7-25-2018

Estimating Adaptation to Climate Change in Groundwater Irrigation

James Keeler
University of Nebraska-Lincoln, keelerjamesb@gmail.com

ESTIMATING ADAPTATION TO CLIMATE CHANGE IN GROUNDWATER IRRIGATION

by

James B. Keeler

A THESIS

Presented to the Faculty of

The Graduate College at the University of Nebraska

In Partial Fulfillment of Requirements

For the Degree of Master of Science

Major: Agricultural Economics

Under the Supervision of Professor Taro Mieno

Lincoln, Nebraska

July 2018

ESTIMATING ADAPTATION TO CLIMATE CHANGE IN GROUNDWATER IRRIGATION

James B. Keeler, M.S.

University of Nebraska, 2018

Advisor: Taro Mieno

Understanding the adaptive capacity of irrigated agriculture, including to what extent producers adjust irrigation choices along the intensive and extensive margins, is vital to the development of accurate and holistic estimates of the impacts of climate change on agricultural production and the sustainability of water-related ecosystem services. This thesis proposes and implements a natural experiment using statistical matching methods to estimate how producers adjust groundwater extraction, irrigated crop acreage, and irrigation technology in response to long-term changes in precipitation and evapotranspiration. Results from groundwater irrigated fields in Kansas suggest that intensive and extensive margin water use adaptations are generally limited in practice, but there is some evidence of adjustments in both crop acreage and mean overall groundwater extraction, particularly for irrigated corn production.

Author's Acknowledgments

This work is a product of the support, guidance, and insights of numerous faculty and peers in the Department of Agricultural Economics at the University of Nebraska-Lincoln. I am indebted to Dr. Taro Mieno who introduced me to research in water resources management and has encouraged me to make my own contributions to the field. I am especially thankful for his willingness to present me with empirical challenges and the freedom to overcome them on my own, while always taking the time to ensure I was headed in a productive direction. Thank you Taro.

I am also grateful for the time and efforts of the other member's of my examining committee, Dr. Nicholas Brozović and Dr. Karina Schoengold. Their advice and suggestions have certainly made this document more palatable and valuable for both myself and the reader, while also providing me with the traction to further develop this work in the future.

I also thank Dr. Simanti Banerjee who granted me the opportunity to study, develop, and design economic experiments which provided a welcome productive outlet at times when my thesis progress was slow going. I am also thankful to Kara Heideman for assistance with navigating graduation requirements and everything in between.

Finally, I am sincerely grateful for the unyielding support of my family. I thank my father and mother, Jim and Laurie Keeler, for encouraging me to believe in myself. I also treasure my sister Shannon's dependable home-cooked meals and conversations over the last two years. Without their support I would not be where I am today, nor as confident and prepared for the trials to come.

Thank you all.

Grant Information

This research was supported by USDA NIFA Water CAP Award number 2015-68007-23133.

Contents

List of Figures

List of Tables

1 Introduction

Access to groundwater has driven agricultural development in semi-arid climates across the globe, particularly for communities above the High Plains Aquifer in the U.S. Midwest. Over the last half-century, groundwater withdrawals have sheltered agricultural producers in the region from variations in surface and atmospheric conditions that would otherwise constrain and jeopardize their productivity. Increasing demand for this fundamental function of groundwater, as well as the price of withdrawal being below social cost, has led to over extraction of the resource and threatened the long-term sustainability of irrigated agriculture as aquifers are depleted and water tables fall (Lin Lawell, 2016).

Climate change is expected to exacerbate this issue by further augmenting the variability of precipitation and soil moisture. As climate patterns shift, rainfall events are expected to be more intense and less frequent (Trenberth, 2011). More intense precipitation, concurrent with increasing evapotranspiration as average temperatures rise, will result in less water naturally infiltrating and being retained in soil. Dry spells and droughts are also expected to increase in frequency and duration, exposing crops to longer periods of extreme heat and water stress.

As uncertainty surrounding the natural provision of effective precipitation accumulates, agricultural producers may respond by extracting groundwater more intensively, applying more irrigation water per acre. Such a response along the intensive margin could be beneficial for agricultural productivity and food security in the short-run, but unsustainable in the long-run as groundwater sources are further depleted and natural recharge is diminished (Meixner et al., 2016). Alternatively, producers may respond to shifting climate patterns by managing available precipitation and other water resources more frugally. This can be accomplished by adjusting irrigated acreage, switching to less-water intensive crops and varieties, investing in more efficient irrigation technologies, and more. These extensive margin adjustments are typically more expensive and slower than intensive margin responses, but are potentially more sustainable in the long-run for regions where groundwater is increasingly scarce. Producer responses on the intensive and extensive margin, alone or in conjunction, may mitigate some of the expected negative impacts of climate change, a process defined in the literature as adaptation. Understanding the adaptive capacity of irrigated agriculture, including to what extent producers adjust their behaviors through various mechanisms and channels, is crucial for developing accurate and holistic estimates for the impacts of climate change on agricultural production and natural resource management over both short and long-run time horizons.

Past literature presents evidence that producer adaptation is limited in practice, suggesting some profoundly pessimistic outcomes for an agricultural sector open and exposed to climate change (Schlenker and Roberts, 2009; Fisher et al., 2012; Burke and Emerick, 2016). Potential explanations for observations of limited adaptation range from physical and institutional constraints on producer adaptability to inherent flaws in the estimation strategies of previous studies. First, producers may face binding constraints which severely limit their capacity to engage in adaptation, e.g. prohibitively expensive technologies or groundwater extraction and land use regulations. Second, recognizing the need for adaptation is a challenge given the fairly slow pace of climate change relative to yearly growing seasons. For producers that rely on long-term moving averages of past climate trends to guide their decisions for current and future growing seasons, a lot of time can pass before the impacts of climate change on environmental conditions become apparent. Cognitive biases common to decision making under uncertainty, like anchoring and representativeness heuristics, further compound this issue. Extensive noise and fluctuations in short-term weather time series can also dampen potential signals for the need to adapt (Burke and Emerick, 2016). Lastly, past empirical approaches to identifying climate change impacts and producer responses in observational data may suffer from endogeneity and misspecification problems, consequently underestimating adaptation.

Given the uncertainty and ambiguity surrounding how irrigated producers respond to climate change in practice, a study of the scope and velocity of adaptation through irrigation choice on both the intensive and extensive margins could clarify how a significant population of the agricultural sector has or has not adapted to climate change. The objective of this thesis is to identify and quantify to what extent irrigated producers in Kansas have adapted their irrigation behavior to changes in climatic conditions over the past couple of decades. If producers do appear to engage in adaption, than a second objective is to determine whether they adjust by applying irrigation water more intensively or frugally.

This thesis' main contribution is the development and implementation of a natural experiment which identifies how crop producers respond to long-term variation in irrigation demand attributed to climate change. As far as the author, it is the first study to directly investigate adaptation responses in irrigation choice. The identification strategy combines both statistical matching methods and fixed effects models with a large panel data set to estimate causal effects of exposure to climate change on observed irrigation behavior along both the intensive and extensive margins. Annual groundwater extraction and irrigated field characteristics from 1991 to 2014 are retrieved from Kansas' WIMAS database. Daily observations of weather variables are extracted from NASA's Daymet gridded data set. Irrigated fields which experienced a meaningful increase in mean irrigation demand (evapotranspiration less precipitation) from 1996 to 2009 are assigned to a treated group. Other fields with covariates that match with the treated group are assigned to a control. Differences in irrigation application and irrigated crop acreage between treated and control fields after exposure to climate change are estimated empirically using differences-in-differences models.

Estimates of climate change treatment effects suggest that adjustment in irrigation choice along both the intensive and extensive margins is scarce in Western Kansas. Irrigation applications and field characteristics were not significantly different before and after exposure to a measure of climate change. However, there is some evidence that producers reduced acreage for crops that are naturally vulnerable to short-term water stress, such as corn, and increased acreage for crops that are less vulnerable, like alfalfa. Treatment effects were then estimated for a sample restricted to susceptible corn production, which also exhibited limited adaptation on each margin. This thesis also estimates the effects of climate change on mean overall groundwater extraction regardless of realized in-season weather conditions using recovered model residuals differenced by treatment status. Estimated differences in intensive margin residuals by experimental group show that corn producers adapted to climate change by reducing their mean overall groundwater extraction without making extensive changes to irrigation technology. It must be noted that these results may not be generalizable to irrigated fields outside of Western Kansas, a concern that arises from using a matched sample. This issue is discussed further in the limitations subsection.

The rest of this thesis is organized as follows. Section 2 reviews the literature for climate change impacts and adaptation in the agricultural sector, including a theory informed critique of past empirical approaches used to identify producer adaptation. Sections 3 through 7 introduce the natural experiment used in this study and detail its implementation, including the necessary assumptions for robust estimation of treatment effects. Section 8 presents results and interprets treatment effects. Section 9 discusses findings, limitations, and extensions. Section 10 concludes.

2 Climate Change Impacts and Adaptation in Agriculture

Past literature has broadly iterated on empirical approaches for identifying adaptation to climate change and its impacts on agricultural production (Mendelsohn et al., 1994; Deschênes and Greenstone, 2007; Burke and Emerick, 2016). Earlier studies focus almost entirely on outputs of production, while more recent investigations bring attention to how adaptations in input use, like irrigation, may stabilize or mitigate reductions in crop yield, revenues, and land values (Schlenker et al., 2005, 2006; Oehninger et al., 2018).

Mendelsohn et al. (1994) apply a ”Ricardian” approach which employs a reduced-form hedonic model to estimate the impact of changes in seasonal weather conditions on cropland prices. Their method improves upon earlier agronomic production function approaches which observe changes in optimal production levels derived from crop yield or revenue maximization problems before and after simple updates to weather inputs like temperature and precipitation. While production function estimates are effective in identifying the agronomic impacts of climate change on agriculture, they ignore any decisions and actions taken by producers to mitigate these impacts through various adaptation mechanisms. The Ricardian approach attempts to remedy this by assessing how changes in weather conditions affect a measure which captures a larger subset of adaptation responses: agricultural land values. Land values embody both changes in agronomic conditions and producer adaptations to climate change like switching crops and investing in newer technologies, in addition to wholesale changes in land use like transitioning to livestock production or exiting the agricultural industry entirely. Mendelsohn et al.'s (1994) cross-sectional estimates, arguably the first to capture adaptation responses, exhibit significant variance in predicted impacts of climate change. Their results for a model emphasizing grain production, which favors cooler temperatures, imply a significant decrease in land values. Results which emphasize high-value per acre crops that favor hotter climates, imply a smaller increase in land values, which contrasts with the less optimistic expectations of preceding studies (Rosenzweig and Parry, 1994).

Mendelsohn et al. (1994) has since served as a benchmark for following literature, motivating improvements in identification strategies which potentially abate misspecification and bias in the original estimates of climate change impacts and adaptation. Schlenker et al. (2005, 2006) extend and supplant the Ricardian approach with estimates which are more robust to sources of endogeneity (measurement error, missing covariates) and rectify conflicting observations for the direction of climate change impacts and adaptation. Yet, given the Ricardian approach is cross-sectional, these earlier studies inadequately control for some of the time-invariant factors (soil and producer characteristics) which affect adaptation responses and outcomes, leading to a well-documented endogeneity problem.

More recent studies adopt panel approaches which include geographical unit (field, county, state, etc.) by year fixed effects. Deschênes and Greenstone (2007) generate estimates of the relationship between annual variation in weather conditions and agricultural yields or profits using a nation-wide panel with observations down to the county level. To predict the impacts of climate change on agriculture, they apply their estimated coefficients for annual weather effects to expectations of climate change derived from global climate models and scenarios, some of which forecast increases in temperature and precipitation of 5°F and 8% respectively. The same method and climate scenarios are also used in cross-sectional approaches to predict agricultural impacts, extrapolating producer responses into the future, but each study applies the method to different sets of dependent variables which account for narrower or wider ranges of adaptation strategies.

Deschênes and Greenstone (2007) find, based on their preferred estimates, that climate change will increase overall agricultural profits by approximately 4%, though the confidence interval for their estimate suggests that impacts are relatively small in expectation. Similar estimates for agricultural yields are not economically or statistically significant. Their findings may indirectly imply either that producers engage in consistent and effective adaptation to annual variations in weather, or that positive impacts in some regions of the U.S. will counteract negative impacts elsewhere. Their results conclude by addressing the fact that their best available forecasts for climate change are spatially uniform, imparting an assumption that either climate change will increase both temperature and precipitation equally across the nation or that all producers in a state experience uniform changes in climate. As temperature and precipitation have conflicting effects on crop yields, their null result could also be attributed to the agronomic impacts of climate change simply balancing out. Evidence from climate change trends in Kansas refute both assumptions, exhibiting significant variation in climate trends across the state, at least over short-run periods.

Fisher et al. (2012) brings attention to concerns that uniform climate model predictions fail to identify climate change as Deschênes and Greenstone (2007) suggest, in addition to reporting anomalies in the original data and code used in Deschênes and

Greenstone (2007)'s analysis. They discover extensive measurement error wherein missing weather observations and climate forecasts take an unreasonable baseline of zero. These errors induce excess random noise in annual weather conditions which attenuate Deschênes and Greenstone (2007)'s climate change impact estimates (bias towards zero). Fisher et al.'s (2012) corrected results suggest significant and large damages to agricultural production owing to climate change, with 11% and 37% reductions in corn yields and agricultural profits, respectively. Deschênes and Greenstone (2012) admit to internal validity concerns in their original estimates and subsequently provide corrected estimates which agree with Fisher et al.'s (2012) pessimistic outcome, suggesting present value discounted total losses in the agricultural sector of $164 billion from 2010-2099 with limited producer adaptation.

