Journal of Environmental Economics and Management 91 (2018) 150–165

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Adaptation to an irrigation water restriction imposed through local governance☆

Krystal M. Drysdale [a], Nathan P. Hendricks [b,*]

[a] City of Phoenix, United States
[b] Department of Agricultural Economics, Kansas State University, Manhattan, KS 66506, United States

ARTICLE INFO

Article history:
Received 3 July 2017
Revised 9 July 2018
Accepted 10 August 2018
Available online 16 August 2018

Jels:
Q15
Q25
Q30

ABSTRACT

We estimate how farmers adapted to a water restriction imposed through local governance. The restriction imposed a uniform quota on water use with a 5-year allocation and allowed trading of the quota within the restricted area. Our analysis exploits unique micro-level data on irrigated water use, irrigated acreage, and crops. We use a difference-in-differences econometric strategy that also includes farmer-time fixed effects to estimate the response to the restriction, where we exploit water rights between 2 and 5 miles of the policy boundary as a control group. Results indicate that farmers reduced water use by 26% due to the policy with most of the response due to reductions in water use intensity on the same crops rather than through reductions in irrigated acreage or changes in crops. The results imply that the short-run welfare impact of the policy was smaller than a policy that reduces irrigated acreage.

1. Introduction

Depleted aquifers due to irrigation extraction are a major concern globally. Using data from NASA's GRACE satellite, Famiglietti (2014) found that groundwater is being depleted in the largest global agricultural zones. Scanlon et al. (2012) show significant groundwater depletion in the southern portion of the High Plains Aquifer (e.g., Kansas, Oklahoma, Texas, and New Mexico) as well as the Central Valley in California. The open-access nature of aquifers is often cited as a reason for excessive depletion. Property rights to extract groundwater according to prior appropriations are currently in place in Kansas, yet these property rights have rarely been exercised to reduce the pumping of junior rights. Instead, one area in Kansas reduced groundwater extraction through a restriction starting in 2013 by forming a Local Enhanced Management Area (LEMA). In this paper, we estimate how farmers adapted to this water restriction.

There is a large literature in economics that studies the optimal extraction of groundwater. Gisser and Sanchez (1980) found that the gains from management of an aquifer may be small. However, other studies relax some restrictive assumptions in Gisser and Sanchez (1980) and found larger gains from management (Koundouri, 2004; Brozović et al., 2010; Guilfoos et al., 2013). Lin Lawell (2016) also documents several reasons why farmers may extract groundwater at a faster rate than is dynamically

☆ This material is based upon work supported by the National Science Foundation under Award No. EPS-0903806 and matching support from the State of Kansas through the Kansas Board of Regents, the USDA National Institute of Food and Agriculture, Agricultural and Food Research Initiative Competitive Program, Agriculture Economics and Rural Communities, grant # 2017-67023-26276, and the Arthur Capper Cooperative Center. We thank anonymous reviewers, Brian Briggeman, Bill Golden, Taro Mieno, Karina Schoengold, and participants at the Agricultural and Applied Economics Association annual meeting for helpful comments.

* Corresponding author. author.

E-mail address: nph@ksu.edu (N.P. Hendricks).

https://doi.org/10.1016/j.jeem.2018.08.002
0095-0696/

optimal.

Recent articles indicate that there could be significant gains to management of the High Plains Aquifer in Kansas. Pfeiffer and Lin (2012) estimate that groundwater extraction is significantly larger due to spatial externalities. Guilfoos et al. (2016) use a simulation model of different policies that reduce water extraction and find larger gains to management in the LEMA area than other regions in northwest Kansas. Edwards (2016) estimate that land values increased more in counties with greater common pool externalities—due to varying hydrology—in response to the introduction of Groundwater Management Districts (GMDs) in Kansas. These GMDs implemented some rules to reduce common pool externalities—such as well-spacing requirements and closure to new drilling—but did not impose any restrictions on existing water rights. The GMDs were formed in 1972, but concerns about aquifer depletion persist. GMDs argued they did not have the authority required to impose restrictions needed to properly manage the aquifer leading to the passage of legislation in 2012 to authorize LEMAs.

Users of a common pool resource take action to reduce extraction when the benefits exceed the costs (Coase, 1960; Demsetz, 1967). The costs include not only the short-run loss in returns from reduced withdrawals but also the transaction costs of taking action (Wiggins and Libecap, 1985; Ayres et al., 2018). Water rights in Kansas follow the prior appropriations doctrine, so a property right structure exists to reduce groundwater extraction—senior water rights can file a claim that junior water rights are impairing the use of their right. Burness and Quirk (1979) and Libecap (2011) note that applying prior appropriations may not be economically efficient due to large transaction costs of trading water rights. Another approach to reduce extraction is to impose a restriction through local governance or collective action (Ostrom, 1990).

Our paper does not address the important question of transaction costs of trading under prior appropriations, but our estimates of how producers adapted to the restriction sheds light on the loss in short-run returns. An upper bound estimate of the cost of applying prior appropriations is the loss due to a reduction in irrigated acreage by junior water right holders (i.e., assuming transaction costs prevent any trades). Farmers in the LEMA could reduce their water use by reducing irrigated acres or reducing water application on the same irrigated area. To the extent that farmers adapt to the restriction through means other than reduced irrigated acreage, the short-run loss in net returns from the restriction must have been smaller than this upper bound estimate of the cost of applying prior appropriations.

We use a variant of a difference-in-differences estimation strategy where our econometric strategy exploits changes in irrigation behavior inside the policy boundary compared to behavior of water rights between 2 and 5 miles from the policy boundary for those farmers that have water rights both inside the LEMA and in the control group. We decompose the total effect on irrigation water use into responses along the extensive margin (changes in irrigated acreage), direct intensive margin (changes in intensity holding cropping pattern constant), and an indirect intensive margin (changes in intensity due to changes in cropping patterns). We find that the water restriction reduced water use by 26% compared to a counterfactual scenario where no restriction had been put in place. Most of the response occurred along the direct intensive margin (21% reduction), with a smaller but significant response along the indirect intensive margin (4.5% reduction) and an insignificant response along the extensive margin.

Unfortunately, we cannot precisely estimate the short-run welfare impact of the water restriction because we do not have field-level yield or cost of production data to understand how the reduced water use impacted yields and costs. We are able to infer that the loss in welfare is less than $28.08 per irrigated acre because the welfare loss must be smaller than the loss of reducing irrigated acreage since farmers chose to adapt mostly through reduced intensity. It is also noteworthy that minutes from recent public meetings indicate that most producers supported a new LEMA proposal that would create a 5-year allocation for 2018–2022.[1] Farmers in the region would only support a renewal of the quota if they perceived any losses in the short run do not exceed the gains in the long run.

Our work is most closely related with Smith et al. (2017). They study the impact of a self-imposed water tax on irrigation in a district in Colorado. In their case, the district imposed the tax in an effort to avoid state regulations on pumping. Smith et al. (2017) also use a difference-in-differences methodology and find that a tax of $75/acre-foot of water reduced extraction by 33%. They also find that most of the response occurs due to reduced irrigation intensity along with limited changes in cropping patterns and irrigated acreage. A key difference between our methodology and Smith et al. (2017) is that we are able to also control for farmer-time fixed effects. We show that omitting these farmer-time fixed effects leads us to overstate the impact of the policy on water use. We also do not have the presence of surface water allocations in our context that complicate the analysis in Smith et al. (2017) and we estimate an event study to understand if the response changed over time.

In another related study, Hornbeck and Keskin (2014) estimate adaptation of agriculture to water availability. They exploit the introduction of new technologies after World War II that allowed farmers to extract groundwater from the High Plains Aquifer. Hornbeck and Keskin (2014) find that counties over the High Plains Aquifer initially adapted by expanding irrigated acreage with little changes in total farmland acreage and then in later periods also began expanding total farmland. They also find that farmers over the High Plains Aquifer adopted higher-valued crops that were more sensitive to droughts. The results in our study indicate that in the short run, farmers adjusted to a water restriction primarily by reducing the intensity of irrigation with relatively small changes in cropping patterns. Therefore, farmers have responded to the LEMA water restriction mostly by

[1] See the formal review and advisory committee reports available at http://agriculture.ks.gov/divisions-programs/dwr/managing-kansas-water-resources/local-enhanced-management-areas/lists/lemas/sheridan-county-6-lema.

exposing themselves to greater production risk.

Several previous studies estimate the response of irrigation water to price along the intensive and extensive margins. Hendricks and Peterson (2012) and Pfeiffer and Lin (2014b) study irrigation in Kansas and both find that most of the demand response is along the intensive margin with relatively small changes in water use due to changes in cropping patterns. In contrast, Schoengold et al. (2006) find that nearly half of the demand response is due to changes in cropping patterns and technology in California. Moore et al. (1994) estimate water demand for the western U.S. and find that essentially all of the demand response is due to changes in cropping patterns or changes in irrigated acreage. Regions other than Kansas may see a greater response due to changes in cropping patterns due to the difference in irrigated crops of the region.

