



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

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

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

Help ensure our sustainability.

Give to AgEcon Search

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

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

Estimating Agricultural Acreage Responses to Input Prices: Groundwater in California

Andrew W. Stevens, Mississippi State University, stevens@berkeley.edu

***Selected Paper prepared for presentation at the 2017 Agricultural & Applied Economics Association
Annual Meeting, Chicago, Illinois, July 30-August 1***

Copyright 2017 by Andrew W. Stevens. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

Estimating Agricultural Acreage Responses to Input Prices: Groundwater in California*

Andrew W. Stevens[†]

May 24, 2017

Abstract

How does agricultural land use respond to variation in the price of agricultural water? Answering this question is difficult in California where there is no well-functioning market for water. To overcome this challenge, I use variation in groundwater depth over space and time to proxy for the price of water. This makes sense in a setting where groundwater pumping is unregulated, meaning the effective price of pumped groundwater is the energy cost to pump it. I construct a panel of agricultural fields in Fresno County, California from 2008 to 2016, and estimate a fixed effects model to estimate groundwater depth's effect on transition probabilities between different categories of land cover. I find that groundwater depth reduces the likelihood that parcels will be planted to an annual crop, but increases the likelihood of fallowing land. Groundwater depth seems to have a less profound effect on choosing to plant perennial crops.

*I am grateful to David McLaughlin for many helpful conversations, as well as to Nathan Hendricks for information on data sources. Any remaining errors are my own.
Preliminary draft. Please do not cite without permission.

[†]Department of Agricultural Economics, Mississippi State University, Box 5187, Mississippi State, MS 39762. E-mail: stevens@berkeley.edu

1 Introduction

Water is arguably the most important input in California agriculture, and its importance has been highlighted by the recent drought. Farmers and researchers both have long been interested in the marginal value of agricultural water and its impact on production. However, due to a patchwork of legal doctrines, historic water rights, and the absence of any reliable market for agricultural water, estimates of water's value in California agriculture have been challenging to come by (Buck et al., 2014). However, all agricultural producers in California have the option to pump groundwater as a source of last resort, and this pumping is currently unregulated. Producers who rely on groundwater use energy (electricity or fuel) to pump water up from an underlying aquifer. Therefore, the cost structure for groundwater is straightforward: the deeper the well, the more expensive the water.

In this paper, I exploit the insight that groundwater depth is an effective proxy for agricultural water costs on farms where groundwater pumping occurs. I use panel data on groundwater levels and field-specific land covers to estimate the effects of groundwater depth (and by extension the price of water) on land use decisions. I demonstrate that deeper groundwater levels decrease the likelihood of land being covered in annual crops, and increase the likelihood of land being left fallow or in grassland.

I am not the first to tie groundwater levels to water costs; authors of previous studies have had the same insight (Schoengold and Sunding, 2014; Green et al., 1996). However, I add to the extant literature by using groundwater's physical characteristics as a source of plausibly exogenous variation. The classic simplification that an aquifer is like a bathtub ignores important hydrological facts. In particular, lateral groundwater movement is slow and leads to a non-uniform water table over space. Thus, even though the entire central valley of California is part of a large aquifer system, different regions face differing well depths at any particular point in time. Simultaneously, lateral groundwater flow ensures that the groundwater depth at any one point is the result of aggregate groundwater pumping in the surrounding area, rather than the private pumping of a single landowner.

Using three distinct datasets, I compile a balanced panel of over eight thousand agricultural fields in Fresno County for the years 2008 through 2016. (See figure 1 for a map of Fresno County within California.) For each parcel of land, I observe that year’s land cover and a measure of groundwater depth from a nearby (less than five miles away) well. I then estimate an econometric model of the effect of groundwater depth on land cover that includes fixed effects for both individual parcels and different years. This approach controls for any time-invariant characteristics of individual parcels as well as any widely shared annual shocks to either groundwater levels or land cover.

Figure 1: Fresno County, CA



Note: This figure displays the extent of Fresno County within the state of California. Notably, it contains a significant portion of the agriculturally productive California Central Valley.

My identification strategy relies on the assumption that, conditional on the included fixed effects, variation in groundwater depth is as good as random. This is, perhaps unintuitively, a credible assumption in this setting. In particular, since aggregate regional pumping deter-

mines groundwater levels, and individual pumpers' impacts on aggregate pumping are quite small, it makes sense that observed groundwater levels are not determined by own-parcel land cover choices.