An interesting question with implications for agricultural policy and the potential effectiveness of adaptation strategies is how these damages from climate change are expected to be distributed across time. Are they more or less damaging in the short-term than in the long-term? Guiteras (2007) administers Deschênes and Greenstone (2007)'s methodology to a panel of major agricultural districts in India, estimating damages separately for medium-term (2010-2039) and long-term (2070-2099) periods. These estimates are not subject to the same measurement errors as Deschênes and Greenstone (2007)'s, given that Guiteras' uses a separate study area and data set. Guiteras finds that climate change impacts are significantly more severe over the long-run (25% decrease decrease in major crop yields) than over the medium-run (4.5% to 9% decrease). Schlenker and Roberts (2009) provide an explanation for the skewed distribution of estimated climate change impacts, finding that temperature effects on crop yields are nonlinear. They show that rising temperatures increase yields initially, but quickly lead to sharp decreases in yields once temperatures surpass extreme heat thresholds which differ between crops. Schlenker and Roberts incorporate these thresholds into a panel analysis which suggests a 30 to 46% reduction in yields, even when using more optimistic climate change forecasts. In terms of adaptation, Guiteras shows that climate change impacts may be more severe for developing countries which historically lack the resources to support a full range of adaptation strategies. Schlenker and Roberts suggest that agricultural production has not adapted to nonlinear temperature effects over the last century.

Estimates of climate change impacts generated by the above panel approaches share in common a dire outlook for agriculture. Predicted impacts are severe, implying that producers either do not engage in adaptation, or historical and modern adaptation strategies are impotent. However, these estimates are inherently flawed in that they do not account for long-term adaptation strategies. Panel studies identify how weather variation in a single growing season affects producer behavior, which is unlikely to account for more involved (requiring significant investments in time and money) responses like switching crops, adjusting acreage, and updating irrigation technologies. A single year of drought or extreme heat does not typically evoke expedient and drastic action, but a consistent trend over multiple years might. Accordingly, panel approaches restricted to annual variation likely overestimate the impacts of climate change as they do not explicitly account for the full range of adaptation strategies available to producers.

Contemporary studies attempt to remedy this flaw in the panel approach by exploiting variation in weather conditions over longer time horizons (Dell et al., 2012; Burke and Emerick, 2016). Burke and Emerick (2016) introduce a ”Long Differences” approach which attempts to identify adaptation and climate change impacts by estimating producer responses to changes in climate trends compounded over multiple decades. Their approach is unique in that adaptation is observed ex post instead of being derived from extrapolations into the future based solely on short-term producer responses. Thus, they essentially approach the problem from a different direction. While previous studies estimate the impacts of climate change to infer whether or not adaptation takes place, Burke and Emerick observe to what extent producers engage in adaptation after long-term exposure to climate change, which then conditions their impact estimates. Yet, their findings are no less pessimistic.

Burke and Emerick generate long difference estimates using differences in average county-level yields and weather conditions between 1978-1982 and 1998-2002 periods, then compare these estimates with those from a standard panel model like the ones used by Deschênes and Greenstone (2007). By comparing panel estimates, which only capture short-term responses, with long differences estimates, which reasonably account for the full range of responses, Burke and Emerick quantify the proportion of climate change adaptation explained by long-term adjustments. They find that long-term adaptation at best mitigates less than half of decreases in corn yields, but more often long-term adjustment is non-existent. They conclude that adaptation to climate change is limited in practice and predict approximately a 15% loss in annual corn yields by 2050. Burke and Emerick's Long Differences approach effectively combines the advantages of both the Ricardian and panel approaches, accounting for the full range of adaptation strategies while also tempering confoundedness from unobserved time invariant factors. Their approach, one of the more recent innovations in the adaptation literature, lays the foundation for this study's methodology. Potential issues and complications with the long differences approach and estimates, which this study endeavors to ameliorate, are discussed in the next subsection.

All of the approaches discussed so far have been employed almost exclusively to investigate the impacts of climate change on outputs of agricultural production e.g. yields and profits. Few studies have explored adaptations in input use, particularly irrigation. For example, Burke and Emerick focus solely on crop yields in counties east of the 100th meridian where precipitation is substantial and consistent enough that irrigation is not a necessity. Schlenker et al. (2005) provide evidence that dry-land and irrigated production exhibit statistically significant structural differences in response to climate change, suggesting that models of climate change impacts must account for each type of production separately. However, data for irrigated land and production was not readily available at the time of their study. The recent ongoing work of Oehninger et al. (2018) amends this gap in the literature by applying the panel approach to groundwater extraction data from irrigated fields in Kansas (the same source as this thesis). They contribute a model of producer irrigation choice which includes both intensive and extensive margin components, from which the total marginal effect of climate conditions on groundwater extraction is derived. Their panel estimates suggest that irrigated producers respond most to shifts in three year monthly precipitation means and increases in the frequency of extreme heat days, when nonlinear temperature effects on yields are at their most severe. While their estimates of producer responses to climate and weather conditions are valuable for understanding the composition of irrigation choice, Oehninger et al. do not extrapolate these effects to predict climate change impacts and infer the adaptive capacity of irrigated agriculture. Their estimates are also subject to the same concerns inherent to the panel approach as discussed earlier, thus they do not adequately capture adaptation to climate change in irrigation.

In summary, the climate change impacts and adaptation literature consists of a diverse set of perspectives and approaches to understanding how climate change will interact with the complex natural and human systems which comprise the agricultural sector. Each of these studies has incrementally improved upon approximations and estimates of how producers respond to long-term shifts in climate patterns, yet exact identification is still an elusive and moving target. This thesis further contributes to this venture with a focus on adaptation in irrigation choices.

2.1 Identifying Adaptation: A Critique of the Long Differences Approach

The first step to identifying the occurrence and extent of adaptation to climate change is to actually define what adaptation is in the context of agricultural production. While a standardized definition doesn't exist, preceding studies share a similar perspective that adaptation in the agricultural sector encompasses how producers adjust their production inputs, choices, and other behaviors in an effort to reduce both the expected and realized negative impacts of climate change. Given this definition, adaptation in agricultural outcomes can theoretically be identified by observing changes in producer behaviors and outcomes after exposure to climate change, holding all else equal. In practice, it is impossible to observe the complete set of actions producers take to mitigate the effects of climate change. Empirical approaches, like the Long Differences approach introduced in Burke and Emerick (2016), circumvent this by assuming that impacts of climate change are first realized in annual weather conditions, which then serve as the catalyst for all producer adaptation responses. If this assumption holds, then producer adaptation is approximately identified by solely estimating how climate change augments the relationship between changes in annual weather conditions and adjustments in producer behavior and outcomes. Empirical estimates generated from weather variables are also arguably more robust to confounding given that weather conditions are plausibly exogenous.

To further investigate how adaptation can be identified empirically, consider two population regression models of producer response to weather in two periods, before (b) and after (a) climate change:

$$y_b = \beta_b + \alpha z_b + \mu_b \tag{1}$$

$$y_a = \beta_a + \lambda z_a + \mu_a \tag{2}$$

The dependent variable y measures either an agricultural output (yield, revenue, etc.) or input use (irrigation, fertilizer, acres, etc.) affected by weather. In keeping with Burke and Emerick's notation, z represents mean short-term climate or weather conditions (e.g. temperature) in a period. For clarity, the quadratic term typically included with meteorological independent variables is omitted, which has no bearing on the following analysis. λ and α capture the relationship between weather conditions and the producer outcome of interest, i.e producer response to weather, in each period. In this case, adaptation to climate change is identified by $\lambda - \alpha$, the difference in producer response to annual weather conditions before and after climate change. If y is groundwater applied per acre, then this difference is expected to be positive (negative) if producers adapt to climate change by increasing (decreasing) irrigation on the intensive margin, and zero otherwise.

Empirically estimating $\lambda - \alpha$, particularly over the long-run, is a challenge given the slow pace of climate change over the past couple of decades relative to its expected acceleration in the future. We simply haven't observed agricultural outcomes (y_a) and weather conditions (z_a) "after climate change", and consequently cannot estimate (2) and λ directly for long-term changes in climate. The solution used in the Ricardian and panel approaches, whose objectives are more so the identification of climate change impacts than adaptation, is to estimate producer response to weather before climate change (α) then plug in climate model predictions as a substitute for z_a, extrapolating α into the future to forecast the impacts of climate change on the agricultural sector (y_a). As α is only estimated from short-term annual variation in weather conditions, the Ricardian and panel model specifications omit long-term adaptations in producer response to weather, potentially overestimating the impacts of climate change.

Burke and Emerick's Long Differences (LD hereafter) approach attempts to leverage this omission as a means to explicitly identify and quantify long-term adaptation in agricultural production. They observe producer outcomes and annual weather conditions in two five year panels separated by a two decades long period. The "long differences" in both mean producer outcomes and weather conditions between periods are then used in the following sample regression model:

$$\hat{y}_a - \hat{y}_b = \beta_1 + \beta_2(\hat{z}_a - \hat{z}_b) + e_a - e_b \tag{3}$$

Which can be rewritten to match Burke and Emerick's notation as:

$$\Delta\hat{y} = \beta_1 + \beta_2\Delta\hat{z} + \Delta e \tag{4}$$

Burke and Emerick argue that β_2, which they denote as the differential climate trend, estimates how producers respond to variations in weather conditions over a longer period wherein climate change has likely occurred. Thus, long-term adaptations to climate change are embedded, along with short-term producer responses, in the differential climate trend and $\mathbb{E}(\hat{\beta}_2) = \lambda$. By differencing LD estimates of β_2 and panel estimates of α that only capture short-term trends, Burke and Emerick contend that the LD approach approximately identifies $\lambda - \alpha$ and quantifies to what extent producers adapt to climate change over the long-run.

Burke and Emerick's main argument for identification hinges on the assumption that $E(\beta_2) = \lambda$. Differencing population models (2) and (1) reveals that this necessary assumption does not hold in theory and that the LD estimating equation is potentially misspecified.

$$\begin{aligned} y_a - y_b &= \beta_a - \beta_b + \lambda z_a - \alpha z_b + e_a - e_b \\ &= \beta_1 + \lambda z_a - \alpha z_b + \lambda z_b - \lambda z_b + \upsilon \\ &= \beta_1 + \lambda(z_a - z_b) + (\lambda - \alpha) z_b + \upsilon \end{aligned} \tag{5}$$

This result is reached by adding a pair of terms, which by themselves sum to zero, to the differenced population regression model, then rearranging and combining like terms to recreate the LD specification and derive the differential climate trend. Following these steps, it becomes apparent that the LD estimating equation omits a relevant factor, $(\lambda - \alpha)z_b$, which is present in the population regression model. Thus, the sample estimate of β_2 is potentially misspecified with $\mathbb{E}(\hat{\beta}_2) = \lambda + (\lambda - \alpha)\frac{cov(z_b,(z_a - z_b))}{var((z_a - z_b))}$. If the long difference in weather conditions $z_a - z_b$ is not independent of short-term variation in weather z_b, then the necessary assumption for the LD estimate to identify long-term adaptation is violated. Since $z_a - z_b$ is clearly a function of z_b, identification eludes the LD approach. It appears that β_2 actually identifies something in between α and λ, suggesting that the LD approach underestimates producer responses to weather conditions after long-term climate change. This provides an alternative explanation for the limited adaptation observed in Burke and Emerick (2016). Estimates of adaptation generated by the matching with differences-in-differences approach proposed in this thesis are not subject to this same specification error, as adaptation is identified using estimated treatment effects in a natural experiment setting.

3 Estimating Adaptation In Irrigation: A Natural Experiment

This thesis follows Burke and Emerick (2016) and preceding adaptation literature in seeking to identify ex post causal effects of climate change on agricultural production outcomes and adaptation using observational data, albeit with agricultural inputs instead of outputs. While methodologies and approaches vary, each subject to different empirical concerns related to eliminating endogeneity and misspecification errors, the common objective is to approximate a randomized experiment which compares the outcomes of treated and control groups of agricultural producers, where the treatment is exposure to climate change. Burke and Emerick suggest that "an ideal but impossible experiment would observe two identical Earths, gradually change the climate on one, and observe whether outcomes diverged between the two". This example essentially describes the comparison between populations models discussed above, just in an experimental setting. However, this design is not feasible in practice because of the fundamental problem of causal inference with observational data (Holland, 1986). For each unique experimental unit, an irrigated field in this study, only a single outcome can ever be observed. Agricultural producers are heterogeneous, as is exposure to climate change over time, thus locating identical treated and control subjects to develop direct counterfactuals which identify adaptation is impossible. These same identification issues render estimates and predictions of treatment effects and climate change impacts from cross-sectional and panel models subject to omitted variable concerns, also discussed above.

Statistical matching methods present an alternative approach to improving the identification of adaptation responses to climate change in a natural experiment. Matching methods serve to further replicate the randomized experimental design by generating treatment and control groups from ”matched” experimental units who are relatively similar in type and characteristics, but not necessarily identical. This approximates a solution to the causal inference problem by generating a counterfactual, the matched control subjects, to compare with the outcomes of the treated. In essence, the matching procedure allows us to observe both $y_a - y_b$ and $z_a - z_b$ and estimate $\lambda - \alpha$ as a treatment effect.