Our work is also related to previous literature that estimates the optimal response to limited water availability using conceptual models or numerical simulations. English (1990) and Wang and Nair (2013) study the economic conditions for optimizing irrigation water use along the extensive and intensive margins with limited water availability. Peterson and Ding (2005) and Foster et al. (2014) model the optimal response to limited water availability due to a limited instantaneous rate of extraction due to groundwater depletion. Graveline and Merel (2014), Foster et al. (2014), and Foster et al. (2017) model the optimal response to a quota on groundwater extraction. A key advantage of our econometric framework is that we estimate how farmers actually responded to a water restriction.

2. Background

Kansas enacted the Water Appropriation Act in 1945, establishing the doctrine of prior appropriations. Pumping of groundwater requires authorization through an approved water right. A water right for irrigation has five key components: (1) priority date, (2) maximum rate of diversion, (3) maximum annual quantity, (4) location of the point of diversion, and (5) place of use (Rogers et al., 2013). The authorization of water rights increased rapidly from the 1960s to the 1980s, but then slowed. The area included in our study is currently closed to new appropriations because of the impact new appropriations would have on existing water right holders.

In 2012, new legislation in Kansas granted Groundwater Management Districts (GMDs) the power to originate their own localized water conservation management plans which are then legally enforced by the state. A LEMA is approved by a vote of the GMD board of directors whose members are elected water right owners (usually farmers). Farmers in a portion of Sheridan and Thomas counties, located in the northwestern corner of Kansas, participated in the process to impose a restriction on themselves by forming a Local Enhanced Management Area (LEMA). The 99 square mile area maintains 185 water rights for irrigation. The restriction was imposed by providing each water right with a five-year allocation of 55 inches per authorized acre effective 2013–2017 (KDA, 2013), representing roughly a 20% reduction of extraction from the historical average for the area. Authorized acres were defined as the maximum number of acres irrigated in a single year between 2007 and 2010.

Providing farmers with a five-year allocation allows them to allocate their water use between years (i.e., intertemporal trading), reducing the potential negative impacts of a drought in a year of the allocation period. Farmers could meet the quota by reducing water use on the same irrigated acreage, reduce irrigated acreage, or some combination of the two. Farmers with multiple water rights could pool their allocation across water rights (i.e., transfers within the same owner). Furthermore, water right owners could transfer their allocation between owners within the LEMA boundary. According to correspondence with the manager of GMD 4, 64% of water rights pooled their allocation. There was only one transfer of an allocation across water rights that occurred in the last year of the five-year period to balance an account. The five-year allocation did not permanently alter the underlying water rights within the LEMA. Farmers were still not allowed to pump more than their authorized annual quantity in any given year. Furthermore, if a LEMA is not approved for a subsequent period, then authorized quantities automatically revert to the underlying water rights.

The restriction was imposed on all water right holders inside the policy boundary even if a water right owner did not support the restriction. Meters are required of all wells in GMD 4 and water right owners are required to report their annual water use. The Kansas Department of Agriculture enforces water rights and also enforces the LEMA allocation. The LEMA order designated a stiff penalty of \$1000 per day of pumping in excess of the allocation if the amount pumped was less than 4 acre-feet above the allocation. Pumping greater than 4 acre-feet in excess of the allocation results in a 2-year suspension of the water right in addition to the fine.

There were 13 meetings—beginning in 2008—of water right owners within the LEMA discussing a potential restriction. And although there was no formal vote among water right owners to approve the LEMA, there were informal votes at the meetings. The minutes of the final meeting on May 9, 2012 states the following.

> “It was the consensus of the group that the proposal as modified during the meeting be written up and presented to the GMD 4 board for adoption and subsequent submission to the chief engineer on their behalves. While few of the consensus decisions recorded during this meeting were unanimous, this record is deemed to reflect the majority consensus of the participants.”

3. Conceptual model

We use the methodology applied in the water demand literature to decompose price response (Moore et al., 1994; Schoengold et al., 2006; Hendricks and Peterson, 2012) in order to examine the margins of response to a quantity restriction. Following

Hendricks and Peterson (2012) we decompose the effect of a quantity restriction into the following three margins of adjustment: (i) extensive (changes in irrigated acreage), (ii) direct intensive (changes in irrigation intensity for a given land use), and (iii) indirect intensive (changes in land use).

Assume there is a representative irrigator such that his or her water demand for a particular well is subject to a water quota denoted q. We assume that the quota constraint is binding. Irrigators choose the optimal irrigated acreage $a(q)$ and the optimal applied water intensity in inches/acre $w(q)$ given the quota. Total water use is written as

$$D(q) = a(q)w(q). \tag{1}$$

Differentiating equation (1) and dividing by $D(q)$ gives the relative change in water use due to a change in the quota. The total change in water is decomposed into two effects

$$\frac{D'(q)}{D(q)} = \frac{a'(q)}{a(q)} + \frac{w'(q)}{w(q)}, \tag{2}$$

where $\frac{a'(q)}{a(q)}$ represents the extensive margin effect and $\frac{w'(q)}{w(q)}$ represents the total intensive margin effect.

The total intensive margin effect can be further decomposed. We let $s_j(q)$ denote the share of irrigated acreage for each crop $j = 1, \ldots, J$ and assume an interior solution for each crop. The average applied water per acre is written as

$$w(q) = \sum_{j=1}^{J} s_j(q)w_j(q). \tag{3}$$

Differentiating $w(q)$ and multiplying by $q/(w(q))$ converts to relative effects such that:

$$\frac{w'(q)}{w(q)} = \left[\sum_{j=1}^{J} s_j\,(q)\,w'_j\,(q) + \sum_{j=1}^{J} s'_j\,(q)\,w_j\,(q)\right]\frac{1}{w\,(q)}. \tag{4}$$

The first term on the right-hand side of equation (4) is the response along the direct intensive margin, which represents the change in water intensity due to less water application per acre while holding constant crops. The second term is the response along the indirect intensive margin, which represents the change in water intensity because farmers change crops (i.e., switch to a less water intensive crop).

Farmers adjust their water use along the margin that gives the smallest reduction in returns in order to achieve the required reduction in water use (see Wang and Nair (2013) for a discussion of the economics of limited irrigation). Assuming that irrigated acreage has constant returns to scale, then farmers find it optimal to reduce water use at the extensive margin if a 1% reduction in water use intensity reduces irrigated returns by more than 1%. If we observe farmers reducing irrigated acreage due to the restriction, then we can infer that the short-run welfare effect is equal to the difference in irrigated and nonirrigated rental rates times the reduction in irrigated acres. If we observe farmers reducing water use along the intensive margin, then we can infer that the short-run welfare effect is less than the welfare effect of reducing irrigated acres. Of course, farmers may find it optimal to adapt by a combination of reduced intensity and reduced irrigated acreage.

4. Data

We create a panel data set for irrigated water rights across a 10-year period (2007–2016). Because the LEMA is not tied to county boundaries but rather an area identified with critical aquifer concerns, it is important to construct the data set at a micro level instead of county level aggregates. Kansas law requires all water right holders to report annually on irrigation water use and data are collected in the Kansas Water Rights Information System Database (WRIS). Because of this, we are able to quantify reported water use data at each irrigator's well (termed a "point of diversion"). Each point of diversion has an associated water right number. Often a water right is associated with a single point of diversion, but in certain cases multiple points of diversion have the same water right number. Because the water use restriction was placed on water rights, we aggregate all the data to the water right level. The WRIS data provide us information on total water withdrawals, irrigated acreage, and crop type. The WRIS data also contain an identification number for the person who filed the water use report, which we refer to as the "farmer." The observations were identified spatially according to the location of the points of diversion associated with each water right.

The specific crops considered in this analysis include alfalfa, corn, sorghum, soybeans, and wheat, with two additional categories identified as multiple and other. The category for other irrigated uses includes fruits, vegetables, sunflowers, pasture, cotton, turf grass, barley, oats, rye, and dry beans. Additionally some reporting merely indicates that "multiple" crops were grown, but not which crops were specifically grown. The Kansas data also do not indicate the number of acres planted to each crop nor how the water was distributed to each crop when multiple crop types were reported. We assume that if k crops were grown, the proportion of the field in each crop is simply $1/k$.

Water rights subject to the LEMA quantity restriction beginning in 2013 were identified from official Kansas Department of Agriculture data. We create a control group that includes all water rights between 2 and 5 miles of the LEMA boundary (Fig. 1). We do not include water rights less than 2 miles from the LEMA boundary due to concerns about spillover of the water restriction—water rights just outside the LEMA boundary may have changed their water use due to changes in their aquifer

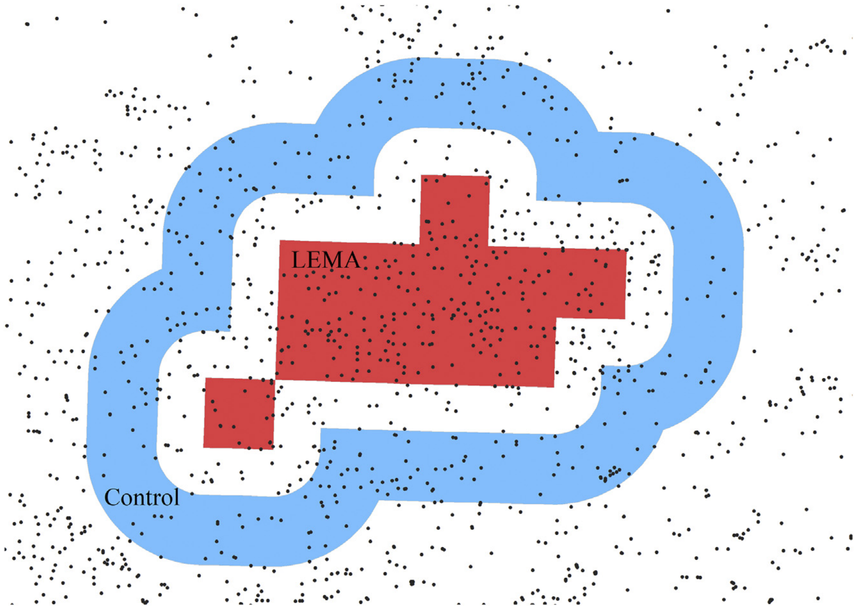

Fig. 1. Points of Diversion located inside the LEMA and the control group.

conditions as a result of the restriction inside the LEMA. We discuss the rationale for our construction of the control group in greater detail in the econometric model section.

