Essays in Water Economics

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Dirk van Duym

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Thomas Holmes, Adviser

August, 2018

Acknowledgements

I am grateful to my adviser Tom Holmes, for invaluable advice and guidance. I am very thankful to Jim Schmitz and Joel Waldfogel, for their constructive discussions and constant support. I am also appreciative of Ellen McGrattan, for her big-picture advice and for serving on my dissertation committee. Speaking with water and agriculture experts Andrew Brait, John Cain, Joe Del Bosque, Carl Hauge, Nathan Hendricks, Richard Howitt, Dave McLaughlin, Josue Medellin-Azuara, Cannon Michael, Abran Padilla, Jon Parker, Jeff Peterson, Laney Siegner, Andrew Stevens, Daniel Sumner, David Sunding, and Jeff Yeazell was essential in being able to understand such a complex subject. Additionally, I benefited immensely from discussions about this work with Joaquin Garcia-Cabo, Bill Walsh, Brian Albrecht, Kailin Clarke, Aradhya Sood, and the participants of the University of Minnesota Applied Microeconomics workshop. I acknowledge the financial support of the Warwick Fellowship and University of Minnesota Graduate Research Partnership Program Fellowship. Finally, I thank Katie Siegner for her passion and enthusiasm, which inspire me every day.

Dedication

To my parents, Anne and Andy, who have supported me unconditionally along every step leading to this point. I am so lucky.

Abstract

This dissertation investigates the issue of water policy in California, focusing on allocation rules in the agricultural sector and how they affect farmer crop choice and water usage. Chapter 1 includes background on California agriculture and the surface water and groundwater allocation system in the state, and a strategy for generating novel estimates of groundwater pumping and its externalities. Then a structural model is estimated and used in counterfactual simulations of new water policy regimes. Chapter 2 extends the analysis in a theoretical context by studying the dynamic farmer decision to grow a permanent crop. The model is solved and the equilibrium studied, the key question being how permanence interacts with water-intensity in a setting with externalities. This chapter also shows the characteristics of an equilibrium involving cycles in the share of farmers who grow water-intensive crops.

Contents

List of Tables

List of Figures

Chapter 1

Water Policy and the Common Pool: Examining Crop Choice in California

1.1 Introduction

Throughout the world, scarcity of water is an important policy issue. In California, it has been on the forefront after a five-year drought that was the worst in 1200 years (Griffin and Anchukaitis, 2014). This paper develops an analysis of water policy in California, using a structural model with parameters pinned down by high-resolution micro data from a large county in the Central Valley. A novel feature of the paper is the way that it integrates an analysis of the market for surface water (e.g. from rivers or reservoirs) and the use of groundwater, extracted from the large aquifer underlying the Central Valley. Surface water property rights are fairly well-defined, but the use of groundwater is a classic Tragedy of the Commons problem, first described by (Hardin, 1968). Even with well-defined surface water rights, the incentives for a market to arise in trade of these claims is limited if it is possible to pump groundwater at low private cost (albeit with potentially high social cost). The paper focuses on agricultural water consumption, which accounts for the vast majority of water used in California. I model the crop choice decisions of farmers. In particular, the model takes account of how access to water feeds into the decision of farmers to plant high or low water intensity crops, and then how these choices feed back into water market. I simulate the effects of water regulation that puts caps on the amount groundwater that can be pumped, considering both caps with no trade in pumping rights (analogous to a recently proposed policy) as well as caps where pumping rights are tradable. I evaluate how regulation of groundwater impacts the market for surface water, and the decisions of which crops get planted where.

A central issue addressed in the paper is determining the extent to which groundwater is indeed a common pool, leading to the associated market failure. Whether this market failure exists depends on specific hydrological features, namely: when a farmer pumps groundwater, where does the level of the water table decline? If pumping affects only a farmer's own water level, then there is no externality. On the other hand, if the effect of local pumping is experienced largely by other landowners, then farmers will have an incentive to over-use water.

These considerations lead to two empirical questions that I address in this paper. First, to what extent is groundwater common property? To make progress on this issue, the paper begins with field-level crop choice data, and combines this with information about water requirements by crop and detailed micro data on surface water allocations to construct estimates of groundwater pumping at the field level. The groundwater withdrawal estimates are then linked to measurements on depth of groundwater at neighboring wells, and these data are used to infer how pumping by one farmer affects his or her own groundwater access and the access of the neighbors. I find that there is significant distance decay, in that an individual farmer's extractions have a greater effect on groundwater levels within a half-mile than levels farther away. However, the cumulative effects of farmers farther out are very large, so that an individual farmer pumping groundwater bears little of the social cost. This finding motivates a modeling choice to make the effect of pumping on neighboring levels entirely as an external cost.

Second, how does access to groundwater affect crop choice, and thus total water usage? In other words, do farmers, who do not bear the full social cost of their groundwater pumping, use more water than they otherwise would? During the sample period that I examine, there was a large change in groundwater depth, going from an average of 194 feet in 2007 to an average of 271 feet in 2015. Furthermore, there is cross-section variation, as changes in depth were not evenly distributed. I use a differences in differences approach, and show that fields experiencing relatively sharper declines in ground water levels were more likely to shift towards less water-intensive crops.

I then build these two estimated behavioral relationships into a structural model, where farmers make a crop choice decision given various characteristics including access to water. Conditional on crop choice, the farmers obtain the requisite water in a cost minimizing way, which in practice typically is to consume the endowment of surface water and then get the residual from groundwater. In taking the model to the data, I first plug in the spatial interaction estimates discussed above, connecting one farmer's pumping decisions to groundwater levels of neighboring farmers. Next I pick parameters to match the substitution patterns for how the water-intensiveness of crops chosen depends upon water access. When evaluating various groundwater policies, I simulate the planting decision for each farmer, and then can calculate how crop choice affects groundwater levels throughout the county. I also determine the equilibrium price of surface water and what causes it to move around.

One important trend in the data is an increase in the percentage of farmers choosing to grow "permanent" crops, like almond and pistachio trees. This is important because with a crop locked in over a long time scale, it is more difficult to respond to shocks. Annual crops allow for more flexibility in times of drought, either through moving towards less water-intensive crops, or letting land lie fallow. There is a corresponding lack of flexibility when the share of permanent crops grows higher. Then farmers must draw down the common aquifer whenever surface water supplies are low. In Chapter 2 of this dissertation, I examine how policies that deal with the common pool problem, e.g. optimal Pigouvian taxes, interact with decisions of farmers to select permanent versus annual crops. In this paper, to keep the analysis tractable, I take as exogenous the decision to be permanent or annual, and focus on the crop choice decisions made by farmers planting annual crops. Nevertheless, the issue of permanent versus annual is important, because the higher the permanent share, the harder it is to respond to supply shocks.

The sustainability of groundwater usage is a major policy concern in California. Of late, the state has experienced extreme volatility in surface water availability. The historic five-year drought, with costs estimated in the billions (Howitt et al., 2015) was followed this past winter by the wettest year on record. During the past few years, farmers used so much groundwater that several small towns saw their domestic wells run dry. Some areas are still far from recovery, and have had to get water tanks or bottled water trucked in on an ongoing basis. Subsidence (the collapse of the land surface due to less groundwater below) is endangering the area around the California Aqueduct, which transports over half of Los Angeles's water supply from the Sacramento-San Joaquin Delta to southern California. It has already decreased the capacity of several other major canals. Such volatility is only expected to increase with climate change, which will deliver both more extreme drought and more intense precipitation events in the years to come.

In response to these issues, California passed the landmark Sustainable Groundwater Management Act (SGMA), the first statewide attempt to regulate groundwater. The basic idea is to bring the state's various groundwater basins into long-term sustainability. The definition of sustainability in SGMA is somewhat nebulous, but it is generally understood by farmers and water managers to mean that the long-run average recharge of groundwater aquifers has to be equal to the long-run average pumping from the aquifers. Regulators and water users will be working in the coming months and years to write the binding rules of the legislation at the local level. At this point, it is unclear what these local management plans will entail, but possibilities include groundwater taxes, restrictions on groundwater pumping, well drilling limits, and groundwater markets.

This paper is at the crossroads of several strands of literature. Burness and Quirk (1979), Schoengold et al. (2006) and many others have written about water rights and usage, and the consequent impacts on agriculture. Gisser and Sanchez (1980), Tsur and Zemel (1995), and Famiglietti (2014) have examined groundwater specifically, noting how excessive groundwater usage in times of drought can cause adverse consequences like land subsidence. Holmes and Lee (2012) have explicitly modeled crop choice in a spatial framework, using detailed crop choice data. And importantly, the concept of the commons has been studied extensively. Almost fifty years ago, Hardin (1968) described the behavior of private agents overusing a shared resource as the tragedy of the commons. Elinor Ostrom won a Nobel Prize for describing how agreements could be made to manage common resources in a more efficient manner, specifically focusing on groundwater basins in (mostly urban) southern California (Ostrom et al., 1994). This body of research has significant bearing on decisions about California's groundwater management systems today, discussed in depth by Ayres et al. (2017). The closest papers to this one are probably Pfeiffer and Lawell (2012), who estimate spatial externalities in groundwater pumping in Kansas, and Guilfoos et al. (2013), who also study Kern County groundwater management, but in an agent-based-modeling framework. Neither of these papers consider surface water or crop permanence. (Franklin et al., 2017) studies permanence in Australian agriculture with a vintage capital model, but does not touch on externalities.

The rest of the paper is structured as follows: Section 1.2 describes the data I use and shows some key facts. Then in Section 1.3 I propose a structural model of crop choice and water usage, and discuss assumptions I make, and show equilibrium outcomes. Next, in Section 1.4 I estimate several important parameters. In Section 1.5 I estimate the full model and evaluate its fit. Section 1.6 discusses counterfactuals, and then I conclude.

1.2 Data and Key Facts

In this section, I describe the various types of data I use, showing some key facts that motivate the modeling choices made in the following section. The data fit nicely into three separate categories: surface water, crops, and groundwater, and wells. After going through each of these four categories, at the end of the section I show groundwater levels impact farmers, both through effects on well failure and effects on crop choice.

1.2.1 Surface Water

Data on water use by individual farmers is hard to come by, since they are under no obligation to report what they use. To work around this issue, it is useful to discuss the water "supply chain." In most of the Central Valley, surface water is delivered to farmers by water or irrigation districts. These are relatively small, quasi-governmental entities that were formed by landowners to hold water rights, maintain infrastructure, and acquire water from state and federal dam projects. Farmers pay their water district to deliver this surface water. Essentially, a water district is a hybrid of a farmer cooperative and a water wholesaler. While district-farmer level transaction data is not available, districts are required to publicly submit Agricultural Water Management Plans in which they report total water supplies, prices, and allocation rules. State law dictates that prices must be at cost, but these costs vary significantly across districts, depending on the original source of the water and the expense of moving it to farmers. Most districts allocate their surface water so that each acre of farmland gets the same amount (not dependent on crop or owner). Drawing from these water plans, I collect data on district acreage and total surface water allocated by district and by year to determine how much surface water each parcel owner can use.

Most districts hold long-term contracts for a specified amount of water every year from one of the two major dam and aqueduct projects in California, the State Water Project and the Central Valley Project. However, actual district allocations are usually less than the contracted amount. This can happen either because of inadequate precipitation and snowmelt, or because of environmental regulations that require enough water to be left to protect fish habitat and to keep salinity levels sufficiently low. Figure 1.1 shows the average surface water allocation by year, averaged over all the districts. The units are acre-feet per acre. An acre-foot is the amount of water it takes to cover one acre of land with one foot of water, or about 325,900 gallons. This is roughly how much a family of four uses for domestic purposes in a year.

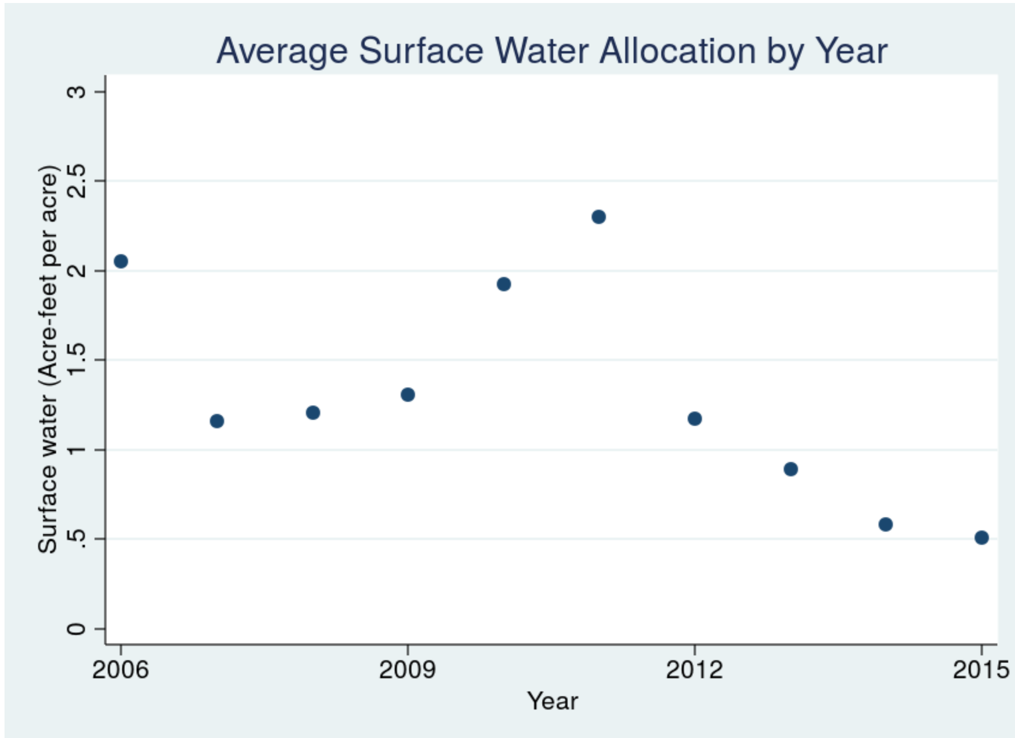

Figure 1.1: Average Surface Water Allocation by Year

As I will show in the next few sections, total cropland and total water used do not vary much with surface water availability. So in effect, surface water purchases by farmers are an infra-marginal transfer. They are happy to pay the (relatively low) cost for whatever is available in a given year, and then pump groundwater to satisfy the water requirement of the crop that they choose to grow. Anecdotally, this is supported by the amount of lobbying done by agricultural interests to loosen environmental restrictions on surface water deliveries.

Figure 1.2 plots per-acre surface water deliveries by the eighteen districts in Kern County. Data is available from 2006 to 2015. Note the relatively significant heterogeneity between districts, and the spikes during the wet years of 2006 and 2011. Some districts follow similar trends as they get their water from the same source.

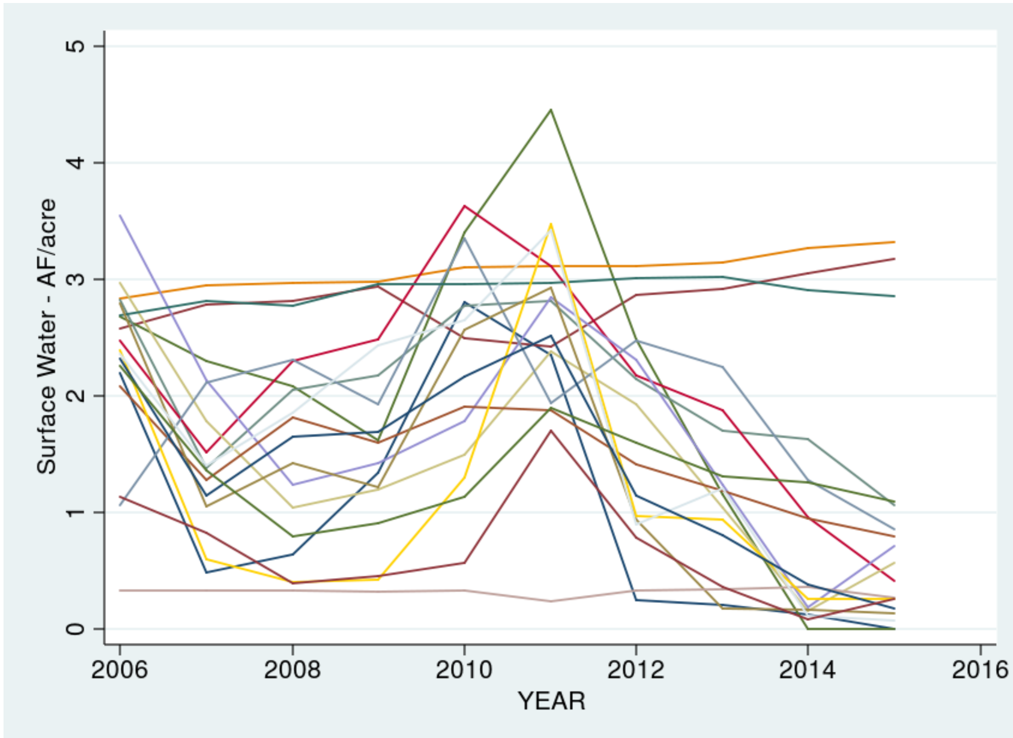

FIGURE 1.2: Variation in District Surface Water Allocations

1.2.2 Crops

The source of data on crop choice is from the Kern County Department of Agriculture and Measurement Standards (County of Kern, 2016). It consists of the crop grown on each parcel of land in the county, from 2000 through 2015. Unfortunately no data is available on crop yields at anything more granular than the county level. The Kern County data is gathered through pesticide reports farmers are required to submit. While the data I use is only available in one county, Kern County is perennially in the top 5 of agricultural counties in the state, and is the 20th largest county by area in the United States. At the southern end of the Central Valley, it is among the driest of the counties of the state's principle agricultural area. The county averages about 6 inches of rainfall per year, and most of it occurs outside of the growing season. Farmers in the county are very reliant on water imports from wetter parts of California.