While my analysis ignores surface water, this omission will bias my findings toward zero and leave me with conservative estimated effect sizes. I plan to incorporate data on surface water rights in future iterations of this work.

Previous literature on water resources in California agriculture has focused in large part on the adoption of efficient irrigation technologies. Caswell and Zilberman (1986) developed a seminal theoretic framework relating land quality, well depth, electricity costs, and irrigation efficiency to technology adoption and production decisions. Dinar (1994) further explored such issues and widened the framework to include groundwater quality and other important agricultural characteristics. Green et al. (1996) applied microparameters at the field level to expand the empirical understanding of technology adoption behaviors. Unlike previous work that has focused on irrigation efficiency, this work instead explores how variations in (implicit) water prices affect crop choices and production decisions.

The paper is organized as follows. In section 2, I describe the sources of my data and provide summary statistics. In section 3, I explain my empirical strategy and econometric specification. In section 4, I present and discuss my results, and finally in section 5, I conclude.

2 Data

I utilize data from three main sources. First, I use the Cropland Data Layer (CDL) to determine land cover and crop choice. Next, I use Common Land Units (CLUs) to determine individual agricultural field boundaries. Finally, I use data from the California Department of Water Resources to determine the depth to groundwater at various monitored wells. I describe each of these data sources in the following subsections. Several of these descriptions

borrow heavily from Stevens (2015).


2.1 Cropland Data Layer

The Cropland Data Layer (CDL) is a raster dataset of landcover in the United States collected and maintained by the National Agricultural Statistics Service (NASS) of the USDA. A satellite records the electro-magnetic wavelengths of light reflected from different points on the earth's surface and uses a ground-tested algorithm to assign each pixel a single land-cover type for the year. Pixels measure 30 meters by 30 meters, except for years 2006–2009 when pixels measured 56 meters by 56 meters.¹ The CDL provides remarkably high-resolution land cover data and is able to distinguish between many different types of vegetation. Figure 2 displays the CDL for Fresno County in 2016. Within the agricultural region of the county, the most prominent land covers can be summarized as follows: grey pixels represent developed (urban) areas, purple pixels represent grapes, dark green pixels represent almonds, red pixels represent cotton, and pink pixels represent alfalfa.

One problem with using raw CDL data is that a 30 meter by 30 meter pixel is likely not the appropriate unit of analysis. Rather, economists are more interested in observing field-level crop choices. Additionally, while CDL data are quite accurate for primary row crops (Boryan et al., 2011), it is apparent that individual pixels are frequently mis-measured. For instance, upon visual inspection of a CDL image, it is not uncommon to observe what is clearly a large field of more than 100 pixels planted to one crop, with one or two pixels somewhere in the field reported as another crop. If analysis is conducted at the pixel level rather than the field level, such mis-measurements become a large concern. To address this concern, I exploit Common Land Unit data to construct field-level crop cover observations.

¹Data collection for the CDL began in the late 1990s in only three states. Data collection for California began in 2007.

Figure 2: Cropland Data Layer (CDL) – Fresno County, 2016



Note: This figure plots land cover for 30 meter by 30 meter pixels across Fresno County for the year 2016. Prominent land covers in the agricultural region of the county include grey for developed (urban) areas, purple for grapes, dark green for almonds, red for cotton, and pink for alfalfa. Source: NASS.

2.2 Common Land Unit


According to the Farm Service Agency (FSA) of the USDA, a Common Land Unit (CLU) is “an individual contiguous farming parcel, which is the smallest unit of land that has a permanent, contiguous boundary, common land cover and land management, a common owner, and/or a common producer association” (Farm Service Agency, 2012). Practically, a CLU represents a single agricultural field. Polygon shapefiles of CLUs are maintained by the FSA, but are not currently publicly available.

I obtain CLU data for California from the website GeoCommunity (<http://www.geo-comm.com>). These data contain shapefiles from the mid 2000s, before CLU data were removed from the public domain. In this research, I implicitly assume that individual CLUs do not change

over time: a reasonable assumption given the FSA definition. In reality, the FSA does adjust individual CLU definitions on a case-by-case basis if necessary, but I assume these adjustments to be negligible as in previous similar studies (Hendricks et al., 2014).