We construct variables for weather conditions using the data from PRISM (Parameter-elevation Relationships on Independent Slopes Model). PRISM estimates weather conditions on a 4 km grid (approximately 2.5 miles). Importantly, this spatial resolution is sufficiently small to control for differences in weather between the treatment and control groups. Since 2002, PRISM uses the precipitation measured from radar to give the most accurate spatial distribution possible. PRISM data provide daily measures of minimum temperature, maximum temperature, and precipitation. From these values, we construct cumulative annual precipitation and cumulative reference evapotranspiration (ET).

5. Summary statistics and visual analysis

5.1. Summary statistics

Our analysis contains 3592 observations (water right-year pairs) of which 2123 are before the LEMA went into effect. Of the 3592 observations, there are 1797 observations in the control group and 1795 observations inside the LEMA boundary. There are 378 water rights and 202 farmers included in the analysis.

Table 1 reports means of the variables in the control group and the LEMA in the pre-treatment period. Most irrigated acreage was dedicated to corn before the restriction inside the LEMA (0.71). Other irrigated acreage was spread between soybeans (0.15), wheat (0.02), sorghum (0.003), alfalfa (0.01), other crops (0.01), and multiple unknown crops (0.09). Annual precipitation in the region averages about 22 inches and reference evapotranspiration is 48 inches. Fig. 2 shows annual precipitation during our study period. There was a drought in 2012 and large rainfall in 2009 and 2011. The years following the water restriction (2013–2016) all had rainfall close to average.

5.2. Graphical analysis

Next, we explore the data by presenting a simple difference-in-differences (DID) analysis by aggregating the data for water rights inside the LEMA boundary and those in the control group for years prior to the water restriction and after. Simply comparing before and after inside the boundary is likely to be misleading because the effect of the water restriction is confounded with the effect of market conditions and weather. Comparing the outcome variable in the LEMA boundary to the control group is also misleading if unobserved characteristics are systematically different inside the LEMA boundary. The DID framework estimates the effect of the LEMA by exploiting the change in the outcome variable in the control group as the counterfactual for how the outcome variable would have changed inside the LEMA boundary after the policy came into effect. Our econometric model in the next section provides a more robust estimate of the causal effect of the LEMA by also controlling for farmer-year fixed effects, but the graphical analysis in this section gives an intuitive visual representation of the data.

Table 1
Summary statistics of pre-treatment sample means.

Variable	Mean of Control	Mean of LEMA
log Applied Inches	6.94	7.39
log Acres Irrigated	4.50	4.82
log Intensity	2.44	2.57
Crop		
Alfalfa	0.03	0.01
Corn	0.62	0.71
Sorghum	0.01	0.003
Soybeans	0.14	0.15
Wheat	0.04	0.02
Other Crops	0.05	0.01
Multiple Unknown	0.10	0.09
Weather		
Precipitation (inches)	21.87	21.57
Evapotranspiration (inches)	48.60	48.48

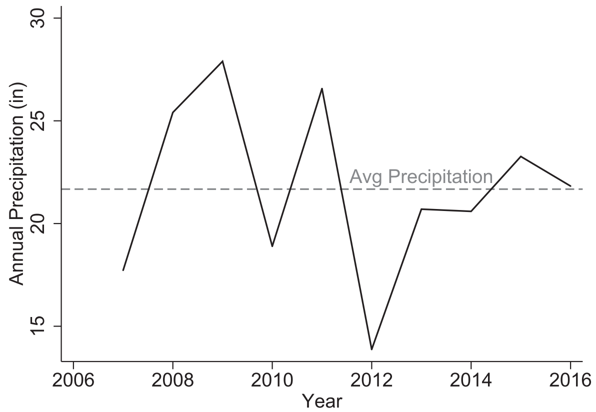

Fig. 2. Annual precipitation in the region.

Fig. 3 shows the aggregate DID results for total irrigated acres (extensive margin response) and water use intensity (intensive margin response). The red line shows the average outcome before and after the restriction inside the LEMA boundary. The blue line shows the outcome in the control group. The dashed red line shows the counterfactual change inside the LEMA, which equals the change in the outcome in the control group applied to the LEMA. The difference between the dashed red line and the solid red line represents the impact of the restriction based on the usual DID parallel trend assumption. Fig. 3 shows that irrigated acres decreased slightly inside the LEMA and acres remained about constant in the control group, implying a 6% decrease in irrigated acreage due to the restriction. Applied water intensity decreased substantially inside the LEMA while it remained about constant in the control group resulting in a 28% decrease in intensity.

Fig. 4 shows how farmers changed crops in response to the restriction. The y-axis in each plot shows the share of irrigated acres for the respective crop averaged across water rights. For example, the share of irrigated acres planted to corn was 0.71 on average before the policy was implemented inside the LEMA. It is also important to note that the scale of the y-axis for corn acres ranges from 0 to 0.80 whereas the y-axis for all of the other crops ranges from 0 to 0.16 because corn production is so prevalent in the region.

Fig. 4 indicates that the share of irrigated acres planted to corn decreased by 0.15 due to the restriction. The decrease in corn acreage was offset by an increase in the share of acres planted to unknown multiple crops (0.08), sorghum (0.03), wheat (0.02), soybeans (0.01), and other crops (0.01). The switch in crop types does not result in a substantial reduction in water use. The average post-policy intensities inside the LEMA are 10.7 in/acre for corn, 9.1 in/acre for multiple crops, 7.7 in/acre for sorghum, 4.3 in/acre for wheat, 9.8 in/acre for soybeans, and 5.4 in/acre for other crops. Multiplying the crop-specific intensity by the change in share of acres gives a reduction in water use per irrigated acre of 0.41 inches per water right. Each water right irrigated about 129 acres after the policy, giving a reduction in water use due to crop switching of roughly 53.9 acre-inches per water right. Therefore, crop-switching alone (i.e., the indirect intensive margin effect) only decreased water use by roughly 3% since the average pre-treatment water use was 1841 acre-inches per water right.

Fig. 5 shows the change in intensity for specific crops. Two or more crops were planted for some water rights but we do not know how much water was applied to each crop. Therefore, Fig. 5 shows the average intensity in cases where a single crop was produced using the entire water right. We only show results for corn, soybeans, and unknown multiple crops because there were a small number of water rights planted completely to the other types of crops. Farmers in the control group decreased

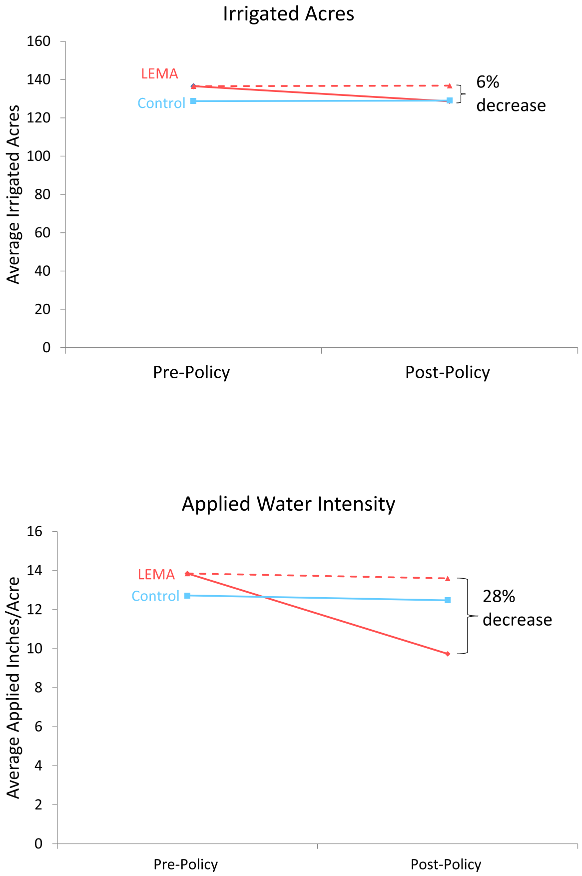

Fig. 3. Difference-in-difference results for average irrigated acreage and applied water intensity.

water intensity for corn, but increased intensity for soybeans and multiple crops. Farmers within the LEMA decreased irrigation intensity subtantially for each crop. The restriction resulted in a decrease in intensity for corn by 22%, soybeans by 28%, and multiple crops by 31%. The decrease in intensity for these crops resulted in a substantial reduction in overall water use since these three categories represent roughly 85% of irrigated acreage in the LEMA after the policy was implemented.