This paper adopts the matching approach with a large data set of groundwater withdrawals for crop production in Kansas, primarily in the western portion of the state which lies above the High Plains Aquifer. Irrigated producers in the region are exposed to different intensities of climate change over space and time. Evapotranspiration (ET), which measures how much water a crop needs given heating conditions, and precipitation are the fundamental weather variables which determine irrigation choice. Their difference composes a direct measure of irrigation demand, i.e. how much water a producer should apply to maximize their crop yield given the natural provision of precipitation. In Kansas, climate change has led to increases in evapotranspiration as growing season temperatures rise, while the impact on the frequency and intensity of rainfall events varies spatially with more arid regions becoming drier (Southern and Western Kansas) and others getting wetter (Northern and Eastern Kansas). In combination, shifting ET and precipitation patterns alter the irrigation demand gap across the state.

The purpose of this thesis is to exploit variation in the intensity of climate change impacts on irrigation demand across Kansas to generate reliable counterfactuals, comparing the irrigation behavior of matched producers with contrasting exposures to climate change (i.e treatment). To conceptualize this approach in the context of groundwater extraction in Kansas, consider the following model of producer response to irrigation demand with (c) and without (t) exposure to climate change:

$$GW_c = \beta_c + \zeta ID_c + \mu_c \tag{6}$$

$$GW_t = \beta_t + \tau ID_t + \mu_t \tag{7}$$

Here, GW is groundwater extraction and ID is irrigation demand. In theory, the difference $\tau - \zeta$ captures how producers adjust their groundwater extraction choice in response to changes in the mean and variability of irrigation demand induced by climate change, defined as adaptation. In practice, we cannot observe states c and t for the same producer due to the fundamental problem of causal inference. Without this counterfactual, it is impossible to meaningfully estimate $\tau - \zeta$. A second best solution is to compare the responses of distinct sub-groups of producers who were exposed to different intensities of climate change (treated and control). This solution requires that ζ, the response to irrigation demand without climate change, be identical for both treated and control groups, as any pre-existing differences in groundwater extraction would be erroneously attributed to producer adaptation embedded in $\tau - \zeta$. Given the large set of water use observations for Kansas, this requirement is approximated by using matching methods to filter sub-groups to the extent that producers in the treated and control groups should exhibit nearly identical values of ζ, without significantly reducing statistical power. If matching is implemented with care and transparency, then $\tau - \zeta$ arguably identifies producer adaptations to long-term variation in irrigation demand. This approach is not solely applicable to adjustments in groundwater extraction on the intensive margin, but also to adjustments in crop choice, irrigation technology, and other water-used related adaptation strategies.

Following the matching procedure, this thesis estimates the $\tau - \zeta$ treatment effect for different climate change treatment definitions using a basic difference-in-differences (DID) panel analysis with fixed effects. The DID specification allows for additional control of pre and post treatment differences between control and treated groups. Binary treatments, which have received the most attention in the econometric and matching literature, are used to simplify the estimation procedure.

The following sections discuss in detail (i) the groundwater extraction and weather data, (ii) treatment definitions and matching implementation, (iii) econometric specification for estimating treatment effects and adaptation responses along the intensive and extensive margins.

4 Data

Observations of individual groundwater well extraction in Kansas are supplied via the Water Information Management and Analysis System (WIMAS), a joint project between the Kansas Geological Survey and the Kansas Department of Agriculture, Division of Water Resources. WIMAS is publicly available and a familiar source of GIS data in literature studying the economics of groundwater extraction (Hendricks and Peterson, 2012; Pfeiffer and Lin Lawell, 2014). The unit of observation is a water right for groundwater irrigation covering a single field. Water rights holders in Kansas are required to report metered water use annually in acre-feet. For parity with weather variables, water use is converted to acre-inches (multiplied by 12). In addition, irrigated producers report irrigated acreage, crop grown, and irrigation technology for each water right owned.

Since producers often allocate portions of an irrigated field and water right to separate crops and irrigation technologies, exact acreage allocation is included in their annual report. Observation at the water right or field level enables this study to account for more variation in extensive margin outcomes of climate change adaptation in the form of irrigated acreage and irrigation technology choices. The collection of water rights from WIMAS composes an unbalanced panel of groundwater irrigation in Kansas from 1991 to 2014. As water rights alternate between inactive and active, are transferred between producers, or have wells decommissioned or installed, observations for some years may be missing. Yet, the majority of water rights are observed through the full 24 year period, with a median of 22 years.

Daily observations of total precipitation, minimum and maximum temperature, vapor pressure, solar radiation, and duration of daylight for each water rights' geographic coordinates are queried from NASA's publicly available Daymet data set. These weather parameters are interpolated across 1 km x 1 km square grids spanning the entirety of North America. Daymet climatological data and algorithms were developed specifically for agricultural applications like crop modeling and other ecosystem process simulations, and provide the parameters and inputs required to calculate evapotranspiration. This study calculates reference evapotranspiration (ET_0) using the standardized ASCE Penman-Monteith equation for a shortgrass reference crop. The Daymet data set does not include a measure of daily wind speed, which is assumed to be a constant 3.25 m/s, noting that wind speed is negligible in ET_0 calculations at the annual level. Irrigation demand, the variable used to capture climate change in this study, is calculated as the difference between ET0 and total precipitation for each day in inches. All weather parameters are aggregated to annual growing-season measurements, which are merged with annual water use from WIMAS for each water right. For the purpose of this study, the growing season is assumed to begin April 1st and end on September 30th for all of Kansas and for each year in the 24 year period.

After merging WIMAS and Daymet data sets and removing observations with missing crop and technology information, the unbalanced panel contains 16,315 unique fields. Figure 1 maps out all fields included in the study, before applying the matching process. Figure 2 visualizes the mean annual trend in irrigation demand over the climate change treatment period for each of these fields. Table 1 presents summary statistics for all relevant panel covariates and dependent variables (intensive and extensive margins) for the 1991 to 1995, 1996 to 2009, and 2010 to 2014 periods. Table 2 presents the same summary statistics but restricted to fields primarily producing corn (share of irrigated acres > 0.5).

5 Treatments

One of the first steps in designing any experiment, randomized or natural, is to define a treatment which identifies the causal effects of interest. For identifying adaptation responses in irrigation, an appropriate treatment should account for structural shifts in irrigation demand associated with a warming climate. Irrigation demand is preferred because it captures both evapotranspiration and precipitation shocks, which often have opposing effects on producer behavior, in a single measure to define a binary treatment. This thesis uses a positive linear trend in irrigation demand over a 14 year ”in-treatment” period to indicate exposure to unfavorable climate change. Specifically, fields which experienced an annual increase in irrigation demand above $\frac{1}{14}$ of an inch from 1996 to 2009 are assigned to the treated group, while those which were not exposed to an increase in irrigation demand are assigned to the control. Fields with a trend between zero and $\frac{1}{14}$ of an inch are dropped from the analysis to create some degree of separation between the experimental groups. The choice of the $\frac{1}{14}$ of an inch cutoff level is entirely subjective, chosen for ease of interpretation (adds up to a 1 inch average increase in irrigation demand over the in-treatment period) and generating large sample groups with well balanced covariate distributions from the matching process. The span of the in-treatment period is also subjective, chosen to allow for equal length pre and post treatment periods and provide enough time for long-term adaption responses to arise. Thus, the complete 1991 to 2014 study period is separated into three stages: pre-treatment from 1991 to 1995 (5 years), in-treatment from 1996 to 2009 14 years (14 years), and post-treatment from 2010-2014 (5 years). Observations during the in-treatment period are only used in assigning the irrigation demand treatment and are dropped in the analysis.

A second treatment is also assigned which mirrors the irrigation demand treatment, but is restricted to fields which primarily grew corn throughout the duration of the study period. Corn is one of the more water intensive and abundant cash crops historically grown in Kansas and across the rest of High Plains aquifer. This restriction should illuminate how a significant subpopulation of producers engage in potentially different climate change adaptation strategies. For clarity, the first treatment is labeled the ”All Crops” treatment, while the second is labeled ”Corn Restricted”. No change is made to the actual treatment definition, as exposure to climate change is still represented by a $> \frac{1}{14}$ mean annual increase in irrigation demand, but fields are rematched. Denoting each sample as a unique treatment serves as reminder for the reader that the Corn Restricted matched sample is not a subset of the All Crops matched sample.

6 Matching

Once fields are assigned to treatment and control groups, the matching procedure ensures their post-treatment outcomes are comparable enough to plausibly identify adaptation responses. As noted in an extensive body of literature, matching methods only generate valid counterfactuals and identify the causal effects of treatment when certain assumptions hold, conditional on the exogenous assignment of treatment and homogeneity of treated and control groups prior to treatment (Imbens and Wooldridge, 2009). These assumptions reduce to ensuring that treated and control groups of producers differ only in climate change exposure. This section first argues that each of these necessary assumptions is valid in the context of climate change adaptation and groundwater extraction in Kansas, then discusses the specific implementation of the matching process used in this thesis.

First, the unconfoundedness assumption requires that exposure to climate change and treatment is independent of all unobservable factors or covariates which influence both irrigation behavior and adaptation, in essence approximating random treatment assignment wherein selection bias is not a concern. If an important covariate correlated with treatment assignment is omitted in the matching of treated and control groups, then separating the true treatment effect or adaptation response from the impact of the omitted factor would be impossible. Treatment assignment as defined in the previous section is plausibly unconfounded or exogenous because irrigation demand, precipitation, and ET are approximately random. Producers have no impact on observed weather conditions, nor their exposure to climate change unless they were to move or shut down their operations. These more extreme adaptation strategies are outside the scope of this study. The unconfoundedness assumption cannot be tested directly, but results from an indirect test with placebo treatment assignment are included as an additional counterfactual.

The second necessary assumption is the Stable Unit Treatment Value Assumption (SUTVA). SUTVA requires that irrigation behavior and adaptation responses for each producer are not influenced by their neighbor's exposure to climate change. This assumption is more challenging to argue for, given the spatial externalities of groundwater extraction. However, well spacing is often regulated to the extent that externalities from neighboring groundwater extraction are modest and relatively homogeneous across an aquifer (Pfeiffer and Lin, 2012). Therefore, SUTVA is assumed to hold over the study period and area.

A final, though not entirely exhaustive, necessary assumption requires that all irrigated fields in the matched sample have a positive likelihood of being exposed to climate change, typically measured with a propensity score. In simpler terms, if an irrigated field has zero likelihood of being exposed to climate change, then it cannot serve as a valid counterfactual for identifying adaptation and should not be matched with any treated fields. This ”overlap” assumption is approximated by selecting treated and control groups whose covariate distributions overlap to a significant degree, referred to as balance in the matching literature.

Given that the unconfoundedness and SUTVA assumptions hold, the only objective of the matching procedure as implemented in this thesis is to generate balanced treatment and control groups so that the overlap assumption holds, approximating the requirement that ζ in the groundwater extraction model is nearly identical between groups. Since the set of important pre-treatment covariates is relatively small in the context of the natural climate change experiment, irrigated fields are matched according to their full set of observed covariates, instead of just their propensity scores. In general, covariate matching methods try to minimize the difference or distance between each field's individual vector of covariates and the vector of covariate means for the sample distribution, typically denoted as the mahalanobis distance in the literature. If the matching process is effective, then treatment and control fields should exhibit a significant degree of parity in climate conditions and production characteristics during the pre-treatment period, which is essential for approximating the overlap assumption and ensuring that estimates of $\tau - \zeta$ identify the causal effects of the climate change treatment. In addition, estimates generated from covariate matched samples are also typically more robust to confounding when unobserved covariates are correlated with those that are observed because the matching process controls for them to some extent as well (Stuart, 2010).

The matching procedure implemented in this thesis follows the genetic multivariate matching approach introduced by Diamond and Sekhon (2013) and carried out it in R using the regenoud and MatchIt packages. Their iterative method searches for successively better balanced covariate distributions, matching each treated field with the statistically nearest control field. Thus, the genetic method preserves all treated observations, but trims outlying control observations. Sekhon's algorithm generalizes a measure of malahanobis distance by incorporating weights for each covariate to optimize balance. Weights are applied to fields in the control group and serve as a condensed estimate of how well a control field matches with the mean treated field. These weights are included in subsequent regressions to improve specification. The genetic approach also attempts to balance quadratic and interaction terms for each unique covariate. This covariate matching process is applied separately for both the All Crops and Corn Restricted treatments as each treatment generates treatment and control groups with different covariate distributions that need to be rebalanced.

Tables 3 and 4 present covariate distribution summary statistics for the All Crops and Corn Restricted treatments, respectively. All covariates listed were used in the matching process. Both tables report treated and control group covariate means during the pre-treatment period before and after applying genetic matching to trim the control group for each treatment. The last column of each table presents a measure of the percent improvement in mean differences between treated and control groups after matching. Note that covariate means are quite similar between each group after matching, evidence that covariate distributions are relatively well balanced and unlikely to violate the overlap assumption. For visual inspection, Figures 3 compares the pre-treatment distribution of precipitation for treatment and control groups (All Crops treatment), before and after matching. Additionally, Figures 4 and 5 present maps of fields included in each matched sample separated by treatment status.

7 Econometric Specification

Estimates of adaptation to climate change on the intensive and extensive margins are derived by comparing pre and post treatment differences in irrigation behavior and land use between treated and control groups, approximately identifying $\tau - \zeta$ for various irrigation-related dependent variables. Matching with binary treatments simplifies the estimation, permitting the application of difference-in-differences (DID) regression models using a panel of post and pre treatment observations. The empirical strategy explicitly accounts for fundamental differences in adaptation on each margin.