As mentioned above, crops vary in their water needs. These values are known as crop evapotranspiration, or ETc, and are reported for Kern County by University of California agricultural and soil scientists, as well as by districts in their water management plans. The units are acre-feet per acre.

Since I have reliable data on acreage planted (described below) as well as crop water needs, it's not necessary to have data on water consumed at the individual level, as it can be easily backed out. In fact, this may even be an advantage: papers that look at smaller areas with data on water distributed to a farmer cannot distinguish if that water was actually used to grow crops, or continued on as runoff. This is generally seen not as waste, but rather as a positive for groundwater replenishment.

In Figure 1.3, I include a screenshot of one year of the geospatial version of this data, overlaid with the water districts of Kern County (in tan). Each small green shape is a field, averaging about 100 acres. The largest fields are 1 square mile, or 640 acres. The definition of a field is a contiguous piece of land, owned by the same person and on which a single crop is planted. About 0.5% of fields are not in a water district, in which case they are solely reliant on groundwater pumping.

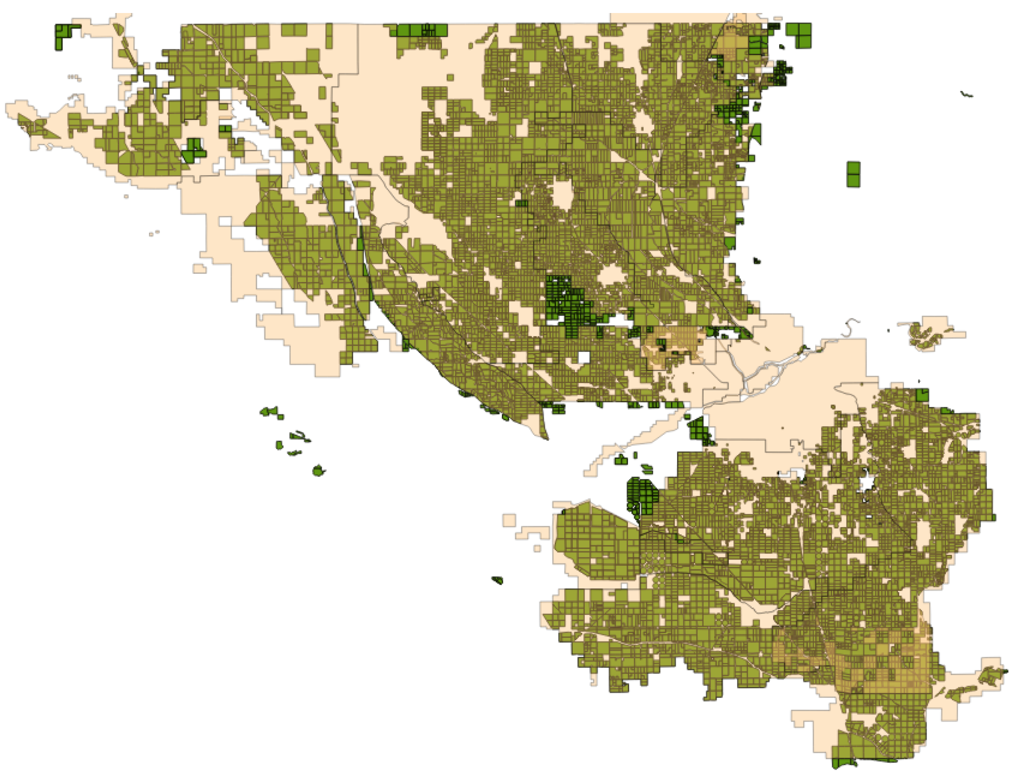

FIGURE 1.3: Kern County Water Districts and Agricultural Parcels

TABLE 1.1: Kern County Crops, Acreage (000s), Water

Crop	Acres 2000	Acres 2015	% change	Water need (AF/acre)
Almonds	125.4	222.5	77	3.5
Pistachios	46.5	124.7	168	3.6
Grapes	110.6	109.4	0	2.4
Wheat	65.5	80.9	24	1.7
Alfalfa	90.9	72.7	-20	4.1
Oranges	48.4	37.9	-22	3.5
Corn	28.9	33.2	15	2.6
Carrots	52.4	30.6	-42	1.5
Cotton	202.4	28.3	-86	2.6
Potatoes	31.2	25.7	-18	2.8
Tomatoes	11.2	18.3	63	1.7
Other perm.	32.4	56.3	74	3.5
Other field	60.3	50.5	-16	2.3
Other veg.	43.6	29.1	-33	1.9
Total crops	962.5	926.1	-4	2.77 $\rightarrow$2.97
Fallow	64.8	95.4	47	0
Total	1027.3	1021.5	0	2.60 $\rightarrow$2.69

The 144 different crops grown in Kern County can be divided into two groups: annual crops, for which planting decisions are made every year, and permanent (or perennial) crops, which usually produce for several decades. Permanent crops (the most prominent of which are almonds, pistachios, and grapes) are significant investments for a number of reasons. First, they take several years to start producing, meaning farmers must forgo revenues during their infancy. Second, they need water every year. Annual crops can be fallowed if surface water allocations are low, if pumping costs are high, or if water sale opportunities arise. With a permanent crop, farmers no longer have flexibility to pump less groundwater during dry years, or to sell their allocation to other farmers or domestic users. Third, for the most part, permanent crops require more water than annual crops such as wheat, peppers, tomatoes, etc. Table 1.1 shows the principal crops grown and their water requirements. Note that due to the change in composition of crops, the average water requirement increases 7% from 2.77 to 2.97 acre-feet per acre. If fallowed land is included in the average, it is instead a 4% increase.

In Kern County, there is a clear increase in the prevalence of permanent crops over time, shown in Figure 1.4. In 2000, the first year detailed data is available, they accounted for 35% of acreage, while in 2015 that number was up to 54%. This has led to what is known in the water policy arena as a "hardening" of water demand. These permanent crop growers cannot adjust their water demand in the face of drought and declining groundwater levels. They are essentially forced to keep pumping groundwater, drilling new million-dollar wells if necessary.

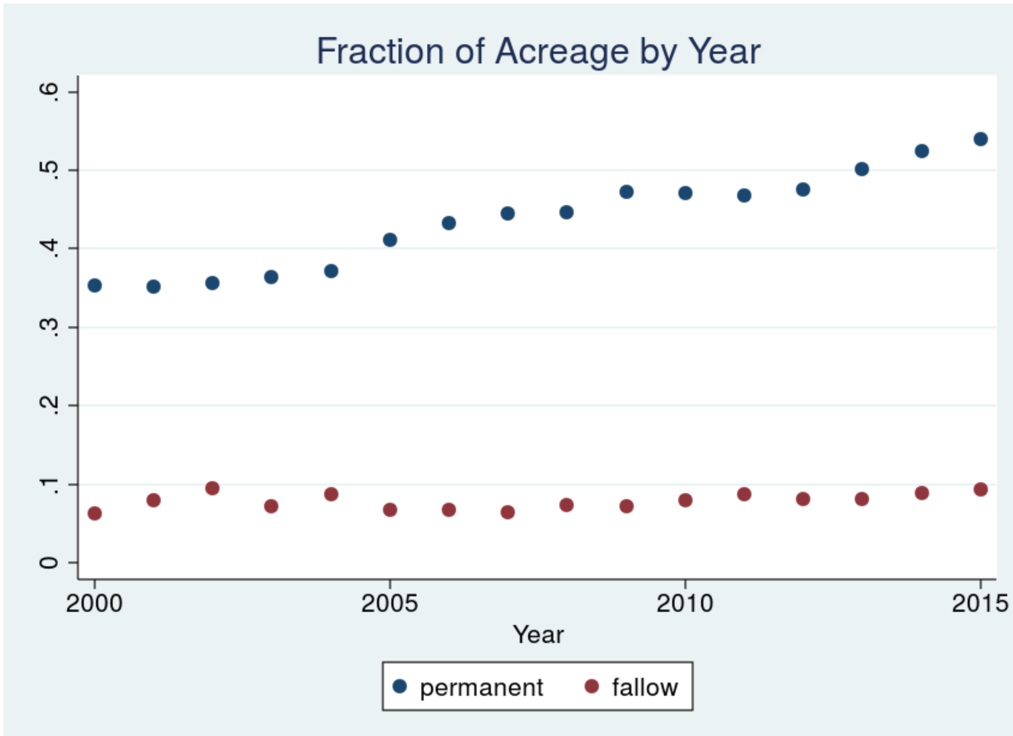

FIGURE 1.4: Fraction of Acreage by Year

As their numbers grow, mostly due to increased demand for nuts, as has happened in the past decade, the problem is exacerbated. On the other hand, annual farmers are more flexible. They can react to a drought by switching to a less water-intensive (or more profitable) crop, or in extreme, fallowing their land and not using any groundwater at all. Perhaps surprisingly, the share of fallowed land in Kern County does not fluctuate very much. Fallowing has benefits for all the farmers who are tapping the aquifer. I further elucidate this point using the model in Section 1.3.

1.2.3 Groundwater

Data on surface water is relatively good compared to groundwater data. In most of the Central Valley there are no regulations requiring groundwater pumping to be reported. This is one of the main reasons that evaluating groundwater policy has been difficult. The only restrictions on groundwater extractions are that they have to be used on the overlying land and not sold for use elsewhere. Most if not all farms have their own wells. Water districts and the state Department of Water Resources also have monitoring wells, for which panel data about groundwater levels is available. The source of this data is California State Groundwater Elevation Monitoring (CASGEM). Most of these wells are measured twice per year: in the early spring, once the wetter months pass; and in the fall, after most groundwater pumping and the dry season has occurred. The observations are not granular enough to be able to get a reliable estimate of individual farmer withdrawals, but clear patterns exist in the aggregate.

To get a sense of the overall trend, I run a panel regression of depth to groundwater on a set of month and year dummies. Figure 1.5 plots those year effects. Note the overall trend that coincides with the increase in permanent crops. It is also interesting to see the uptick after the last wet year before the drought in 2011, and the sharp decrease during the recent drought.

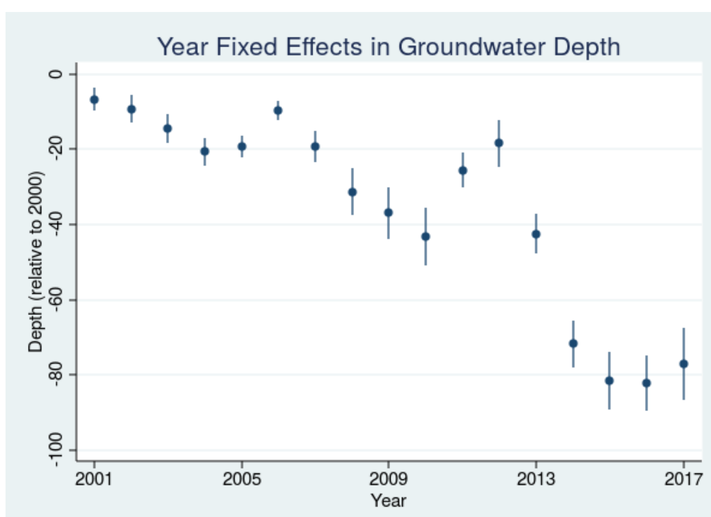

FIGURE 1.5: Year Fixed Effects in Groundwater Depth, Kern County

Since information on most pumping (from active private wells) either is not publicly available or does not exist at all, this is not perfect, but it provides evidence that groundwater is used as a buffer stock. When more surface water is delivered, groundwater levels rise. There is also a clear downward trend in depth to groundwater as more land has been planted with permanent crops. Famiglietti (2014) also shows evidence of decreasing groundwater availability using NASA data. Satellites detect changes in gravitational pull above the aquifer, showing changes in total aquifer volume. These data can be a bit misleading, since much of the groundwater reserves are miles deep, too far to ever be of use. Some of this groundwater can be too salty as well. These satellites are also used to determine the extent to which subsidence has occurred. When too much groundwater is pumped out of certain types of aquifers, the soil layers can compact, causing everything above to drop, up to around two feet per year. The worst areas are slightly north of Kern County. Some subsidence is temporary, and wet years can solve the problem. However, when the soil has more clay content, subsidence can be irreversible. In this case, pumping too much groundwater can permanently decrease the aquifer's storage capacity.

Figure 1.5 hides large cross-sectional variation in depth to groundwater. Figure 1.6 shows this heterogeneity in 2015. The lightest color is where the groundwater level is only ten feet below ground, and the darkest is the deepest, at 600 feet. I create this figure by interpolating over depth measurements from the spring of 2015[1] Much of the difference across the county is due to land elevation changes as the edge of the valley approaches on the south and the east. I will later on look to estimate the sensitivity of water usage to groundwater levels. This variation in depth to groundwater is not enough to make the case, as soil or other unobserved farmer or geographic characteristics could explain any cross-sectional variation in water usage. However, I make use of this heterogeneity and how it changes over time in Subsection 1.4.3.

[1] I use the ordinary kriging method preferred by hydrologists Delhomme (1978). Results are similar with the Triangular Irregular Network and Inverse Distance Weighted methods.

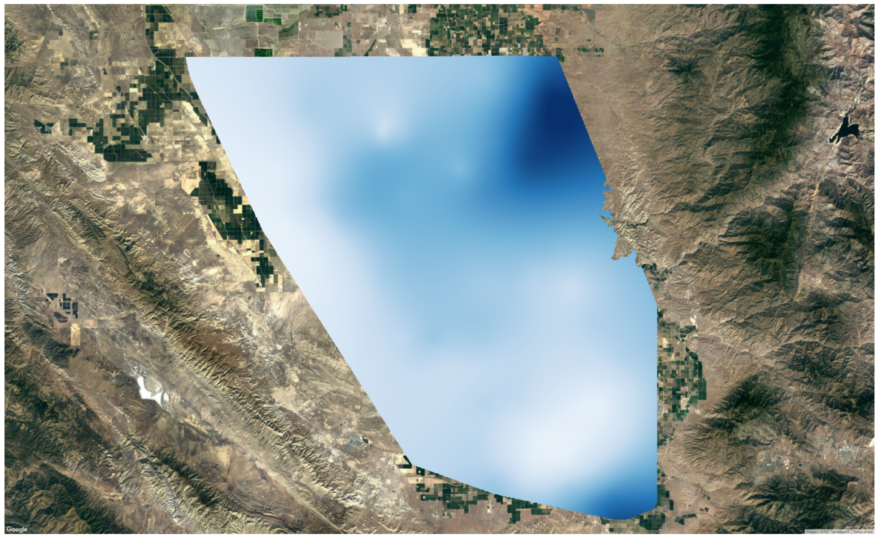

Figure 1.6: 2015 Groundwater Depth (darker shows deeper)

1.2.4 Well Failures

Data on well failures is not available in Kern County. However, Tulare County, which borders Kern County to the north, kept a detailed record during the drought years of 2013-2015 of how often wells were failing due to dropping groundwater levels. Along with these records, I obtain data on groundwater levels from the public monitoring wells in Tulare County over this time period, from the CASGEM program described above. This allows me to track how declines in the aquifer level are correlated with the number of wells failing. The data are shown in Figure 1.7. Well failures can go negative when levels rise and wells start working again.

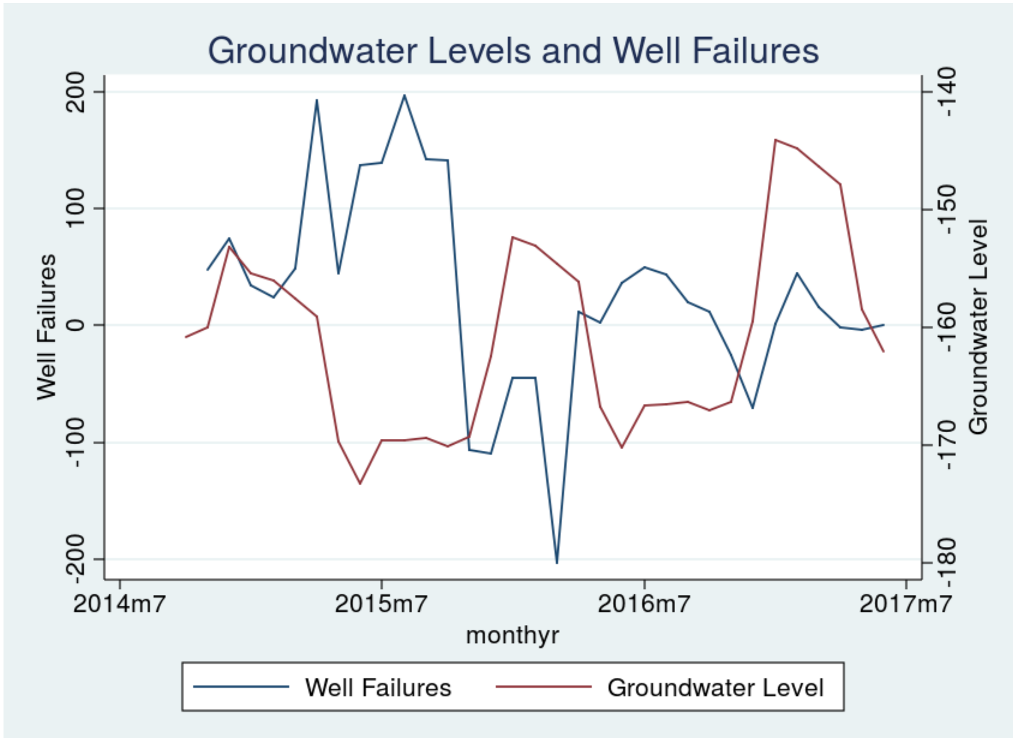

Figure 1.7: Groundwater Levels and New Well Failures, Tulare County

I simply regress the number of well failures by month on the average groundwater level by month. From here, I find that a groundwater level drop of one foot is associated with 3.4 more wells failing per month. According to a state groundwater assessment, Tulare County has about 20,000 domestic wells California State Water Resources Control Board (2016). This translates to an increased probability of needing a well in a year of .002 per extra foot of depth.