6. Econometric model

For our preferred estimate of the impact of the LEMA water restriction we extend the difference-in-differences model to include farmer-time fixed effects and weather controls. Our regression of interest is written as

$$\ln(y_{ift}) = \alpha_i + \lambda_{ft} + \beta D_{ift} + \boldsymbol{\theta} \boldsymbol{Z}_{ift} + \varepsilon_{ift} \tag{5}$$

where $\ln(y_{ift})$ is the natural log of the outcome (i.e., total water use, irrigated acreage, or water use intensity) for water right i of farmer f in year t, α_i are water right fixed effects, λ_{ft} are farmer-time fixed effects, $\boldsymbol{Z}_{ift}$ is a vector of weather controls, and D_{ift} is equal to 1 if the LEMA water restriction is applied to the water right and 0 otherwise. We cluster standard errors by the water right to allow for heteroskedastity across water rights and autocorrelation of the errors for a particular water right over time.

The stable unit treatment value assumption (SUTVA) for our analysis is that the outcomes for water rights in our control group are not affected by the treatment status of water rights within the LEMA. This assumption is violated if reduced withdrawals inside the LEMA boundary affect aquifer conditions in the control group, and thus affect water use decisions in the control group. SUTVA is least likely to hold for water rights nearest the LEMA boundary, therefore we do not include water rights within 2 miles of the boundary in the control group. Aquifer conditions for water rights further than 2 miles from the LEMA boundary are unlikely to be affected by reduced withdrawals inside the LEMA in an economically meaningful way due to the slow lateral flows of the aquifer. Stotler et al. (2011) analyze the hydrology of a different area of Thomas County—one of the counties in the LEMA—and find that long term flow of the aquifer takes about 15–20 years per mile.[2] Pfeiffer and Lin (2014b) estimate the

[2] The analysis in Stotler et al. (2011) examines the effect of a potential reduction in water use in four townships. Stotler et al. (2011, p. 116) conclude “the first and greatest effects of either conservation or depletion will be experienced in the immediate area.”

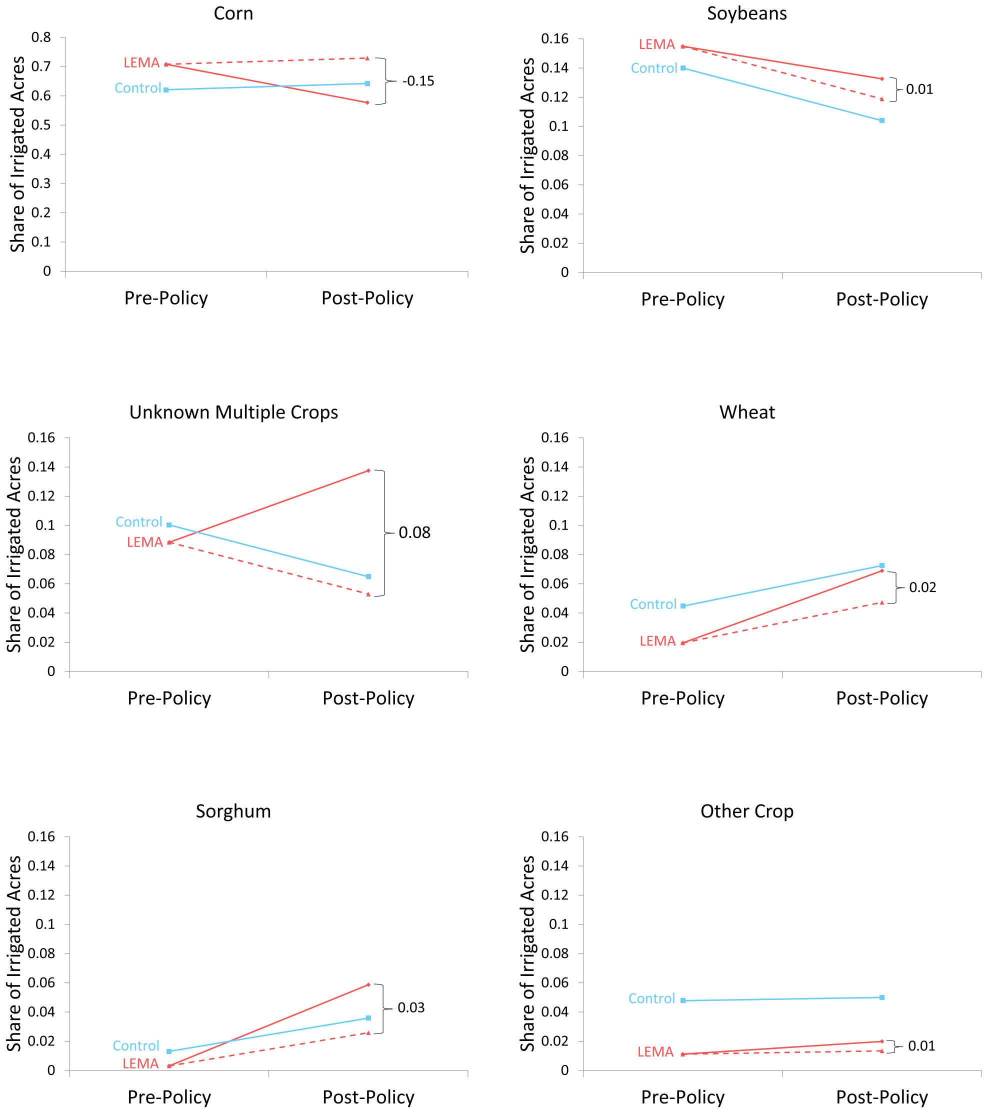

Fig. 4. Difference-in-difference results for share of irrigated acres by crop.

impact of neighbors' pumping on the water table between years in Kansas and find little impacts from pumping more than two miles away. Our results are also robust to using different distances to construct the control group as we discuss in the results section.

One potential concern with our estimates is that the policy boundary could be endogenously determined (i.e., the boundary was formed to include those farmers that were willing to reduce their water use). Although the LEMA was in part farmer-led, the resulting defined boundary was formally decided externally, by the advisory board which consisted of not only residents, owners, and operators but also representatives from the Division of Water Resources, Kansas Department of Agriculture, and Groundwater Management District 4. The boundary was formally defined as those sections whose saturated thickness of groundwater had decreased by 9% or more between 1997 and 2006. Therefore, farmers did not necessarily select into the water restriction.

Furthermore, including farmer-time fixed effects reduces the concern that farmers inside the LEMA boundary may be systematically different from farmers in the control group and thus may have changed water use over time differently than those in the control group. For example, farmers inside the LEMA may be adopting conservation technologies faster or slower than those in the control group. Or farmers inside the LEMA may be more or less responsive to changes in crop prices over time. The direction of such bias is unclear, *a priori*, but including farmer-time fixed effects reduces this concern. Intuitively, including farmer-time fixed effects means that we exploit changes in outcomes for farmers who manage water rights both inside and outside the LEMA because the effect of the water restriction is not identified for farmers that only manage land inside the LEMA boundary. One potential concern with including farmer-time fixed effects is if farmers that had land inside and outside the LEMA

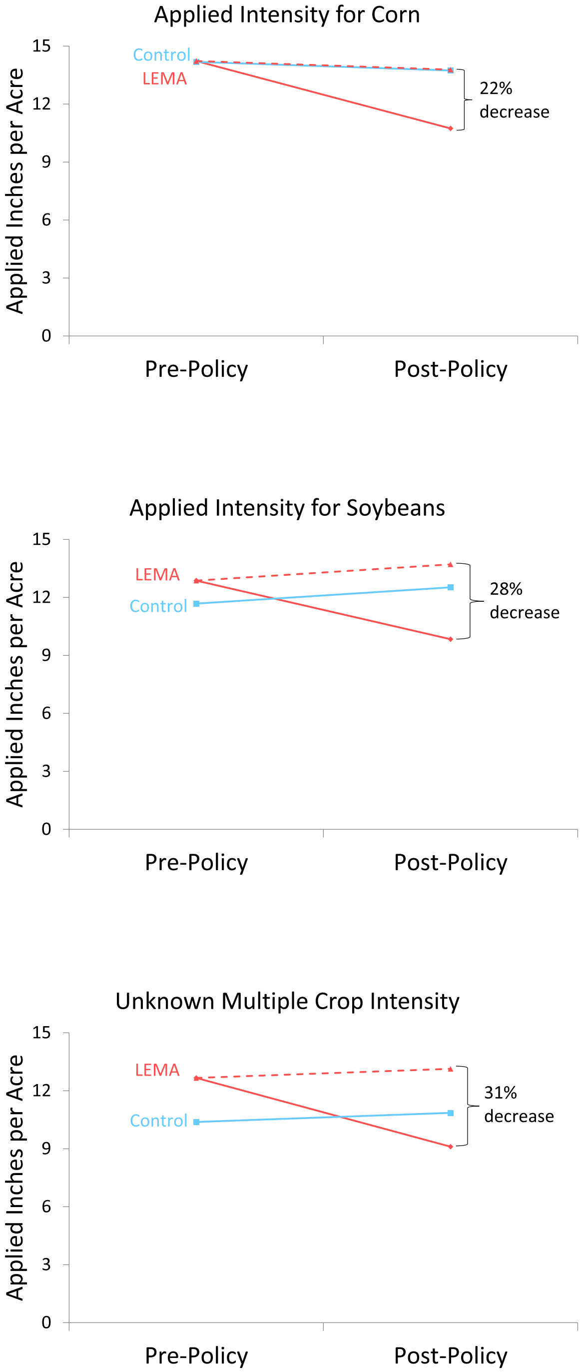

Fig. 5. Difference-in-difference results for average applied water intensity by crop.