I overlay the CDL raster data with CLU polygons as shown in Figure 3. Upon visual inspection, the fit is quite good: CLU boundaries line up with crop changes in the CDL, CLU boundaries largely do not exist for non-agricultural areas, and geographical features such as waterways are visible. One concern is that many CLUs are quite small, and this is particularly pronounced in areas near urban sprawl. Therefore, to maintain confidence that the fields I study are actually “fields” in the way we think of them, I drop all CLUs from my dataset with areas of less than 5 acres.

Figure 3: Common Land Units – Fresno County

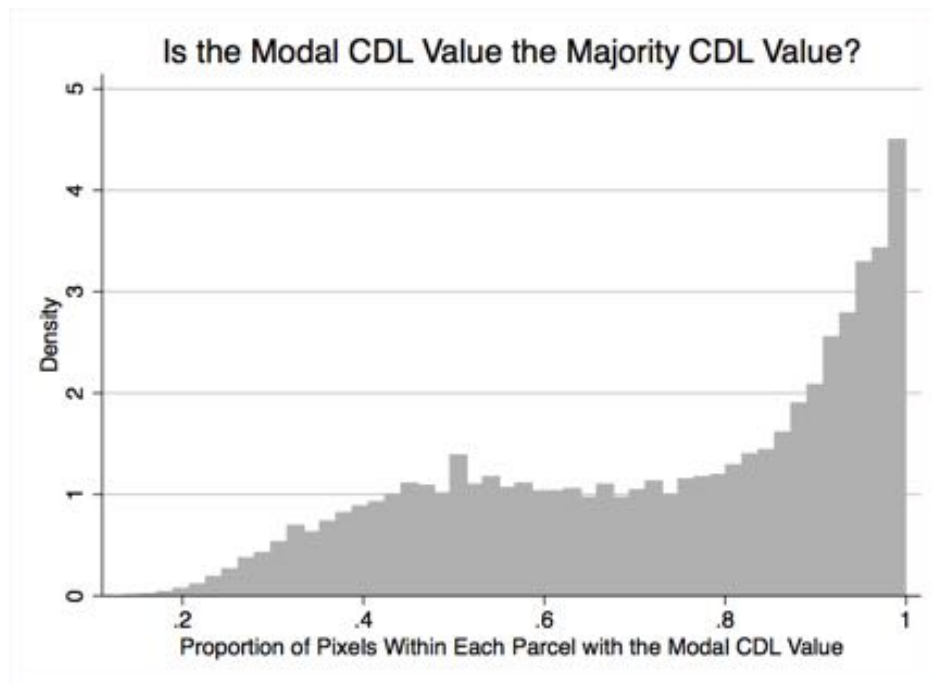


Note: This figure plots Common Land Unit outlines over the CDL data reported in figure 2. Source: GeoCommunity.

To assign each CLU a single crop cover, I follow Stevens (2015) and calculate the modal

value of the raster pixels contained within each CLU polygon. I then assign that modal value to the entire CLU. This procedure enforces the assumption that each field (CLU) is planted to a single crop. However, this is not strictly true. Figure 4 reports the proportion of modal values within each CLU in my final dataset. While many fields are dominated by their modal CDL value, there are some for which the modal CDL value is a minority value. Future work will assess the sensitivity of my results to this measure.

Figure 4: Modal CDL Values



Note: This figure plots a histogram of all CLU parcels in my final dataset, and reports the proportion of CDL pixels in each parcel that share the modal CDL value.


Finally, for each CLU polygon, I construct a centroid for the field. I then use these CLU centroids to calculate distances from each field to nearest well in my data.

2.3 Groundwater Depth

I obtain data on groundwater depth from the California Department of Water Resources. Specifically, I begin with the universe of well-measured groundwater levels available as of

March 2017.² I then restrict my data to only those wells in Fresno County that have at least annual readings dating back to at 2007. This leaves me with forty-seven (47) unique wells as shown in figure 5. I then calculate an annual average groundwater depth for each well, leaving me with a balanced panel of forty-seven wells with annual observations from 2007 to 2016. Notably, these wells include those in the CASGEM program (California Statewide Groundwater Elevation Monitoring) as well as other wells that voluntarily report data.

Figure 5: Well Locations



Note: This figure plots the location of the forty-seven wells used in my analysis. Source: CA Department of Water Resources.


