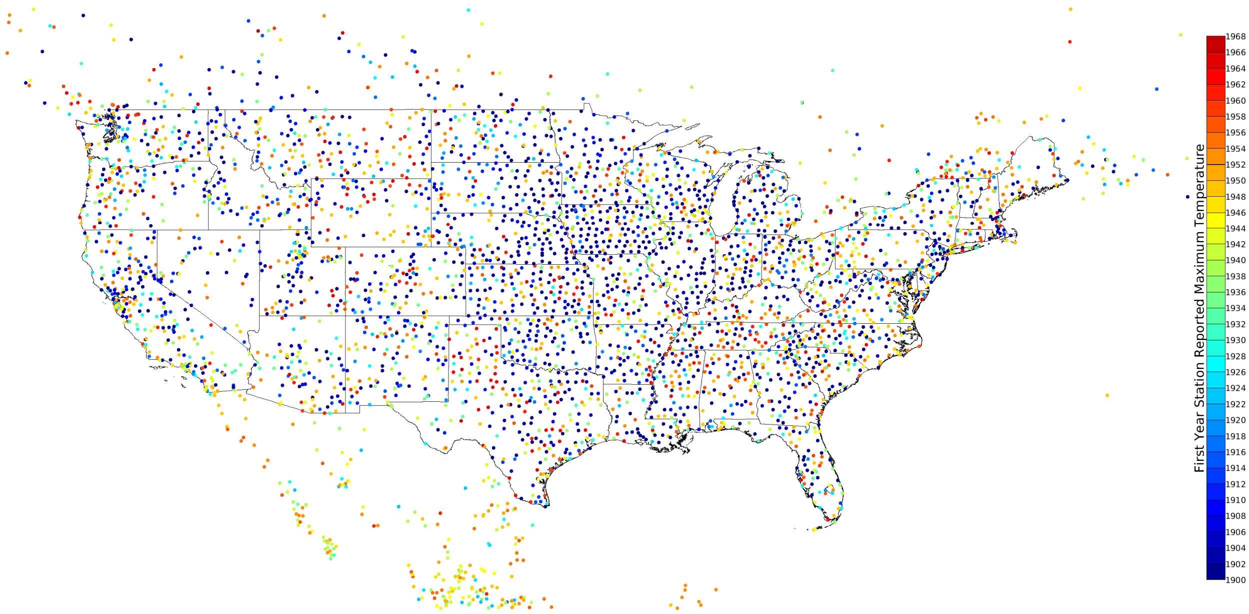

Figure A2: *Location of station and first year station reported data for maximum temperature*

Notes: Figure displays the location of the weather stations used in the interpolation as circle, which are colored by the first year the station reported data (see right legend). Outlines of the 48 states in the contiguous US are added in black. Some stations are located in Canada or Mexico.

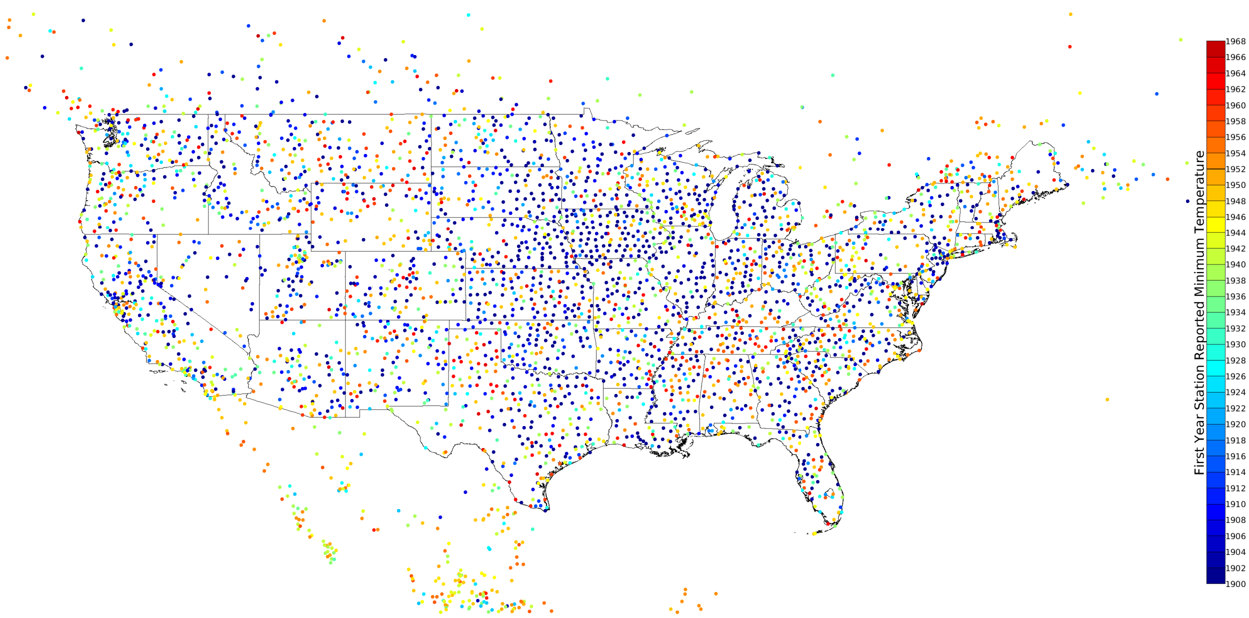

Figure A3: *Location of station and first year station reported data for minimum temperature*

Notes: Figure displays the location of the weather stations used in the interpolation as circle, which are colored by the first year the station reported data (see right legend). Outlines of the 48 states in the contiguous US are added in black. Some stations are located in Canada or Mexico.