responded by reallocating water use or production of corn into the control region. This type of behavior would imply some sort of constraint—such as a credit constraint—such that farmers could only produce a fixed amount of corn and they reallocate the location of the corn due to the restriction. We think that these constraints are unlikely to have a significant impact on our results because farmers tend to make irrigation decisions based on field-specific conditions. Furthermore, this concern would likely lead us to overestimate the importance of crop switching in reducing water use (i.e., the indirect intensive margin) but we find small impacts due to crop switching. Our main results are also supported by falsification tests to alleviate concerns that we find a spurious effect of the water restriction if the control is not a valid counterfactual.

The coefficient on a binary variable when the dependent variable is in log form cannot be directly interpreted as the relative change due to the treatment. Kennedy (1981) shows that transforming the coefficient estimate as $e^{\beta - 0.5Var(\beta)} - 1$ provides an estimate of the relative change in the outcome due to the treatment, where Var(β) is the variance of the respective coefficient. We include this correction for all estimates in the article.

6.1. Decomposing the marginal effects of water use

We estimate the total relative change in water use (β^{aw}), extensive margin response (β^{a}) and the intensive margin response (β^{w}) from the following regressions:

$$\ln(a_{ift}w_{ift}) = \alpha_i^{aw} + \lambda_{ft}^{aw} + \beta^{aw}D_{ift} + \boldsymbol{\theta}^{aw}\boldsymbol{Z}_{ift} + \varepsilon_{ift}^{aw}, \tag{6}$$

$$\ln(a_{ift}) = \alpha_i^{a} + \lambda_{ft}^{a} + \beta^{a}D_{ift} + \boldsymbol{\theta}^{a}\boldsymbol{Z}_{ift} + \varepsilon_{ift}^{a}, \tag{7}$$

$$\ln(w_{ift}) = \alpha_i^{w} + \lambda_{ft}^{w} + \beta^{w}D_{ift} + \boldsymbol{\theta}^{w}\boldsymbol{Z}_{ift} + \varepsilon_{ift}^{w}, \tag{8}$$

where $\ln(a_{ift}w_{ift})$ is the log of total water use, $\ln(a_{ift})$ is the log of irrigated acres and $\ln(w_{ift})$ is the log of applied water per acre. In cases where total water use was zero, we set total water use and acres irrigated equal to one before taking the log.[3] The total effect on water use in response to the policy is approximately equal to the sum of the extensive and intensive margin responses ($\beta^{aw} \approx \beta^{a} + \beta^{w}$).[4]

We decompose the intensive margin response into a direct and indirect effect by either holding land use constant or omitting land use (Hendricks and Peterson, 2012).[5] The direct intensive margin is estimated from the following regression:

$$\ln(w_{ift}) = \alpha_i^{ww} + \lambda_{ft}^{ww} + \beta^{ww}D_{ift} + \boldsymbol{\theta}^{ww}\boldsymbol{Z}_{ift} + \rho\boldsymbol{S}_{ift} + \varepsilon_{ift}^{ww} \tag{9}$$

where $\boldsymbol{S}_{ift}$ is a vector of variables indicating the share of irrigated acreage for each crop. The indirect intensive margin response is recovered by subtracting the indirect intensive margin from the total intensive margin response, $\beta^{ws} = \beta^{w} - \beta^{ww}$. We estimate the standard error of the indirect intensive margin response using a cluster bootstrap routine with 1000 replications.

6.2. Event study

We examine how the effect of the water restriction varied over time using an event study methodology. For each outcome of interest, we estimate the regression

$$\ln(y_{ift}) = \alpha_i + \lambda_{ft} + \sum_{m=2008}^{2016} \beta_m D_{ift}^{m} + \boldsymbol{\theta}\boldsymbol{Z}_{ift} + \varepsilon_{ift}, \tag{10}$$

where $D_{ift}^{m} = 1$ if the LEMA restriction was applied to the water right and $t = m$. Therefore the estimate for β_m indicates how water use changed for water rights inside the boundary compared to the change in the control group in year m compared to a baseline year of 2007.

The event study helps us to validate our identification strategy since we should see no significant difference between irrigators in the LEMA and in the control group before the policy was implemented. In other words, β_m should equal 0 for years 2008–2012. The results are also insightful to understand how farmers adapted to the water restriction over time.

7. Econometric results

Table 2 reports the set of decomposed marginal effects of water use from the preferred fixed effects regression with the LEMA effect adjusted for the log-linear form where the total, extensive, intensive, and direct intensive estimates are derived by estimating equations (6)–(9). We find that the water restriction imposed by the LEMA reduced total water use by 26.4% compared to water use in the control group. Most of the reduction in water use (25.5% reduction) occurred at the intensive margin rather than the extensive margin (2.2% reduction), and the extensive margin response is statistically insignificant.

---

[3] Results are not sensitive to using different values to replace zero. Our preferred specification gives a total elasticity of −0.264 and an extensive margin elasticity of −0.022. Replacing 0 with 0.1 gives a total elasticity of −0.270 and an extensive margin elasticity of −0.029. Replacing 0 with 2 gives a total elasticity of −0.263 and an extensive margin elasticity of −0.020.

[4] The only reason that $\beta^{aw} \neq \beta^{a} + \beta^{w}$ in our empirical results is that we have a different number of observations in the regressions due to some observations with zero water use. We cannot estimate equation (8) for observations with no water use.

[5] Intuitively, omitting land use as in equation (8) gives the change in intensity due to changes in intensity holding constant land use (direct intensive) and changes in intensity through changes in crops (indirect intensive) because it does not hold constant crops. Hendricks and Peterson (2012) show this more formally by applying the omitted variable bias formula.

Table 2
Decomposed estimates of water use response to restriction.

Variable	Total	Extensive	Intensive	Direct	Indirect
LEMA policy effect	−0.264**	−0.022	−0.255**	−0.209**	−0.045*
	(0.050)	(0.057)	(0.043)	(0.050)	(0.024)
Crop					
Alfalfa				0.085	
				(0.088)	
Corn				0.087	
				(0.045)	
Sorghum				−0.293**	
				(0.086)	
Soybeans				−0.045	
				(0.047)	
Wheat				−0.506**	
				(0.112)	
Other Crop				−0.381**	
				(0.146)	
Weather					
Precipitation	−0.049**	−0.014	−0.035**	−0.041**	
	(0.017)	(0.016)	(0.012)	(0.013)	
Evapotranspiration	−0.248	−0.312	0.002	0.045	
	(0.282)	(0.306)	(0.087)	(0.079)	
Observations	3592	3592	3411	3247	
R^2	0.624	0.546	0.712	0.747	

Note: Standard errors in parentheses are clustered at the water right level. LEMA effect estimates are adjusted for the log-linear correction. All specifications include water right and farmer-time fixed effects. The standard error for the indirect intensive margin is computed using a cluster bootstrap. * and ** denote significance at the 5% and 1% levels.

Estimates conditional on the cropping type give the direct intensive margin response. We estimate that conditional on the same cropping type, farmers reduced water use by 20.9%. Therefore, the change in water use due to changes in crops was only a 4.5% reduction.

The results in Table 2 indicate that farmers in the LEMA found it optimal to reduce water use mostly by reducing the intensity of irrigation. As discussed in our conceptual model, this implies that adjusting water use along the intensive margin gives a smaller reduction in net returns than adjusting along the extensive margin. Assuming that irrigated acreage has constant returns to scale, then our results imply that a reduction in water use of 26.4% was achieved while reducing irrigated net returns by less than 26.4%. This result implies diminishing returns to water application and is consistent with numerical simulations of Foster et al. (2014) and Wibowo et al. (2017) that farmers respond to reduced water availability by first reducing water use at the intensive margin. Note, however, that if the restriction became increasingly stringent, farmers could switch to reducing water use along the extensive margin rather than the intensive margin if the loss in crop yields from decreased intensity begins to decrease rapidly due to further reductions in intensity.

Other coefficient estimates from the regression specification are reasonable. Irrigation water use on corn is 8.7% larger than when unknown multiple crops are produced, but not a statistically significant difference. Water use is also similar for alfalfa and soybeans. Water use on sorghum, wheat, and other crops are significantly smaller. While we do not know what is included when the farmer lists multiple unknown crops, the regression results indicate that switching to multiple crops from corn does not give significant reductions in water use. The coefficients on precipitation for total water use and intensity are negative as expected. The effect of changes in precipitation from year-to-year are captured by the farmer-year fixed effects, so there is relatively little variation in precipitation remaining to exploit. Therefore, it is not surprising that the coefficients on precipitation are small and the coefficients on ET are not significant.