2.4 Final Dataset and Summary Statistics

To construct the final dataset for use in my econometric analysis, I restrict my sample to only those CLU parcels within five (5) miles of a well. Figure 6 plots this subset of parcels. This

²See <http://www.water.ca.gov/waterdatalibrary/groundwater/index.cfm>

sample restriction prevents me from attributing groundwater readings from too far away to a particular field that may experience different local groundwater levels due to slow lateral groundwater flow. Subsequent work will further test the effects of using different cut-off points. I then match each CLU parcel to its nearest well, and use the annual readings from that well as a proxy for that parcel's true (unobserved) groundwater depth.

Figure 6: Final Dataset



Note: This figure plots the forty-seven Fresno County wells used in my analysis, as well as the Fresno County parcels no more than five miles from these wells. These are the parcels included in my econometric analysis.

Next, I classify each CLU parcel's land cover into one of seven categories: annual crop, perennial crop, water, developed (urban), forest or wetland, fallow or grassland, and missing or undefined. Then, for each year, I determine a parcel's land cover category in the previous year. This ultimately gives me a balanced panel of 8,804 agricultural fields with annual land cover observations from 2008 to 2016.

Table 1 summarizes the annual percentage of CLU parcels in each category of land cover

Table 1: Annual Aggregate Land Cover, Percent of Total

Land Cover	Year								
	2008	2009	2010	2011	2012	2013	2014	2015	2016
Annual Crop	32.16	36.74	36.95	35.09	33.41	32.61	29.40	22.57	23.36
Perennial Crop	39.97	25.41	29.04	41.98	41.30	46.10	45.26	45.90	46.43
Water	0.49	0.70	0.45	0.58	0.65	0.68	0.66	0.65	0.62
Developed (Urban)	2.27	2.67	1.90	2.04	1.98	1.93	2.92	2.76	2.70
Forest or Wetland	0.31	1.43	1.31	0.08	0.12	0.05	0.05	0.06	0.03
Fallow or Grassland	24.65	32.89	30.21	20.09	22.41	18.49	21.57	27.93	26.70
Missing or Undefined	0.16	0.15	0.14	0.14	0.14	0.14	0.15	0.14	0.14

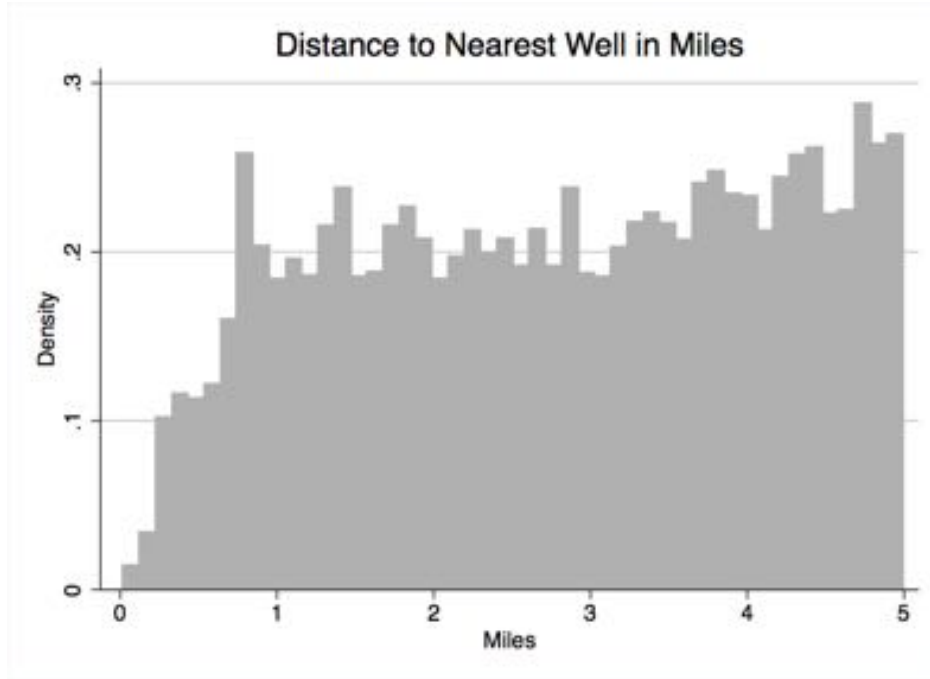
Note: This table records the proportion of CLU parcels in my final dataset with each of the eight categories of land cover for each year between 2008 and 2016. Note that the most common categories are clearly annual crops, perennial crops, and fallow or grassland.

from 2008 to 2016. The overall proportion of observations in each land cover category is relatively stable over time, but displays non-negligible year-to-year variation. The three most common land cover categories are annual crop, perennial crop, and fallow or grassland. Annually, more than ninety-five percent of CLU parcels are in one of these three categories. Therefore, in my subsequent analysis, I will focus on land use transitions between these three categories.