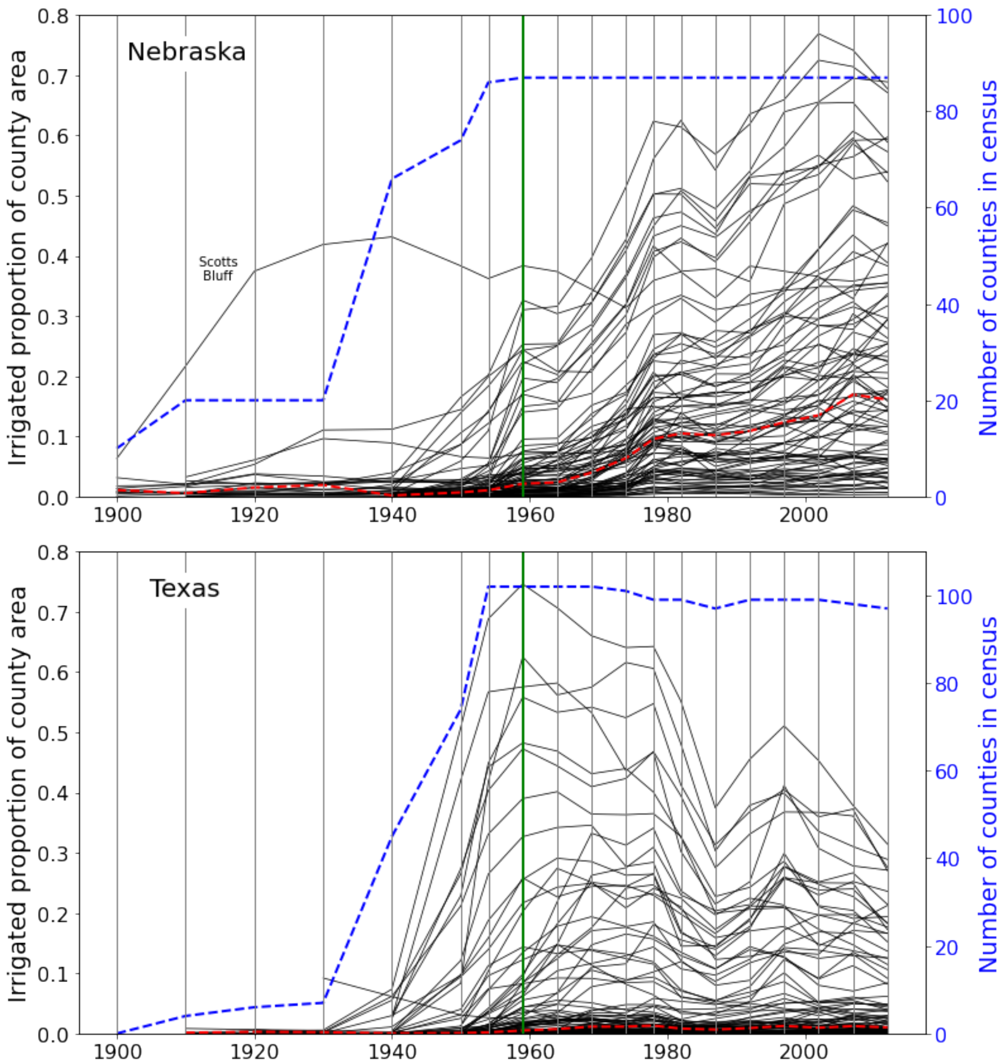

Figure A4: *Evolution over 1900-2012 of the proportion of irrigated area for each county in Nebraksa (**top**) and Texas (**bottom**) that falls within the area of interest (continuous black lines, **primary axis**) and total number of counties with available information about irrigated acres (dashed blue line, **secondary axis**).*

Notes: The median proportion of irrigated county area is indicated by the dashed red line. Vertical lines represent years with agricultural census, among which the year 1959 (in green) marks the start of the period of interest.