Table 3 shows results for the effect of the restriction on cropping patterns. Each column in Table 3 represents a separate regression with the share of irrigated area for the respective crop as the dependent variable and the regression specifications are the same as the previous regressions. The policy mainly caused farmers to switch from corn to sorghum, other crops, and unknown multiple crops.[6] Note that the coefficient in the corn equation likely overstates the reduction in corn area because corn production is likely one of the crops planted when the farmer indicated multiple crops. While the impact of switching crops on total water use is small—as indicated in Table 2—there was still significant changes in cropping patterns that have important implications for agribusiness firms in the local community. Corn production uses more fertilizer and chemical inputs and produces a greater volume of grain than the alternative crops. Therefore, less corn production in the short run is likely to have negative impacts on local agribusiness firms in the short run. If the water restriction was implemented on a larger scale, the reduction in corn area could also have negative impacts on cattle feeding and ethanol plants in the region.

[6] The effect of the policy on the area of unknown multiple crops is significant at the 10% level with a p-value of 0.053.

Table 3
Effect of water restriction on share of irrigated acres in each crop.

	(1) Alfalfa	(2) Corn	(3) Sorghum	(4) Soybeans	(5) Wheat	(6) Other	(7) Multiple
LEMA policy effect	−0.024	−0.162**	0.085**	−0.027	0.011	0.044*	0.073
	(0.017)	(0.052)	(0.028)	(0.046)	(0.041)	(0.022)	(0.038)
Observations	3248	3248	3248	3248	3248	3248	3248
R^2	0.652	0.583	0.593	0.542	0.613	0.566	0.582

Note: The dependent variable is the share of irrigated acres planted to the crop indicated in the respective column heading. Standard errors in parentheses are clustered at the water right level. All specifications include water right fixed effects, farmer-time fixed effects, precipitation, and evapotranspiration as controls. * and ** denote significance at the 5% and 1% levels.

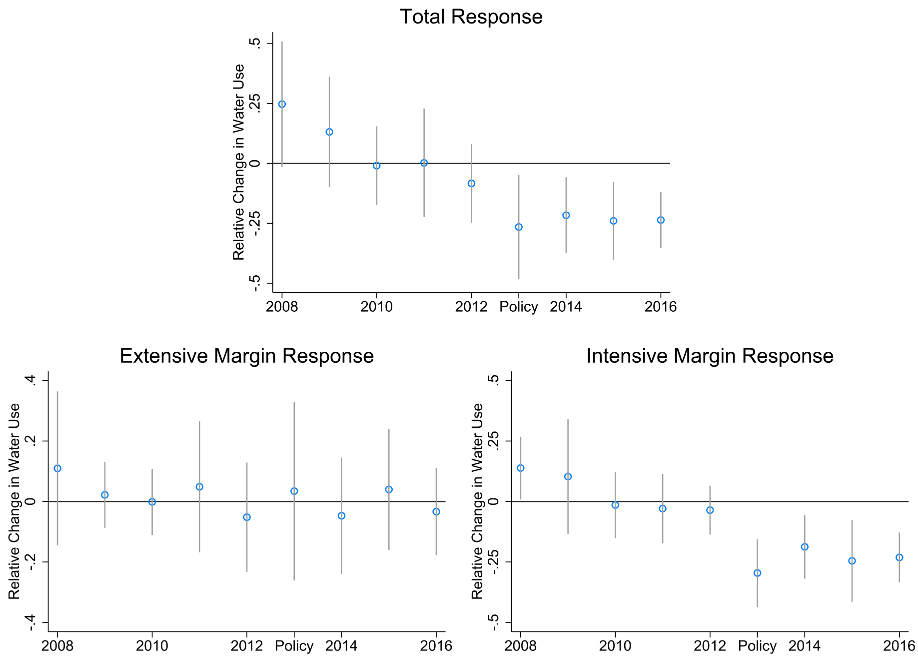

Fig. 6. Event study results.

7.1. Event Study Results

Fig. 6 presents results from the event study using equation (10). The coefficients for pre-policy impacts are not significantly different from zero supporting the validity of our identification strategy. The total change in water use due to the restriction is similar in each period after the policy came into effect. There appears to be a slightly larger response to the restriction in the first year—perhaps because farmers wanted to provide a buffer of water savings to use in future years—but the difference in response is not significantly different from other years after the policy. We also estimated the direct and indirect intensive margin responses with the event study (results not reported) and there was little change in the response along the indirect intensive margin. Results from the event study do not provide evidence that farmers changed their adaptation strategy as they learned over time how to manage the restriction.

7.2. Heterogeneous response by size of irrigators

A concern of some farmers prior to the LEMA being implemented was that a subset of farmers with multiple water rights inside the LEMA might have a less binding allocation because they could pool their allocation across constrained and unconstrained wells and effectively have a less binding allocation (KDA, 2013). For example, if a well had a well capacity (i.e., the rate at which water can physically be extracted) that limited extraction so that the quota was not constraining, then the farmer could reallocate the unused authorized quantity from this water right to other water rights in the LEMA that had a better well capacity. This concern is similar to the concern raised by Levinson (1997) and Oates et al. (1989) that trading can undermine the reduction in use by a quota when some users are unconstrained by the quota. In the LEMA, farmers with more water rights would be more likely to exploit the pooling of allocations with low and high capacity wells, so that farmers with more water

Table 4
Effect of water restriction by number of water rights inside LEMA.

Variable	Total	Extensive	Intensive
LEMA policy effect	−0.265**	−0.083	−0.214**
	(0.083)	(0.107)	(0.062)
LEMA × Large	−0.016	0.124	−0.110
	(0.140)	(0.146)	(0.103)
Observations	3592	3592	3411
R^2	0.624	0.546	0.712

Note: Standard errors in parentheses are clustered at the water right level. LEMA effect estimates are adjusted for the log-linear correction. All specifications include water right and farmer-time fixed effects. * and ** denote significance at the 5% and 1% levels.

Table 5
Total effect of restriction with alternative specifications.

Variable	(1)	(2)	(3)	(4)
LEMA policy effect	−0.408**	−0.371**	−0.315**	−0.264**
	(0.041)	(0.043)	(0.049)	(0.050)
Precipitation		−0.048**		−0.049**
		(0.007)		(0.017)
Evapotranspiration		0.002		−0.248
		(0.055)		(0.282)
Fixed Effects	Water Right Time	Water Right Time	Water Right Farmer-Time	Water Right Farmer-Time
Observations	3592	3592	3592	3592
R^2	0.137	0.148	0.622	0.624

Note: The dependent variable in each column is the total water use of the water right in a given year—the same as the total response in Table 2. Standard errors in parentheses are clustered at the water right level. LEMA effect estimates are adjusted for the log-linear correction. * and ** denote significance at the 5% and 1% levels.

rights have a smaller reduction in use due to the LEMA.[7]

We evaluate this concern by including an interaction term between the treatment variable D_{it} and a dummy variable that indicates if a farmer operated 2 or more water rights inside the policy boundary. The results that include the interaction term are reported in Table 4. The coefficient on the interaction term (second row) is statistically insignificant for all margins of response. We also conducted estimations by defining large with different levels of well ownership (e.g., greater than 3 water rights) and found that this did not change the significance of the estimates or the interpretation of the results.

7.3. *Alternative specifications*

If we assume that unobserved differences across farmers are constant across time, then the addition of a farmer specific fixed effect would suffice. However, in our fixed effects model this is already accounted for in the water right fixed effect α_i. We suggest the need to additionally account for unobserved differences across farmers that are varying through time.

Table 5 shows the importance of inclusion of additional controls in the model. Column (1) shows the DID estimates that only include a water right (α_i) and time (δ_t) fixed effects without weather controls. Column (2) additionally includes a set of weather controls that vary across water rights and time. Column (3) includes the set of farmer-time fixed effects (λ_{ft}) but omits the weather controls. Column (4) represents the preferred model with both the farmer-time fixed effects and weather controls.

Including farmer-time fixed effects reduces the estimated impact of the policy. Omitting the farmer-time fixed effects indicates the LEMA policy reduced water use by 40.8%, whereas our preferred estimate is only 26.4%. These results suggest that omitting farmer-time fixed effects introduces downward bias in our estimator for the impact of the policy on water use. We conduct a statistical test for omitted variable bias (see Pei et al., 2017, p. 17), where we effectively compare the coefficient on the treatment variable when we include farmer-time fixed effects versus only time fixed effects and reject the null hypothesis of no difference (p-value = 0.08). This result supports including farmer-time fixed effects to account for changes in management practices over time for different farmers and to allow common shocks to have differing impacts across farmers. Adding the weather controls has a smaller impact on the estimates—compare columns (1) and (2) or columns (3) and (4). The small impact of the weather controls is due to the fact that the policy area is relatively small with little variability in weather across the region for a given year and the impact of weather across years is captured in the time fixed effects.

[7] The ability to transfer authorized quantities between water rights before the LEMA was difficult and most farmers were not seriously constrained by the authorized quantity of the water right, so this type of behavior was not a concern prior to the LEMA restriction.