Figure 7 presents a histogram of the distance of each CLU parcel in my dataset to its nearest well. The distribution of distances is roughly uniform except for distances under one mile, which are less prevalent. This is encouraging evidence that distance-to-well is unlikely to drive my results in any systematic way.

Figure 8 summarizes groundwater depth readings over time for the forty-seven wells in my dataset. Several observations are worth noting. First, there is a wide range of groundwater depths within Fresno county, even in a single year. In 2015, for instance there is a nearly 500 feet difference between the deepest groundwater level and the most shallow, while the average depth is around 175 feet. Second, there is meaningful year-to-year variation in groundwater levels: the average annual depth fluctuates between about 150 and 175 feet. Third, from 2011 to 2016, the figure shows groundwater depth increasing for many wells. This fits with

Figure 7: Distance to Nearest Well




Note: This figure plots a histogram of all CLU parcels in my final dataset, and reports the distance from each parcel centroid to its nearest well in my dataset.

anecdotal observations that farmers relied on increased groundwater withdrawals during these years as California experienced a prolonged drought.

Finally, table 2 summarizes the unconditional probabilities of CLU parcels transitioning between annual crops, perennial crops, and fallow or grassland between any two years. Notably, this table does not control for any possible determinants of these transitions, and merely summarizes my dataset. In my empirical exercise, I will be estimating how groundwater depth affects these transitions probabilities.

Figure 8: Groundwater Depth Over Time



Note: This figure plots annual summaries of the groundwater depths measured at each of the forty-seven wells in my dataset.

Table 2: Unconditional Land Cover Transition Probabilities

Previous Land Cover	Current Land Cover		
	Annual Crop	Perennial Crop	Fallow or Grassland
Annual Crop	75.39	9.28	15.00
Perennial Crop	6.32	84.31	8.22
Fallow or Grassland	16.83	15.05	66.01

Note: This table records the unconditional probability of a CLU parcel having a particular land cover given its previous land cover. I focus on the three most common land covers: annual crop, perennial crop, and fallow or grassland. All numbers are percentages.

3 Empirical Strategy

My goal is to estimate the effect of groundwater depth on the probability that land cover transitions between any two particular categories. Conceptually, increased groundwater

depth results in more expensive water if that water is pumped from aquifers. Therefore, we would expect relatively deeper groundwater levels to cause farmers to transition from relatively-more-water-intensive land uses to relatively-less-water-intensive land uses. Between annual crops, perennial crops, and fallow or grassland, the third category is clearly the least water-intensive. Thus, we expect deep groundwater levels to increase transitions to fallow or grassland.

It is less clear, however, whether annual or perennial crops as a category are more water-intensive. A relevant concern here is the option value involved in this trade-off. For instance, an almond farmer with a relatively young orchard of trees has a strong incentive to keep her trees watered, even in a drought. However, at some point, an old and less productive orchard becomes less lucrative to water than an annual crop that doesn't require as much water. On the other hand, a farmer who currently farms an annual crop may balk at investing in a perennial crop when groundwater levels are sufficiently deep. In short, deep groundwater levels are likely to increase annual crop cover. However, it is unclear what effect they would have on perennial crop cover.

To estimate groundwater depth's effect on land cover transitions, I estimate the fixed effects model specified in equation 1 on different subsets of my data. In this specification, $LandCover_{it}$ is a dummy variable for a land cover category such as annual crop or perennial crop. Subscript i indexes different CLU parcels and subscript t indexes year. The variable $GroundwaterDepth_{it}$ represents the groundwater depth in feet as measured at the nearest well to field i in year t . I include a constant term β_0 , a CLU parcel fixed effect α_i , and a year fixed effect γ_t . The error term ε_{it} is clustered at the CLU parcel level to allow for correlation in a single field's land cover decisions over time.

$$LandCover_{it} = \beta_0 + \beta_1 GroundwaterDepth_{it} + \alpha_i + \gamma_t + \varepsilon_{it} \quad (1)$$

To clarify how I implement my empirical strategy, consider the following example. To determine the effect of groundwater depth on the transition probability from annual crop

cover to perennial crop cover, $LandCover_{it}$ would be the dummy variable $Perennial$ that takes on a value of 1 for parcel i in year t if it is in a perennial crop land cover in year t . I then estimate specification 1 on all observations in my data for which the dummy variable $Annual_{prior}$ is equal to 1. That is, I estimate the specification only on those parcel-year observations for which the land cover category in the previous year was annual crop.