	(1)
	failures
level	-3.354*
	(1.561)
N	32

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.2: Effect of Depth to Groundwater on Water-Intensity of Crop Choice

1.3 Model

In this section, I lay out the theoretical model of farmer crop choice and water usage, and the evolution of groundwater levels.

There are N farmers who each own one parcel of land, indexed by i and located in a water district d. Farmers live for T years, and discount the future at rate β. They can grow one of three crops, $j = h$ (water-intensive annual crop), $j = l$ (low water annual crop) or $j = p$ (permanent crop), which I describe in the next subsection. They can also choose to fallow their land and use no water. Their land has an inherent vector of soil qualities γ_{ij}, where γ_{ij} is soil quality.

Importantly, the production function is fixed proportions. Each crop j requires a fixed amount of water per period, denoted w_j. Given that amount of water is used, production is equal to $f_j(\gamma_{ij})$. Crop output prices p_{jt} and other input costs x_j are exogenous.

The water that farmers use in crop production can either be surface water, w^s, or groundwater w^g. They are perfect substitutes in the production function. A farmer with land in district d is allocated surface water $\bar{w}^s_{dt}$. This is a district-specific shock, as different water districts have access to different sources of water. It can be obtained for a price of p^s. (Note that $\bar{w}^s_{dt}$ and p^s will be endogenized in the counterfactuals, when groundwater is regulated and there is then more of an incentive for a robust surface water market to exist.) Groundwater w^g_{it} has to be pumped from the underlying aquifer, according to the following (private) cost function, where L_{it} is the depth to groundwater for farmer i at time t:

$$c(w^g_{it}, L_{it}) = \beta_0 + \beta_1 L_{it} w^g_{it}. \tag{1.1}$$

The parameters β_0 and β_1 in the groundwater cost function are meant to capture well fixed costs and the cost of energy needed to pump one acre-foot of groundwater a lift of one foot. Groundwater at a particular location evolves according to the law of motion:

$$-L_{it+1} = -L_{it} + \alpha r_{it} - \theta \sum_{k} e^{-\rho d_{ik}} w_{kt}^{g}, \tag{1.2}$$

where r_{it} is recharge of the aquifer, w_k^g is withdrawals at a distance d_{ik}, and ρ is a parameter meant to capture the extent to which the aquifer is a common pool. This functional form allows the depth to the aquifer to be more affected by groundwater pumping from nearby than farther away. Writing the law of motion in this way is a strong assumption, as it is imposing a functional form on the flow of groundwater. To study this, hydrologists use Darcy's Law, which is a general formulation of fluid dynamics in porous media. The formulation above works in a similar way, with the exponential decay parameter ρ standing in for unknown soil parameters and flow rates in the aquifer. It also has the attraction of being quite tractable, allowing me to study the evolution of groundwater levels with simulated groundwater pumping decisions in the model.

1.3.1 Crop Choice

I make the assumption that exogenously, some farmers grow permanent crops and some grow annual crops. Without loss of generality, let the permanent crop farmers be indexed by $i = 1, ..., P$ and the annual crop farmers be indexed by $i = P + 1, ..., N$. Each year, farmers with annual crops are able to choose between a higher-water intensity crop, a lower-water intensity crop, and fallowing their land ($j \in \{h, l, f\}$). Suppose that farmers take groundwater evolution as given, which I show evidence of in Subsection 1.4.2. Then I can write the annual farmer maximization as a series of static problems. For any t, their problem is:

$$\max_{j_t, w_{it}^g} \; [p_{jt} f_j(\gamma_{ij}, w_j) - p_t^s \bar{w}_{it}^s - c(w_{it}^g, L_{it}) - x_j + \varepsilon_{ijt}] \tag{1.3}$$

$$s.t. \; w_j = \bar{w}_{it}^s + w_{it}^g.$$

The permanent crop farmer's problem is:

$$\max_{w_{it}^g} \; [p_{pt} f_p(\gamma_{ij}, w_p) - p_t^s \bar{w}_{it}^s - c(w_{it}^g, L_{it}) - x_p] \tag{1.4}$$

$$s.t. \; w_p = \bar{w}_{it}^s + w_{it}^g.$$

Given the assumptions of the model, it is only farmers with annual crops whose decisions are affected by crop price changes, surface water allocations, and groundwater costs. The other way of saying this is that water demand is perfectly inelastic for farmers who have a permanent crop. While clearly not globally true, it is reasonable that on the margin, almond and pistachio growers do not reduce irrigation when water becomes more expensive.

Due to the fixed proportions production function, choosing crop amounts to choosing water consumption as well. Here, when fallowing, farmers sell their surface water allocation, meaning the second term in the objective function is positive. They also draw a type-one extreme value shock for each of their three possible crop choices. Re-writing the objective function as $\delta_{ij} + \varepsilon_{ij}$ leads to the logit formula for the probability that a farmer grows a particular crop.

$$Pr(j) = \frac{exp(\delta_{ij})}{\sum_k exp(\delta_{ik})} \tag{1.5}$$

This formulation allows annual crop growers to make yearly decisions to react to changing surface water allocations, groundwater levels, and crop prices. These reactions can serve to mitigate the common pool problem–depletion of the aquifer can result in switches to crops that use less water, or even fallowing, rather than permanent crop growers who only serve to make the problem worse.

If the exponential decay parameter ρ is large, a larger fraction of the effects of groundwater pumping are felt by the pumpers themselves. That is, farmers are affecting their own future supplies more than with a small ρ. So they have the incentive to pump less groundwater today in order to keep their groundwater level from dropping as much in the future. In this case, it may not be a good assumption to have farmers take groundwater levels as given. When farmers do take into account their own effect, I can no longer write out the maximization as a static problem. The choice to use more groundwater today must be weighed against more costly groundwater tomorrow. However, in Subsection 1.4.2 I provide evidence that ρ is sufficiently small for me to ignore this effect. In Kern County, farmers are small relative to total groundwater pumping. This assumption can also be interpreted behaviorally: farmers are short-sighted, which seems unlikely; or, farmers believe that their effect on future groundwater levels are negligible, which is likely the case.

1.3.2 Equilibrium in Surface Water

The production function f is fixed proportions, so given any crop choice j, farmers optimize by obtaining total water w_j. In most of California, there is no robust market for surface water. On occasion, surface water transactions are made, but only after significant transaction costs. These come in the form of rules by government at all levels, transportation costs, and frequent legal challenges. Part of the reason why a market has not developed is thanks to the existence of the outside option of groundwater. The misallocation problems that a robust surface water market would solve are not as significant when groundwater is available to farmers without access to surface water.

However, as groundwater regulations are put into place, surface water misallocation becomes more severe. This is simply due to the outside option of groundwater no longer being as reliable. Thus, it is important to allow for a surface water market to arise when groundwater usage is restricted, which is what I lay out in this subsection.

It is useful to go through the intuition of how this surface water market matters. In the status quo, to satisfy crop water demand, farmers use up their surface water allocation and then turn to groundwater for the rest of their water need. So except for in the wettest and driest years, most farmers use both sources of water. When a surface water market arises, this is no longer the case. Farmers use whichever source of water is cheaper, and since marginal costs are constant for a given farmer within a given year, this means each farmer uses only one type of water. Farmers for whom groundwater is more expensive will buy surface water on the market, and farmers for whom groundwater is cheaper sell their surface water allocations and rely entirely on groundwater.

In 1.6, I define the excess demand for surface water for farmer i. This is surface water demand (which depends on crop choice, price of surface water, and depth to groundwater) less the initial surface water allocation (t suppressed):

$$ED_i(j_i, p^s, L_i, \bar{w}_d^s) = w_i^s(j_i, p^s, L_i) - \bar{w}_i^s. \tag{1.6}$$

A farmer k with high groundwater costs consumes her surface water allocation $\bar{w}_k^s$, and then looks to buy more surface water on the market to satisfy crop need. But her groundwater pumping is $w_k^g = 0$. A farmer m with low groundwater costs sells his surface water allocation $\bar{w}_m^s$ on the open market, and then pumps groundwater w_m^g to satisfy crop need.

To find the equilibrium price of surface water, p^s, I search over a grid of possibilities and calculate decisions by solving (1.3) and (1.4). Then I update, and this process continues until the market clearing condition (1.7) for surface water is satisfied.

$$\sum_i ED_i = 0. \tag{1.7}$$

Qualitatively, this delivers the pattern that when groundwater levels drop, annual crop farmers move toward less water-intensive crops, and the market-clearing price of surface water increases.

This model can generate several features that we see in the data: a severe drought decreases groundwater levels and increases well construction, and (maybe surprisingly) annual farmers rarely choose to fallow. I match the responsiveness of annual crop farmers' water choices to the responsiveness I find in the data in Subsection 1.4.3. This allows me to do counterfactual simulations with new groundwater policies.

1.3.3 Boundary Conditions

In this subsection, I discuss assumptions that must be made in order for the above equilibrium to exist.

With a market for surface water, Equation 1.7 must be satisfied for an equilibrium to exist (with an interior solution). Consider low surface water shock realizations in the status quo environment (no groundwater restrictions). Equation 1.7 will be satisfied with large amounts of groundwater pumping. However, in an environment with groundwater quantity caps, this may not work. Due to the assumption that permanent crop acreage (and thus permanent crop total water need) is fixed, excess demand may be positive for all possible values of the price of water. The interpretation is that with extremely low surface water shock realizations (and strict enough groundwater restrictions), the total amount of surface water and groundwater available is not enough to satisfy the water needs of the stock of permanent crops. So even with 100% of annual crop farmers choosing to fallow and sell all their water, excess demand is still positive and an equilibrium does not exist. Taking the model literally, I would have to make an assumption about how profits are affected if half of permanent crops die off, or are watered less-intensively, etc.

More formally: there exists an equilibrium price such that excess demand for water is equal to zero if:

$$\sum_{i=1}^{N} acre_i(\bar{w_i^s} + \bar{w^g}) \geq \sum_{i=1}^{P} acre_i * w_p. \tag{1.8}$$

The good news is that the conditions for this lack of equilibrium to occur are quite unlikely, as only about half of acreage is for permanent crops. In any case, in the simulations in Section 1.6, I preclude this possibility of nonexistence, by only drawing surface water shocks from a distribution with a lower bound such that Equation 1.7 is satisfied with equality. In practice, even with relatively stringent groundwater restrictions, this lower bound is lower than any average surface water realization in the past twenty years, making this a rather innocuous assumption.

It is possible for there to be a corner solution where the equilibrium price of surface water is equal to zero. This happens when surface water availability is high enough that no groundwater needs to be pumped. These kind of realizations can happen in extremely wet years, which while rare, have occurred in the past. I do not preclude this possibility in the simulations.

1.4 Empirical Results

In this section, I go directly to the data to estimate several important parameters of the model. Before that though, I must obtain granular estimates of groundwater usage. Then I can go after the essential question of figuring out the physical size of the groundwater pumping externality: how much does groundwater pumped out of the aquifer at one location affect the depth to groundwater at another? After that, I show how changes in depth to groundwater affect the probability that groundwater wells fail, necessitating costly new construction.

1.4.1 Groundwater Usage

With all of the data described in Section 1.2, I can solve for the per-acre amount of groundwater used on each field from the following equation:

$$w_i^g = w_j - w_i^s, \tag{1.9}$$

where w_i^g is groundwater used on field i, w_i^s is surface water used on field i, and w_j is the crop evapotranspiration (ETc), or water requirement, of crop j.

This is an important, rather stark assumption, so it is instructive to go through the implications carefully.

The first is that each crop uses a specific amount of water, the crop evapotranspiration. It is probably a reasonable approximation, and is the approach used by most water districts to estimate water used. Some advances have been made using newer methods that rely on satellites that detect moisture, but that is beyond the scope of this paper. I need to make this assumption in order to back out groundwater usage at a granular level. It does remove the importance of the decision to install water-saving irrigation technology. There is an extensive literature on the move from flood or sprinkler irrigation toward drip irrigation. There is no doubt that advanced irrigation technology allows farmers to purchase less water from their district supplier to grow a given amount of crop. However, much of the formerly "wasted" water from flood irrigation techniques would seep down to recharge groundwater.

Also, while detailed data is not available, from references in district Agricultural Water Management Plans it appears that the vast majority of Kern County farmers already use efficient drip irrigation.

The second is that when surface water allocations are diminished, groundwater must fill in the gap. Anecdotally, this is exactly how farmers react, and it is consistent with how constant total water usage has been over time. During extremely wet years like 2017 (for which data is not available), enough surface water is available for everyone to not have to use groundwater. This did not happen except in 2011 in two districts in the years of my sample.

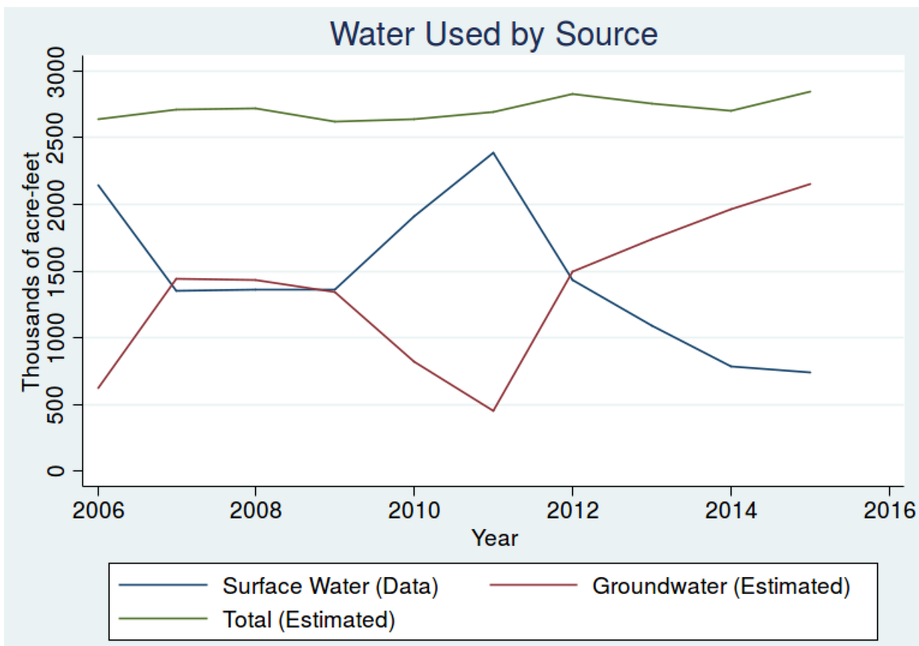

FIGURE 1.8: Water Used by Source

1.4.2 Spatial Externalities

It is important to figure out to what extent the aquifer that farmers tap into is a common pool. This depends on the flow of groundwater below the surface–how quickly the aquifer re-equilibrates. Some work, like the seminal Gisser and Sanchez (1980) paper, simply assume one extreme possibility: the "bathtub model," where groundwater withdrawals at any point affect the aquifer equally across space. All the way at the other extreme

Year	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015
Mean drop	3.3	34.9	25.6	14.4	-0.1	-18.8	16.7	31.2	22.9	14.0
Sample size	249	280	195	235	200	233	163	201	140	170

Table 1.3: Groundwater Level Changes by Year

would be totally restricted flows, as if each farmer had her own personal tank of groundwater that only she could access. In this extreme, each farmer's pumping only affects her own groundwater levels, erasing the rationale for groundwater management policies. If the aquifer below Kern County were like this second extreme, there would be no spatial interaction between groundwater pumpers.

There is strong evidence from a wide-range of disciplines that subsurface flows do occur, leading to spatial externalities. In principle, Darcy's Law answers the question, but it requires complete data on highly localized factors like the entire composition of soils well below the surface, the exact timing, rate, and location of groundwater pumping, and groundwater elevations at every point in the aquifer.

Instead, I use an econometric approach informed by Darcy's Law to estimate the extent of this spatial externality. The simple idea is to compare changes in groundwater levels to estimates of groundwater pumping in an explicit spatial framework. One challenge is that aquifer recharge is notoriously difficult to determine. To deal with this, I take advantage of California's seasonality. Very little rainfall occurs in the summer months: about 1.5 inches on average, and little of this seeps into the groundwater. Compare this to the typical crop water requirement of 36 inches. Aquifers are recharged either by excessive irrigation, or from natural sources during the wetter winter months.

So I restrict attention to wells that have measurements at the beginning and end of the primary groundwater pumping season (March and October of the same year).

First, in Table 1.3 I provide information on well level drops, by year. A negative value means levels rose. The effects of the droughts of 2007-2009 and 2012-2015 can be seen simply in these summary statistics.

The next step is to obtain the exact location of the groundwater that is pumped. To do this, I assume that groundwater used on a particular field is pumped from a well on that field. There are some instances of groundwater being supplied by the water district, but in general this is rare. So using Equation 1.9 in Subsection 1.4.1, I have an estimate for how much, and exactly where, groundwater is pumped in Kern County. While this paper focuses on agricultural groundwater use, it is necessary to include groundwater pumping from cities and towns as well. More is described about this data and process in the appendix.