Table 6
Falsification tests for the total effect of restriction.

Variable	(1)	(2)	(3)	(4)
False policy effect	−0.065	−0.075	−0.118	−0.110
	(0.088)	(0.083)	(0.069)	(0.074)
Observations	3673	2123	2123	2123
R^2	0.758	0.610	0.611	0.610

Note: Standard errors in parentheses are clustered at the water right level. LEMA effect estimates are adjusted for the log-linear correction. All specifications include water right fixed effects, farmer-time fixed effects, precipitation, and evapotranspiration as controls. ∗ and ∗∗ denote significance at the 5% and 1% levels.

Table 7
Robustness checks.

Specification	Total	Extensive	Intensive	Direct	Indirect
Control Group 1–4 miles	−0.302∗∗	−0.040	−0.281∗∗	−0.249∗∗	−0.032∗
	(0.037)	(0.046)	(0.033)	(0.037)	(0.016)
Control Group 2–7 miles	−0.296∗∗	−0.055	−0.264∗∗	−0.230∗∗	−0.034
	(0.047)	(0.057)	(0.043)	(0.046)	(0.019)
Control Group 3–6 miles	−0.299∗∗	−0.021	−0.281∗∗	−0.251∗∗	−0.030
	(0.050)	(0.060)	(0.049)	(0.054)	(0.024)
Exclude Extreme Precip Years	−0.277∗∗	−0.034	−0.257∗∗	−0.211∗∗	−0.046
	(0.058)	(0.065)	(0.048)	(0.058)	(0.026)

Note: Standard errors in parentheses are clustered at the water right level. LEMA effect estimates are adjusted for the log-linear correction. All specifications include water right fixed effects, farmer-time fixed effects, precipitation, and evapotranspiration as controls. The standard error for the indirect intensive margin is computed using a cluster bootstrap. ∗ and ∗∗ denote significance at the 5% and 1% levels.

7.4. *Falsification tests*

Table 6 reports results from falsification tests. The purpose of the falsification tests is to evaluate whether our identification strategy gives us a spurious result or represents a true causal impact. The falsification tests create false "treatments" and we estimate the model to see if we obtain an insignificant effect as expected.

Results in column (1) of Table 6 compare water rights within 2–5 mile of the LEMA to water rights within 5–8 miles of LEMA and creates a false treatment in 2013–2016 in the 2 to 5 mile area. Results in column (1) omit the observations from inside the LEMA boundary. We could obtain a false effect in column (1) if water rights in different regions change water use differently in the post-policy period. Results in column (2) restrict the original data to contain only years between 2007 and 2012—prior to the actual LEMA policy in 2013—and create a false policy effect in 2009 inside the LEMA boundary. Column (3) is the same as column (2) except the false policy effect is in 2010 and column (4) creates a false policy effect in 2011. A significant result in columns (2)–(4) would suggest that water rights in the LEMA change water use over time differently than in the control group even in the absence of the policy. We estimate statistically insignificant results in all falsification tests, giving validity to our identification strategy.

7.5. *Robustness checks*

Table 7 shows results using different control groups as a robustness check. Each coefficient in the table is from a separate regression. The columns indicate the response to the LEMA policy along the different margins of response. The rows represent different definitions of the control group. The first three rows use different distances from the LEMA to define the control group. The last row removes observations with the largest annual precipitation (i.e., 2009) and the smallest precipitation (i.e., 2012) from the data. Estimates of the change in water use due to the policy and that most of the response was along the intensive margin are robust across all specifications.

8. Implications and conclusion

We provide robust econometric estimates of the impact of a water restriction in Kansas imposed through local governance. Our results show that farmers reduced water use by 26%. Most of the reduction in water use occurred due to reducing the intensity of water use given the same cropping patterns. However, there were also small, but significant, reductions in water use due to changes in cropping patterns. We find no statistically significant effect of the restriction on irrigated area.

In August 2017, the LEMA was renewed for another five-year allocation of 55 inches per authorized acre for the period 2018–2022. The new LEMA order allows farmers to carry-over a maximum of 5 inches per authorized acre of unused allocation from the initial allocation period. Farmers likely anticipated the potential to carry-over unused allocation based on previous meetings so farmers may have reduced water use more in the first 5-year period than they will in subsequent periods in order to accumulate a greater stock of allocated water. Another behavior that we may see in 2017—the last year of the first allocation—is that farmers with more than 5 inches of unused allocation will use more water in 2017 than previous years to avoid effectively losing allocated water assuming that the carry-over amount was known during the 2017 irrigation season.

The fact that most of the response to the water restriction occurred along the intensive margin is highly relevant to policymakers. Water rights in Kansas are assigned according to the prior appropriation doctrine. According to prior appropriations, a junior water right may not use water if it impairs the use of water by senior water rights. Defining an "impairment" is difficult in groundwater systems, but a strict application of prior appropriations would eliminate irrigation on the most junior water rights in the area. An upper bound estimate of the economic cost of applying the prior appropriation doctrine (i.e., the cost with no trading) to achieve the same reduction in water use is the welfare loss of reducing irrigated area by 26%.[8] Since farmers responded to the restriction primarily along the intensive margin, our econometric results imply that the LEMA was able to achieve the reduction in water use with a smaller loss in welfare than reducing irrigated acreage by 26%.

We can get a rough estimate of the short-run welfare impact of reducing irrigated acreage as the difference between irrigated and nonirrigated cash rental rates. Recent cash rental rates in Thomas County provide a reasonable estimate of rental rates in the LEMA region.[9] Average (2013–2017) cash rental rates in Thomas County were $167/acre for irrigated and $59/acre for nonirrigated according to National Agricultural Statistics Service data. Assuming that land converted from irrigated to nonirrigated has average productivity, a reduction in irrigated acreage would result in an average welfare loss of roughly $28.08 per acre initially irrigated ($(167-59) \times 0.26$). This provides an estimate of the short-run cost to the government of an annual water right buyout program and an upper bound on the short-run welfare loss from applying prior appropriations.

The spillovers of irrigation on other sectors of the economy have also received substantial interest (e.g., Hornbeck and Keskin, 2015). Our econometric results have important implications for the impact of the LEMA on agribusiness firms that sell inputs to farmers. Unfortunately, we do not have data on the use of inputs, but we can still get a rough estimate using our estimates of changes in cropping patterns and information on expenditures from Kansas State University irrigated crop budgets. One challenge in obtaining a good estimate is that we do not know which crops were planted when the farmer indicated multiple crops. For the sake of getting a rough estimate of changes in expenditures, we allocate multiple crops as follows: 50% corn, 20% soybeans, 10% wheat, 10% sorghum, and 10% other.[10] We calculate the change in expenditures for fertilizer, herbicide, insecticide, and seed using the pre-treatment average share of each crop and the change in the share of each crop from Table 3. Our results indicate a 14% reduction in chemical and seed expenditures on irrigated land within the LEMA.[11] These calculations assume that expenditures for corn did not change due to the restriction but farmers may have reduced their expenditures along with the reduction in intensity (e.g., a lower seed population and less fertilizer). Therefore, the reduction in expenditures may have been larger than our estimate. So although changes in cropping patterns had a relatively small impact on water savings, the change did have a substantial impact on chemical and seed expenditures.[12] It is also worth noting that the reduction in fertilizer and chemical usage could have benefits through reduced water pollution.

The effectiveness of the LEMA at achieving a reduction in groundwater extraction with at least general support among farmers is noteworthy. Pfeiffer and Lin (2014a) find that previous cost-share programs to upgrade to more efficient irrigation technologies may have actually increased groundwater extraction. In principle, a tax on water could achieve a similar reduction in water use (e.g., Johansson et al., 2002), but is likely to be highly unfavorable to farmers unless tax revenues are returned to farmers through lump sum transfers. Using a demand elasticity of −0.5 as estimated by Mieno and Brozović (2016), achieving a 26% reduction in water use would impose a tax burden on farmers of roughly $16.30/acre ($17.78/acre-foot).[13] Using a demand elasticity of −0.1 as estimated by Hendricks and Peterson (2012), the tax policy would impose a tax burden of roughly $81.51/acre ($88.92/acre-foot). These predicted taxes are comparable to the tax of $75 per acre-foot imposed in the Rio Grande Water Conservation District studied by Smith et al. (2017). This tax burden would be in addition to any reduction in returns due to reduced water use. Although the overall welfare impact on society is similar with the tax or quota, the distributional consequences are dramatically different. The tax provides benefits to taxpayers at a cost to producers, whereas the quota provides no

[8] Water right trading is minimal in this region of Kansas because the legal requirements to transfer water create large transaction costs. If transaction costs of water trading were reduced, then an application of prior appropriations could achieve the same result as the LEMA, at least in theory, but the distributional impacts would have been very different—farmers with junior water rights would have incurred all the loss in the short run instead of all farmers sharing the loss.

[9] There are 8 water rights in Thomas County in the LEMA out of a total of 786 water rights in Thomas County. Therefore, the LEMA should have a negligible impact on average cash rental rates for Thomas County.