On the intensive margin, producers irrigate to close the gap between crop water requirements and precipitation during the growing season. The model specification accounts for this behavior by including interactions with treatment indicators for in-season precipitation and controls for ET0:

$$
\begin{aligned}
y_{it} = & \beta_1 Preci_{it} + \beta_2 Trt_{it} \times Preci_{it} + \beta_3 Post_{it} \times Preci_{it} \\
& + \delta_1 Trt_{it} \times Post_{it} \times Preci_{it} + \beta_4 ET0_{it} + \alpha_{ip} + \sigma_t + e_{it}
\end{aligned} \tag{8}
$$

Again, the unit of observation i is a field endowed with a water right for irrigation spanning most or all of the 1991 to 2014 period, with year observed denoted by t. Total acre feet of groundwater pumped and acre feet applied per irrigated acre each serve as a measure of intensive margin response (y) to weather conditions. Trt is 1 if the field experienced an average increase in irrigation demand of more than $\frac{1}{14}$ of an inch during the 1996 - 2009 treatment period, 0 if the trend in irrigation demand was non-positive. $Post$ is 1 for years in the 2010 - 2014 post-treatment period, 0 otherwise. $Preci$ and $ET0$ are mean in-season precipitation and reference evapotranspiration (inches), respectively. Precipitation and ET are included as separate independent variables, as opposed to being combined in a measure of irrigation demand, to control for their separate and unique impacts. For example, producers may weight ET and precipitation differently when making irrigation decisions, assigning more importance or confidence to one over the other. If the model was estimated solely for irrigation demand, then estimates could be biased according to the same argument for misspecification as in the LD approach. Individual responses to weather conditions (i.e. coefficients) for treatment and control groups in pre and post periods are estimated using the within transformation for each 5 year pre and post treatment period. Given this two-way fixed effects approach, intercepts for treatment status and period are perfectly collinear with the field-period and year fixed effects (α_{ip} and σ_t) and are consequently dropped from the specification.

δ_1 captures the difference in intensive margin response to annual in-season precipitation between treated and control groups after treatment (i.e. the average treatment effect on the treated or $\tau - \zeta$ in the earlier model of groundwater extraction), as long as the parallel trend assumption holds. The matching procedure and fixed effects approach control for most pre-treatment and time-invariant factors which could violate the parallel trends assumption. Figures 6 graphs mean trends in total and per acre water use by treatment status. Simple visual inspection shows that each variable exhibited relatively parallel trends from 1991 to 2014. Thus, δ_1 can be interpreted as intensive margin adjustment or adaptation to the negative impacts of climate change on irrigation demand. A similar estimate for response to ET could be constructed with additional DID interaction terms, but the main specification favors only using ET as a control and interpretation of its impact is ignored in the results. This is due to measurement error concerns, which are discussed further in a later section.

Also note that other observed covariates that are typically correlated with both water applied per acre and weather conditions, like crop type and irrigation technology, are not included in the intensive margin specification. With a typical cross-section or panel approach, omitting these covariates could severely bias estimates of adaptation. However, this is not the case for approaches that approximate the design of a randomized experiment wherein post-treatment bias is a concern (Montgomery et al., 2018). In fact, including any of these covariates would actually bias the treatment effect estimate δ_1, essentially altering its interpretation to the extent that it no longer identifies adaptation. This is because almost all of the observed covariates are also outcomes of the treatment and including them in the DID specification needlessly soaks up variation in the intensive margin response which should be attributed to the treatment alone. For example, exposure to climate change can induce changes in crop choice post-treatment. If crop choice was included as a control, then δ_1 would only capture the partial effect of climate change on intensive margin responses which is not attributed to adjustments in crop choice. The objective of this study is to estimate the full effect of climate change separately for each margin, thus observed covariates which are also outcomes of the treatment are not included in the estimating equations, but do serve as separate dependent variables. Precipitation and ET are also arguably subject to post-treatment bias, as they are directly impacted by climate change, but are necessary for structurally modeling the intensive application of irrigation water.

Estimation of treatment effects on the extensive margin is considerably less complex, as the specification only requires the treatment period interaction term:

$$w_{it} = \beta_1 Trt_{it} \times Post_{it} + \alpha_i + \sigma_t + e_{it} \tag{9}$$

The treatment effect β_1 is independent of in-season weather conditions due to the defining characteristics of extensive margin irrigation responses, that they occur over long periods and in response to long-term trends in weather or climate forecasts. Therefore, the irrigation demand treatment alone reasonably identifies adaptation to climate change in land and technology use. The set of extensive margin dependent variables (w_{it}) includes irrigated and crop acreage, as well as acres allocated to specific types of irrigation systems. The interpretation of the treatment effect is standard and estimates the difference in treated and control groups' extensive margin quantities. The intensive and extensive margin models only differ in their estimating equation and set of dependent variables. Only field and year fixed effects are included for the extensive margin specification. The matching procedure is applied equally to both models and estimates are generated from the same treatment and control groups.

The primary objective of both models is to test the hypothesis that producers adapt to climate change over the long-run, which is accomplished with t-tests for the statistical significance of treatment effects for each dependent variable and on each margin. Following findings from recent studies of DID estimation in practice, special consideration must be given to eliminate bias in standard errors due to serial correlation and clustering (Bertrand et al., 2004; Abadie et al., 2017). In particular, treatment assignment in this study could be clustered at a level above the individual field as trends in irrigation demand are potentially shared between fields within the same section, township, county, or groundwater management district. This is the experimental design issue discussed in Abadie et al. (2017). If this is the case, then standard errors would underestimate the standard deviation of coefficients, potentially leading to erroneous inferences which suggest that producers adapt to climate change when in reality they do not.

Given that caution in standard error estimation is warranted, this study conducts hypothesis testing with more conservative cluster-robust standard errors. Deciding which level of clustering is appropriate is slightly more challenging, requiring a balance between eliminating bias and using too few clusters. Clustering by groundwater management district could be unnecessarily conservative as there are only 6 groups (5 actual districts plus the group of fields that aren't included in a district), and irrigation demand would reasonably be expected to vary substantially within these geographically large groups. In contrast, clustering by PLSS section would have little effect on bias as a single field can cover an entire section. Thus, the choice is between clustering by county or township. As precipitation can vary significantly within a single county, clustering by township is arguably a more appropriate choice. Therefore, all statistical inference in this study is done with standard errors clustered by PLSS township.

7.1 Intensive Margin Residual Analysis

While the intensive and extensive margin models provide estimates of climate change adaptation through direct adjustments to irrigation and land use choice, they ignore more subtle, long-term adaptation mechanisms wherein producers apply water more efficiently without making drastic and expensive changes to production. For example, producers may shift their irrigation schedule to apply more water during cooler nights and cloudy (but dry) days when evaporation is reduced and more water infiltrates into the soil. In this case, the producer reduces their irrigation per acre over the long-term regardless of in-season precipitation and evapotranspiration, though scheduling would still be correlated with weather conditions. Observations for factors like irrigation scheduling are not readily available for the Kansas study area. But because these factors are relatively time-invariant, they are likely embedded in group fixed effects.

For the intensive margin model, these subtle adjustments in total water use are captured by field-period and year fixed effects (α_{ip} and σ_t), as well as in idiosyncratic error (e_{it}). These fixed and omitted factors accumulate in the constant/intercept term estimated by (8), which differences out of the model as a casualty of the within transformation. However, estimates of intercepts for individual-period combinations can be recovered by partialing out the estimated coefficients in (8):

$$\upsilon_{ip} = y_{it} - \hat{\beta}_1 Preci_{it} - ... - \hat{\beta}_4 ET0_{it} \tag{10}$$

Here, $\upsilon_{ip} = \alpha_{ip} + \sigma_t + e_{it}$ is the recovered intercept for each individual field-period combination. Again, y_{it} is a vector of water applied per acre observations. The fixed effects are of particular interest in the analysis, as they capture subtle adaptation behavior like adjustments in irrigation scheduling. Comparison of estimated intercept means for treatment and control groups before and after treatment identifies changes and adaptations in water use regardless of in-season precipitation and ET. $E[\upsilon_{treated,pre}] - E[\upsilon_{treated,post}]$ is the change in intensive margin residual for treated fields, while $E[\upsilon_{control,pre}] - E[\upsilon_{control,post}]$ is the same measure for the control group. Thus, $E[\upsilon_{treated,pre}] - E[\upsilon_{treated,post}) - (E[\upsilon_{control,pre}] - E[\upsilon_{control,post}])]$ measures the treatment effect, i.e the difference in intensive margin residuals between treated and control fields after exposure to climate change. This difference is estimated using another differences-in-differences model with field and period fixed effects which includes the recovered intercept estimates for each field-period combination from (8) as the dependent variable:

$$\upsilon_{ip} = \theta(Trt_{ip} \times Post_{ip}) + \phi_{ip} + \mu_{it}$$

The intensive margin residual differences-in-differences are estimated for the same matched samples used in the intensive and extensive margin models for both the All Crops and Corn Restricted treatments.

8 Results

Estimates of treatments effects are generated from matched samples as discussed above. For the All Crops treatment, 1,019 control fields are matched with 5,065 treated fields. The matched sample for the Corn Restricted treatment is considerably smaller, with only 52 control fields and 225 treated fields. This significant reduction is a result of both the aggressiveness of the genetic matching algorithm (configured with defaults) and that only fields which made no changes to crop choice over the entire 14 year in-treatment period are included, a substantial restriction. Estimates are first presented for the All Crops treatment, with Corn Restricted results to follow. Analysis for each treatment begins with estimates of adaptation on the intensive margin and concludes by exploring the distribution of intensive margin model residuals before and after climate change. In general, results show that irrigation behavior after exposure to increases in irrigation demand is not significantly different from irrigation behavior before exposure, implying that producer adaptation on the intensive and extensive margins is minimal. However, estimates suggest that producers may respond to some extent by reducing acres used to produce crops which are particularly sensitive to water and heat stress. Results from a placebo treatment are also put forward to serve as a robustness check and indirect test of the unconfoundedness assumption.

8.1 All Crops Treatment

Tables 5 through 10 report results for the All Crops treatment.

8.1.1 Adaptations on the Intensive Margin

Table 5 presents estimates of the differences in intensive margin response to mean in-season precipitation between treatment and control groups before and after treated fields are exposed to more than a $\frac{1}{14}$ of an inch annual increase in irrigation demand on average. The two columns report regression results for different dependent variables, each measuring a different intensive margin outcome. The first column uses total volume of groundwater applied (acre-inches), while the second uses volume of groundwater applied per acre irrigated (acre-inches per acre).

The point estimates for precipitation and its interaction (the first four estimates in a column) each have a typical differences-in-differences interpretation. The first coefficient captures the baseline response to precipitation of the control group during the pre-treatment period. The second estimates the pre-existing difference in response to precipitation between the treated and control fields. The third is interpreted as the change in the control group's response from the pre to post treatment periods. Lastly, the fourth coefficient estimates the treatment effect δ_1. Complete interpretations are provided only for the volume per acre results (2nd column), as they are nearly identical to those for the total volume regression results. The results for both dependent variables are also quite similar, differing primarily in the magnitude of their point estimates, which is expected given their definitions.

Estimates show that on average 0.25 fewer acre-inches of groundwater per acre was applied to control fields for every additional inch of annual in-season precipitation. This baseline estimate of intensive margin response to precipitation is statistically significant and reassuring, as its direction agrees with the common sense expectation that producers substitute irrigation for precipitation. The point estimate for pre-existing differences between treated and control is not statistically significant, which is consistent with the use of a matched sample wherein most pre-existing differences have been eliminated. For the estimated average treatment effect on the treated, the null hypothesis, which states that the difference in intensive margin response to annual in-season precipitation between treated and control groups after treatment is zero, cannot be rejected. Thus, the results do not suggest that crop producers have adjusted their irrigation choice along the intensive margin in response to climate change.

8.1.2 Adaptations on the Extensive Margin

Tables 6 to 8 present estimates of differences in extensive margin responses for the All Crops treatment. As discussed earlier, the econometric specification for the extensive margin is much simpler, including only a single independent variable to capture the average treatment effect. However, given the wider range of adaptation strategies and outcomes on the extensive margin, the model is estimated for a larger set of dependent variables, 7 in total. Table 6 reports the estimated treatment effect on total acres irrigated in a field. Table 7 uses field acres allocated to each of Kansas' five most abundant crops (Wheat, Corn, Soybeans, Sorghum, Alfalfa) and the last accounts for acres producing all other crop types. For all dependent variables, the treatment effects (Trt*Post) are interpreted as the post-treatment differences in treated and control fields' total irrigated acreage and acreage specifically allocated for each unique crop type.

In theory, producers may reduce total irrigated acres in response to climate change as a means to reduce total crop water needs, and subsequent groundwater extraction, while still ensuring that each individual plant receives enough soil moisture to maximize yield. This is the textbook example of adaptation on the extensive margin, i.e. adjusting the quantity of production units requiring inputs instead of adjusting the quantity of inputs applied per unit. In practice, Table 6 presents evidence that fields exposed to a long-term increase in irrigation demand did not appear to adjust acres irrigated downward, or even at all, given that the estimated treatment effect is not statistically significant.