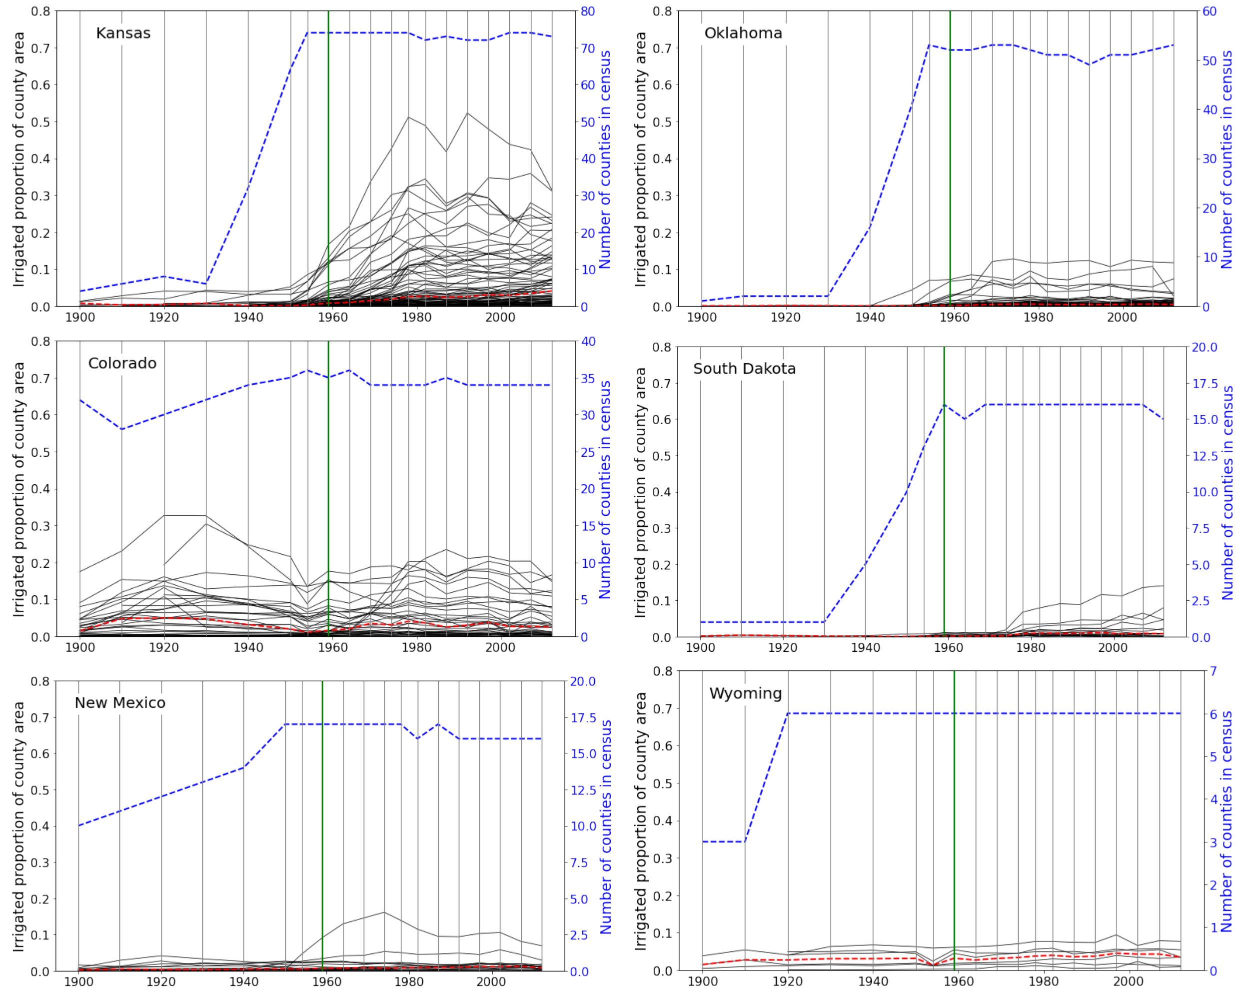

Figure A5: *Evolution over 1900-2012 of the proportion of irrigated area for each county in Kansas, Oklahoma, Colorado, South Dakota, New Mexico and Wyoming that falls within the area of interest (continuous black lines,* ***primary axis****) and total number of counties with available information about irrigated acres (dashed blue line,* ***secondary axis****).*

Notes: The median proportion of irrigated county area is indicated by the dashed red line. Vertical lines represent years with agricultural census, among which the year 1959 (in green) marks the start of the period of interest.

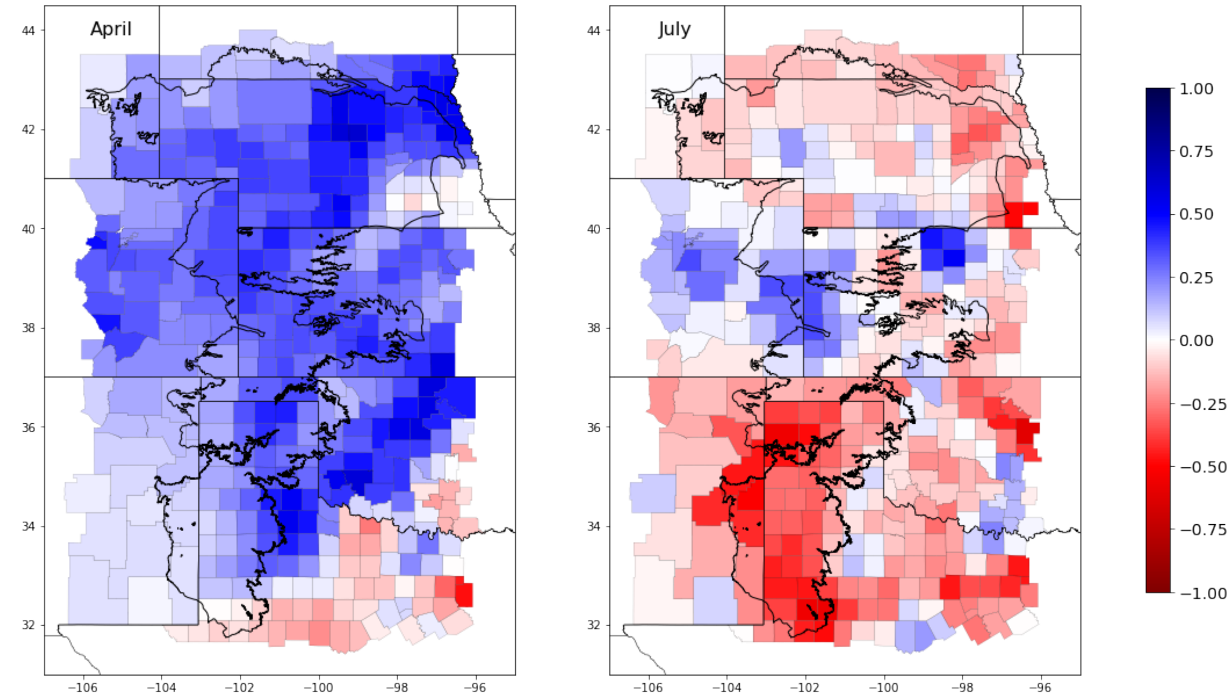

Figure A6: *Observed trends in total rainfall* ($\hat{\rho}_{im}$) *in the 393 counties* (i) *of the area of interest for months* (m) *of April* (***left***) *and July* (***right***).

Notes: Counties having experienced an average increase/decrease in rainfall over the 1959-2019 period are colored in blue/red respectively. The color bar is on a common scale and given in units of mm/y.