In order to consider β_1 as a causal effect in my regressions, I rely on the identifying assumption that groundwater depth is as good as random after accounting for parcel and year fixed effects. More precisely, I assume groundwater depth for a particular field is uncorrelated with the error term ε_{it} after accounting for α_i and γ_t . This assumption would clearly be incorrect if groundwater were a private good – that is, both excludable and rival. Instead, groundwater is a common pool resource: rival but not perfectly excludable. Any one farmer’s groundwater depth is ultimately determined by the aggregate pumping of those farmers nearby, and any one farmers’ contribution to the aggregate pumping is assumed to be small enough to be insignificant. In other words, I identify β_1 using deviations from annual location-specific average groundwater levels, which I assume to be as good as random and driven by idiosyncratic aggregate pumping levels.

Currently, my empirical analysis ignores access to surface water for irrigation. Water rights in California are certainly important, and are expected to have a large impact on both land use decisions and the choice to pump groundwater. However, because California follows the appropriative doctrine for water rights, these rights are legally tied to individual parcels of land (Wilkinson, 1992). Therefore, parcel fixed effects should capture the overall effect of having access to some level of water rights. Additionally, unobserved surface water use would bias my estimates toward zero insofar as a farmer with no need to pump groundwater would not change their land use decisions at all in response to changes in groundwater levels. Future work will explore more deeply the interaction between surface water rights and the effects estimated here.

4 Results

Since over ninety-five percent of my observations fit into three land cover categories, I focus my analysis on transition probabilities between these three categories: annual crop, perennial crop, and fallow or grassland. This leads me to estimate specification 1 nine times to fill a 3×3 transition matrix.

To begin, I report the estimated $\hat{\beta}_0$ coefficients from these nine regressions in table 3. Table 3 should be considered as a companion to table 2 in that they both report transition probabilities between different land cover categories. However, table 3 controls for parcel and year fixed effects, resulting in “conditional” transition probabilities. The three largest differences between the two sets of transition probabilities are that, after controlling for fixed effects, (1) annual crop cover is more likely after annual crop cover, (2) annual crop cover is more likely after fallow or grassland cover, and (3) fallow or grassland cover is less likely after fallow or grassland cover.

Table 3: Conditional Land Cover Transition Probabilities

Previous Land Cover	Current Land Cover		
	Annual Crop	Perennial Crop	Fallow or Grassland
Annual Crop	85.72	8.76	5.51
Perennial Crop	5.65	81.84	10.66
Fallow or Grassland	25.25	14.19	57.18

Note: This table records the conditional probability of a CLU parcel having a particular land cover given its previous land cover, controlling for field fixed effects and year fixed effects. Specifically, this table reports the values of $\hat{\beta}_0$ estimated by running equation 1. I focus on the three most common land covers: annual crop, perennial crop, and fallow or grassland. All numbers are percentages.

Next, table 4 reports the effects of groundwater depth on the transition probabilities contained in table 3. Each of these reported coefficients can be interpreted as the effect of an additional foot of groundwater depth on the relevant transition probability. For instance,

consider a parcel that had an annual crop land cover in the previous year (i.e. look at the first row of table 4). Increasing the groundwater depth for that parcel by 100 feet would decrease the likelihood that parcel would have an annual crop land cover this year by 6.1% (column one), and increase the likelihood the parcel would be fallow or grassland this year by 5.6% (column three).

Table 4: Effect of Groundwater Depth (feet) on Transition Probabilities

Previous Land Cover	Current Land Cover		
	Annual Crop	Perennial Crop	Fallow or Grassland
Annual Crop	-0.061*** (0.009) $n = 25,795$	0.003 (0.004) $n = 25,795$	0.056*** (0.008) $n = 25,795$
Perennial Crop	0.005* (0.003) $n = 30,964$	0.019*** (0.005) $n = 30,964$	-0.018*** (0.004) $n = 30,964$
Fallow or Grassland	-0.062*** (0.008) $n = 19,706$	0.006 (0.006) $n = 19,706$	0.065*** (0.010) $n = 19,706$