Results for adjustments in crop acreage are mixed, with acreage for corn and alfalfa exhibiting the only treatment effects which are statistically significantly different from zero. Corn is a substantially water-intensive crop, typically requiring irrigation from a reliable source like groundwater to be profitable on land in more arid climates like

Western Kansas. Increases in irrigation demand combined with decreases in water tables escalate the costs and risks of corn production, potentially to the extent that it becomes profitable for some producers to switch to a less water intensive crop. Switching to a crop that is typically irrigated from a different water source, like surface water, could also plausibly increase profitability given vast shifts in climate patterns, as would ceasing crop production entirely. However, the dataset used in this study does not include surface water irrigated or fallowed fields, so more drastic adjustments on the extensive margin are not observed. The point estimate for corn acreage suggests that treated fields on average put approximately 21.4 fewer acres into corn production than control fields. This result agrees with the expectation that water-intensive corn production should decrease on lands where the frequency and effectiveness of precipitation diminishes with climate change. By contrast, treated fields on average put 7.6 more acres into alfalfa production. This result is surprising as alfalfa is also considered a fairly water-intensive crop, often more so than corn. However, alfalfa is typically more resilient to increases in the number of continuous days without precipitation given its deeper root system and other agronomic characteristics. Thus, producers may switch to alfalfa primarily as an adaptation to increases in the variability of irrigation demand caused by climate change, and not necessarily in response to shifts in mean weather conditions. Alternatively, producers may be exploiting the fact that alfafla uses less water at earlier growth stages to produce approximately the same yield as in later growth stages. Therefore, perennial alfalfa production has an additional margin or adaptation strategy that annual crops lack. It is important to note that the empirical strategy as implemented does not directly identify switches from one specific crop to another, thus we cannot infer from these estimates that producers switch directly from corn to alfalfa.

Table 8 reports results for adjustments in acres irrigated by different irrigation technologies. Technologies are grouped into four categories: Flood only, Flood and/or Center Pivot, Center Pivot Drop, and Other. Flood includes only acreage irrigated entirely by surface flood. Flood and/or Center Pivot includes acres irrigated by a combination of surface flood and standard center pivot, as well as those irrigated entirely by center pivot alone. The exact ratio of flood and pivot acres is not reported in the WIMAS database. Center Pivot Drop includes all acres irrigated entirely by center pivot with drop nozzles. The Other category accounts for less-common irrigation technologies and practices, including sub-surface drip and sprinkler. The first three columns of table 8 are ordered from least to most water use efficient. In theory, producers may adapt to climate change by switching to a more efficient irrigation technology, ensuring that more water gets to the plant without having to increase groundwater extraction. Upgrading irrigation technology is typically a very expensive endeavor, however it may become financial viable if increases in irrigation demand lead to rising crop prices.

Point estimates for irrigation technology investment treatment effects suggest that treated fields irrigated 13.5 fewer acres with surface floods and 23.6 more acres with some combination of flood and center pivot on average compared to control fields during the post-treatment period. The estimate for center pivot drop surprisingly suggests that treated fields used the more efficient technology on 13.8 fewer acres. However, none of these point estimates are statistically significant, implying that irrigation technology choice per acre was not adjusted in response to increasing irrigation demand. Similarly, the treatment effect for the compilation of all other observed irrigation technologies is insignificant. The summary statistics reported in Table 1 show that mean acreage allocated to these other technologies diminished over the study period, but results from the extensive margin model suggest that this trend is unrelated to changes in irrigation demand.

8.1.3 Intensive Margin Residuals

Table 9 presents estimated residuals recovered from the volume per acre intensive margin model containing estimates of field-period fixed effects and idiosyncratic errors, distributed by treatment status. Applying another DID model with fixed effects to estimate differences in residuals across these groups is certainly overkill, but provides a clustered standard error needed for robust statistical inference without much hassle. The estimated average treatment effect on intensive margin residuals is reported in Table 10 and suggests that mean irrigation per acre conditional on in-season precipitation and ET does not differ between treated and control fields with the All Crops irrigation demand treatment. Though the sign of the point estimate agrees with expectations that exposure to climate change causes producers to make subtle adjustments in groundwater extraction, like rescheduling irrigation to save water, the estimated effect is not statistically significantly different from zero.

8.2 Corn Restricted Treatment

Tables 11 through 15 report results for the Corn Restricted treatment. Treatment effects and coefficients share the same interpretation with the those from the All Crops treatment. The only differences are the restricted sample of matched fields and that treatment effects on crop acreage obviously cannot be estimated as included fields are limited to producing corn for the entire study period. As referenced earlier, corn is a groundwater irrigated crop with significant production in Kansas and other states which lay above the High Plains Aquifer. The restricted treatment attempts to highlight and amplify adaptation strategies beyond switching crops, which results from the All Crops treatment showed was the favored response to climate change for producers in Kansas over the 1996-2009 period.

8.2.1 Adaptations on the Intensive Margin

Table 11 reports results for the intensive margin response model using the restricted sample. Corn fields appear to be slightly less responsive to increases in precipitation, with only 1.79 fewer acre-inches of water applied per acre for each additional inch of precipitation. Corn is generally more susceptible to short-term water stress than other crops like wheat, soybeans, sorghum, and alfalfa. Therefore, risk-averse corn producers may be less responsive to precipitation because they are ”hedging their bets” and irrigating with less regard to realized and expected precipitation. The timing of when corn yields are most susceptible to stress could also play a role. The estimated treatment effects for identifying intensive margin adaptations are not statistically significant for the Corn Restricted sample. No difference in irrigation application response to precipitation for treated and control corn fields is observed.

8.2.2 Adaptations on the Extensive Margin

Table 12 reports the estimated treatment effect for irrigated acreage, the only remaining observed non-tech extensive margin response when switching crops is restricted. The result shows that exposure to a long-term positive trend in irrigation demand has no appreciable effect on acres irrigated as the point estimate is neither statistically or economically significant. Compared to other crops produced in Kansas, the more water-intensive corn crop arguably has more to benefit from adjustments in irrigated acreage made to ensure each individual plant's water requirements for maximizing yield are attained when water supplies are limited. Thus, the null result for irrigated acreage adaptation is fairly consistent amongst more and less water-intensive units of production.

Table 13 presents treatment effects for each of the four categories of irrigation technology used to irrigate corn in Kansas. The point estimates are considerably smaller for the Corn Restricted sample compared to the results from the All Crops sample. This could be attributed to diminishing marginal returns to irrigation efficiency and the fact that most fields producing corn at the start of the study period were already irrigated with center pivots (see Table 2) and are unlikely to downgrade to surface flood. The directions of each treatment effect agree with expectations from theory, but again none are statistically significant.

8.2.3 Intensive Margin Residuals

Tables 14 and 15 present Corn Restricted estimates of intensive margin residuals following the same setup as for the All Crops Treatment. They are the only estimates from the Corn Restricted treatment to exhibit significant differences associated with treatment status. Results suggest that treated fields on average applied 2.3 fewer acre-inches of water per acre regardless of in-season precipitation and ET compared to control fields after treatment. Caution is warranted in directly interpreting differences in residuals, but the estimated treatment effect suggests that corn producers engage in more subtle adaptation strategies, like shifting irrigation schedules, to meet increasing crop water needs without extracting more groundwater. Given that producers do not switch crops in the Corn Restricted treatment, this result serves as an example of how restricting the sample to production of a single crop amplifies the identification of alternative adaptation responses.

8.3 Robustness Checks

To indirectly test whether the treatment definition and assignment used in this study reliably identifies differences in irrigation behavior with exposure to climate change, a placebo treatment which should not elicit changes in behavior is assigned to all irrigated fields observed in pre and post treatment periods. With the placebo treatment, fields are randomly assigned (by coin flip) to treated and control groups. The matching process is reapplied and treatment effects are re-estimated following the exact same identification strategy used in the All Crops and Corn Restricted treatments. The re-estimated treatment effects then serve as reasonable counterfactuals to those estimated using the original irrigation demand treatments. Estimates for the placebo treatment are presented in Tables 16 through 21. Results show that fields treated with the placebo do not respond differently than control fields with regard to most intensive and extensive margin irrigation choices. The placebo treatment effect for alfalfa acreage is statistically different from zero, but the point estimate is not economically significant. Baseline irrigation per acre responses to precipitation and ET are statistically significant and the directions of their point estimates agree with expectations, showing that the model of intensive margin irrigation behavior is consistent with a larger sample of irrigated fields. The only anomaly is the statistically and economically significant treatment effect for irrigated acreage. However, this result is not entirely damning because a single iteration of random treatment assignment does not preclude spurious correlation with a randomly observed covariate. This anomaly would only be a concern if it is consistent across thousands of iterations of random treatment assignment. For the most part, results from this single iteration of the placebo treatment support the reliability of the identification and modeling strategies implemented in this thesis.

9 Discussion

The above results tend to agree with findings from studies discussed earlier which suggest that producer adaptation to climate change is limited. While these studies only investigated adjustments in agricultural yields and profits, it's not particularly surprising that water use would exhibit a similarly null result, as any major adjustments in irrigation behavior would have likely had an effect on irrigated yields. Potential explanations for the estimated deficit of adjustments in water use also overlap with those for agricultural outputs, as both channels for adaptation rely on producers recognizing the need to respond to climate change and having the capacity and right incentives to act. Additionally, the decomposition of water use into separate margins and outcomes permits a more in-depth analysis of alternative explanations unique to irrigated agriculture. See Burke and Lobell (2010) for an extensive review of the factors which influence to what extent producers adapt to climate change through adjustments in irrigation and other production choices.

Adaptation in irrigation choice along the intensive margin could be suppressed for the Kansas study area due to a salient set of physical and institutional constraints on groundwater extraction. In terms of physical constraints, the depletion of portions of the High Plains Aquifer in Western Kansas, which may accelerate with concurrent decreases in natural recharge due to climate change (Scanlon et al., 2012; Meixner et al., 2016), is a significant threat to groundwater irrigated crop production in the region. Water tables are projected to further decline over the next 15 to 20 years in the same areas where saturated thickness has already diminished considerably (Steward et al., 2013).

Increasing groundwater extraction may not be a feasible adaptation strategy in these areas as pumping costs increase and groundwater wells inch towards a ceiling on well yields, eventually losing the capacity to maintain stable and sufficient extraction rates.

Groundwater sustainability concerns given past and current extraction trends have provided the impetus for local regulations and the establishment of Kansas' Groundwater Management Districts (GMDs) in 1972. The presence of a regulator and expectations of future regulation alone may have dissuaded producers in the region from relying on any remaining capacity to increase groundwater extraction in response to increases in irrigation demand. In addition, the structure and implementation of water rights across the state has potentially disincentivized adjustments in groundwater extraction following exposure to climate change. In particular, Kansas has historically practiced ”use it or lose it” water rights wherein producers forfeit any portion of their groundwater extraction allocation that is not used over a multi year period. Thus, producers have an incentive to irrigate to their full irrigation allocation, regardless of whether or not effective precipitation is sufficient to meet crop water needs. Producers are also not allowed to carryover any remaining allocation amounts to future periods, intensifying the perception that any unused allocation is ”lost”. Considering this implementation of water rights, producers over the study period and area faced strict institutional constraints which likely limited their capacity to adapt on the intensive margin. In fact, incentives may have aligned in a manner that encouraged producers to maintain the same level of groundwater extraction throughout the study period, solely in accordance with the allocation level established in their water right.

While these physical and institutional constraints are likely to limit adjustments in groundwater extraction along the intensive margin, they may have the opposite effect on extensive margin adjustments wherein producers can potentially adapt to climate change while sustaining the allocation level established in their water right. Results show that producers favored the extensive margin strategy of adjusting crop acreage according to how different crops vary in susceptibility to short-term water stress, effectively reducing their vulnerability to increases in the variability of precipitation. Yet, this adjustment was limited to only a pair of crops, and producers did not appear to reduce irrigated acreage or upgrade to more efficient irrigation technologies, implying that adaptation along the extensive margin is also relatively scarce in Western Kansas. As discussed by Burke and Lobell (2010) and Burke and Emerick (2016), this apparent dearth in acreage and technology adjustments could be attributed to conflicting incentives in producer profitability and risk management. For instance, the price of water intensive crops may increase with regional changes in climate, potentially making these crops more profitable even if producers cannot fully satisfy their crop water needs and experience some losses in yield. Similarly, the significant cost to upgrade irrigation technologies may exceed their yield loss mitigating benefits. Thus upgrading may not be profitable, even with increases in prices for water intensive crops. Producer capacity to upgrade irrigation technologies could also be constrained by field characteristics, particularly for those with hilly terrains. Finally, government programs like subsidized crop insurance, may disincentivize adaptation along the extensive margin because insured producers are required to maintain the same irrigated acreage throughout the entire growing season or risk receiving a smaller insurance payment in the event of a yield loss. Insured producers could also conceivably respond to climate change induced yield loss in the short-term by making insurance claims instead of engaging in costly adaptation along the extensive margin. However, this strategy is only viable if insurers fail to accurately predict the effects of climate change and do not update insurance policies accordingly.