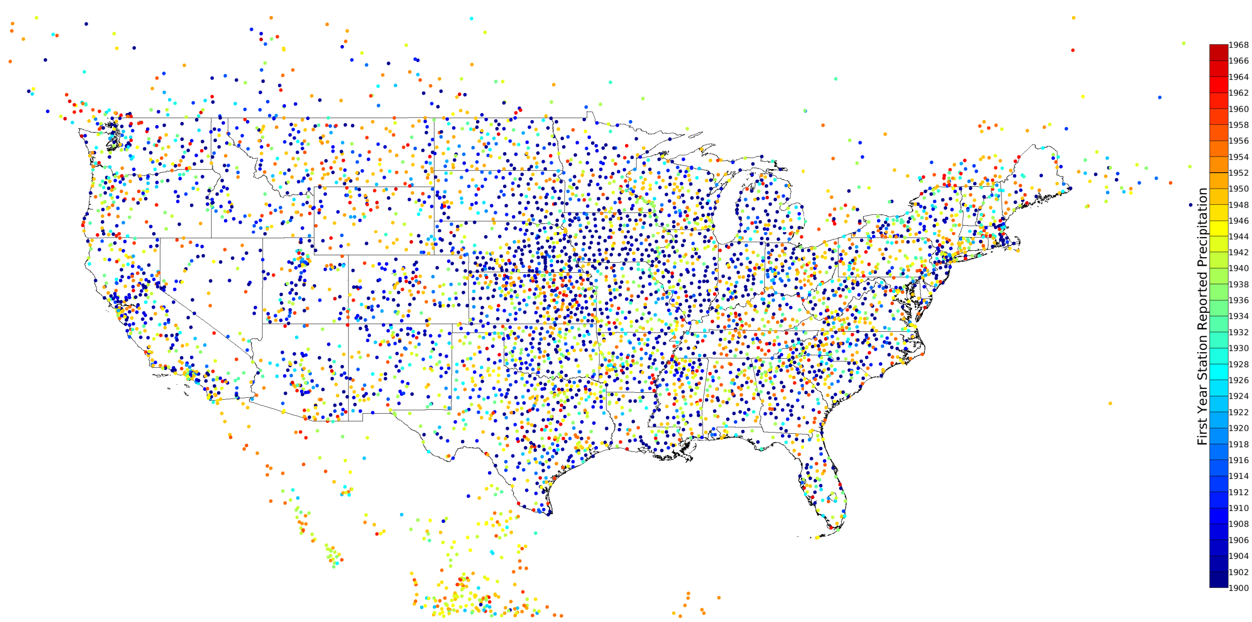

Figure A7: *Location of station and first year station reported data for precipitation*

Notes: Figure displays the location of the weather stations used in the interpolation as circle, which are colored by the first year the station reported data (see right legend). Outlines of the 48 states in the contiguous US are added in black. Some stations are located in Canada or Mexico.