I assume the following functional form for the groundwater level law of motion from year to year, where i denotes a well, k denotes groundwater users, d_{ik} is the distance from well i to user k in miles, w^g_{kt} is groundwater pumped by k in year t, and r_{it} is rainfall in year t (which happens after groundwater pumping has finished):

$$-L_{it+1} = -L_{it} + \alpha r_{it} - \theta \sum_k e^{-\rho d_{ik}} w^g_{kt}. \tag{1.10}$$

Since recharge happens after the growing season, I can write the law of motion from the beginning to the end of a growing season in a simpler form, where $drop_{it}$ is the drop in groundwater levels from March to October at well i in year t.

$$drop_{it} = \theta \sum_k e^{-\rho d_{ik}} w^g_{kt} \tag{1.11}$$

To identify the exponential decay parameter ρ, I can compare the effects of pumping from different distances. This is similar to the work of Pfeiffer and Lawell (2012), with adjustments due to not having data on every agricultural well as they do.

My approach is the following: I consider circles of various sizes around the 502 wells in the data. As an approximation, pumping is assigned to seven discrete distance bins, or rings[2], around each well. I assume that fields are uniformly distributed within each ring, and assign the distance to be the average within that ring. I then add up the groundwater pumped within the boundaries of the smallest circle, and then each ring moving outward. The smallest circle has a radius of .5 miles, and each successive circle has double the radius.

[2]The smallest region is a circle of radius 0.5 miles. The next five rings have radii of 1, 2, 4, 8, and 16 miles. The largest ring includes the entire rest of the county.

Parameter	Drop	Drop
ρ	.153***	.201***
	(.035)	(.046)
θ	.00014***	.00022**
	(.00005)	(.00008)
N	2,066	1,992

Table 1.4: Decay Results

In addition, it is important to account for the presence of artificial groundwater recharge through irrigation runoff. This happens more when farmers use flood irrigation than drip irrigation, but without data on irrigation technology, I simply assume 10% of applied water ends up recharging groundwater. Note that this can increase groundwater levels during a wet year when surface water is all that is needed.

In Appendix A., I use a similar procedure to estimate decay using a more flexible approach where I do not assume the exponential decay functional form, and obtain similar results.

I run a non-linear least squares regression to estimate the exponential decay parameters ρ and θ. The results are in Table 1.4. The first column includes all available data. The second column removes observations on the border of the study area, where groundwater pumping to the north is unobserved.

Armed with the estimate for ρ, I can calculate more readily interpretable statistics. An important figure is what fraction of groundwater changes are due to activity from close by. The expression for this is:

$$\frac{e^{-\hat{\rho}d_1}\bar{w}_1^g}{\sum_k e^{-\hat{\rho}d_k}\bar{w}_k^g}, \tag{1.12}$$

where I use the estimated $\hat{\rho} = .201$, and the upper bar indicates the mean. This is important because much more groundwater pumping happens farther away than nearby. So even with high rates of spatial decay, nearby pumping could account for just a small portion of level changes. This is indeed the case. On average, 0.7% of changes at a particular well are due to pumping from within .5 miles of the well. The mean number of fields within that radius is 6.5, so a particular farmer's contribution is typically even smaller.

The takeaway from this result is that the pumping of a particular farmer is not an important determinant of that farmer's groundwater level. This evidence shows that even for farmers who value future availability of groundwater, there is very little incentive to cut back on their pumping today.

1.4.3 Response to Groundwater Levels

In this section, I aim to determine the cost of dropping groundwater levels. My strategy is to look at how sensitive crop-water usage is to groundwater levels. Under my assumptions, permanent crop growers are unable to adjust at all, and annual crop growers can only do so by adjusting their crop choice. The empirical question is how much they actually do so. I start off with a difference-in-difference strategy, building off of Figure 1.9. The idea is to compare the crop-water decisions of farmers who saw the worst groundwater level drops over my sample period with those whose levels rose or stayed relatively constant. Figure 1.9 shows this variation from 2007 to 2015 (the darkest green is a 200 foot drop, and the darkest purple is a 150 foot rise[3]). In Appendix C., I use a multinomial logit model of crop choice to answer the question in a different way.

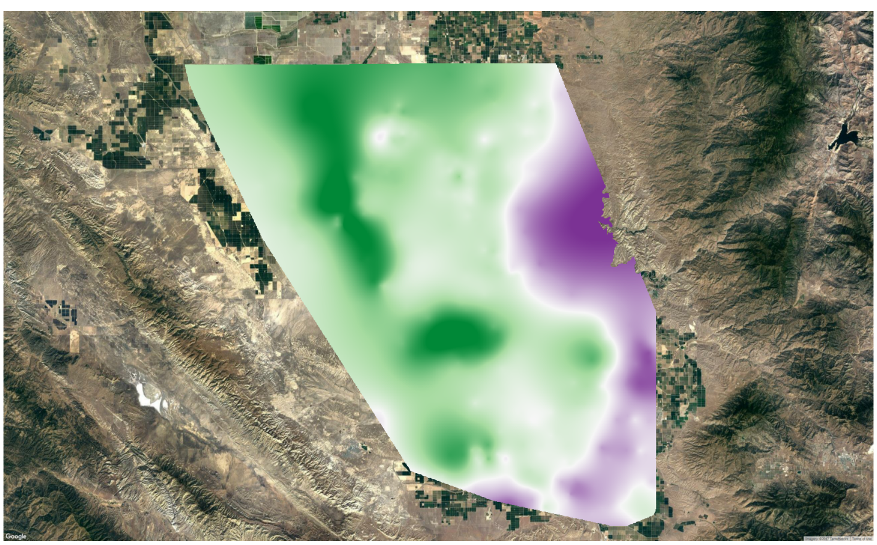

FIGURE 1.9: Groundwater Level Changes

[3]Note that most of the areas where levels rose are actually in non-agricultural areas where, presumably, little groundwater pumping occurs.

	Worst level changers	Best level changers
$L_{2015} - L_{2007}$	-156	50
w^s_{2007}	0.90	0.82
w^s_{2015}	0.31	0.33
$\Delta\bar{w}$	-5.0%	$+0.7\%$
$\Delta Pr(alfalfa)$	$-.01$	$+.03$
$\Delta Pr(fallow)$	$+.07$	$-.02$

Table 1.5: Comparing Behavior Over Time

In Table 1.5, I break out the areas where groundwater changes were largest during my sample period of 2007 to 2015. This comparison clearly illustrates the effect of groundwater levels on crop-water intensity, controlling for farmer and soil characteristics. I restrict attention to those farmers who are more able to switch their crop choice by removing farmers with tree and vine crops from the sample. A distinction should be made that it is not too costly in and of itself to switch away from a permanent crop; it just does not happen very often since they require the significant investments of the first few years of no production. I report the mean groundwater level change for the two groups, as well as their surface water allocations (in acre-feet/acre) in 2007 and 2015, which were quite similar. The fourth row is the crop-water intensity. This is mostly driven by changes in the probability of choosing alfalfa (about .25 overall) and the probability of fallowing (about .12 overall).

Alternatively, I estimate a series of specifications, with the following general form:

$$w_{it} = \alpha w^s_{it} + \beta L_{it} + \gamma w_{it-1} + dummies + \varepsilon_i, \tag{1.13}$$

where the dependent variable is the crop-water usage of farmer i in year t. The main coefficient of interest is β. My preferred specification is (3) in Table 1.6. The estimated β can be interpreted in the following way: a 1-foot drop in groundwater level is associated with annual crop farmers choosing a crop that uses .00131 less acre-feet per acre. On average, groundwater levels dropped 80 feet over my data period. This implies that annual crop farmers decreased their water usage by 4.97%.

	(1) Water Need	(2) Water Need	(3) Water Need	(4) Water Need
Surface Water	-0.00356	-0.0260**	0.0330***	-0.00517
	(0.00778)	(0.00975)	(0.00924)	(0.00766)
Depth to Groundwater	-0.00232***	-0.00234***	-0.00131***	-0.000945***
	(0.0000599)	(0.0000605)	(0.0000706)	(0.0000485)
Lagged Water Need				0.629***
				(0.00426)
Constant	2.877***	2.842***	2.753***	1.033***
	(0.0171)	(0.0244)	(0.0239)	(0.0228)
Observations	35122	35122	33664	35122
YearDummies		Yes	Yes	Yes
TownDummies			Yes	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 1.6: Effect of Depth to Groundwater on Water-Intensity of Crop Choice

1.5 Model Estimation and Parameters

Many of the parameters in the model to be used for counterfactual analysis have already been estimated, or will be taken from the literature. However, I estimate the model's scaling parameter using indirect inference. The main thrust of how this works is the following: first, I run specific regressions in the data. Then, I pick model parameters (in addition to using parameters that are estimated outside of the model), and simulate the model. Next, using the model-generated data, I run the same regressions that I ran on the real data, and then update the chosen model parameters and continue until the identified coefficients are the same using the model-generated data as they are using real data.

The main model parameter that I am estimating using this method is the variance of the unobserved term in the annual crop farmer's choice problem. This is important, as it governs how sensitive crop choice probabilities are to changes in for example, groundwater levels.

In order to construct as close to the same regression as I can in the model-generated data, I of course only use decisions made by annual crop farmers. In the real data, this involves

	Data Water Need	Model Water Need
Surface Water	0.0330***	0.224***
	(0.00924)	(0.00315)
Depth to Groundwater	-0.00131***	-0.00131***
	(0.0000706)	(0.0000205)
Constant	2.753***	2.278***
	(0.0239)	(0.00798)
Observations	33664	119220

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.7: Regressions on Data and Simulated Data

Annual decisions	Data	Model
Share water-intensive	.255	.234
Share fallow	.101	.045

Table 1.8: Model Fit: Untargeted Moments

dozens of different crops with a range of water-intensities that serve as the dependent variable in the regression in Subsection 1.4.3. In the model-generated data, there are only three choices for annual crops, water-intensive, low-water, and fallow, so I use the average water requirements for each of the first two, and obviously zero for the third.

I choose specification (3) in Table 1.6, exactly matching so that the coefficient on depth to groundwater is identical.

Note that the model matches the water-intensive share quite well, while slightly under-estimating the share that go fallow. The model also cannot replicate the lack of a strong relationship between surface water availability and choice of water-intensity.

Parameter	Meaning	Source	Value
ρ	Decay	Estimated	.201
θ	Scaling decay	Estimated	.00022
κ	New well probability	Estimated	.002
F	Well fixed cost	Literature	100000
β	Discount factor	Literature	0.95

Table 1.9: Other Model Parameters

1.6 Counterfactual Simulations

I can now think about what the effects of new water policies would be, given the estimates above. Specifically, it is important to have nailed down the responsiveness of annual crop farmers in terms of the water-intensity of their crop choice, as well as the extent to which there are spatial externalities in groundwater usage.

The Sustainable Groundwater Management Act allows for local flexibility in what groundwater regulations will look like. The model is well suited to examine different possible groundwater policy regimes. Perhaps the most interesting is groundwater taxes. While not very common in the western US, they have been recently tried in Colorado. A relevant point of comparison is the groundwater taxation regime implemented in the San Luis Valley of Colorado, documented in Smith et al. (2017). They find that farmers, acting collectively, decreased groundwater usage by 33% by implementing a $75 per acre-foot tax on groundwater pumping.

In order to evaluate groundwater taxation policy using my estimated model, I need to make some assumptions about how surface water reallocations can work. This is very important for the results. The issue is that surface water allocations observed in the data may be a result of market forces. There are significant transaction costs, but (while rare) transactions do occur.[4] When groundwater is restricted, it is important to allow farmers to buy surface water instead of paying exorbitant costs for groundwater, or letting their crops die. In this case, it is certainly true that surface water transactions would occur. To keep the comparisons fair, I allow for surface water markets at all tax levels, as described in Subsection 1.3.2.

[4] conversations with Andrew Bell, Buena Vista Water Storage District

1.6.1 Taxing Groundwater Usage

The first counterfactual simulations are with Pigouvian taxes on groundwater usage. I simplify matters by only considering per-acre groundwater pumping taxes that are uniform across farmers and across time. The former is quite reasonable, as it is unlikely that taxes would vary by farmer characteristics. (Note that they do in the first-best policy solution, as different farmers have different social costs of groundwater pumping, due to spatial heterogeneity. However, geographically heterogeneous groundwater taxes would be extremely difficult to implement in a political economy sense.) Keeping constant groundwater taxes over time is a stronger assumption. However, good arguments can be made in both directions: groundwater taxes should be higher in dry years, since externalities are higher; but groundwater taxes should be higher in wet years, so that farmers are incentivized to save groundwater for the dry years. In light of these arguments, it seems reasonable to examine taxes that are constant over time as well. Results in Chapter 2 show that results in terms of water-intensity are similar with the two types of taxes.

In order to find this constrained-optimal tax, I simulate out over a twenty-year period. I impose that groundwater taxes are equal for everyone, across time and space.

More formally, in the simulations annual farmers solve:

$$\max_{j_t, w_{it}^g, w_{it}^s} \; [p_{jt} f_j(\gamma_{ij}, w_j) - p_t^s w_{it}^s - c(w_{it}^g, L_{it}) - \tau w_{it}^g - x_j + \varepsilon_{ijt}] \tag{1.14}$$

$$s.t. \; w_j = w_{it}^s + w_{it}^g$$

$$-\bar{w}_{it}^s \leq w_{it}^s \leq w_j.$$

The permanent crop farmer's problem is:

$$\max_{w_{it}^g, w_{it}^s} \; [p_{pt} f_p(\gamma_{ij}, w_p) - p_t^s w_{it}^s - c(w_{it}^g, L_{it}) - \tau w_{it}^g - x_p] \tag{1.15}$$

$$s.t.\ w_p = w^s_{it} + w^g_{it}$$

$$-\bar{w}^s_{it} \leq w^s_{it} \leq w_j.$$

These are very similar to what is outlined in the model section, but a few key distinctions should be made. One is obviously that groundwater is more costly, in that it is no longer just the private cost determined by $c(w^g_{it}, L_{it})$, but there is an additional per-acre-foot tax of τ. Two, there is an additional constraint on the choice of surface water, w^s_{it}. This is just saying that farmers can sell up to their endowment of surface water, and can buy up to their crop need. Three, the price of surface water p^s_t is endogenous, found as described in Subsection 1.3.2. And four, perhaps most importantly, is the issue of who and where the permanent crop farmers are. Their presence remains exogenous, as there is no decision about permanence in the model, but I take their location and crop choice directly from the Kern County crop data. This means that the simulations are more accurate in determining from where exactly groundwater is likeliest to be pumped. Identifying these patterns helps to predict how groundwater levels evolve into the future, as I still use the estimated relationship in the counterfactual simulations in Section 1.6.

Next comes the process of comparing between the different outcomes driven by different options for the tax. I compute welfare as the sum of three components: 1) the sum of discounted profits for permanent crop farmers, 2) the sum of discounted realized profits for annual crop farmers given their crop choices, and 3) the sum of discounted tax revenues.

Restricting attention to constant groundwater taxes, I find that the optimal Pigouvian tax on groundwater pumping is $85 per acre-foot. This is a large increase in the cost of groundwater pumping, but it only increases aggregate discounted welfare by 0.5%. This is due in large part to the inflexible permanent crop farmers. However, if we measure the welfare effects relative to the status quo, where there is no market for surface water, the increase is 6.3%.

The reason welfare goes up at all is that groundwater levels are not depleted as much, benefits that compound into the future. Under the optimal tax, define the average groundwater pumped as $\tilde{w}_t^g$ acre-feet per acre. This is 15% lower than the status quo. This quantity will be used in the following section. It is important to note that my results do not take into account the benefits of improved groundwater levels on municipal water users, or on the decreased probability of land subsidence.

It is illustrative to compare my simulated results with the Colorado experimental results. With the same $75 tax on groundwater pumping (inefficiently low in my case), I find that farmer groundwater pumping is less responsive, only decreasing by 14%. This is at least partly due to the fact that permanent crops are much more common in California.

1.6.2 Quantity Regulations on Groundwater Usage

In this subsection, I look at regulations that act directly through quantities, rather than prices. It does seem more likely that SGMA will institute quantity rules than price-based policies.

A point of clarification should be made. In general, surface water availability is quite volatile. This means that from year to year, groundwater usage swings wildly as well. In this case, a quantity regulation that does not vary with surface water availability will clearly not be optimal. This is because a constant cap on groundwater has no chance to both keep farmers afloat in the driest years (providing buffer value) while also having any bite in most years. So in order to make a fair comparison between price and quantity regulations, I set the cap on groundwater availability to be equal to the average equilibrium quantity $\tilde{w}_t^g$ pumped under the optimal Pigouvian tax, found in the previous subsection. This is in the spirit of introductory models of negative externalities, when comparisons can be made between surplus with a tax, with command and control, and with cap and trade.

Command and Control

In this counterfactual, the per-acre quantity of groundwater is restricted. In essence, a farmer i now has two water allocations: $\bar{w}_{it}^s$ and $\tilde{w}_{it}^g$. Both are indexed by i, but more practically, surface water allocations can differ by district and groundwater allocations are likely to be constant across the county. The surface water allocation is tradable, and the groundwater allocation is not. (Note: with the fixed-proportions production function, having neither allocation tradable would mean permanent crops would die in dry years. Since this rarely happens, it seems reasonable to allow for tradable surface water.) So the usual inefficiencies of command and control are mitigated, as high marginal product producers can purchase surface water on the open market. Still, there is a type of "use it or lose it" inefficiency with groundwater.

More formally, in these simulations annual farmers solve:

$$\max_{j_t, w_{it}^g, w_{it}^s} \; [p_{jt} f_j(\gamma_{ij}, w_j) - p_t^s w_{it}^s - c(w_{it}^g, L_{it}) - x_j + \varepsilon_{ijt}] \tag{1.16}$$

$$s.t. \; w_j = w_{it}^s + w_{it}^g$$

$$w_{it}^g \leq \tilde{w}_{it}^g$$

$$-\bar{w}_{it}^s \leq w_{it}^s \leq w_j.$$

The permanent crop farmer's problem is:

$$\max_{w_{it}^g, w_{it}^s} \; [p_{pt} f_p(\gamma_{ij}, w_p) - p_t^s w_{it}^s - c(w_{it}^g, L_{it}) - x_p] \tag{1.17}$$

$$s.t. \; w_p = w_{it}^s + w_{it}^g.$$

$$w_{it}^g \leq \tilde{w}_{it}^g$$

$$-\bar{w}_{it}^s \leq w_{it}^s \leq w_j.$$

Note the extra quantity constraint that groundwater pumping has to be less than or equal to the cap $\tilde{w}_{it}^g$, found in Subsection 1.6.1.