While results show that adaptation is limited along the intensive and extensive margins, the outlook is more optimistic for adjustments in mean overall groundwater extraction. In particular, corn producers exposed to increases in irrigation demand appear to apply irrigation per acre more efficiently to meet crop water requirements, regardless of realized in-season precipitation and ET. Corn production is a microcosm for such adjustments in mean irrigation applications because optimal corn growth and development requires a consistent supply of water. As climate change decreases the frequency of effective precipitation and limits increases in groundwater extraction, corn producers are compelled to make the most of every drop of water they can muster from their current water supply. Interestingly, this result suggests that corn producers were not initially operating at the efficient frontier of their crop-water production functions before exposure to climate change, as they wouldn't have had the capacity to apply irrigation water more efficiently otherwise. Thus, corn producers appear to have had additional opportunities for adaptation along the extensive margin beyond adjusting acreage choice. Restricting the treatment only to producers that kept producing corn after climate change highlights these additional opportunities. While these findings only apply to estimates from corn production, it is reasonable to expect similar subtle groundwater extraction adjustments for crops that are more resilient than corn in the short-term, but will eventually succumb to the effects of climate change. Finally, adaptations in total irrigation regardless of short-term weather conditions may not be subject to the same constraints as adaptations in acreage and technology choice. For example, rescheduling irrigation to cooler and less windy days is a much cheaper approach to improving irrigation efficiency than installing another center pivot. If subtle reductions in mean overall groundwater extraction are relatively inexpensive and benign in terms of yield loss, such adjustments could potentially explain why alternative adaptations on the intensive and extensive margins were not observed over this short study period relative to the slow pace of climate change.

9.1 Limitations and Extensions

It is vital to note that these results are subject to important caveats and limitations that arise from the matching with DID model approach, particularly concerning some trade offs between internal and external validity. This thesis has argued in favor of using matching methods which approximate a randomized experiment to generate estimates of adaptation which are reasonably more robust to bias and specification error, i.e more internally valid. However, this supposed increase in internal validity comes with a significant cost to external validity because the matching procedure introduces sampling bias. After the matching procedure trims the control group to include only irrigated fields that are nearly identical to the mean treated field, the matched sample no longer captures all the heterogeneity present in the true population of interest. Estimates generated from the matched sample are then only applicable to irrigated fields that share the same covariate distribution as the treated fields in Western Kansas. As irrigated field characteristics vary considerably not only within Kansas but across the U.S., the treatment effects estimated by the matching approach are likely not generalizable to larger, more interesting populations. This concern prohibits this thesis from forecasting the adaptation conditioned impacts of climate change on irrigation, as in previous studies. In contrast, Burke and Emerick (2016) use a sample containing nearly all corn and soybean producing counties east of the 100th meridian. While this thesis showed that Burke and Emerick's estimates are not internally valid, they are certainly more externally valid and generalizable to a larger population of interest.

Another important limitation related to external validity concerns whether or not a positive annual trend in irrigation demand over a 14 year period adequately represents the effects of long-term climate change. Climate change is a complex process which this thesis simplified to a single binary treatment based solely on how climate change effects the difference between mean annual precipitation and reference evapotranspiration. This definition completely ignores mean preserving changes in the variability and intensity of weather conditions, potentially underestimating both the effects of climate change and subsequent adaptations in irrigation behavior. The finding that producers switched to alfalfa, a crop that is water intensive but less vulnerable to increases in the variability of precipitation, suggests that producers may be more responsive to changes in weather variability than changes in changes in the mean, potentially highlighting the inappropriateness of the binary treatment definition. A continuous or multivalued treatment variable could potentially capture the effects of changes in both the central tendency and variability of weather conditions.

Another argument for the inappropriateness of the irrigation demand treatment in the context of climate change is that the change in mean irrigation demand is economically insignificant. For irrigated fields to be considered treated, they must have been exposed to more than a 5% increase in mean irrigation demand over a 14 year long period. On an annual basis, this amounts to a relatively modest $\frac{1}{14}$ inches per year. One could argue that this is such a small amount that producers may fail to recognize the trend, thus explaining why most of the observed treatment effects are insignificant. However, this problem is unrelated to changes in irrigation demand not adequately representing climate change, but rather due to the considerably short treatment period relative to the pace a climate change. The 14 year treatment period likely only captures short-term climate change, not long-term. This limitation is imposed by the availability of groundwater extraction data which, as of this thesis, is only available from 1991 to 2014 for Kansas. A future study could reestimate irrigation demand treatment effects with updated data that spans a longer period. These updated estimates could then be compared with this study's original estimates as an alternative approach to identifying long-term adaptation.

While the matching procedure and treatment definition as implemented are likely deficient in externally validity, this thesis has argued that the quasi-experimental approach produces more internally valid estimates of adaptation to climate climate change in irrigation. However, as with all natural experiments, there are likely many unobserved factors that have the potential to confound estimates of treatment effects, particularly those that are non-randomly distributed between treatment and control groups. Figures 4 and 5 clearly show that treated and control groups for both the All Crops and Corn Restricted treatments are not randomly distributed across Kansas. Specifically, control fields are clustered around the northwestern and central portions of the state, while almost all treated fields are located in the southwest. Any factors that disproportionately impact irrigation behavior in one region compared to its neighbors could confound estimates of adaptation responses. Notable examples of such factors in Kansas include disparities in aquifer saturated thickness and water tables between regions, variation in groundwater management regimes, and the impact of large-scale irrigation on local and regional climates. In particular, the control and treated groups are mostly situated in separate Groundwater Management Districts which differ in their management objectives and the reliability of their annual groundwater extraction reports. In addition, the covariate balance statistics in tables 3 and 4 show that treated fields tended to irrigate more than matched control fields in the the pre-treatment period. Given recent evidence that large-scale irrigation may induce changes in local evapotranspiration and precipitation rates (Szilagyi, 2018), the gap in the intensity and scope of irrigation between groups could lead to differences in local climates for treated and control fields which could confound the irrigation demand treatment over the study period. The magnitude and direction of bias generated by these omitted factors are difficult to estimate given ambiguity in how each factor interacts with irrigation choice. If these interactions are economically significant in the context of groundwater irrigated production in Kansas, then adaptation estimates may lack internal validity unless omitted factors are substantially correlated with observed covariates used in the matching process.

Suggested extensions of this study's approach and findings involve the development and implementation of treatments that better represent the effects of climate change, as well as applying these treatments to additional producer outcomes and adaptation strategies. An alternative climate change treatment that is defined with regard to how producers perceive and recognize the need for adaptation, instead of solely capturing the occurrence of climate change, would better illuminate long-term strategies that only prevail after severe climate change damages are realized. In addition, this study has ignored other channels for adaptation in irrigation like rotating crops to increase soil water storage, switching to drought tolerant crop varieties, or moving planting dates. Factors which influence the profitability of certain adaptation strategies, like harvest prices and technology costs, are also omitted. Thus, a future study could apply a more well-defined long-term climate change treatment with a wider range of agricultural outcomes, choices, and factors to improve the identification of producer adaptation to climate change.

10 Conclusion

This thesis develops and implements a natural experiment to generate estimates of producer adaptation in irrigation choices along the intensive and extensive margin. This experimental setting is both easy to replicate and extendable to future studies with potentially more appropriate treatment definitions and data sets spanning a longer time period and larger geographic area. Results generally agree with the pessimistic outlook of previous studies which find that producer adaptation to climate change is limited in practice. Plausible explanations for the estimated dearth in water use adjustments include groundwater scarcity and regulation constraints, the misalignment of incentives and profitability concerns with the need to adapt, or that producers simply fail to recognize that the climate is changing.

Yet, there is evidence that some adaptation is occurring in irrigated corn production as acreage planted with this notably water intensive crop dropped following increases in irrigation demand. Results also suggest that groundwater extraction regardless of realized precipitation and evapotranspiration decreased in irrigated corn fields that ”survived” the climate change treatment. This result illuminates how the need for adaptation in irrigation varies with field and irrigation system characteristics. As climate change continues to increase the variability of weather conditions, more vulnerable subpopulations of irrigated producers are of course going to adapt before less vulnerable subpopulations. While this finding is not groundbreaking or unexpected, it reiterates the fact that caution is warranted in estimating the impacts of climate change using approaches that pool subpopulations and implicitly assume that producers engage in adaptation uniformly over space and time. Identifying and quantifying long-term adaptation and climate change impacts comprehensively requires observing irrigation and other production outcomes over substantially long periods of time wherein the long-term effects of climate change are realized. Of course, this comprehensive approach doesn't produce impact estimates in reasonable enough time to meet the needs of policymakers and stakeholders. However, the experimental setting proposed and implemented in this thesis approximates the comprehensive approach more so than the Ricardian and panel approaches.

Quantifying adaptation in irrigation, approximately or comprehensively, serves to inform agricultural stakeholders, as well as natural resources managers who are tasked with securing the sustainability of ecosystem services, of some of the potential impacts of climate change that have not yet been realized to a significant extent in the agricultural sector. Adaptation in irrigation is likely a double-edged sword as increases in groundwater extraction may mitigate yield losses over the short-term, while potentially inducing severe negative impacts on the health of regional water supplies over the long-run. These impacts could cascade to other sectors of the economy as natural steam flows diminish, flood cycles shift, watershed biodiversity is threatened, and more. However, if producers adjust their irrigation behavior in ways that satisfy lasting increases in irrigation demand without extracting more groundwater, then many of these water-related negative impacts of climate change adaptation may never come to fruition. Thus, the results of this thesis can be interpreted as evidence that adaptations to climate change in irrigation may not have unintended side effects on ecosystem services over the short run as producers initially adapt by reducing mean overall groundwater extraction and making other adjustments on the extensive margin. But if producers eventually reach the efficient frontier of their crop-water production functions and extensive margin strategies are extremely limited over the long run, than a conflict may arise between sustaining irrigated agricultural production and protecting water-related ecosystem services.

11 Bibliography

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. (2017). When Should You Adjust Standard Errors for Clustering? NBER Working Papers 24003, National Bureau of Economic Research, Inc.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1):249–275.

Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from us agriculture. *American Economic Journal: Economic Policy*, 8(3):106–40.

Burke, M. and Lobell, D. (2010). *Food Security and Adaptation to Climate Change: What Do We Know?*, pages 133–153. Springer Netherlands, Dordrecht.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Deschênes, O. and Greenstone, M. (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review*, 97(1):354–385.

Deschênes, O. and Greenstone, M. (2012). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Reply. *American Economic Review*, 102(7):3761–73.

Diamond, A. and Sekhon, J. S. (2013). Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. *Review of Economics and Statistics*, 95(3):932–945.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., and Schlenker, W. (2012). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment. *American Economic Review*, 102(7):3749–60.

Guiteras, R. (2007). The impact of climate change on indian agriculture. Working paper.

Hendricks, N. P. and Peterson, J. M. (2012). Fixed effects estimation of the intensive and extensive margins of irrigation water demand. *Journal of Agricultural and Resource Economics*, pages 1–19.

Ho, D. E., Imai, K., King, G., and Stuart, E. A. (2011). MatchIt: Nonparametric preprocessing for parametric causal inference. *Journal of Statistical Software*, 42(8):1–28.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of economic literature*, 47(1):5–86.

Kansas Geological Survey and Kansas Department of Agriculture (2018). Water information management and analysis system. http://hercules.kgs.ku.edu/geohydro/wimas/index.cfm.

Lin Lawell, C.-Y. C. (2016). The management of groundwater: Irrigation efficiency, policy, institutions, and externalities. *Annual Review of Resource Economics*, 8(1):247–259.

Mebane, Jr., W. R. and Sekhon, J. S. (2011). Genetic optimization using derivatives: The rgenoud package for R. *Journal of Statistical Software*, 42(11):1–26.

Meixner, T., Manning, A. H., Stonestrom, D. A., Allen, D. M., Ajami, H., Blasch, K. W., Brookfield, A. E., Castro, C. L., Clark, J. F., Gochis, D. J., Flint, A. L., Neff, K. L., Niraula, R., Rodell, M., Scanlon, B. R., Singha, K., and Walvoord, M. A. (2016). Implications of projected climate change for groundwater recharge in the western united states. *Journal of Hydrology*, 534:124 – 138.

Mendelsohn, R., Nordhaus, W. D., and Shaw, D. (1994). The impact of global warming on agriculture: A ricardian analysis. *The American Economic Review*, 84(4):753–771.

Montgomery, J. M., Nyhan, B., and Torres, M. (2018). How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science*.

Oehninger, E. B., Lin Lawell, C.-Y. C., and Springborn, M. R. (2018). The effects of climate change on agricultural groundwater extraction. Working paper.

Pfeiffer, L. and Lin, C.-Y. C. (2012). Groundwater pumping and spatial externalities in agriculture. *Journal of Environmental Economics and Management*, 64(1):16 – 30.

Pfeiffer, L. and Lin Lawell, C.-Y. C. (2014). The effects of energy prices on agricultural groundwater extraction from the high plains aquifer. *American Journal of Agricultural Economics*, 96(5):1349–1362.

R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rosenzweig, C. and Parry, M. L. (1994). Potential impact of climate change on world food supply. *Nature*, 367:133–138.

Scanlon, B. R., Faunt, C. C., Longuevergne, L., Reedy, R. C., Alley, W. M., McGuire, V. L., and McMahon, P. B. (2012). Groundwater depletion and sustainability of irrigation in the us high plains and central valley. *Proceedings of the National Academy of Sciences*, 109(24):9320–9325.

Schlenker, W., Hanemann, W. M., and Fisher, A. C. (2005). Will u.s. agriculture really benefit from global warming? accounting for irrigation in the hedonic approach. *American Economic Review*, 95(1):395–406.

Schlenker, W., Hanemann, W. M., and Fisher, A. C. (2006). The impact of global warming on u.s. agriculture: An econometric analysis of optimal growing conditions. *Review of Economics and Statistics*, 88(1):113–125.

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to u.s. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37):15594–15598.

Steward, D. R., Bruss, P. J., Yang, X., Staggenborg, S. A., Welch, S. M., and Apley, M. D. (2013). Tapping unsustainable groundwater stores for agricultural production in the high plains aquifer of kansas, projections to 2110. *Proceedings of the National Academy of Sciences*, 110(37):E3477–E3486.

Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statist. Sci.*, 25(1):1–21.

Szilagyi, J. (2018). Anthropogenic hydrological cycle disturbance at a regional scale: State-wide evapotranspiration trends (1979–2015) across nebraska, usa. *Journal of Hydrology*, 557:600 – 612.

Thornton, P., Thornton, M., Mayer, B., Wei, Y., Devarakonda, R., Vose, R., and Cook, R. (2017). Daymet: Daily surface weather data on a 1-km grid for north america, version 3.

Trenberth, K. E. (2011). Changes in precipitation with climate change. *Climate Research*, 47(1/2):123–138.

12 Tables & Figures

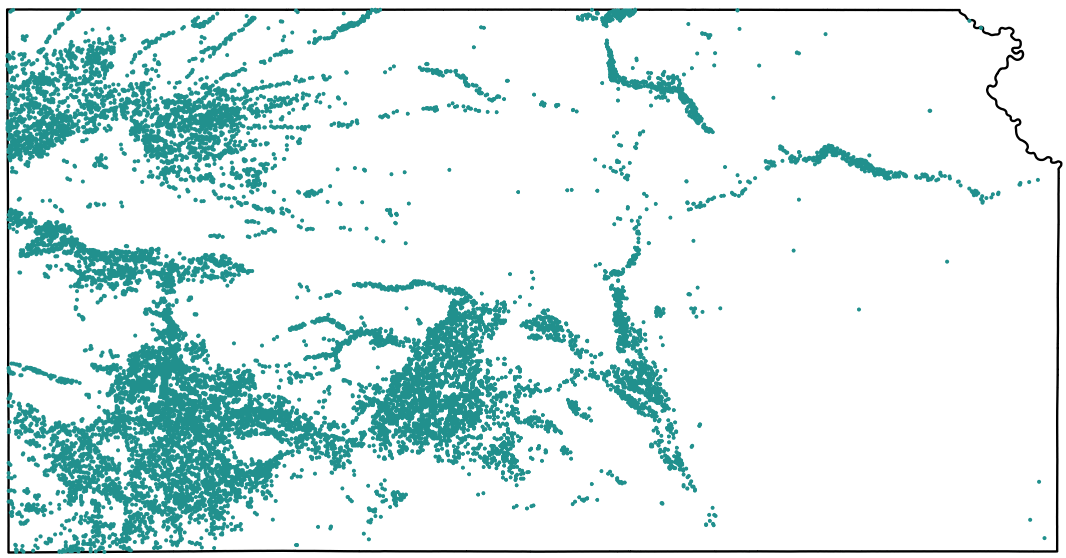

Figure 1: Spatial Distribution of Observed Irrigated Fields

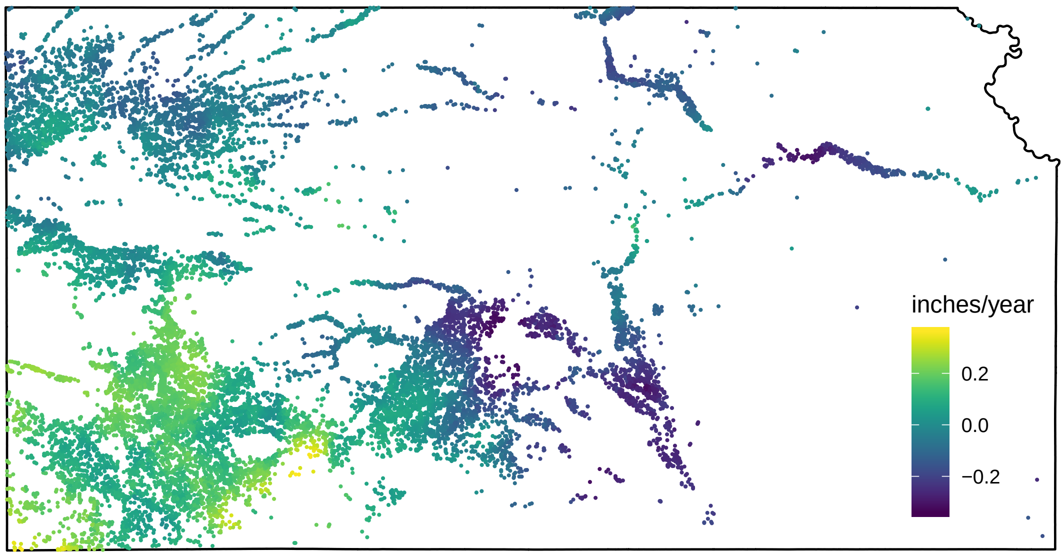

Figure 2: Irrigated Fields by Annual Irrigation Demand Trend 1996-2009

Table 1: Summary Statistics by Period - All Irrigated Fields

	1991-1995	1996-2009	2010-2014
Variables	mean (sd)	mean (sd)	mean (sd)
Weather (inches)			
Reference ET	38.37 (3.071)	39.9 (3.609)	42.98 (4.719)
Precipitation	18.06 (5.21)	16.88 (5.355)	15.08 (5.966)
Irrigation Demand	20.31 (7.939)	23.03 (8.602)	27.9 (10.4)
Intensive Margin (acre-inches)			
Total Water Applied	2030 (1744)	1839 (1403)	1935 (1468)
Water Applied Per Acre	14.18 (7.389)	13.29 (6.419)	14.19 (6.751)
Extensive Margin (acres)			
Irrigated	139.3 (86.51)	138.8 (87.01)	135.7 (85.61)
Corn	104 (53.55)	109.3 (59.12)	109.1 (61.08)
Wheat	95.87 (54.51)	93.75 (57.06)	87.4 (53.86)
Soybean	75.36 (43.08)	82.91 (43.88)	87.14 (46.72)
Alfalfa	99.76 (52.65)	103 (54.18)	103.7 (54.1)
Sorghum	79.41 (50.26)	76.55 (59.27)	82.09 (53.18)
Other Crop	143.8 (97.84)	144.8 (102.1)	141.8 (99.63)
Surface Flood	128.8 (97.5)	96.62 (90.2)	67.02 (59.21)
Flood + Center Pivot	143 (69.99)	152.2 (85.59)	150.3 (103.9)
Center Pivot w/ Drop Nozzles	134.2 (74.06)	141.5 (70.68)	137.8 (71.72)
Other Tech	101.2 (100.3)	63.45 (73.3)	71.9 (74.11)

Table 2: Summary Statistics by Period - Irrigated Fields Primarily Producing Corn

	1991-1995	1996-2009	2010-2014
Variables	mean (sd)	mean (sd)	mean (sd)
Weather (inches)			
Reference ET	37.9 (3.044)	39.65 (3.545)	42.72 (4.662)
Precipitation	18.7 (5.326)	17.06 (5.37)	15.22 (5.889)
Irrigation Demand	19.19 (8.028)	22.59 (8.562)	27.5 (10.26)
Intensive Margin (acre-inches)			
Total Water Applied	1959 (1507)	1882 (1248)	1979 (1321)
Water Applied Per Acre	14.63 (7.015)	14.23 (6.03)	15.12 (6.431)
Extensive Margin (acres)			
Irrigated	131.4 (71.01)	133.2 (71.89)	130.7 (71.24)
Surface Flood	114.6 (84.13)	88.95 (69.59)	64.84 (47.85)
Flood + Center Pivot	139.8 (61.17)	146 (72.57)	139.4 (96.07)
Center Pivot w/ Drop Nozzles	139.5 (67.01)	140 (66.05)	136 (64.56)
Other Tech	129.4 (72.41)	115.7 (94.97)	108.2 (66.82)

Table 3: Covariate Balance Summary - All Crops

	Before Matching			After Matching			
Covariates	Mean Trt	Mean Cntrl	Diff.	Mean Trt	Mean Cntrl	Diff.	% Imprv.
Weather (inches)							
Reference ET	39.87	37.01	2.86	39.87	39.47	0.4	86.01
Precipitation	15.85	20.14	-4.29	15.85	15.92	-0.07	98.37
Intensive Margin (acre-inches)							
Total Water Applied	2905	1210	1695	2905	2428	477	71.86
Extensive Margin (acres)							
Irrigated	174.3	104.9	69.4	174.3	164.8	9.5	86.31
Extensive Margin (shares)							
Corn	0.33	0.46	-0.13	0.33	0.33	0	96.83
Wheat	0.16	0.06	0.1	0.16	0.16	0	98.82
Soybean	0.03	0.11	-0.08	0.03	0.02	0	94.95
Alfalfa	0.12	0.06	0.05	0.12	0.12	0	97.67
Sorghum	0.07	0.09	-0.01	0.07	0.07	0	85.02
Other Crop	0.29	0.22	0.07	0.29	0.3	0	93.45
Surface Flood	0.4	0.35	0.05	0.4	0.39	0.01	76.21
Flood + Center Pivot	0.53	0.54	-0.02	0.53	0.54	-0.01	22.88
Center Pivot w/ Drop Nozzles	0.05	0.07	-0.02	0.05	0.05	0	91.63
Other Tech	0.02	0.04	-0.02	0.02	0.02	0	93.49

Table 4: Covariate Balance Summary - Corn Restricted

	Before Matching			After Matching			
Covariates	Mean Trt	Mean Cntrl	Diff.	Mean Trt	Mean Cntrl	Diff.	% Imprv.
Weather (inches)							
Reference ET	39.53	37.21	2.32	39.53	39	0.53	77.16
Precipitation	16.02	19.59	-3.57	16.02	16.34	-0.32	91.04
Intensive Margin (acre-inches)							
Total Water Applied	2941	1531	1410	2941	2642	299	78.79
Extensive Margin (acres)							
Irrigated	146.2	113.5	32.7	146.2	134.2	12	63.3
Extensive Margin (shares)							
Corn	0.92	0.85	0.07	0.92	0.95	-0.03	57.77
Wheat	0.02	0.01	0	0.02	0.01	0.01	-198.8
Soybean	0.02	0.06	-0.04	0.02	0.02	0	94.45
Alfalfa	0	0	0	0	0	0	50.28
Sorghum	0.01	0.02	-0.01	0.01	0	0.01	46.12
Other Crop	0.04	0.06	-0.02	0.04	0.03	0.01	43.81
Surface Flood	0.34	0.26	0.08	0.34	0.3	0.04	47.11
Flood + Center Pivot	0.61	0.63	-0.02	0.61	0.66	-0.05	-126.8
Center Pivot w/ Drop Nozzles	0.04	0.09	-0.06	0.04	0.03	0.01	89.62
Other Tech	0.01	0.01	0	0.01	0.01	0	2.771