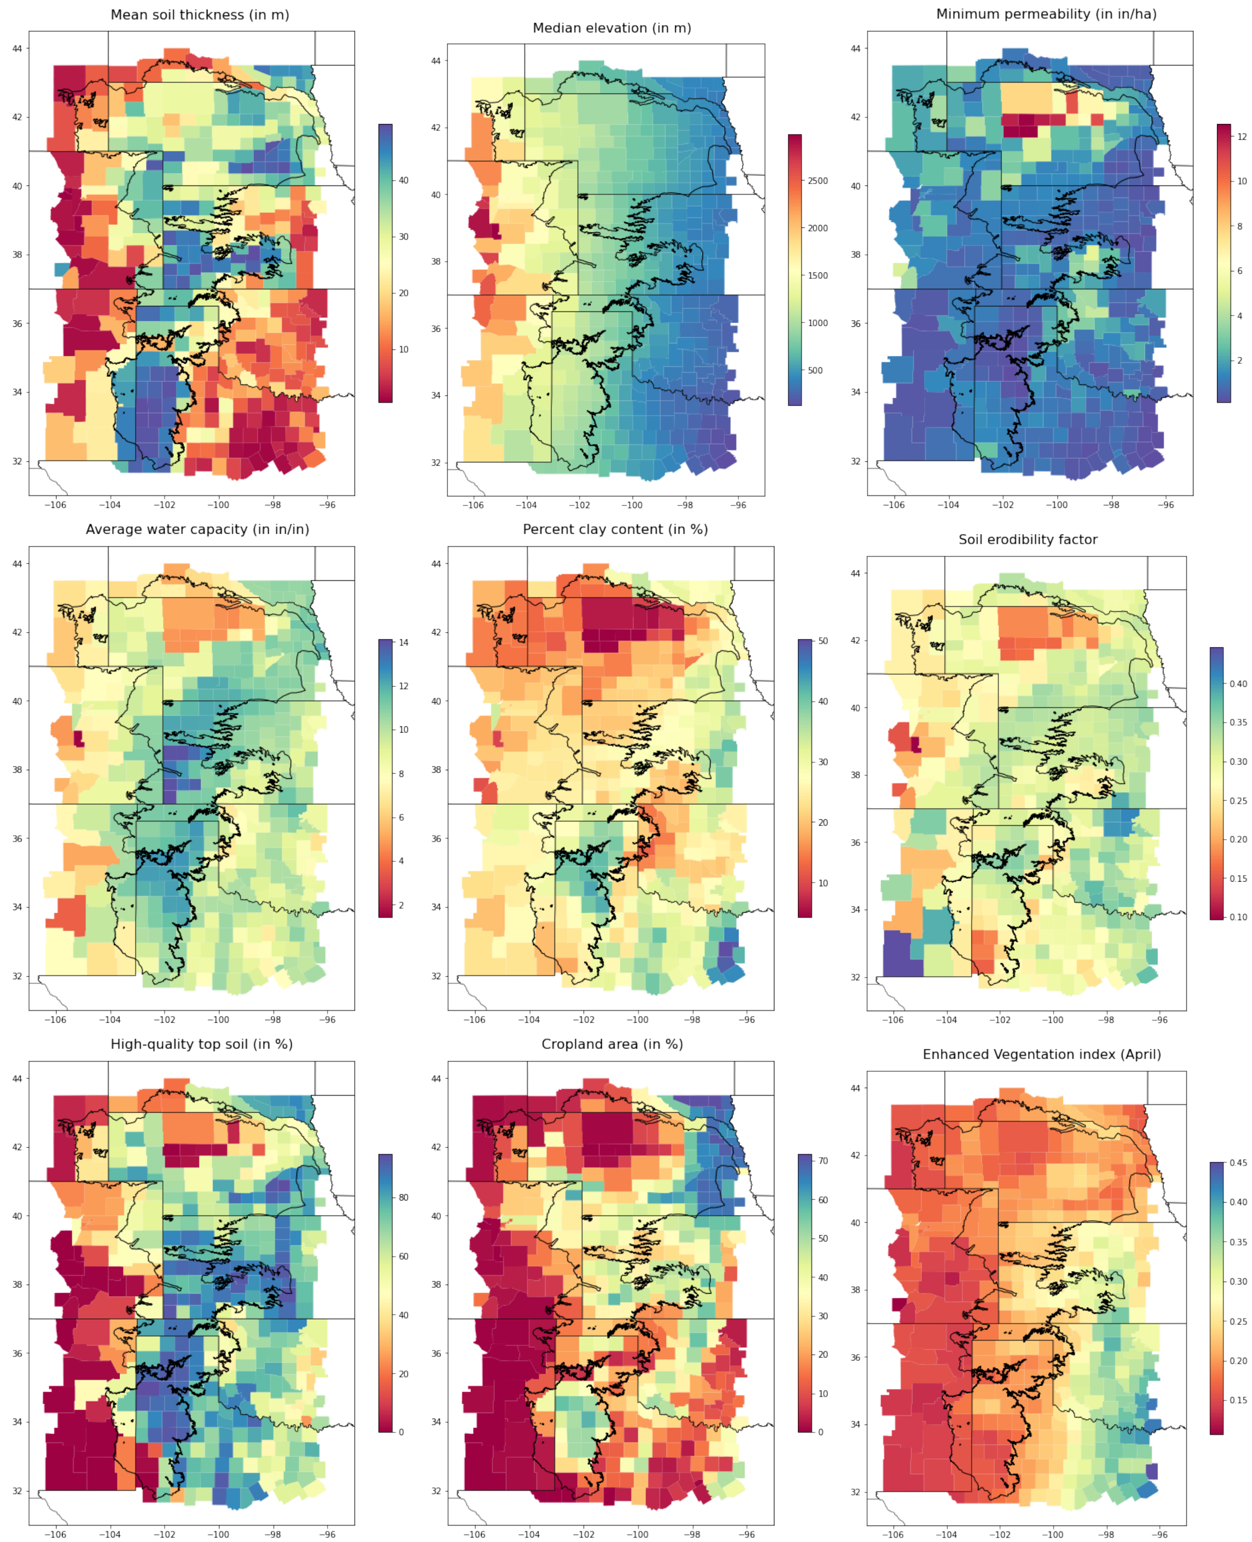

Figure A8: *Map of soil controls, namely soil thickness, elevation, permeability, water capacity, clay content, erodibility, top-soil quality, cropland area, EVI for the 393 counties in the area of interest.*

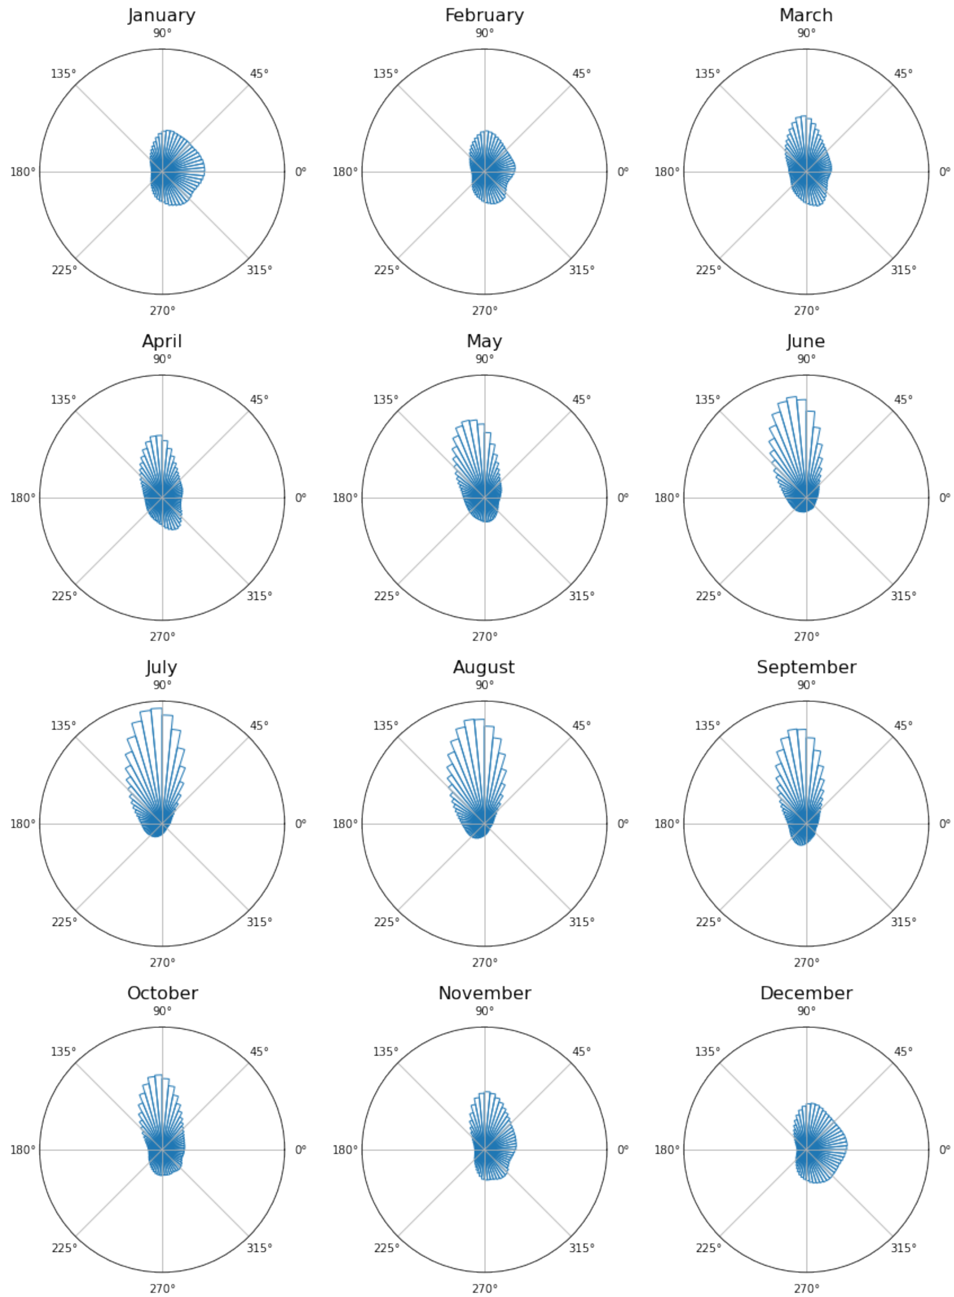

Figure A9: *Polar distribution of hourly wind directions for counties in the area of interest for each month of the year over the 1979-2019 period.*

Notes: Contrary to the usual convention for wind roses, we orient directions in agreement with (not against) the wind flow.

