Effects of Economic and Climate Factors on Central Arizona Agricultural Water Use

EFFECTS OF ECONOMIC AND CLIMATE FACTORS ON CENTRAL ARIZONA

AGRICULTURAL WATER USE

by

Brian McGreal

A Thesis Submitted to the Faculty of the

DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS

In Partial Fulfillment of the Requirements

For the Degree of

MASTER OF SCIENCE
WITH A MAJOR IN APPLIED ECONOMICS AND POLICY ANALYSIS

In the Graduate College

THE UNIVERSITY OF ARIZONA

2021

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Master's Committee, we certify that we have read the thesis prepared by: Brian McGreal
titled: Effects of Economic and Climatic Factors on Central Arizona Agricultural Water Use

and recommend that it be accepted as fulfilling the thesis requirement for the Master's Degree.

Bonnie G Colby
_______________________ Date: Dec 7, 2021
Bonnie Colby

George Frisvold
_______________________ Date: Dec 7, 2021
George Frisvold

_______________________ Date: Dec 7, 2021
Daniel Scheitrum

Final approval and acceptance of this thesis is contingent upon the candidate's submission of the final copies of the thesis to the Graduate College.

I hereby certify that I have read this thesis prepared under my direction and recommend that it be accepted as fulfilling the Master's requirement.

Bonnie G Colby
_______________________ Date: Dec 7, 2021
Bonnie Colby
Thesis Committee Chair
Agricultural & Resource Economics

Acknowledgements

This work would not have been possible without the help and support of my thesis advisor, Dr. Bonnie Colby, as well as the other members of my thesis committee, Dr. George Frisvold and Dr. Daniel Scheitrum. In addition to these committee members, various students and faculty at the University of Arizona's Department of Agricultural and Resource Economics have provided ideas and feedback invaluable to this study. These include Dr. Russ Tronstad, Dr. Satheesh Aradhyula, Dr. Jeffrey Michler, Dr. Anna Josephson, Dr. Gary Thompson, Dr. Na Zuo, and Dr. Tauhidur Rahman, as well as Danielle Buhrow, Ashley Kerna Bickel, and Hannah Hansen. This project was funded through a grant jointly issued by the National Science Foundation and the US Department of Agriculture. Much of the research presented in the coming pages was conducted alongside a team of students and faculty from the University of Arizona's hydrology department, including Dr. Laura Condon and students Danielle Tadych and Matthew Ford. Dr. Craig Wissler from the University of Arizona's Department of Natural Resources and the Environment taught me to use QGIS software for the purposes of incorporating remote sensed data into this study. Dr. Susanna Eden from the University of Arizona's Water Resources Research Center provided me with an in-depth education on the history of water management in Arizona. In addition to those mentioned above, a number of people from various state and federal institutions were kind enough to share their expertise with me during this research. These include Tyler Fitzgerald and Thomas Whipple from the Arizona Department of Water Resources, Ken Seasholes from the Central Arizona Project, Ron Klawitter from the Salt River Project, and Dave DeWalt and Patrick Willis from the USDA National Agricultural Statistics Service. Everyone mentioned above contributed greatly to the quality of this work, and all have my utmost appreciation and thanks. Finally, I'd like to thank Pat and Carol McGreal, Penelope Cottrell-Crawford, Caitlin McGreal, Jay Sherman, and Gigi. You guys are the best.

Table of Contents

List of Figures

List of Tables

Abstract

Throughout the arid American West, agriculture is the dominant consumptive use of water, with farming operations dependent on groundwater or surface flows for necessary irrigation. Although surface flows are limited in the region, widespread availability of groundwater has allowed agricultural economies to develop in otherwise dry areas. However, groundwater withdrawals have outpaced natural aquifer recharge throughout much of the twentieth and twenty-first centuries, and as water tables decline in elevation, surface flows are adversely affected.

This study seeks to model the climatic and economic factors that contribute to farmers' water use decisions in central Arizona, a region that has been historically dependent on groundwater to satisfy the water demands of agriculture, urban expansion, and heavy industry. Today, the area's water needs are met through a combination of groundwater, Colorado River water delivered via the Central Arizona Project, and additional surface flows. The modeling approach presented is applicable to a wide range of agricultural communities that are at least somewhat dependent on irrigation for agriculture. This study specifically examines the effect of climatic, economic, and remote sensed land cover variables on water deliveries to and irrigation intensity within irrigation districts in Central Arizona. The study's panel data set is enumerated the level of irrigation districts annually from 2008 to 2019, and incorporates remote sensed land cover data as well as a set of economic variables and climate measures. Econometric analysis finds that climate, the prices of December Cotton Futures, CAP water costs, and fallowed area have significant impacts on water deliveries to irrigation districts. It is also found that climate, the prices of December Cotton Futures, and CAP water costs significantly impact the intensity of irrigation water application (water applied/planted area). Understanding irrigators' water use decisions is useful to those concerned with the impact of water availability on local economies, ecosystems, and aquifers.