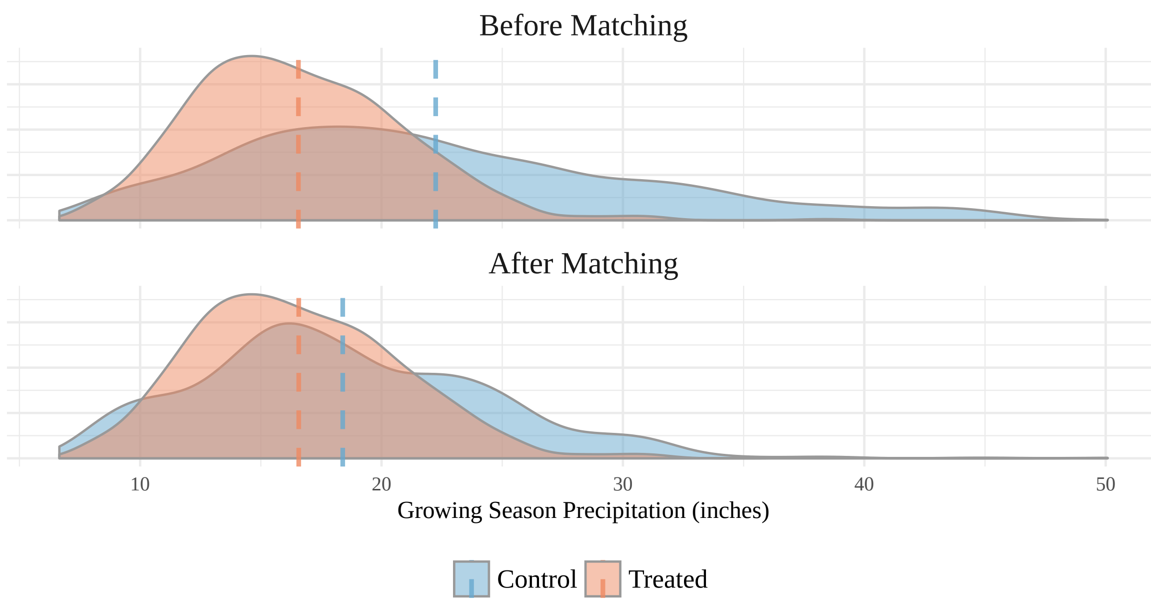

Figure 3: Example Covariate Distributions Before & After Matching - All Crops

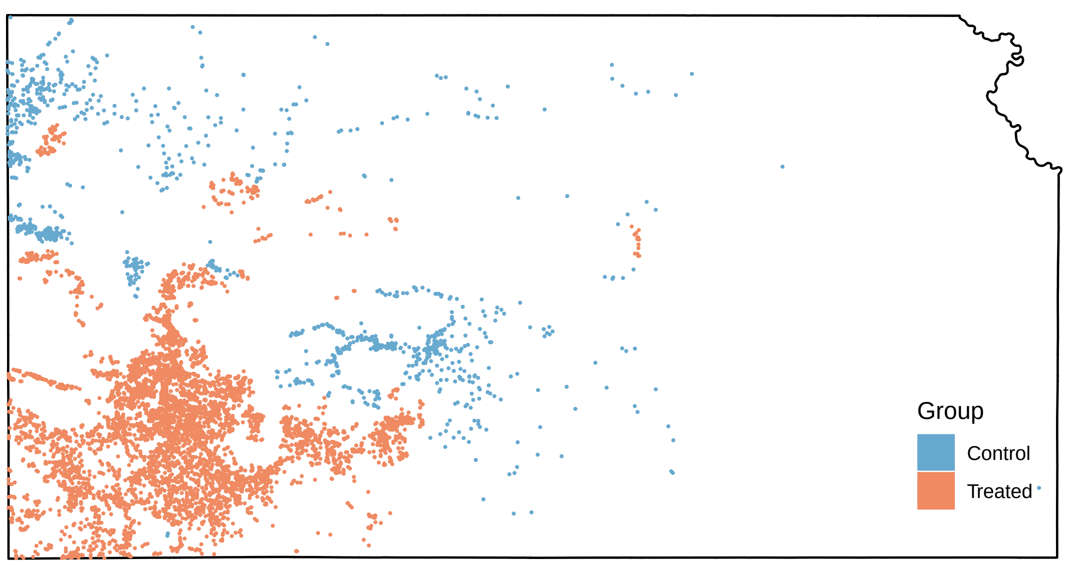

Figure 4: Spatial Distribution of Irrigated Fields by Group - All Crops

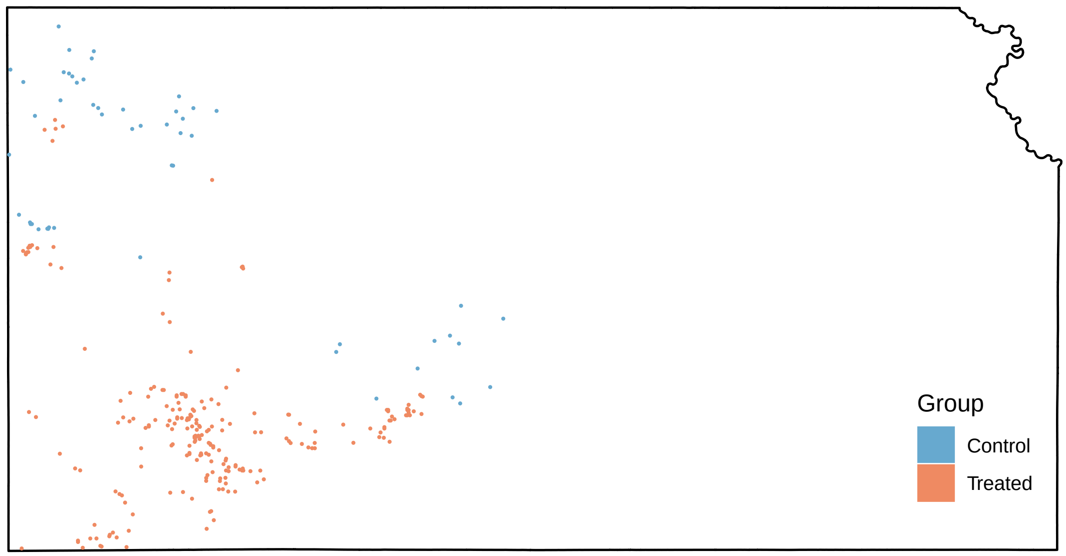

Figure 5: Spatial Distribution of Irrigated Fields by Group - Corn Restricted

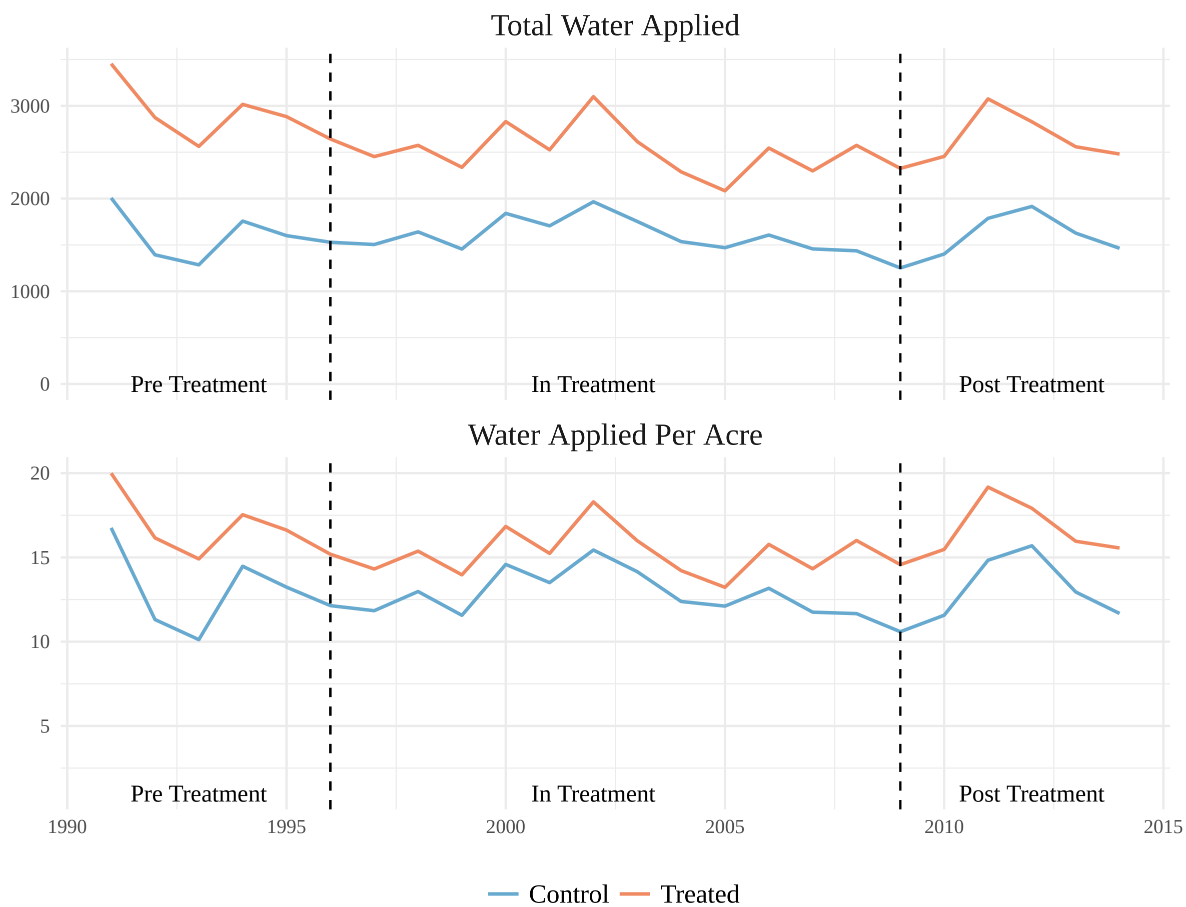

Figure 6: Mean Total and Per Acre Water Use by Group (acre-inches) - All Crops

Table 5: All Crops - Intensive Margin (acre-inches)

	Total Water Use	Water Use Per Acre
Precip	-22.05^{*}	-0.25^{***}
	(13.21)	(0.06)
Trt*Precip	17.10	0.10
	(14.86)	(0.07)
Post*Precip	8.86	0.06
	(17.04)	(0.09)
Trt*Post*Precip	-14.09	-0.05
	(16.28)	(0.08)
ET0	40.93^{***}	0.16^{***}
	(6.53)	(0.04)
Site-Period FEs	Yes	Yes
Year FEs	Yes	Yes
Observations	56,730	56,730
R^2	0.82	0.65
Adjusted R^2	0.77	0.56
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 6: All Crops - Irrigated Acreage

	Irrigated Acreage
Trt*Post	-2.09
	(4.08)
Site FEs	Yes
Year FEs	Yes
Observations	56,730
R^2	0.77
Adjusted R^2	0.74
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 7: All Crops - Crop Acreage

	Corn	Wheat	Soy	Alfa	Sorg	MultiOther
Trt*Post	-21.38^{***}	-0.98	-1.04	7.63^{*}	1.00	11.67
	(6.40)	(3.67)	(1.46)	(4.00)	(1.94)	(8.01)
Site FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	55,775	55,775	55,775	55,775	55,775	55,775
R^2	0.48	0.42	0.23	0.52	0.30	0.53
Adjusted R^2	0.41	0.35	0.14	0.46	0.21	0.47
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 8: All Crops - Irrigation Tech Acreage

	Flood	ComboPivot	CPivotDrop	Other
Trt*Post	-13.52	23.65	-13.80	0.42
	(11.28)	(15.92)	(10.99)	(1.12)
Site FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	55,837	55,837	55,837	55,837
R^2	0.53	0.44	0.64	0.23
Adjusted R^2	0.47	0.37	0.60	0.13
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 9: All Crops Treatment - Mean Overall Groundwater Extraction

	Pre Treatment	Post Treatment
Control	10.66	8.434
Treated	12.43	10.15

Table 10: All Crops - Overall Groundwater Extraction

	Overall Groundwater Extraction
Trt*Post	-0.07
	(0.49)
Observations	12,166
R^2	0.70
Adjusted R^2	0.40
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 11: Corn Restricted - Intensive Margin (acre-inches)

	Total Water Use	Water Use Per Acre
Precip	-15.07	-0.18^{**}
	(11.53)	(0.09)
Trt*Precip	-14.55	-0.04
	(15.03)	(0.10)
Post*Precip	-18.83	0.001
	(14.16)	(0.16)
Trt*Post*Precip	24.50	0.09
	(19.28)	(0.19)
ET0	22.31^{**}	0.22^{*}
	(9.95)	(0.11)
Site-Period FEs	Yes	Yes
Year FEs	Yes	Yes
Observations	2,674	2,674
R^2	0.79	0.62
Adjusted R^2	0.74	0.52
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 12: Corn Restricted - Irrigated Acreage

	Irrigated Acreage
Trt*Post	6.06
	(6.79)
Site FEs	Yes
Year FEs	Yes
Observations	2,674
R^2	0.57
Adjusted R^2	0.52
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 13: Corn Restricted Treatment - Irrigation Tech Acreage

	Flood	ComboPivot	CPivotDrop	Other
Trt*Post	-10.38	9.51	4.47	2.46
	(21.39)	(17.93)	(12.67)	(1.99)
Site FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	2,674	2,674	2,674	2,674
R^2	0.52	0.44	0.75	0.39
Adjusted R^2	0.47	0.37	0.72	0.32
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 14: Corn Restricted - Mean Overall Groundwater Extraction

	Pre Treatment	Post Treatment
Control	12.09	10.46
Treated	15.79	11.87

Table 15: Corn Restricted - Overall Groundwater Extraction

	Overall Groundwater Extraction
Trt*Post	-2.30^{**}
	(1.11)
Observations	554
R^2	0.75
Adjusted R^2	0.49
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 16: Placebo - Intensive Margin

	Total Water Use	Water Use Per Acre
Precip	-12.92^{***}	-0.22^{***}
	(2.13)	(0.02)
Trt*Precip	0.91	0.01
	(1.45)	(0.01)
Post*Precip	5.26^{*}	0.08^{***}
	(2.82)	(0.02)
Trt*Post*Precip	0.73	0.01
	(1.81)	(0.02)
ET0	48.81^{***}	0.27^{***}
	(3.10)	(0.02)
Site-Period FEs	Yes	Yes
Year FEs	Yes	Yes
Observations	107,324	107,324
R^2	0.85	0.68
Adjusted R^2	0.81	0.59
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 17: Placebo - Irrigated Acreage

	Irrigated Acreage
Trt*Post	-12.92^{***}
	(1.60)
Site FEs	Yes
Year FEs	Yes
Observations	107,324
R^2	0.79
Adjusted R^2	0.76
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 18: Placebo - Crop Acreage

	Corn	Wheat	Soy	Alfa	Sorg	MultiOther
Trt*Post	-0.53	0.21	0.14	-1.65^{**}	-0.46	1.14
	(1.26)	(0.81)	(0.54)	(0.69)	(0.51)	(1.57)
Site FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	105,117	105,117	105,117	105,117	105,117	105,117
R^2	0.45	0.41	0.28	0.52	0.28	0.53
Adjusted R^2	0.39	0.34	0.19	0.46	0.20	0.47
Note:				*p<0.1; **p<0.05; ***p<0.01		

Table 19: Placebo - Irrigation Tech Acreage

	Flood	ComboPivot	CPivotDrop	Other
Trt*Post	-0.97	-0.89	0.94	-0.25
	(1.46)	(1.84)	(1.71)	(0.36)
Site FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	105,232	105,232	105,232	105,232
R^2	0.51	0.49	0.63	0.23
Adjusted R^2	0.45	0.42	0.59	0.13
Note:		*p<0.1; **p<0.05; ***p<0.01		

Table 20: Placebo - Mean Overall Groundwater Extraction

	Pre Treatment	Post Treatment
Control	7.101	5.154
Treated	6.993	4.855

Table 21: Placebo Treatment - Overall Groundwater Extraction

	Overall Groundwater Extraction
Trt*Post	-0.19
	(0.15)
Observations	23,111
R^2	0.71
Adjusted R^2	0.41
Note:	*p<0.1; **p<0.05; ***p<0.01