Note: This table reports the effect of an additional foot of groundwater depth on the probability (percent chance) that a CLU has a particular land cover. Specifically, this table reports the values of $\hat{\beta}_1$ estimated by running equation 1 on various subsets of my data. These effects can be directly compared to the conditional transition probabilities reported in table 3. Standard errors are reported in parentheses and are clustered at the CLU level. n reports the number of CLU observations included in each regression. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The results reported in table 4 paint a relatively clear picture that largely matches expectations. Groundwater depth reduces the likelihood that parcels will be planted to an annual crop, and this effect is especially large and statistically significant for parcels that have been recently planted to an annual crop or left as fallow or as grassland. Conversely, groundwater depth increases both the likelihood of fallowing land after growing annual crops and the likelihood of keeping land fallow or in grassland. Groundwater depth seems not to have a profound effect on choosing to plant perennial crops, except as to increase the likelihood that

perennial crops stay planted. This fits with the idea that the dominant force with perennial crops is an option value determination that relies on the large fixed cost associated with many perennial crops.

5 Conclusion

My results support the prediction that farmers, when facing relatively more expensive sources of agricultural water, will transition to less water-intensive land uses. In particular, for an increase in groundwater depth of 100 feet, the likelihood that a parcel previously covered with an annual crop will be fallowed in the next year increases by 5.6 percent. Given that the conditional probability of this land use transition is only 5.5% to begin with, groundwater levels (and hence water costs) can have large and meaningful impacts on land use decisions.

To put my findings into perspective, Martin et al. (2011) notes that each additional 100 feet of groundwater depth requires approximately 0.9 more gallons of diesel fuel to pump an acre-inch of water. Currently, a gallon of diesel costs roughly \$2.50, meaning that an approximately \$27/acre-foot increase in the cost of agricultural water would have similar effects to those I report in table 4.

While encouraging, these preliminary results require additional scrutiny. In future work, I plan to expand the scope of this study to more of California, test various data assumptions, incorporate information on surface water rights, and more rigorously quantify the implied water costs of groundwater depth. Nonetheless, this paper clearly demonstrates that groundwater – as a water source of last resort – can tell us much about the value of agricultural water.

References

- Boryan, Claire, Zhengwei Yang, Rick Mueller, and Mike Craig (2011), “Monitoring US agriculture: the US Department of Agriculture, National Agricultural Statistics Service, Cropland Data Layer Program.” *Geocarto International*, 26, 341–358.
- Buck, Steven, Maximilian Auffhammer, and David Sunding (2014), “Land Markets and the Value of Water: Hedonic Analysis Using Repeat Sales of Farmland.” *American Journal of Agricultural Economics*, 96, 953–969.
- Caswell, Margriet F. and David Zilberman (1986), “The Effects of Well Depth and Land Quality on the Choice of Irrigation Technology.” *American Journal of Agricultural Economics*, 68, 798–811.
- Dinar, Ariel (1994), “Impact of Energy Cost and Water Resource Availability on Agriculture and Groundwater Quality in California.” *Resource and Energy Economics*, 16, 47–66.
- Farm Service Agency (2012), “Common Land Unit (CLU): Information Sheet.” Technical report, USDA.
- Green, Gareth, David Sunding, David Zilberman, and Douglas Parker (1996), “Explaining Irrigation Technology Choices: a Microparameter Approach.” *American Journal of Agricultural Economics*, 78, 1064–1072.
- Hendricks, Nathan P., Aaron Smith, and Daniel A. Sumner (2014), “Crop Supply Dynamics and the Illusion of Partial Adjustment.” *American Journal of Agricultural Economics*, 96, 1469–1491.
- Martin, Derrel L., Tom W. Dorn, Steve R. Melvin, Alan J. Corr, and William L. Kranz (2011), “Evaluating Energy Use for Pumping Irrigation Water.” In *Proceedings of the 23rd Annual Central Plains Irrigation Conference*, Burlington, CO.

Schoengold, Karina and David Sunding (2014), “The Impact of Water Price Uncertainty on the Adoption of Precision Irrigation Systems.” *Agricultural Economics*, 45, 729–743.

Stevens, Andrew W. (2015), “Fueling Local Water Pollution: Ethanol Refineries, Land Use, and Nitrate Runoff.” Working Paper.

Wilkinson, Charles F. (1992), *Crossing the Next Meridian: Land, Water, and the Future of the West*. Island Press, Washington, D.C.