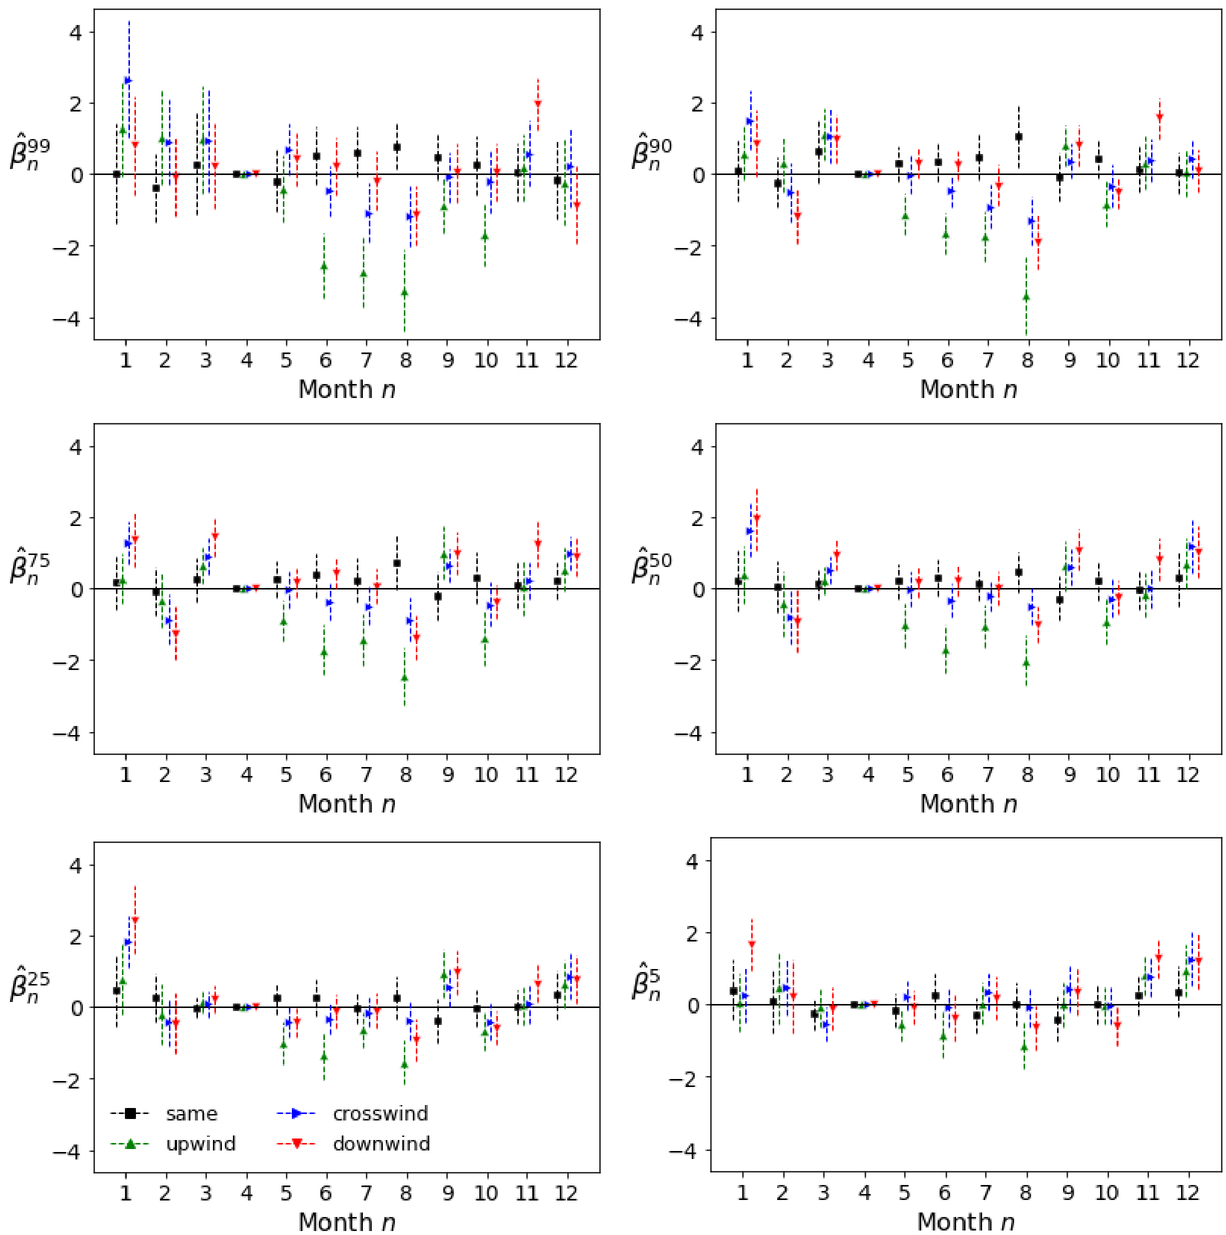

Figure A10: *Effects of irrigation from current (black), upwind (green), crosswind (blue) and downwind (red) counties on the 99th, 90th, 75th, 50th, 25th and 5th temperature percentiles estimated jointly via stand-alone continuous difference-in-differences for each month of the year and with 95% confidence intervals.*

Notes: See specification (4). Units are in °C/(% of irrigated county area). Compare with Figures 3 (different scale) and A11.