Cap and Trade

Here, both surface water and groundwater are tradable. In principle, it would be desirable to have two markets for water. One would be for surface water, as described in previous counterfactuals. The other would be for the right to pump groundwater. In that case, when a transaction occurs, farmer i simply pumps more and farmer j pumps less, and money changes hands. However, then I need to solve for the price of surface water and the price of groundwater claims. For computational purposes, I assume that when a transaction occurs, the groundwater is pumped by the seller, not the buyer. This allows there to just be one equilibrium price of water. It is lower than the market price of water under command and control because groundwater supplies are available on the market as well. Here, what can be sold on the open market is both surface water and groundwater.

Unlike a basic model of externalities, this is also different from the tax counterfactual, since what is allowed to be traded is again different.

More formally, in these simulations annual farmers solve:

$$\max_{j_t, w_{it}^g, w_{it}^s} \; [p_{jt} f_j(\gamma_{ij}, w_j) - p_t w_{it}^s - c(w_{it}^g, L_{it}) - x_j + \varepsilon_{ijt}] \tag{1.18}$$

$$s.t. \; w_j = w_{it}^s + w_{it}^g$$

$$w_{it}^g \leq \tilde{w}_{it}^g$$

$$-\bar{w}_{it}^s - \tilde{w}_{it}^g \leq w_{it}^s \leq w_j.$$

The permanent crop farmer's problem is:

$$\max_{w_{it}^g, w_{it}^s} \; [p_{pt} f_p(\gamma_{ij}, w_p) - p_t w_{it}^s - c(w_{it}^g, L_{it}) - x_p] \tag{1.19}$$

$$s.t. \; w_p = w_{it}^s + w_{it}^g.$$

$$w_{it}^{g} \leq \tilde{w}_{it}^{g}$$

$$-\bar{w}_{it}^{s} - \tilde{w}_{it}^{g} \leq w_{it}^{s} \leq w_{j}.$$

The main differences here are that the price of water in the objective functions is not p_t^s, but rather just p_t, since both types are included. And the third constraint is showing how the amount sold is limited only by the sum of both water allocations.

A good way of illustrating the differences between these allocation systems is by examining the options of a farmer who is considering fallowing, under the three counterfactual regimes. Under command and control, by choosing to fallow the farmer can only sell her surface water allocation. However, this is quite valuable, since permanent farmers will likely be in need of it and not have any other options once they exhaust their own groundwater caps. Under cap and trade, by choosing to fallow the farmer can not only sell her surface water allocation, but also has the option to pump groundwater to sell. Water will not be as pricey as before, since permanent farmers have more options, but fallowing still looks pretty good since she has more water to sell. Under the groundwater tax regime, by choosing to fallow the farmer can only sell her surface water allocation. This does not have very much value, since no permanent farmers will be capped. They will only want to buy her water if their groundwater is more costly. So it makes sense that we see less fallowing in the tax regime, as it is simply less lucrative.

The results from each of the counterfactuals are in Table 1.10. Each row is a different policy environment that is being compared to the status quo, no-trade environment outlined in Section 1.3.

Again, to make a fair comparison between the four, I allow for an endogenous surface water market to arise. In some sense, this cautions against interpreting the full size of the gains too literally: there is no particular reason the market should be able to exist if it does not now. It's simply the case that without the surface water market performing, reasonable analysis of quantity regulations could not be done.

Note how much better annuals do with a quantity restriction. They can force permanents to pay much more for their surface water in these cases, as can be seen in the last column, the market price for water.

	Welfare	Annuals welfare	Permanents welfare	$\bar{L_{21}}$	$\sum w^g$	fallow	$\bar{p^s}$
Market, no tax	+5.8%	+5.6%	+5.9%	+14.8%	−5.2%	+2 pp	164
Optimal tax	6.3	10.4	5.0	19.8	−12.8	+6 pp	227
Command	4.0	29.9	−4.2	13.6	−20.7	+14 pp	436
Cap & trade	6.0	24.8	0.0	15.5	−24.7	+15 pp	312

Table 1.10: Counterfactual results (relative to no-trade status quo)

1.7 Conclusion

This paper studies the complex issue of agricultural water usage in California. Data collected from numerous sources on crop choice, surface water deliveries, groundwater levels, and more are necessary to inform empirical analysis in this paper. Then the relatively simple model laid out above can be an intuitive way to think about agricultural water decisions and policy in California.

After determining parameters governing crop choice, water usage, and the extent to which externalities matter, I use the model to run policy experiments, reflecting the range of implementation options of SGMA. I find that the optimal tax on groundwater pumping is $85 per acre-foot, and that this increases total welfare by 6.3%. However, most of this increase is simply due to there being a market for surface water. Quantity regulations can also increase welfare, but not by as much. A cap and trade regime significantly outperforms command and control.

By using the model to estimate the impacts of various policy options, this paper contributes to the literature on optimal groundwater management, showing how aspects like groundwater hydrology and crop permanence play an important role. While I show that permanence matters, in this paper I do not endogenize the farmer decision to move to growing a permanent crop. This is clearly a very important margin, and deserves further study. In Chapter 2, I study this decision in a similar but simplified context.

Chapter 2

The Common Pool and Socially Excessive Volatility

2.1 Introduction

In this chapter, I illustrate the dynamics of farmers choosing between permanent and annual crops, either of which can be water-intensive. An important feature is that choosing water-intensive crops makes growing those crops more difficult for everyone else. This externality is a key part of the model.

This setting is relevant for the study of farming in California, where a significant portion of land is planted with permanent crops (Johnson and Cody, 2015). These are usually trees, like almonds, pistachios, and citrus fruit. However, vineyards and blueberries are also permanent crops. These crops are on a range of water-intensity, with pistachios being on the high end and grapes on the low end. The same range of water-intensity can be found amongst annual crops. The distinction with annuals is that they are more flexible when conditions change: a farmer can more easily switch to a more or less water-intensive crop the next year. While not in the agricultural context, (Sandholm, 2007) studies a similar setting involving externalities and Pigouvian pricing from a theoretical standpoint. The closest paper to this one is (Franklin et al., 2017), which studies permanence with a vintage capital model focused on Australia. This paper is the first to study the interaction between the choice over water-intensity and the dynamic decision to plant a permanent crop.

This is a significant issue in California, a drought-prone state that has the nation's largest population and its largest agricultural sector, according to the agricultural census (USDA, 2012). Scientists project climate change to impact precipitation, snowpack, and reservoirs in California, potentially increasing extreme weather and drought (Williams et al., 2015).

Why does growing water-intensive crops impose external costs? The model will be general, but the preferred interpretation is because of the lack of property rights for groundwater usage. This draws from the seminal work on the Tragedy of the Commons by (Hardin, 1968). Both groundwater and surface water are used in California agriculture, and most years the proportion is roughly even, with groundwater becoming more important during droughts (Hanak et al., 2016). Generally it is the case that farmers have limited access to surface water, and supplement crop need with additional groundwater pumping. Surface water property rights are generally well-defined, so the model will just have externalities in groundwater usage, and thus only for the water-intensive crops. The externality exists by assumption in this paper, but there is plenty of evidence that these kind of externalities are significant in agricultural contexts (Pfeiffer and Lawell, 2012), (Guilfoos et al., 2013), and Chapter 1. The recent drought led to land subsidence, a phenomenon where the land surface sinks due to huge extractions of groundwater. All of this spurred California to pass the Sustainable Groundwater Management Act, which put into place a framework to begin regulating groundwater usage in the coming years (Kiparsky et al., 2016). These regulations are not yet written, but price-based instruments are on the table (Kiparsky, 2016). In the rest of the paper, I show what these kind of policies might do. Importantly, this takes into account the fact that so much agricultural land is planted with permanent crops, which essentially lock in water demand for years to come.

In the model, crops are either water-intensive or low-water. Much of the literature on these issues classify crops in this way (Cooley, 2015), minimizing the importance of the intensive-margin of water use (for a given crop), as irrigation technology and other water-saving production measures are now widespread in California. That interpretation can be made here as well, but it does not need to be: I retain the flexibility to interpret a particular farmer as choosing between a water-intensive or a low-water version of the same crop, meaning that choice could be made on the intensive margin as well. Each type can be permanent or annual, making four distinct possibilities. Permanent is not necessarily any more water-intensive than annual--it is just a commitment. This commitment is not literally permanent, since removing an almond orchard is (while costly) an option available to farmers. It is best to think of the permanent crop as significantly reducing future flexibility.

It is easy to solve the model if we assume that farmers are not forward-looking, because then in the model permanent crops and annual crops are no different. I show this in the first part of Section 2.3.

Adding forward-looking behavior to the model means that farmers choosing permanent crops are likely giving up flexibility in the future. This flexibility is more valuable when there is more variability in which option will be preferred in the future, either because of idiosyncratic or aggregate shocks. In this paper, I focus on an aggregate shock structure where the state alternates between high and low (wet and dry). This gives rise to an equilibrium featuring cycles in the proportion of farmers who choose water-intensive. As expected, a tax on water-intensive crops leads them to be less attractive in both parts of the cycle (and does so for both the annual and permanent versions). Less obvious however, is that the amplitude of cycles decreases with the tax, with interesting implications on the share that choose permanent: we actually see slightly more permanent crops with a tax on water-intensity, because the cyclicality of the aggregate shock is mitigated.

I show how these results (with Pigouvian taxes) compare to an environment where the tax remains constant over time, as studied in Chapter 1 of this dissertation. We are more likely to see this type of tax in the California setting, and it is thus important to examine how closely this regime approximates the Pigouvian tax environment (where the per-unit tax varies depending on where in the cycle we are). I find that this constant tax delivers similar results in what share of farmers have permanent crops and the average water-intensity, but the aforementioned cycles in the share water-intensive are significantly less affected. Notably for policymakers, I find that both taxes slightly increase the share of farmers that grow permanent crops, as the taxes serve to mitigate the volatility driven by the aggregate shock. This could be a potential unintended consequence of groundwater regulation.

While the back-and-forth structure of the aggregate shock is obviously not a complete picture of reality, it does reflect the pattern of farmers choosing to grow more water-intensive crops during wet periods, which sometime can lock them into those choices for many years, including periods of severe drought. The incentives faced by farmers in this model have not been well-studied in the literature.

The rest of the chapter proceeds as follows. Section 2.2 provides more background and evidence of patterns in the data. In Section 2.3, I lay out the model in detail. First I solve a simpler version, and then add in the components necessary to study the permanence/water-intensity interaction. Then in Section 2.4 I show how to solve and simulate the model, examining characteristics of equilibrium and what they can tell us about this setting. I then discuss robustness and conclude.

2.2 Background and Evidence

California is like few other places in the world in terms of its agricultural productivity. The Mediterranean climate allows for the cultivation of an extremely diverse set of crops, many of which are not possible to grow at any kind of large scale in other parts of the country. Fourteen of these crops see 99% of their production come from California (CDFA, 2017). These include many permanent crops, like nut trees, but also some annuals, such as garlic and melons, which come on a wide range of water-intensity. The most water-intensive crops like alfalfa use about 4 acre-feet of water per acre, and the least, like most grains and vegetables use about 2 acre-feet per acre. An acre-foot is the amount of water it takes to cover one acre of land with one foot of water, or about 325,900 gallons. This is roughly how much a family of four uses for domestic purposes in a year.

Figure 2.1 shows the fraction of acreage devoted to growing permanent crops in Kern County, California, using detailed micro data on yearly crop choice (County of Kern, 2016). Kern County is in an important agricultural area in the Central Valley of California, and is the top-ranked county in the country in terms of agricultural revenue.

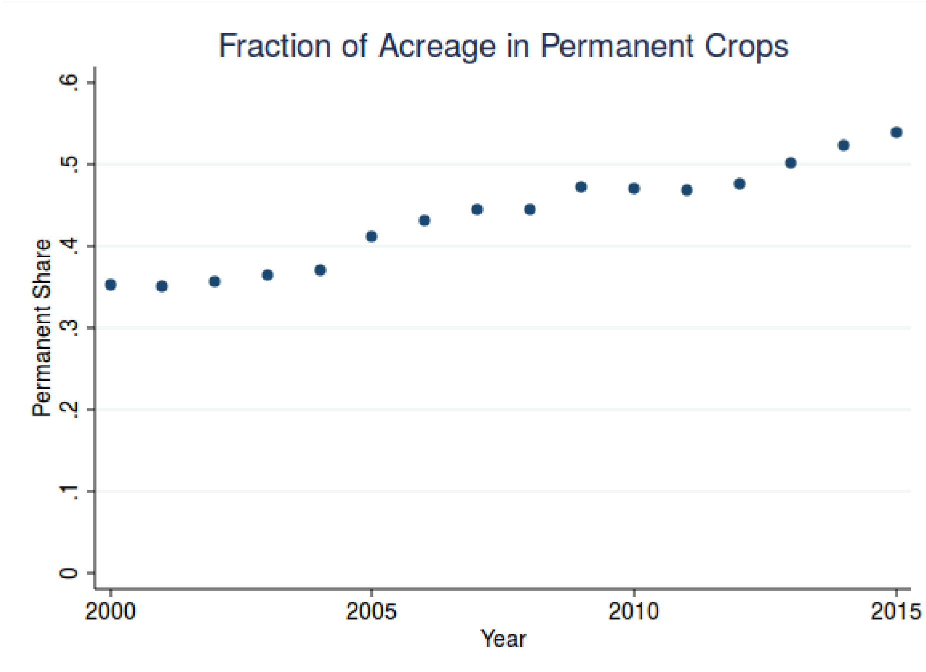

FIGURE 2.1: Permanent Acreage, Kern County

In the state as a whole, there has been relatively little study on the decision to grow permanent crops, but data is available based on agricultural surveys (USDA, 2017). From 2000 to 2014, permanent crop acreage increased 21% to 2.9 million acres. This was despite total farmland dropping by one million acres over the same time period. The gains in permanent acreage mostly came at the expense of crops like cotton and alfalfa, both of which are relatively water-intensive. This is similar to the patterns observed in the Kern County micro data, which is focused on in Chapter 1.

2.3 Model

In this section, I lay out a theoretical model that fleshes out the topics discussed in the introduction.

There are a continuum of infinitely-lived farmers indexed on the unit interval by i. Time is discrete, and the common discount factor is β. Crops are of two distinct water-intensities: $c = 0$ is low-water, and $c = 1$ is water-intensive. There is a permanent and annual version of each. Based on last period's decisions (and a probability λ of no longer being stuck with permanent), a farmer is classified as either initially permanent or initially annual. If she is initially permanent, she is stuck and makes no choice. If she is initially annual, she picks between one of four options, two of which are permanent. The timing of the model works as follows: when making a decision in period t, farmers see their own shocks (idiosyncratic and whether they have reverted, if they were permanent in $t-1$), as well as the distribution of choices in period $t-1$.

Suppose a farmer picks $c=0$ (or is stuck with the low-water permanent option). Then she gets return ϕ plus an idiosyncratic shock explained below. Another interpretation of $c=0$ is that it is the outside option.

If she picks $c=1$, then she gets profit $f(x_t, \xi_t)$ (plus an idiosyncratic shock), where x_t is the period t share that are water-intensive (annual or permanent) and ξ_t is an aggregate shock, like rainfall or surface water availability. Assume $f_x < 0$ and $f_\xi > 0$.

Now let's explain the idiosyncratic shock. Let ε_{ia0t}, ε_{ia1t}, ε_{ip0t} and ε_{ip1t} be the idiosyncratic returns for choosing a combination of $c=0$ or $c=1$ and p or a. Each is distributed type I extreme-value.

Let n°_{at} be the initial annual and n°_{p0t} and n°_{p1t} be the initial permanent of each type at the beginning of period t. Let s_{p0t}, s_{p1t}, and $s_{at} = s_{a0t} + s_{a1t}$ be the fraction with each option based on decisions in period t. Then we obtain

$$n^\circ_{p0t+1} = (1-\lambda)\, s_{p0t} \tag{2.1}$$

$$n^\circ_{p1t+1} = (1-\lambda)\, s_{p1t} \tag{2.2}$$

$$n^\circ_{at+1} = s_{at} + \lambda s_{p0t} + \lambda s_{p1t}. \tag{2.3}$$

2.3.1 Simplest case

Consider the simple case where there is no aggregate shock, no permanence, and $\beta = 0$. I will suppress the t as well. Let x be equilibrium share choosing water-intensive. A farmer

i does this if

$$\phi + \varepsilon_{ia0} \leq f(x) + \varepsilon_{ia1}. \tag{2.4}$$

With the assumption that the idiosyncratic shock is distributed type I extreme value, and since all farmers are identical save for that shock,

$$x = \frac{\exp(f(x))}{\exp(\phi) + \exp(f(x))}. \tag{2.5}$$

Simplifying,

$$x = \frac{1}{\exp(\phi - f(x)) + 1}. \tag{2.6}$$

Under $f'(x) < 0$, it is clear that there exists a unique equilibrium. Let x^e solve the above equation. Due to the externality imposed by water-intensive crops, we are also interested in the social planner's solution, x^*. The externality imposed is $xf'(x)$, the product of how many water-intensive farmers there are and how much each is hurt by the marginal water-intensive farmer. Suppose the water-intensive crop farmers are made to pay this cost they impose on everyone else. Then a farmer i chooses water-intensive if

$$\phi + \varepsilon_{ia0} \leq f(x) + f'(x)x + \varepsilon_{ia1}. \tag{2.7}$$

As above, this implies that

$$x = \frac{\exp(f(x) + f'(x)x)}{\exp(\phi) + \exp(f(x) + f'(x)x)}. \tag{2.8}$$

Assume that $f(x) + f'(x)x$ is decreasing, meaning:

$$f'(x) + f'(x) + f''(x)x \leq 0,$$

or

$$2f'(x) + f''(x)x < 0. \tag{2.9}$$

For example, let $f(x) = \exp(-\delta x)$. Then we have:

$$\begin{aligned} f'(x) &= -\delta \exp(-\delta x) & (2.10) \\ f''(x) &= \delta^2 \exp(-\delta x). & (2.11) \end{aligned}$$

Substituting in, we obtain:

$$2f'(x) + f''(x)x = -2\delta \exp(-\delta x) + \delta^2 x \exp(-\delta x). \tag{2.12}$$

So we require

$$\begin{aligned} 2\delta &> \delta^2 x \\ 2 &> \delta x. & (2.13) \end{aligned}$$

Since $x \leq 1$, this means $\delta < 2$.