1 Introduction

The American West has a water problem. The vast majority of the region is arid or semi-arid, and water scarcity has driven the course of human habitation throughout its history. Irrigated agriculture in the desert southwest appears among indigenous cultures very early in the historical record (Bayman, 2001). The Mormons, the most successful early American settlers in the region, were able to thrive at the base of the Wasatch Mountains due to their implementation of complex community-based irrigation companies (Alexander, 2002). Los Angeles would never have become the sprawling metropolis it is today were it not for the controversial seizure of water rights in the Owens Valley and the construction of a 420 mile aquaduct to deliver Owens River water to the city (Reisner, 1986; Libecap, 2007). The 20^{th} century saw the construction of dozens of federally funded water reclamation projects throughout the region, designed to capture every last drop of rainfall or snowmelt in order to put this water to economic use. The most well known include the Hoover Dam and Glen Canyon Dam on the Colorado River, and the Grand Coulee Dam on the Columbia (US Bureau of Reclamation, 2021). The Colorado River, a vital lifeline in the desert Southwest, has been fully commoditized, its waters statutorily portioned between seven western states and Mexico (Bickel et al., 2019).

Today, a shifting climate is driving less predictable rain patterns, decreased mountain snowpack, and unprecedented summer temperatures throughout the region (MacDonald, 2010). The effects of changing weather patterns and widespread water withdrawals for irrigation have brought us to a point in time where "half of Western U.S. rivers have lost more than 50% of their summer flow (Richter et al., 2019)." In addition to this, since the end of World War II, the West has seen a dramatic increase in population. In the Southwest, the combined populations of the cities of Los Angeles, San Diego, Las Vegas, and Phoenix have increased from 2,409,037 in 1950 to 7,129,411 in 2010 (US Census Bureau, 2021). A direct consequence of these factors is the depletion of water stored at Lake Mead, the reservoir created by the Hoover Dam. In 2019, anticipating water shortages, Arizona, California, and Nevada adopted a Drought Contingency Plan which triggers cuts to their Colorado River water allotments based on water levels in the reservoir (York et al., 2020). The Bureau of Reclamation currently projects water levels in Lake Mead dropping to below an elevation of 1,075 feet in 2021, low enough to trigger the first round of cuts to states' water allotments (Central Arizona Project, 2021a). The bulk of these cuts will be to water intended for agricultural purposes (Ferris and Porter, 2021).

With all this in mind, an out-of-state visitor driving down I-10 from Phoenix to Tucson would perhaps be surprised to see pecan orchards and fields of cotton growing alongside the saguaro and arid mountain ranges of Central Arizona. In fact, since the late 1800s, for-profit irrigated agriculture has flourished in Arizona. The year round sunshine and warm temperatures make the state an ideal place to grow a wide variety of crops as long as water is available. Today, many communities throughout the state are economically dependent on irrigated agriculture. Much of the water for Central Arizona irrigation is sourced either from underground aquifers or from the Colorado River via the Central Arizona Project (CAP) canal network (Lahmers et al., 2018). This makes the notion of Colorado River water cuts especially alarming. In 2021, with the level of Lake Mead sitting at 1,066 feet above sea level, the U.S. Secretary of the Interior declared the first ever Tier 1 shortage along the Colorado. This shortage declaration means that CAP water deliveries in 2022 will be reduced by around 30% (Central Arizona Project, 2021a). At greatest risk of economic loss due to these cuts are agriculturally oriented communities in Pinal and Maricopa Counties (Bickel et al., 2019). Rather than face diminished profits, growers in these areas may choose to offset the reduction in Colorado River water with increased groundwater extraction. While this might maintain local economies in the short run, groundwater supplies do not replenish themselves quickly, and aquifer depletion is known to drastically restructure the hydrologic landscape (Condon and Maxwell, 2019). Increased groundwater withdrawals would at best be a stop-gap solution for these agricultural communities.

This study is concerned with addressing drivers of agricultural water use in Central Arizona. A 2014 Arizona Department of Water Resources (ADWR) report described agricultural water use as accounting for 68% of water use in the state. In 2017, agriculture contributed $23.3 billion to Arizona's economy (Lahmers et al., 2018). As water supplies become more and more scarce, the question of what motivates irrigators' water use decisions becomes increasingly salient. This analysis considers the impact of economic factors, climate, alternative land cover, and regulatory policy on water use. Identifying what motivates increased or decreased agricultural water use has the potential to aide policy makers and water managers in anticipating growers' water needs. This sort of analysis is especially relevant given the looming cuts to Arizona's allotment of Colorado River water.

This study begins with a short history of agriculture in Arizona and brief descriptions of the study area, data used, analysis methods, and findings. Next comes a review of existing economic literature concerned with agricultural water demand. This is followed by the presentation of a set of conceptual models based on the literature designed to inform a later econometric exploration of agricultural water use in Central Arizona. After this comes a discussion of the data and methods employed in the study's econometric analysis, before going on to present findings. The study then concludes by considering the implications of these findings.

1.1 A Short History of Irrigation and Groundwater Management in Arizona

Irrigated agriculture is a fundamental feature of human habitation in the Southwestern United States, and irrigation infrastructure has likely existed in Arizona for over 3,000 years.