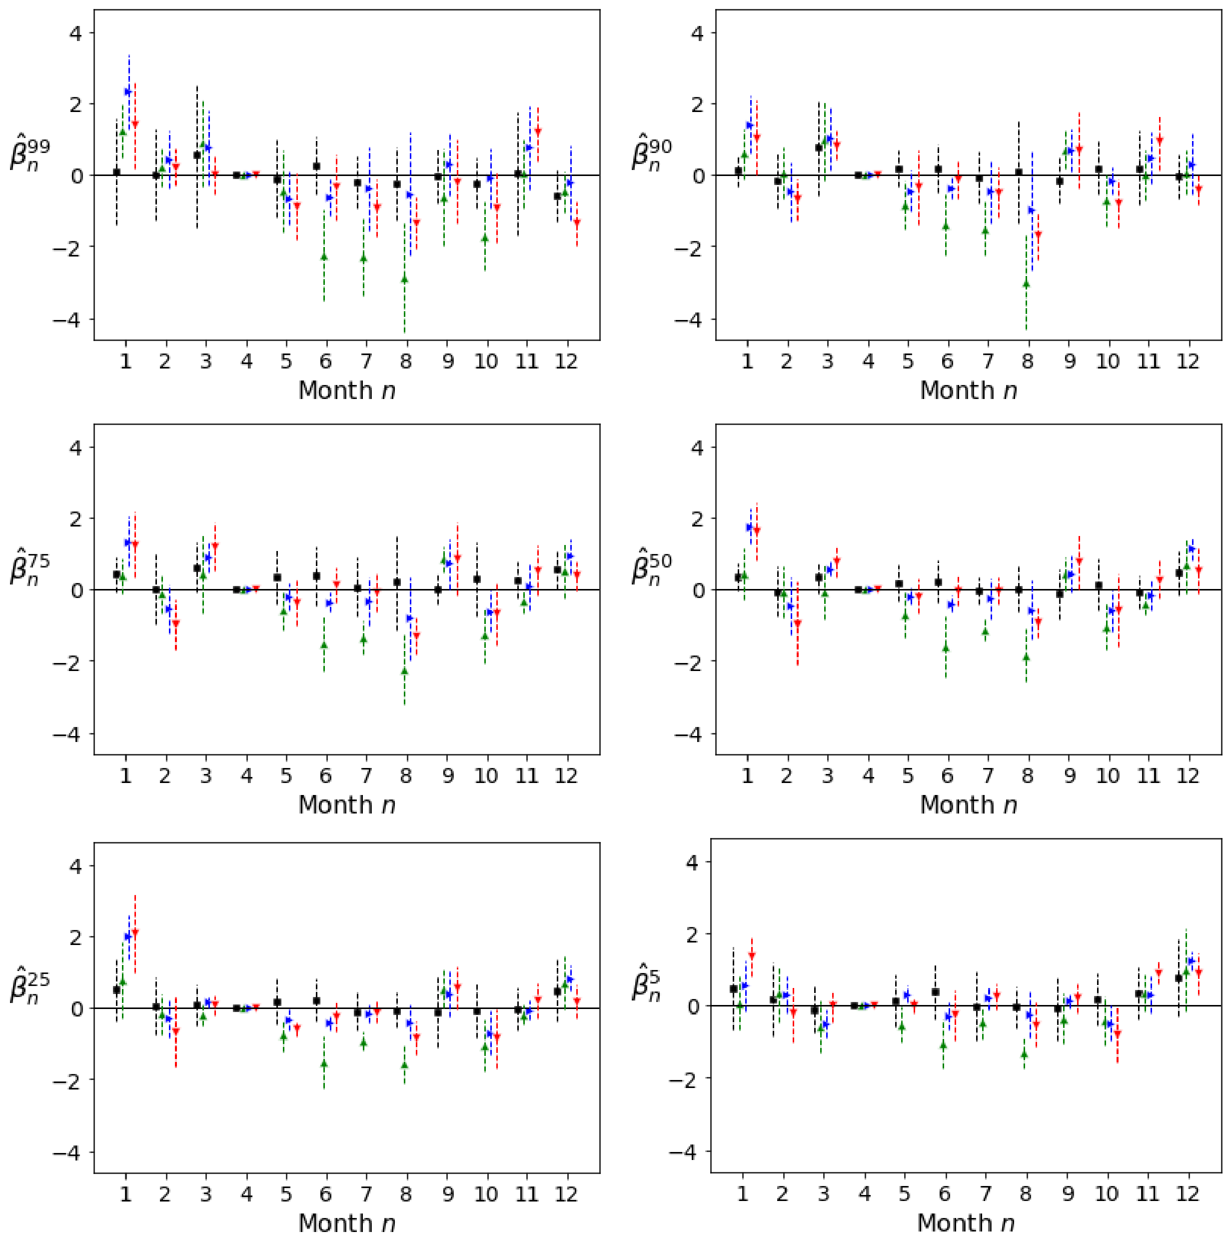

Figure A11: *Effects of irrigation from current (black), upwind (green), crosswind (blue) and downwind (red) counties on the 99^th^, 90^th^, 75^th^, 50^th^, 25^th^ and 5^th^ temperature percentiles estimated jointly via fixed effects model for the panel data in levels for each month of the year and with 95% confidence intervals.*

Notes: See specification described in footnote 23. Units are in °C/(% of irrigated county area). Compare with Figures 3 and A10.

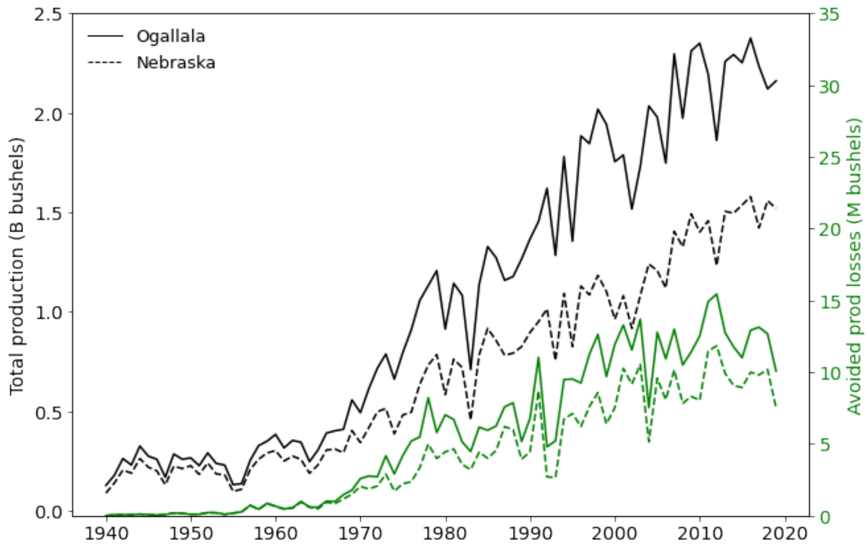

Figure A12: *Total corn production (in billion bushels, black) and avoided corn production losses via the cooling-by-irrigation effect induced by irrigation in upwind counties (in million bushels, green) since 1940, for the region of interest (continuous lines) and in Nebraska alone (dashed lines).*