2.3.2 Adding permanence

Now, add back in the option of growing the permanent crops, but retain $\beta = 0$. In equilibrium we get the same result, due to the following argument. Suppose we take as given x^e from before is what will happen. We can break the water-intensive choices up: choose between permanent water-intensive or annual water-intensive. Suppose the farmer begins as annual. The probability of choosing permanent and water-intensive is

$$s_{p1} = \frac{\exp(f(x))}{\exp(f(x)) + \exp(\phi) + \exp(f(x)) + \exp(\phi)}. \tag{2.14}$$

Then the probability of choosing water-intensive is

$$\begin{aligned} s_{p1} + s_{a1} &= \frac{\exp(f(x)) + \exp(f(x))}{\exp(f(x)) + \exp(\phi) + \exp(f(x)) + \exp(\phi)} \\ &= \frac{\exp(f(x))}{\exp(f(x)) + \exp(\phi)}, & (2.15) \end{aligned}$$

so we get the same as above. Thus, simply adding permanent doesn't change anything. This is unsurprising without forward-looking farmers, with which I augment the model in the next section.

2.3.3 Full model: cyclical aggregate shock

In this subsection, I consider what changes when there is an aggregate shock ξ to the utility of water-intensive crops. This shock will take a very particular form of cycling between high and low every period, analogous to wet conditions and drought conditions. Let $\xi_L < \xi_H$ and assume the shock alternates every period between the two types, with $\xi_t = \xi_L$ in odd periods and $\xi_t = \xi_H$ in even periods. The objective is to characterize the equilibrium and to examine its properties for two cases: first no regulation, and second where farmers pay Pigouvian taxes.

The state variables can be just as easily understood both as last period's distribution of permanent crops, or this period's distribution of farmers who are stuck, since the latter is simply a fraction $1-\lambda$ of the former. For notational simplicity, I write everything in terms of the latter, so the state variables are $\{n^\circ_{p0t}, n^\circ_{p1t}, \xi_t\}$, where n°_{p0t} and n°_{p1t} be the measure of agents of each type at the beginning of the period.

There is a value function associated with each of the following situations: entering a period with a water-intensive permanent crop, defined as $V_{p0}(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t)$; entering a period with the low-water permanent crop, defined as $V_{p1}(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t)$; and entering a period with the ability to choose (either due to having an annual crop the previous period or having the opportunity to switch from a permanent crop), defined as $V_a(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t)$.

$$
\begin{aligned}
V_{p0}(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t) &= \phi + \\
&\quad \beta\left[\lambda V_a(n^\circ_{p0t+1}, n^\circ_{p1t+1}, \xi_{t+1}) + (1-\lambda)V_{p0}(n^\circ_{p0t+1}, n^\circ_{p1t+1}, \xi_{t+1})\right] && (2.16)\\
V_{p1}(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t) &= f(x_t) + \\
&\quad \beta\left[\lambda V_a(n^\circ_{p0t+1}, n^\circ_{p1t+1}, \xi_{t+1}) + (1-\lambda)V_{p1}(n^\circ_{p0t+1}, n^\circ_{p1t+1}, \xi_{t+1})\right] && (2.17)\\
V_a(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t) &= \max\Big\{ V_{p1}(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t), V_{p0}(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t), && (2.18)\\
&\quad f(x_t) + \beta V_a(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t), \phi + \beta V_a(n^\circ_{p0t}, n^\circ_{p1t}, \xi_t)\Big\}
\end{aligned}
$$

Note that the arguments of the value functions on the right-hand side are the next period's state variables. An expectation does not appear, since there is no aggregate uncertainty. Each farmer takes as given what everyone else does, and can forecast next period's state variables perfectly.

Now consider choice behavior, which together with the initial conditions give us distributions over time. Taking as given x_t, farmers that can choose pick between their four options, with the following utilities:

$$\begin{aligned}
U_{a0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= \phi + \beta V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + \varepsilon_{ia0t} && (2.19)\\
U_{a1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= f(x_t) + \beta V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + \varepsilon_{ia1t} && (2.20)\\
U_{p0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= \phi + \beta \Big[(1-\lambda)\, V_{p0}(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + && (2.21)\\
&\quad \lambda V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) \Big] + \varepsilon_{ip0t} &&\\
U_{p1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= f(x_t) + \beta \Big[(1-\lambda)\, V_{p1}(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + && (2.22)\\
&\quad \lambda V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) \Big] + \varepsilon_{ip1t} &&
\end{aligned}$$

The value of having the opportunity to choose is:

$$\begin{aligned}
V_a(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= E\big[\max\{U_{a0t}, U_{a1t}, U_{p0t}, U_{p1t}\}\big] &&\\
&= \log(\exp(U_{a0t} - \varepsilon_{ia0t}) + \exp(U_{a1t} - \varepsilon_{ia1t}) && (2.23)\\
&\quad + \exp(U_{p0t} - \varepsilon_{ip0t}) + \exp(U_{p1t} - \varepsilon_{ip1t})). &&
\end{aligned}$$

Given that opportunity to choose, we can derive choice probabilities based on the distributional assumptions on the idiosyncratic shock.

$$cp_{a0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) = \frac{\exp(U_{a0t} - \varepsilon_{ia0t})}{\exp(V_a(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t))} \tag{2.24}$$

$$cp_{a1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) = \frac{\exp(U_{a1t} - \varepsilon_{ia1t})}{\exp(V_a(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t))} \tag{2.25}$$

$$cp_{p0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) = \frac{\exp(U_{p0t} - \varepsilon_{ip0t})}{\exp(V_a(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t))} \tag{2.26}$$

$$cp_{p1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) = \frac{\exp(U_{p1t} - \varepsilon_{ip1t})}{\exp(V_a(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t))} \tag{2.27}$$

Definition of equilibrium

Definition. Given a set of parameters $\{\lambda, \sigma, \xi_H, \xi_L, \phi\}$, a production function $f(x, \xi)$, preferences, an initial distribution of crops $\{\tilde{s}_{p0}, \tilde{s}_{p1}, \tilde{s}_{a0}, \tilde{s}_{a1}\}$, and idiosyncratic shocks $\{\varepsilon_{ip0t}, \varepsilon_{ip1t}, \varepsilon_{ia0t}, \varepsilon_{ia1t}\}$ for all i and for all t, a **competitive equilibrium** is a set of value functions $\{V_a, V_{p0}, V_{p1}\}$ (Equations 2.16-2.18) such that the choice probabilities $\{cp_{p0t}, cp_{p1t}, cp_{a0t}, cp_{a1t}\}$ (Equations 2.24-2.27) given by the value functions are consistent with maximizing behavior and satisfy the following conditions:

$$\begin{aligned}
s_{p0t} &= n^{\circ}_{p0t} + cp_{p0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * n^{\circ}_{at} \\
&= (1-\lambda)s_{p0t-1} + cp_{p0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * (s_{a0t-1} + s_{a1t-1} + \lambda s_{p0t-1} + \lambda s_{p1t-1}) && (2.28) \\
s_{p1t} &= n^{\circ}_{p1t} + cp_{p1t} * n^{\circ}_{at} \\
&= (1-\lambda)s_{p1t-1} + cp_{p1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * (s_{a0t-1} + s_{a1t-1} + \lambda s_{p0t-1} + \lambda s_{p1t-1}) && (2.29) \\
s_{a0t} &= cp_{a0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * n^{\circ}_{at} \\
&= cp_{a0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * (s_{a0t-1} + s_{a1t-1} + \lambda s_{p0t-1} + \lambda s_{p1t-1}) && (2.30) \\
s_{a1t} &= cp_{a1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * n^{\circ}_{at} \\
&= cp_{a1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) * (s_{a0t-1} + s_{a1t-1} + \lambda s_{p0t-1} + \lambda s_{p1t-1}) && (2.31) \\
x_t &= s_{p1t} + s_{a1t}. && (2.32)
\end{aligned}$$

In equilibrium, we get stationary cycle conditions (after enough periods from the initial states)

$$\begin{aligned}
n^{\circ}_{p0L} &= (1-\lambda)\, n^{\circ}_{p0H} + \left(1 - n^{\circ}_{p0H} - n^{\circ}_{p1H}\right) * cp_{p0H}(n^{\circ}_{p0H}, n^{\circ}_{p1H}, \xi_H) && (2.33)\\
n^{\circ}_{p1L} &= (1-\lambda)\, n^{\circ}_{p1H} + \left(1 - n^{\circ}_{p0H} - n^{\circ}_{p1H}\right) * cp_{p1H}(n^{\circ}_{p0H}, n^{\circ}_{p1H}, \xi_H) && (2.34)\\
n^{\circ}_{p0H} &= (1-\lambda)\, n^{\circ}_{p0L} + \left(1 - n^{\circ}_{p0L} - n^{\circ}_{p1L}\right) * cp_{p0L}(n^{\circ}_{p0L}, n^{\circ}_{p1L}, \xi_L) && (2.35)\\
n^{\circ}_{p1H} &= (1-\lambda)\, n^{\circ}_{p1L} + \left(1 - n^{\circ}_{p0L} - n^{\circ}_{p1L}\right) * cp_{p1L}(n^{\circ}_{p0L}, n^{\circ}_{p1L}, \xi_L) && (2.36)
\end{aligned}$$

To explain the first two equations above, take as given the previous period was a high (wet) period. Of the n°_{p0H} that were initially permanent (i.e., stuck) in the high period, a fraction $1 - \lambda$ of those are initially permanent the next period, which will be an L period. We also have to add on the annuals (given by $1 - n^{\circ}_{p0H} - n^{\circ}_{p1H}$) who choose to become stuck: a fraction $cp_{p0H}(n^{\circ}_{p0H}, n^{\circ}_{p1H}, \xi_H)$ become permanent and low-water, and $cp_{p1H}(n^{\circ}_{p0H}, n^{\circ}_{p1H}, \xi_H)$ become permanent water-intensive.

Also of interest is the fraction that are growing either permanent crop option. These choices serve to lock in farmers for an expected $\frac{1}{\lambda}$ periods. Note that in this model, for simplicity, there is no added benefit to permanent crops, aside from the idiosyncratic shock. Different policy environments or parameters simply lead to a smaller or larger disincentive to choose permanent. The idea is that bigger swings in the aggregate shock, or a larger variance σ in the idiosyncratic shock to utility mean that locking in to a permanent option looks less attractive. This idea is formalized in the following theorem.

Theorem. *For all* $\lambda \in (0, 1)$, *the sum of the choice probabilities for the permanent options,* $cp_{p1t} + cp_{p0t}$, *is less than* $.5$.

Proof. By definition, the sum of the choice probabilities is 1. As described above, the permanent version of a crop provides no benefit compared to the annual version—it simply reduces flexibility. So for all feasible combinations of $\{n^{\circ}_{p0,t}, n^{\circ}_{p1,t}, \xi_t\}$, it must be true that both

$$\begin{aligned}
cp_{p1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &< cp_{a1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t), && (2.37)\\
cp_{p0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &< cp_{a0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t). && (2.38)
\end{aligned}$$

Assume the theorem does not hold. So for some combination of parameters and state variables, $cp_{p1t} + cp_{p0t} \geq .5$. Then by the previous, since the annuals have higher choice probabilities than their permanent counterparts, $cp_{a1t} + cp_{a0t} > .5$. Adding the previous two inequalities gives a strict inequality: $cp_{p1t}+cp_{p0t}+cp_{a1t}+cp_{a0t} > 1$, a contradiction. □

A distinction should be made between choice probabilities and the distribution of choices after each period, both of which vary over time and with the cycle. The two are different due to the permanence aspect of the model. The share of permanent crops can certainly sum to greater than one-half, and is indeed likely for most of the parameter space. An equal distribution amongst the four crops can only be sustained with significantly higher choice probabilities for the annual options.

The process of solving the model also requires finding a fixed point for x_t each period. To study the path from the initial distribution to equilibrium, I allow for t subscripts. Eventually this leads to the high-low equilibrium cycle, as can be seen in the figures in the Appendix.

2.3.4 Pigouvian taxes

Farmers who choose either water-intensive crop are imposing an externality on the other water-intensive crop growers. The Pigouvian tax in period t is

$$\tau_t = f'(x_t, \xi_t)x_t, \tag{2.39}$$

where x_t is the fraction choosing water-intensive crops in period t. This tax varies based on the period, since x_t is not constant. Next we write down the choices if the optimal Pigouvian tax is implemented. Note there is the same tax for permanent as annual.

$$\begin{aligned}
U_{a0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= \phi + \beta V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + \varepsilon_{ia0t} && (2.40)\\
U_{a1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= f(x_t) + f'(x_t)x_t + \beta V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + \varepsilon_{ia1t} && (2.41)\\
U_{p0t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= \phi + \beta \left[(1-\lambda) V_{p0}(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + \lambda V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1})\right] && (2.42)\\
&\quad + \varepsilon_{ip0t} \\
U_{p1t}(n^{\circ}_{p0t}, n^{\circ}_{p1t}, \xi_t) &= f(x_t) + f'(x_t)x_t + && (2.43)\\
&\quad \beta \left[(1-\lambda) V_{p1}(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1}) + \lambda V_a(n^{\circ}_{p0t+1}, n^{\circ}_{p1t+1}, \xi_{t+1})\right] + \varepsilon_{ip1t}.
\end{aligned}$$

We can derive choice probabilities similarly to above. Theorem 1 still goes through, and we will clearly get lower choice probabilities for both water-intensive options.

2.3.5 Average taxes

It is also important to look at a slightly more feasible option for internalizing externalities, even if it is not optimal. That would be holding fixed the tax at the average of the Pigouvian tax at the high and low parts of the cycle. Define this as:

$$\bar{\tau} = f'(\bar{x}_t)\bar{x}_t. \tag{2.44}$$

Utilities, value functions, and choice probabilities are defined analogously to the above. If f is linear in x_t, as it will be in the simulations in Section 2.4, note that this tax will be higher than the Pigouvian tax in dry periods and lower in wet periods.

2.4 Numerical Simulations

In this section, I aim to answer the following questions using a model solved and simulated computationally: how does the equilibrium with the uncorrected common pool problem compare with the Pigouvian tax case? And in particular, what are the differences in degree of lock-in, where more are choosing one of the permanent crops? And how does it compare with the extent to which there is time variation? How are cycles affected?

Assume the following about the aggregate shock:

$$\xi_H = \bar{\xi} + \frac{\Delta}{2} \tag{2.45}$$

$$\xi_L = \bar{\xi} - \frac{\Delta}{2}. \tag{2.46}$$

So the mean is $\bar{\xi}$ and the spread is Δ.

Let's take the following simple functional form for $f(x_t, \xi_t)$:

$$f(x_t, \xi_t) = \xi_t - m x_t. \tag{2.47}$$

The coefficient m is a constant, and recall that x is the fraction that are water-intensive. As an aside, note that in this formulation there is no state variable regarding the extent to which there is an externality, unlike in Chapter 1, where groundwater levels may be extensively depleted over time.

Now consider comparative statics in Δ. What happens when Δ grows? Do we get excessive lock-in of water-intensive crops, when the state is L? If there is no Pigouvian tax, then farmers avoid paying $\frac{\partial f(x_t, \xi_L)}{\partial x} x_t$ tomorrow. With the dynamic nature of the decision, I make use of the cycle structure of the aggregate shock. In equilibrium, there will be a fraction of farmers making each possible choice, due to the extreme-value shocks. There will be a higher fraction choosing water-intensive (both annual and permanent) in wet periods, similarly for less water-intensive in dry periods. To figure out characteristics of the cycle and how they vary with different parameters, I solve the model using value function iteration.

The baseline case ($\Delta = 0$) involves no cycles. Because of the model's simple shock structure (where it never takes the mean value, only the two extremes) I measure the amplitude of the cycle as the difference between the fraction that choose either water-intensive crop in wet periods and the fraction that choose either water-intensive crop in dry periods:

$$amplitude = x_H - x_L = s_{a1H} + s_{p1H} - s_{a1L} - s_{p1L}. \tag{2.48}$$

Table 2.1 shows the parameters used for the numerical simulations. Several parameters are more important, both for the results and for the idea as a whole, so for those I show results for different parameter values and how things change.

Parameter	Value
β	0.95
λ	0.1
m	1
$\bar{\xi}$	1.5
σ	1

Table 2.1: Assumed Parameters

In Table 2.2, I show results of the simulations for three different policy environments: 1) the free market, 2) with a tax exactly equal to the per-period externality imposed on others, and 3) with a uniform tax over time equal to the average externality. I specifically break out the third environment, as it is probably the type of tax that we are most likely to see in the real world. This type of tax is studied in Chapter 1. I break out the results by using different values of Δ, the difference in the attractiveness of the water-intensive crop between wet and dry periods.