[10] This distribution of crops is based roughly on the observed pattern of cropping (74% corn, 17% soybeans, 4% wheat, 3% sorghum, and 1% other), but we assume a smaller share of land planted to multiple crops is planted to corn since farmers likely diversify the cropping pattern in part to reduce water use.

[11] We include alfalfa with other crops and we assume the expenditures on other crops are the average of wheat and sorghum expenditures. If we assume that multiple crops are comprised of less corn (25% corn, 10% soybeans, 25% wheat, 25% sorghum, and 15% other), the reduction in chemical and seed expenditures is 16%. If we assume that multiple crops are comprised of the same pattern as the rest of the irrigated land (74% corn, 17% soybeans, 4% wheat, 3% sorghum, and 1% other), the reduction in chemical and seed expenditures is 12%.

[12] It is also possible that the reduction in chemical usage had some positive environmental impacts, but we do not attempt to quantity these impacts.

[13] We calculate the tax burden as

$$\frac{Relative\,Reduction\,Water\,Use}{Demand\,Elasticity} \times \$/Acre - Inch\,to\,Extract \times Quantity\,Extracted.$$

The only cost of pumping groundwater in Kansas is the cost of extraction. We assume it costs $2.85 per acre-inch to extract the water according to recent irrigated crop budgets for Northwest Kansas by Kansas State University. We assume the quantity extracted is 11 inches per acre as imposed by the restriction.

revenues to taxpayers with a smaller cost to producers. However, a tax could be viewed favorably by farmers if tax revenues were returned to farmers.

Our estimates are highly relevant to ongoing policy discussions in Kansas. GMD 4 in northwest Kansas has approved a LEMA in early 2018 that would cover the entire district—an area the size of roughly 4 counties or 4852 mi^2. GMD 1 in west central Kansas has also been in discussions to form a district-wide LEMA and a county within GMD 1 is seeking a purely voluntary reduction in water use. An area within GMD 3 of Southwest Kansas is also working on proposing a LEMA. Our results are also of interest for the management of other kinds of resources where it is important to understand how users adapt to restrictions on the use of the resource.

References

Ayres, A.B., Edwards, E.C., Libecap, G.D., 2018. How transaction costs obstruct collective action: The case of California's groundwater. J. Environ. Econ. Manag. 91, 46–65.
Brozović, N., Sunding, D.L., Zilberman, D., 2010. On the spatial nature of the groundwater pumping externality. Resour. Energy Econ. 32, 154–164.
Burness, H.S., Quirk, J.P., 1979. Appropriative water rights and the efficient allocation of resources. Am. Econ. Rev. 69, 25–37.
Coase, R.H., 1960. The problem of social cost. J. Law Econ. 3, 1–44.
Demsetz, H., 1967. Toward a theory of property rights. Am. Econ. Rev. 57, 347–359.
Edwards, E.C., 2016. What lies beneath? Aquifer heterogeneity and the economics of groundwater management. J. Assoc. Environ. Resour. Econ. 3, 453–491.
English, M., 1990. Deficit irrigation. I: analytical framework. J. Irrigat. Drain. Eng. 116, 399–412.
Famiglietti, J., 2014. The global groundwater crisis. Nat. Clim. Change 4, 945–948.
Foster, T., Brozovi, N., Butler, A.P., 2017. Effects of initial aquifer conditions on economic benefits from groundwater conservation. Water Resour. Res. 53, 744–762.
Foster, T., Brozović, N., Butler, A.P., 2014. Modeling irrigation behavior in groundwater systems. Water Resour. Res. 50, 6370–6389.
Gisser, M., Sanchez, D., 1980. Competition versus optimal control in groundwater pumping. Water Resour. Res. 16, 638–642.
Graveline, N., Merel, P., 2014. "Intensive and extensive margin adjustments to water scarcity in France's cereal belt. Eur. Rev. Agric. Econ. 41, 707–743.
Guilfoos, T., Khanna, N., Peterson, J.M., 2016. Efficiency of viable groundwater management policies. Land Econ. 92, 618–640.
Guilfoos, T., Pape, A.D., Khanna, N., Salvage, K., 2013. Groundwater management: the effect of water flows on welfare gains. Ecol. Econ. 95, 31–40.
Hendricks, N.P., Peterson, J.M., 2012. Fixed effects estimation of the intensive and extensive margins of irrigation water demand. J. Agric. Resour. Econ. 37, 1–19.
Hornbeck, R., Keskin, P., 2015. Does agriculture generate local economic spillovers? Short-run and long-run evidence from the ogallala aquifer. Am. Econ. J. Econ. Pol. 7, 192–213.
Hornbeck, R., Keskin, P., 2014. The historically evolving impact of the ogallala aquifer: agricultural adaptation to groundwater and drought. Am. Econ. J. Appl. Econ. 6, 190–219.
Johansson, R.C., Tsur, Y., Roe, T.L., Doukkali, R., Dinar, A., 2002. Pricing irrigation water: a review of theory and practice. Water Pol. 4, 173–199.
KDA, 2013. Order of Designation Approving the Sheridan 6 Local Enhanced Management Area within Groundwater Management District No. 4. http://dwr.kda.ks.gov/LEMAs/SD6/LEMA.SD6.OrderOfDesignation.20130417.pdf.
Kennedy, P.E., 1981. Estimation with correctly interpreted dummy variables in semilogarithmic equations. Am. Econ. Rev. 71, 801.
Koundouri, P., 2004. Potential for groundwater management: Gisser Sanchez effect reconsidered. Water Resour. Res. 40.
Levinson, A., 1997. Why oppose TDRs?: Transferable development rights can increase overall development. Reg. Sci. Urban Econ. 27, 283–296.
Libecap, G.D., 2011. Institutional path dependence in climate adaptation: coman's "some unsettled problems of irrigation". Am. Econ. Rev. 101, 64–80.
Lin Lawell, C.Y.C., 2016. The management of groundwater: irrigation efficiency, policy, institutions, and externalities. Annu. Rev. Resour. Econ. 8, 247–259.
Mieno, T., Brozović, N., 2016. Price elasticity of groundwater demand: attenuation and amplification bias due to incomplete information. Am. J. Agric. Econ. 87, 401–426.
Moore, M., Gollehon, N., Carey, M., 1994. Multicrop production decisions in western irrigated agriculture: the role of water price. Am. J. Agric. Econ. 76, 859–887.
Oates, W.E., Portney, P.R., McGartland, A.M., 1989. The *net* benefits of incentive-based regulation: a case study of environmental standard setting. Am. Econ. Rev. 79, 1233–1242.
Ostrom, E., 1990. Governing the Commons. Cambridge University Press, New York, NY.
Pei, Z., Pischke, J.S., Schwandt, H., 2017. Poorly Measured Confounders Are More Useful on the Left than on the Right. NBER Working Paper Series, Working Paper 23232. National Bureau of Economic Research.
Peterson, J.M., Ding, Y., 2005. Economic adjustments to groundwater depletion in the high Plains: do water-saving irrigation systems save water? Am. J. Agric. Econ. 87, 147–159.
Pfeiffer, L., Lin, C., 2014a. Does efficient irrigation technology lead to reduced groundwater extraction? Emperical evidence. J. Environ. Econ. Manag. 67 (2), 189–208.
Pfeiffer, L., Lin, C.Y.C., 2012. Groundwater pumping and spatial externalities in agriculture. J. Environ. Econ. Manag. 64, 16–30.
Pfeiffer, L., Lin, C.Y.C., 2014b. The effects of energy prices on agricultural groundwater extraction from the high Plains aquifer. Am. J. Agric. Econ. 96, 1349–1362.
Rogers, D.H., Powell, G.M., Ebert, K., 2013. Water Primer: Part 5, Water Law. K-State Research and Extension, p. MF3024.
Scanlon, B., Faunt, R., Claudia, C., Longuevergne, L.R., Alley, R.C., McGuire, W.M., McMahon, V.L., Peter, B., 2012. Groundwater depletion and sustainability of irrigation in the US high Plains and Central Valley. Proc. Natl. Acad. Sci. Unit. States Am. 109, 9320–9325.
Schoengold, K., Sunding, D.L., Moreno, G., 2006. Price elasticity reconsidered: panel estimation of an agricultural water demand function. Water Resour. Res. 42.
Smith, S.M., Andersson, K., Cody, K., Cox, M., Ficklin, D., 2017. Responding to a groundwater crisis: the effects of self-imposed economic incentives. J. Assoc. Environ. Resour. Econ. 4, 985–1023.
Stotler, R., Butler Jr., J., Buddemeier, R., Bohling, G., Comba, S., Jin, W., Reboulet, E., Whittemore, D., Wilson, B., 2011. High Plains Aquifer Calibration Monitoring Well Program: Fourth Year Progress Report. KGS Open File Report No. 2011-4. Kansas Geological Survey.
Wang, C., Nair, S., 2013. The economics of deficit irrigation. Nat. Resour. Model. 26, 331–364.
Wibowo, R., Hendricks, N., Kisekka, I., Araya, A., 2017. Using a crop simulation model to understand the impact of risk aversion on optimal irrigation management. Trans. ASABE 60, 2111–2122.
Wiggins, S.N., Libecap, G.D., 1985. Oil field unitization: contractual failure in the presence of imperfect information. Am. Econ. Rev. 75, 368–385.