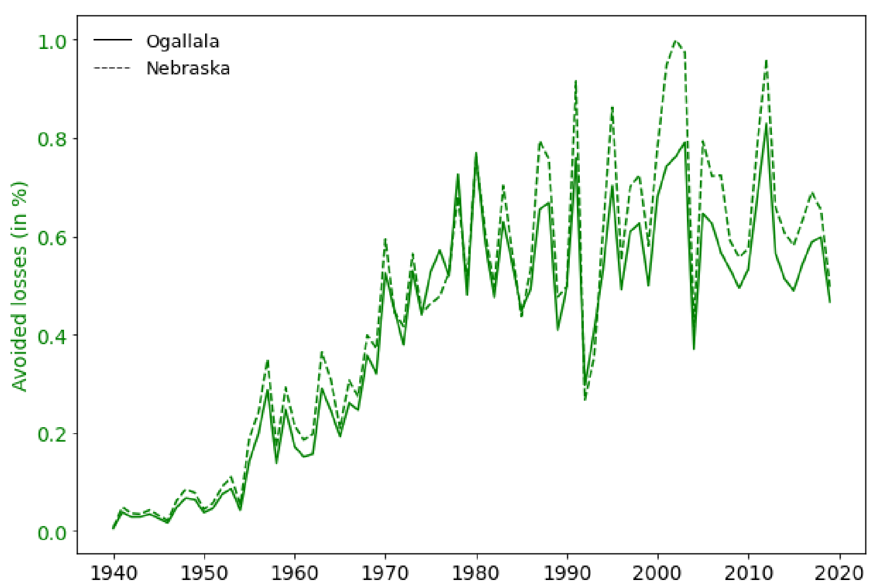

Figure A13: *Proportion of avoided losses in the total corn production via the cooling-by-irrigation effect induced by irrigation in upwind counties for the region of interest (continuous line) and Nebraska alone (dashed line).*

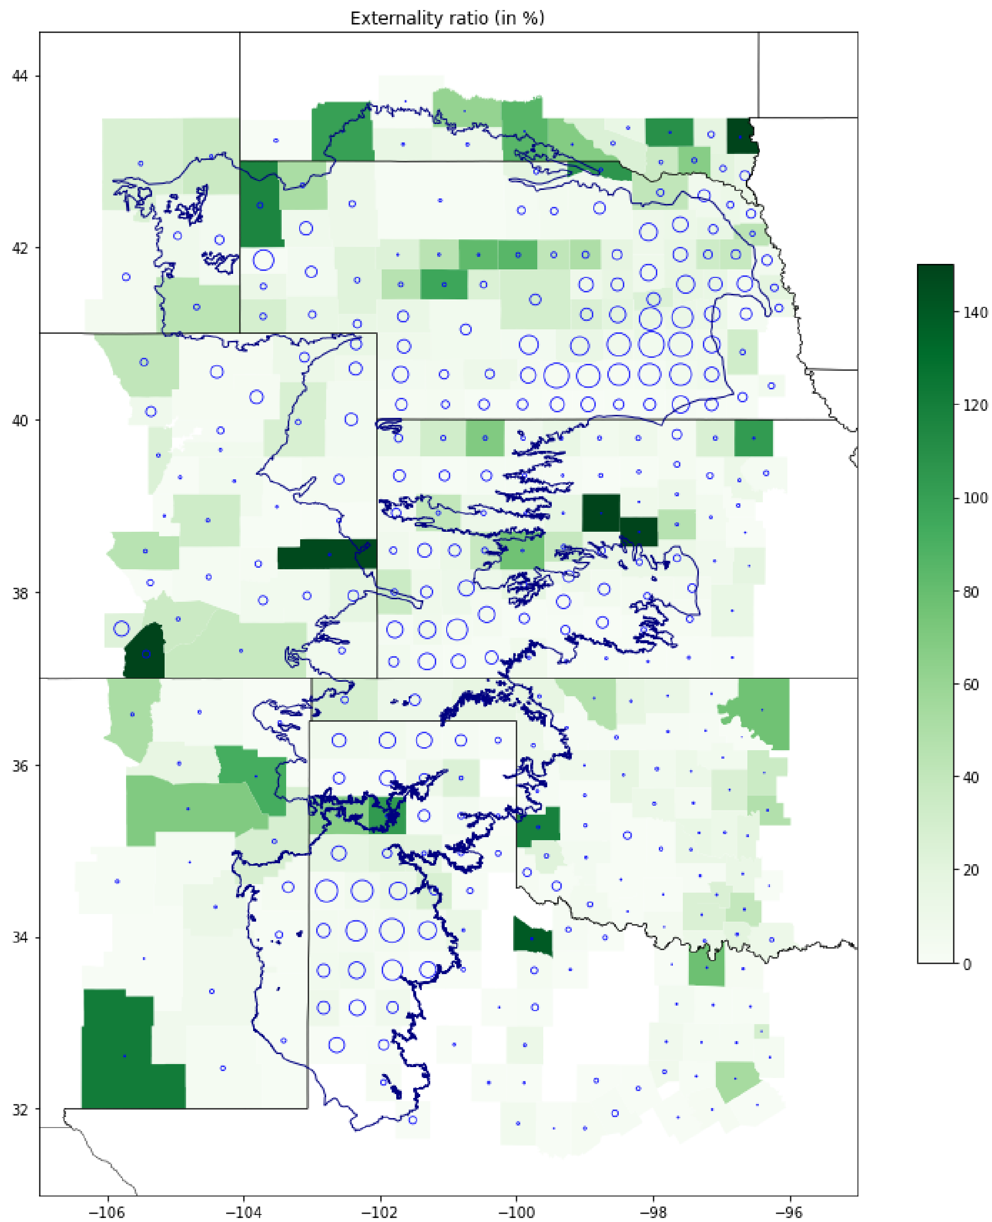

Figure A14: *Ratio of avoided corn production losses due to unintended cooling induced by irrigation in upwind neighbors to avoided corn production losses due to the primary effect of within-county irrigation.*

Notes: Size of blue bubbles is proportional to average irrigation intensity over the 1959-2019 period.