Δ	Variable	Free market	Pigouvian tax	Avg. tax
0	cycle amplitude	0	0	0
	share permanent	.585	.600	.600
.1	cycle amplitude	.001	.009	.009
	share permanent	.584	.600	.599
.5	cycle amplitude	.049	.042	.046
	share permanent	.582	.598	.597
1	cycle amplitude	.098	.085	.093
	share permanent	.577	.594	.592

Table 2.2: Simulation Results: Three Regimes, Four Types of Shocks

Again, the key issue is the extent to which the policy environment affects the aforementioned cycles in the share that are water-intensive, as well as the propensity to choose the permanent crop options.

No matter the initial conditions, the cycle stabilizes after 10-20 periods (this mostly just depends on λ, the probability that permanents can switch next period, which I set to 0.1 in all simulations, and to some extent, on the period 0 distribution across choices, which I assume to be uniform at .25 each).

As might be expected, when Δ is larger, the amplitude of the cycles is larger as well. Note that for all $\Delta > 0$, the environment with the Pigouvian tax on the water-intensive crop leads to cycles with smaller amplitudes, ranging from 12.8% lower with low Δ to 14.7% lower with high Δ. As expected, cycles in the average tax are in between the free market and Pigouvian ones. The magnitudes of these relationships depends on parameters, but the directions do not. Share of permanence is relatively unaffected, but perhaps surprisingly is higher with the taxes. This shows the potential unintended consequences of taxing groundwater usage; it could lead to even more permanent crops and water demand that is less flexible.

However, the total permanence statistic masks a shift in the breakdown of water-intensive versus low-water in that permanence share. This is shown well in the figures in the appendix, which break out the four crop shares.

The pattern can also be seen in Table 2.3. First consider the free market case. Both permanent types have the same equilibrium share no matter what the cycle looks like. This is because, in equilibrium, they give the same return (for both permanent and annual, without considering idiosyncratic shocks).

When we add in a tax on the water-intensive crop, obviously low-water becomes much more attractive and the two permanent options no longer have equal shares. So despite taxation actually increasing the total permanent share, it significantly affects the distribution of that permanence over water-intensity. An interesting note is how similar that effect is with the Pigouvian and average tax regimes.

Δ	Variable	Free market	Pigouvian tax	Avg. tax
0	share permanent water-intensive	.292	.148	.148
	share permanent low-water	.292	.452	.452
.1	share permanent water-intensive	.292	.148	.147
	share permanent low-water	.292	.452	.452
.5	share permanent water-intensive	.291	.147	.147
	share permanent low-water	.291	.451	.451
1	share permanent water-intensive	.289	.146	.145
	share permanent low-water	.289	.448	.447

Table 2.3: Simulation Results: Permanence Distribution

2.5 Robustness

In this section, I discuss how sensitive the above results are to various assumptions.

One key choice is the functional form for the externality. In the analysis above, farmer utility for water-intensive crops decreases linearly with the fraction that also choose water-intensive. The slope doesn't affect the nature of the cycle, apart from making water-intensive crops less attractive in all periods. An alternative functional form would be exponential---say $f(x_k, \xi_k) = \xi_k + \exp(x_k)$. This simply introduces curvature in the size of the externality: one more farmer choosing water-intensive has more of an effect when fewer peers are doing the same than when more peers are. Again, results are not too different with this functional form.

What can make a bigger difference is if the aggregate shock affects the extent of the externality. Note that above I only consider cases where the aggregate shock simply shifts up or down the relative attractiveness of the water-intensive option, by changing Δ, but not the actual mechanism of the externality. This may be a relevant case empirically, as groundwater extractions are especially important in drought years, potentially increasing externalities.

2.6 Conclusion

This paper examines the question of how the choices of water-intensity and permanence interact in a setting with externalities. The model is general and theoretical, but I use the example of agriculture in California, which relies heavily on groundwater. Since rights to use groundwater are not well defined, modeling usage as having externalities on other farmers is a suitable assumption. I show how in the presence of a cyclical aggregate shock, meant to capture wet and dry periods, the equilibrium is characterized by a time-varying fraction of farmers who choose water-intensive crops. By adding in Pigouvian taxes to the model, I illustrate how the free market equilibrium not only has a socially excessive amount of water-intensity, but there is socially excessive volatility (12-15%) in the cycle driven by the aggregate shock. I also study other more feasible tax structures and show that the equilibria closely approximate the socially efficient equilibrium. Future work should look at stochastic aggregate shocks, allow for farmer heterogeneity, and estimate the model using suitable data from California. Endogenizing the size of the externality as well may affect the properties of the equilibrium cycles.

References

Arvin-Edison Water Storage District (2015): "Water Management Plan Update," Tech. rep., Arvin-Edison Water Storage District.

Ayres, A., E. Edwards, and G. Libecap (2017): "How Transaction Costs Obstruct Collective Action to Limit Common-Pool Losses: Evidence from California's Groundwater," NBER Working Paper No. 23382.

Belridge Water Storage District (2015): "Agricultural Water Management Program," Tech. rep., Belridge Water Storage District.

Berrenda Mesa Water District (2015): "Agricultural Water Management Program," Tech. rep., Berrenda Mesa Water District.

Brar, G. S., D. Doll, L. Ferguson, E. Fichtner, C. E. Kallsen, R. H. Beede, K. Klonsky, K. P. Tumber, N. Anderson, and D. Stewart (2015): "Sample costs to establish and produce pistachios," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

Buck, S., M. Auffhammer, and D. Sunding (2014): "Land markets and the value of water: Hedonic analysis using repeat sales of farmland," *American Journal of Agricultural Economics*, 96, 953–969.

Buena Vista Water Storage District (2016): "2015 Agricultural Water Management Plan for the Buena Vista Water Storage District," Tech. rep., Buena Vista Water Storage District.

Burness, S. and J. Quirk (1979): "Appropriative water rights and the efficient allocation of resources," *The American Economic Review*, 69, 163–179.

California State Water Resources Control Board (2016): "Groundwater Ambient Monitoring and Assessment: Domestic Well Project Groundwater Quality Data Report," Tech. rep., California State Water Resources Control Board.

Cawelo Water District (2014): "Cawelo Water District Agricultural Water Management Plan," Tech. rep., Cawelo Water District.

CDFA (2017): "California agricultural statistics review: 2016-2017," Tech. rep., California Department of Food and Agriculture.

Clark, N., C. A. Frate, D. A. Sumner, K. Klonsky, D. Stewart, and C. A. Gutierrez (2016): "Sample costs to establish and produce alfalfa," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

Cooley, H. (2015): "California agricultural water use: key background information," Tech. rep., Pacific Institute.

County of Kern (2016): "Kern County Department of Agriculture and Measurement Standards Spatial Data," .

Delano-Earlimart Irrigation District (2014): "Water Management Plan," Tech. rep., Delano-Earlimart Irrigation District.

Delhomme, J. (1978): "Kriging in the Hydrosciences," *Advances in Water Resources*, 1, 251–266.

Famiglietti, J. (2014): "The global groundwater crisis," *Nature Climate Change*, 4, 945–948.

Faunt, C. C., ed. (2009): *Groundwater Availability of the Central Valley Aquifer, California*, 1766, U.S. Geological Survey.

Franklin, B., K. C. Knapp, and K. A. Schwabe (2017): "A dynamic regional model of irrigated perennial crop production," *Water Economics and Policy*, 3.

Frate, C. A., K. M. Klonsky, and R. L. D. Moura (2013): "Sample costs to produce blackeye beans," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

Gisser, M. and D. Sanchez (1980): "Competition versus optimal control in groundwater pumping," *Water Resources Research*, 16, 638–642.

Goulder, L. and I. Parry (2008): "Instrument Choice in Environmental Policy," *Review of Environmental Economics and Policy*, 2, 152–174.

Griffin, D. and K. J. Anchukaitis (2014): "How unusual is the 2012–2014 California drought?" *Geophysical Research Letters*, 41, 9017–9023, 2014GL062433.

Guilfoos, T., A. Pape, N. Khanna, and K. Salvage (2013): "Groundwater management: The effect of water flows on welfare gains," *Ecological Economics*, 95, 31–40.

Hanak, E., J. Mount, and J. Lund (2016): "California's water," Tech. rep., Public Policy Institute of California.

Hardin, G. (1968): "The tragedy of the commons," *Science*, 162, 1243–1248.

Hendricks, N. P. and J. M. Peterson (2012): "Fixed Effects Estimation of the Intensive and Extensive Margins of Irrigation Water Demand," *Journal of Agricultural and Resource Economics*, 37, 1–19.

Holmes, T. and S. Lee (2012): "Economies of density versus natural advantage: Crop choice on the back forty." *Review of Economics and Statistics*, 94, 1–19.

Howitt, R., D. MacEwan, J. Medellin-Azuara, J. Lund, and D. Sumner (2015): "Economic analysis of the 2015 drought for California agriculture," Tech. rep., UC-Davis Center for Watershed Studies.

Hutmacher, R. B., S. D. Wright, L. Godfrey, D. S. Munk, B. H. Marsh, K. M. Klonsky, R. L. D. Moura, and K. P. Tumber (2012): "Sample costs to produce cotton," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

Johnson, R. and B. A. Cody (2015): "California agricultural production and irrigated water use," Tech. rep., Congressional Research Service.

Kern Delta Water District (2015): "Kern Delta Water District Agricultural Water Management Plan," Tech. rep., Kern Delta Water District.

Kern-Tulare Water District (2016): "Agricultural Water Management Plan," Tech. rep., Kern-Tulare Water District.

Kiparsky, M. (2016): "Unanswered questions for implementation of the Sustainable Groundwater Management Act," *California Agriculture*, 70.

Kiparsky, M., D. Owen, N. G. Nylen, J. Christian-Smith, B. Cosens, H. Doremus, A. Fisher, and A. Milman (2016): "Designing effective groundwater sustainability agencies: criteria for evaluation of local governance options," Tech. rep., UC Berkeley Center for Law, Energy and the Environment.

Koundouri, P. (2004): "Potential for groundwater management: Gisser-Sanchez effect reconsidered Authors Potential for groundwater management: Gisser-Sanchez effect reconsidered," *Water Resources Research*, 40.

Lost Hills Water District (2015): "Agricultural Water Management Program," Tech. rep., Lost Hills Water District.

Medellin-Azuara, J., D. MacEwan, R. E. Howitt, G. Koruakos, E. C. Dogrul, C. F. Brush, T. N. Kadir, T. Harter, F. Melton, and J. R. Lund (2015): "Hydro-economic analysis of groundwater pumping for irrigated agriculture in California's Central Valley, USA," *Hydrogeology Journal*, 23, 1205–1216.

North Kern Water Storage District (2015): "2015 Agricultural Water Management Plan," Tech. rep., North Kern Water Storage District.

Ostrom, E. (1990): *Governing the commons*, Cambridge University Press.

Ostrom, E., R. Gardner, and J. Walker (1994): *Rules, games, and common-pool resources*, University of Michigan Press.

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

Rosedale-Rio Bravo Water Storage District (2017): "2015 Operations Report," Tech. rep., Rosedale-Rio Bravo Water Storage District.

Sandholm, W. (2007): "Pigouvian pricing and stochastic evolutionary implementation," *Journal of Economic Theory*, 132, 367–382.

Schoengold, K., D. Sunding, and G. Moreno (2006): "Price elasticity reconsidered: panel estimation of an agricultural water demand function," *Water Resources Research*, 42.

Semitropic Water Storage District (2013): "Agricultural Water Management Plan," Tech. rep., Semitropic Water Storage District.

Shafter-Wasco Irrigation District (2013): "Water Management Plan," Tech. rep., Shafter-Wasco Irrigation District.

Smith, S. M., K. Andersson, K. C. Cody, M. Cox, and D. Ficklin (2017): "Responding to a Groundwater Crisis: The Effects of Self-Imposed Economic Incentives," *Journal of the Association of Environmental and Resource Economists*, 4, 985–1023.

Southern San Joaquin Municipal Utilities District (2012): "Five Year Update: Agricultural Water Management Plan," Tech. rep., Southern San Joaquin Municipal Utilities District.

Tsur, Y. and A. Zemel (1995): "Uncertainty and irreversibility in groundwater resource management," *Journal of Environmental Economics and Management*, 29, 149–161.

USDA (2012): "Census of Agriculture," Tech. rep., Department of Agriculture.

USDA (2017): *National Agricultural Statistics Service*.

Wheeler Ridge-Maricopa Water Storage District (2015): "Agricultural Water Management Plan," Tech. rep., Wheeler Ridge-Maricopa Water Storage District.

Williams, A. P., R. Seager, J. T. Abatzoglou, B. I. Cook, J. E. Smerdon, and E. R. Cook (2015): "Contribution of anthropogenic warming to California drought during 2012-2014," *Geophysical Research Letters*, 42.

Wright, S., K. Klonsky, and D. Stewart (2015): "Sample costs to produce field corn," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

WRIGHT, S. D., R. HUTMACHER, K. M. KLONSKY, AND R. L. D. MOURA (2013): "Sample costs to produce wheat for grain," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

YAGHMOUR, M., D. R. HAVILAND, E. J. FICHTNER, B. L. SANDEN, M. VIVEROS, D. A. SUMNER, D. E. STEWART, AND C. A. GUTIERREZ (2016): "Sample costs to establish an orchard and produce almonds," Tech. rep., University of California Cooperative Extension Agriculture and Natural Resources.

ZILBERMAN, D., N. MACDOUGALL, AND F. SHAH (1994): "Changes in water allocation mechanisms for California agriculture," *Contemporary Economic Policy*.

Appendix A

A.1 Decay Parameter

For the purposes of estimating the decay parameter above, it is necessary to consider municipal groundwater pumping as well as agricultural pumping. For that question, it is the pumping that matters, not the end use. Agriculture does use most of the water in Kern County, but the population is still relatively high (882,000 in 2015). The largest city, Bakersfield, accounts for about 400,000 of that. Most municipal water allocation is done through water districts as well (usually separate from their agricultural counterparts). Urban Water Management Plans are publicly submitted by these agencies, and have much more detail than the agricultural plans discussed above. There are 13 primary municipal water districts in Kern County. Some of them have access to surface water through state projects, but most rely at least partly on groundwater. Indeed, some of the districts serving smaller towns are entirely reliant on groundwater. This is another reason why dropping levels can be so costly. Groundwater pumping data since 2006 is available directly from district Urban Water Management Plans. The bigger limitation comes in knowing exactly where that pumping occurs. Shapefiles are available for district boundaries, but not for district wells. I make the same assumption as I do for groundwater pumping on agricultural fields: that it occurs at the centroid. This is less reasonable for the larger districts, but it does make sense that they would want to locate their wells as close to the centroid as possible, to minimize transportation costs.

I can then simply merge the municipal pumping data by location with the agricultural pumping data by location. This is what I use, along with the panel of groundwater levels, to estimate the decay in the regressions in Section 4.

In the bathtub model, pumping would be estimated to have the same effect from each of the circles, since distance would not matter.

As a first step, I estimate the following regression:

$$drop_{it} = \beta_1 w^g_{1t} + \beta_2 w^g_{2t} + \ldots + \beta_7 w^g_{7t} + \varepsilon_{it} \tag{A.1}$$

The dependent variable is the change in depth at a particular well from March to October. A drop in the groundwater level is coded as a negative. Each independent variable w^g_{it} is the total groundwater pumped in the ith closest region to the well. I report the results for various specifications, where the regions are defined differently (also see a graphical representation of the results below).

Table A.1: Pumping Effects: Alternative Specifications

	(1) drop	(2) drop	(3) drop	(4) drop	(5) drop
total	-0.0000213***				
	(0.00000126)				
$<$ 4 mi		-0.000205***		-0.0000921	
		(0.0000448)		(0.0000519)	
$>$ 4 mi		-0.0000155***			
		(0.00000189)			
$<$ 2 mi			0.000147		
			(0.000205)		
ring 2-8			-0.000122***		
			(0.0000215)		
$>$ 8 mi			-0.0000114***		
			(0.00000209)		
ring 4-16				-0.0000432***	
				(0.00000656)	
$>$ 16 mi				-0.00000827***	-0.00000890***
				(0.00000246)	(0.00000247)
$<$ .5 mi					-0.000361
					(0.00159)
ring .5-1					0.000238
					(0.000843)
ring 1-2					-0.0000616
					(0.000331)
ring 2-4					0.00000370
					(0.000109)
ring 4-8					-0.000117***
					(0.0000348)
ring 8-16					-0.0000275**
					(0.00000974)
Observations	2066	2066	2066	2066	2066

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

I report five different specifications in Table A.1. The first column includes one independent variable, total groundwater pumped in Kern County. This is essentially a sanity check, but serves as a baseline. Each acre-foot of groundwater pumped is associated with an average drop in the distance to groundwater of 0.0000235 feet. This translates to a drop of 0.1 inches with the pumping required for an average-sized field of almonds (373 acre-feet). This specification ignores any spatial heterogeneity in pumping or its effects.

The second column is the simplest way to check for evidence of spatial heterogeneity in the effects of groundwater pumping. This regression includes two independent variables. For each well, groundwater pumping in the county is divided between that which is extracted from within a four mile radius, and that which is extracted from outside the four mile radius. The estimated coefficients are significantly different. The typical field of almonds is now associated with an average drop of 0.93 inches if it is within four miles, but only 0.08 inches if not. In a qualitative sense, results are not very sensitive to the radius chosen.

The last three columns further break up the groundwater pumping that happens around a well. Breaking the data up this way weakens the statistical relationship. The point estimates are even positive in some cases. However, as we zoom out to the wider rings, the pattern is still there. The farther out that we move, the smaller the effect of pumping on the depth to groundwater. In all specifications, standard errors are clustered at the well-level.

Figure A.1 is a visual representation of the coefficients.

A.2 Surface Water Data

For the districts in Kern County, most submit Agricultural Water Management Plans. From here, I take information on surface water deliveries to farmers. What I do not observe is transactions between farmers in the same district. To the extent to which these are very common, I will be misattributing groundwater pumping to the wrong fields. This should not be a huge problem, as districts are generally relatively small.

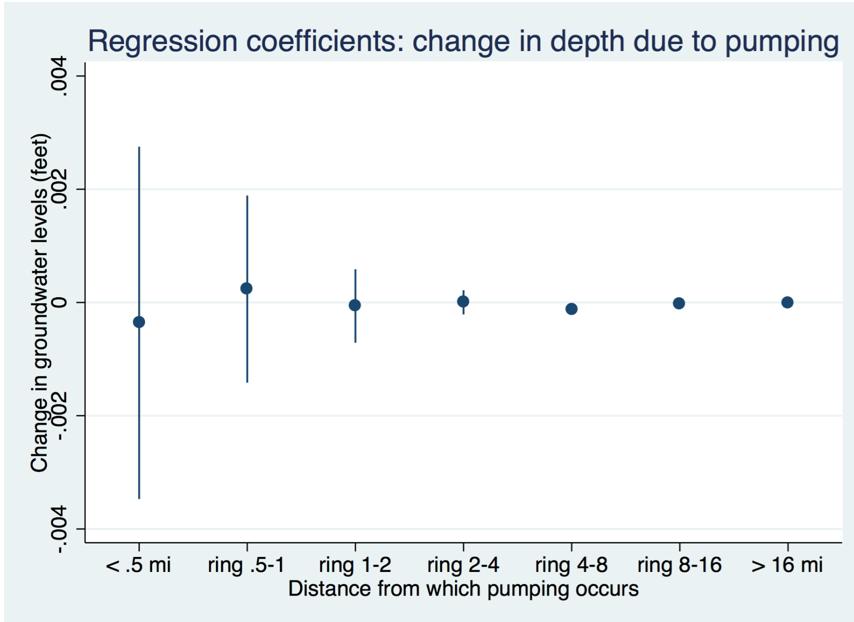

FIGURE A.1: Pumping Effects

A.3 Multinomial Logit

Here I explain how I estimate a multinomial logit model of crop choice, taking into account crop output prices, farmer soil quality, and crop input costs (including water needs).

There are several more data sources that I need in this section. First, soil data comes from SSURGO, a comprehensive national soils database built by the USDA. I use information on the irrigation capability class, suitable for measuring the quality of soils in the San Joaquin Valley, where irrigation is necessary for any significant productivity.

Crop input costs are taken from very detailed crop-specific studies done by University of California agricultural extension programs (Clark et al., 2016). While not done for all crops, they are available for the most prominent options, and are meant to be useful specifically to farmers considering growing these crops. From these studies, I include all costs except for anything related to water (the quantities necessary closely match what I find from water district plans and the Caltech irrigation studies) and land, as the goal is to only look at crop-specific non-water costs. These costs are not available at the yearly level.

Crop output price is from the USDA NASS, and is at the state level. Crop yields are

Table A.2: Multinomial Logit Results

	(1) chosen	(2) chosen	(3) chosen	(4) chosen	(5) chosen
Profits	-0.00960	-0.00660	-0.0124	-0.0105	-0.0129
	(0.0158)	(0.0160)	(0.0162)	(0.0163)	(0.0160)
z	-0.00130**	-0.00122**	-0.00164***	-0.00158***	
	(0.000493)	(0.000445)	(0.000363)	(0.000346)	
Observations	226158	225309	226158	225309	225309
Soil		Yes		Yes	Yes
Location			Yes	Yes	
Level					Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

also available from this source, but county output is published by Kern County, so I use the two to form my revenue estimate.

I also use the previously explained data on surface water allocations, crop-water needs, and groundwater levels to form z_{ijt}, which is meant to capture differing groundwater levels but also the fact that only some farmers actually need to use groundwater, depending on their surface water allocation. Annual crop farmers are assumed to have utility function

$$u_{ijt} = \alpha(p_{jt}f_{jt} - c_j) + \beta_1 z_{ijt} + \beta_2 s_i + \varepsilon_{ijt},$$

where $z_{ijt} = L_{it}(w_j - \bar{w}^s_{it})$.

The two important estimated parameters are α and β_1. The estimate for β_1 is more reasonable, as it is saying that farmers are less likely to choose crops for which groundwater costs are high. The estimate for α is not significantly different from zero. This is essentially saying that controlling for groundwater costs and soil, farmers are no likelier to choose more lucrative crops. There are several reasons that I could be finding this pattern. One unobservable characteristic is that high-value specialty crops might require specific capital (or human capital), and low-value crops such as alfalfa or corn are simply easier to grow. Another is that those provide certain benefits for other parts of the farm: this makes sense in Kern County, as much of the land planted with annual crops is owned by dairies that are growing cattle feed.

In any case, for the comparison to model simulations, I use the reduced-form estimates in the response section.

Appendix B

B.1 Distribution of Crop Choices

In this appendix, I show the distribution across crop choices over time, for a number of scenarios. I show results for three different levels of the aggregate shock, which represents how much worse a dry year is as compared to a wet year. For each shock, I break out results into the free market regime, the Pigouvian tax regime, and the average (or constant) tax regime. In the figures, water-intensive permanent is p1, low-water permanent is p0, water-intensive annual is a1, and low-water annual is a0. By assumption, in each of the nine scenarios I start with an equal distribution of each crop, and simulate out over 100 periods.

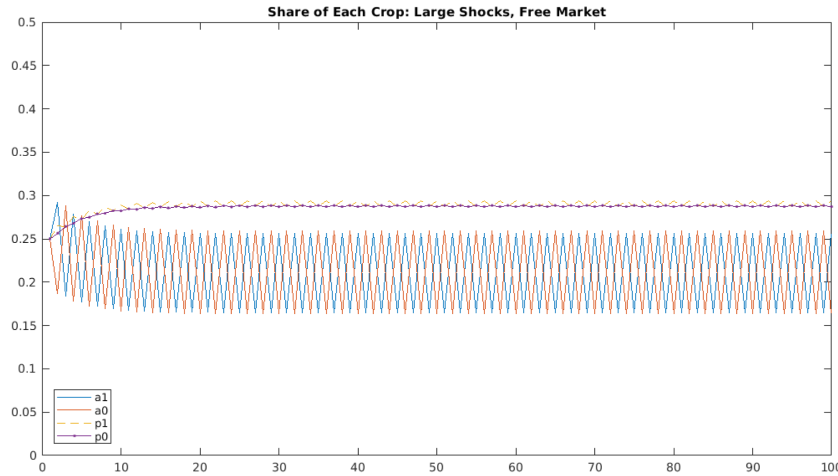

FIGURE B.1: Crop Shares: Large Shocks, Free Market

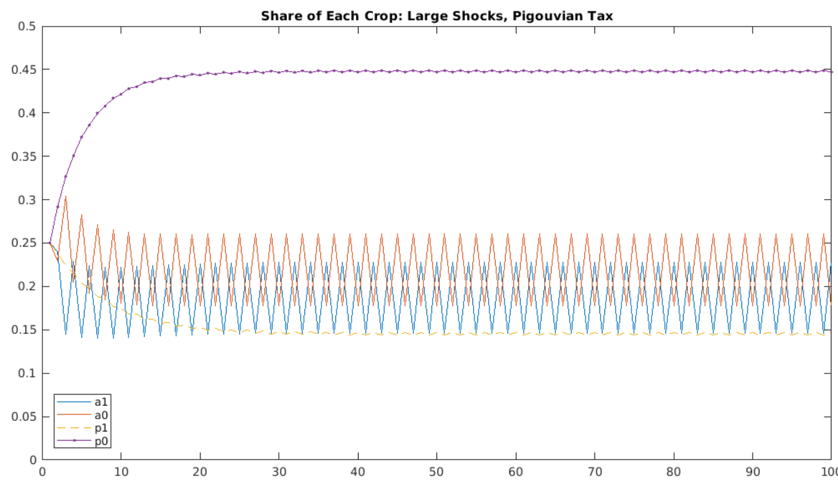

FIGURE B.2: Crop Shares: Large Shocks, Pigouvian Tax

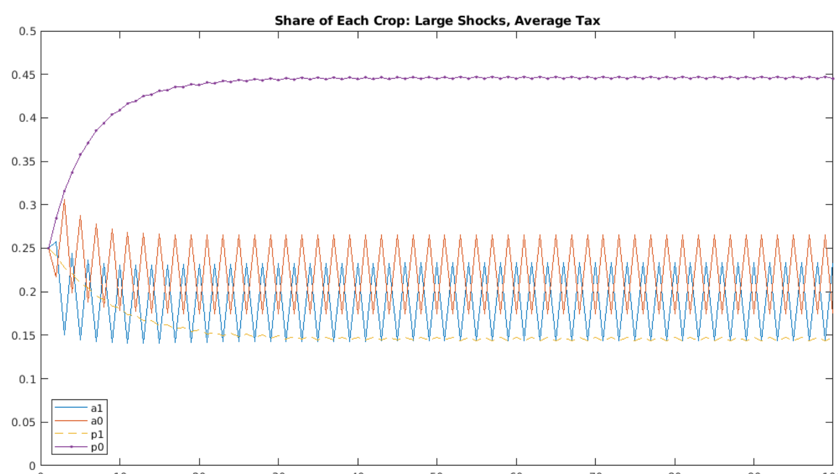

FIGURE B.3: Crop Shares: Large Shocks, Average Tax

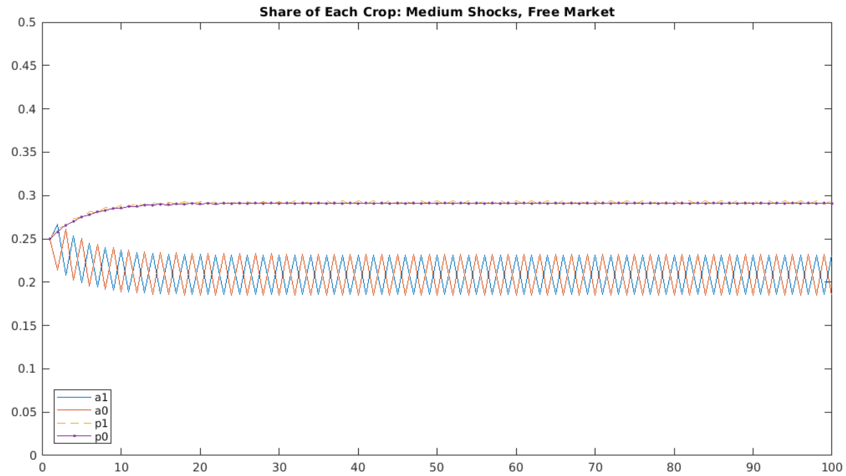

FIGURE B.4: Crop Shares: Medium Shocks, Free Market

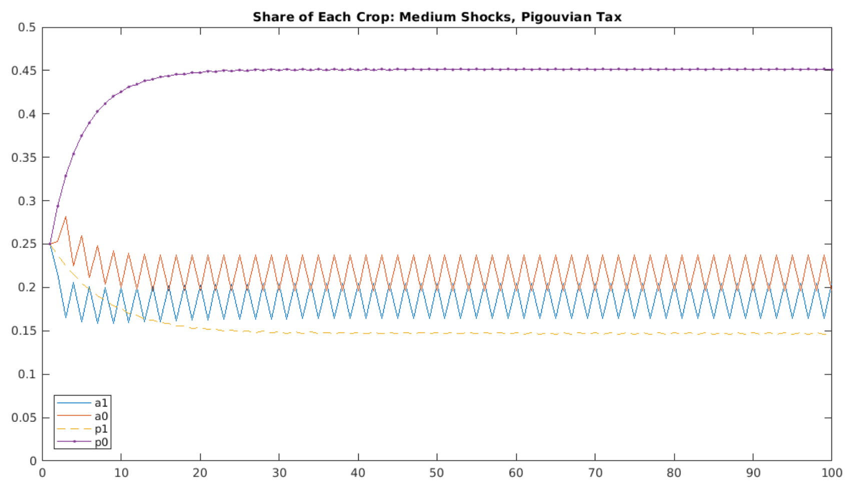

FIGURE B.5: Crop Shares: Medium Shocks, Pigouvian Tax

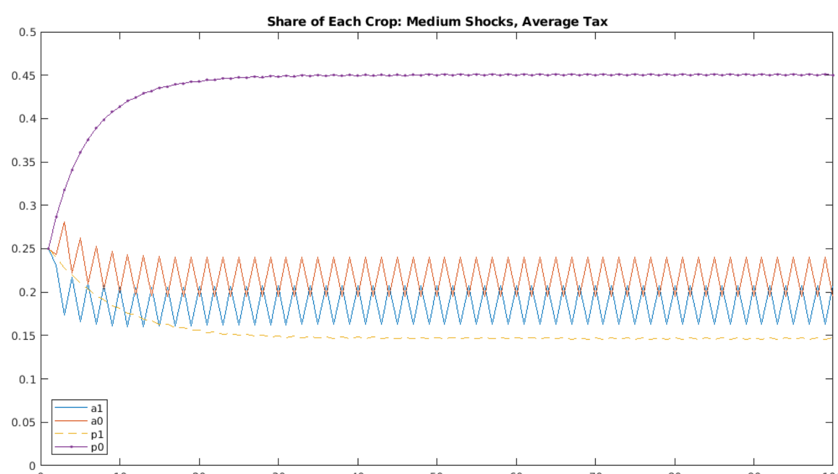

FIGURE B.6: Crop Shares: Medium Shocks, Average Tax

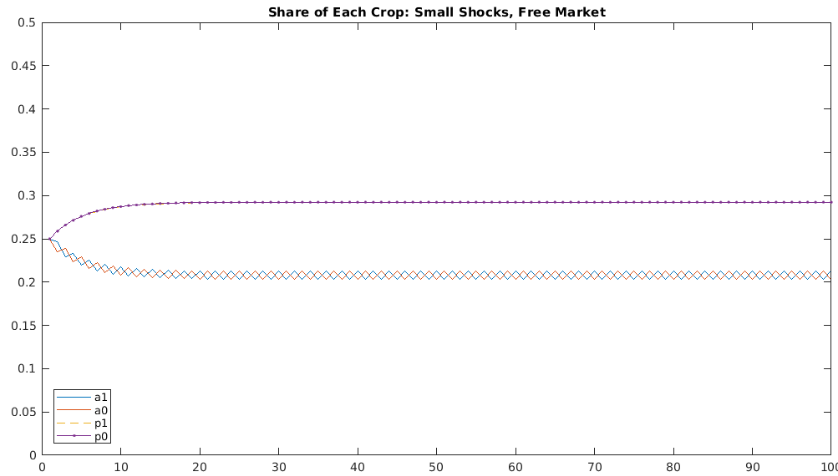

FIGURE B.7: Crop Shares: Small Shocks, Free Market

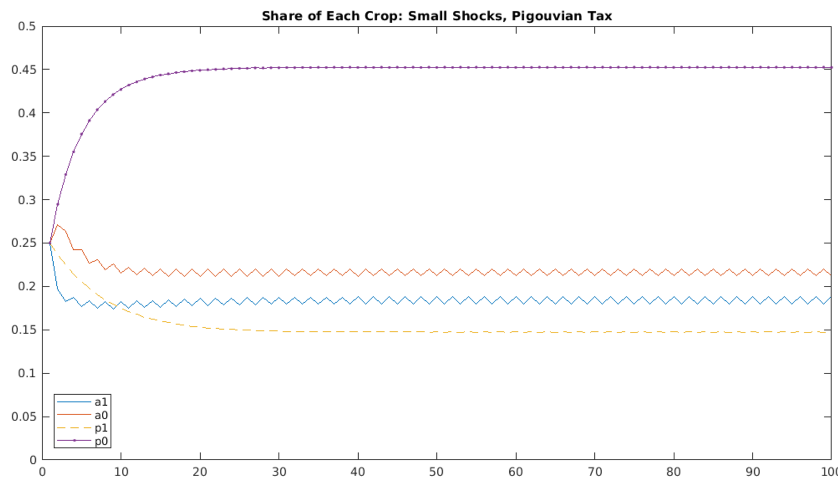

FIGURE B.8: Crop Shares: Small Shocks, Pigouvian Tax

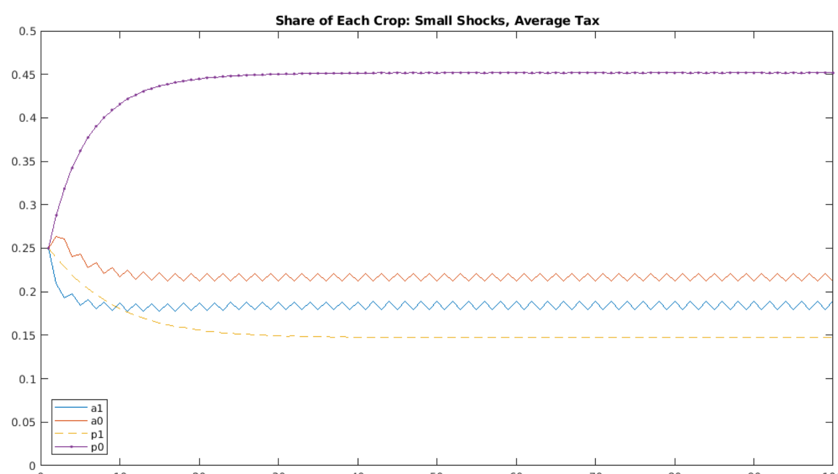

FIGURE B.9: Crop Shares: Small Shocks, Average Tax