Archaeological evidence suggests that indigenous peoples maintained irrigation canals along the Santa Cruz River as early as 1200 BCE in order to grow corn, tobacco, and squash. Between 300 BCE and 1450 CE, the Hohokam constructed complex irrigation networks along the Salt and Gila Rivers in and around what is now the City of Phoenix. By the 1800s, American settlers were expanding on the original Hohokam canals to irrigate fruit trees, alfalfa, and grain in the Salt River Valley. Federally funded dam projects, such as the Roosevelt Dam on the Salt River and the Hoover Dam on the Colorado, would follow in the early 20th century. Around this same time, the first irrigation districts began operation in the state, with the Yuma Auxillary Project beginning water deliveries to member farmers in 1905. In 1917, the Salt River Project (SRP) in Central Arizona would follow suit and begin delivering around a million acre-feet of water to subscribers each year (Lahmers et al., 2018).

Throughout this period, as sophisticated canal systems and massive dam projects captured an ever increasing share of Arizona's scarce surface flows in order to apply them to economic uses, a huge amount of freshwater sat buried underground in essentially untapped aquifers. This groundwater supply, the product of millions of years of natural accumulation, remained almost entirely intact, as the technology necessary to lift quantities capable of sustaining farm level agriculture did not yet exist (Silber-Coats et al., 2017). The 1937 invention of the high-speed centrifugal turbine pump changed this, allowing groundwater dependent agriculture to expand to areas removed from existing canals or streams. As a result, the region saw "...a major boom in irrigated farming [which] drove a tripling of groundwater extraction rates" (McGreal et al., 2021). Irrigation districts were established in areas that previously had no access to surface flows, with water supply portfolios almost entirely built on groundwater (Lahmers et al., 2018). All of this led to a major decline in groundwater levels in many parts of the state, as groundwater was put to use supporting not only agriculture, but also rapidly increasing municipal needs. The year 1968 saw withdrawals from Central

Arizona aquifers outpace natural recharge by an estimated 2.5 million acre-feet (McGreal et al., 2021).

By the 1970s, lawmakers recognized that, absent some statewide management scheme, this vital resource would be at risk of total depletion. 1980 saw the passage of the Arizona Groundwater Management Act, a package of reforms that established ADWR and introduced a multi-tiered groundwater management scheme throughout the state. The areas of greatest concern were designated as Active Management Areas (AMAs) and were subject to the most stringent regulations contained in the act. Initially, AMAs were established in and around the cities of Phoenix, Tucson, and Prescott, as well as in the largely rural but heavily agricultural Pinal County. Most AMAs (with the exception of Pinal) were tasked with the goal of achieving "safe-yield", a balance between groundwater extraction and replenishment, by 2025. These areas contain 80% of Arizona's population, as well as approximately half of the state's irrigated farmland (McGreal et al., 2021). Within the AMAs, the legislation restricted the expansion of irrigated agriculture and required measurement and reporting of groundwater pumping for non-domestic wells. All existing water rights would be fixed at this time, with groundwater rights based on pumping in the five years preceding the Groundwater Management Act's passage. These grandfathered rights to pump water are fairly generous, as groundwater use had been at an all time high prior to the act's passage. Crucially though, the act ensured that, from this point on, new water rights would not be granted within the AMAs, meaning that in these areas there exists a legally mandated upper-bound on groundwater extraction (Megdal, 2012). Additionally, in the years since the act's passage, Arizona's groundwater code has been amended to allow holders of these grandfathered water rights to accumulate "flex credits" by not withdrawing their full allotment of groundwater in a given calendar year. These credits may then be conveyed or sold to another grandfathered water right holder within the same AMA, with some restrictions (Arizona Revised Statutes, 2016).

Around this same time, another potential lifeline for Arizona's imperiled aquifers was in the offing. Approved by the US Congress as part of the Colorado River Basin Project Act of 1968, the Central Arizona Project (CAP) was intended to deliver water from the Colorado River to the population centers of Central Arizona via a system of pumping stations and canals. CAP would allow the AMAs more flexibility in the effort to achieve "safe-yield" conditions, as water users could now choose to forego pumping any amount of their groundwater allotments in exchange for CAP water (McGreal et al., 2021). Administered by the US Bureau of Reclamation, the project would begin water deliveries in the late 1980s, with delivery costs structured only to cover expenses (Central Arizona Project, 2016). In the years to come, the introduction of water banking schemes (the Underground Water Storage and Recovery Act of 1986 and the Underground Water Storage and Replenishment Act of 1994) would assure that any disused CAP water could be stored in aquifers until some later date (Colby, 2016).

As a result of the 1980 Groundwater Management Act, irrigation districts in Central Arizona must annually report all water received and delivered to ADWR. It is this requirement which makes this study of irrigators' water use decision making possible. Today, irrigation districts' water portfolios are largely composed of some mixture of groundwater, CAP deliveries, "in-lieu" water (referring to CAP water delivered in exchange for growers limiting their groundwater use), and surface flows (Fleck, 2013). A more detailed discussion of specific irrigation districts in my study area follows in the next section.

1.2 Study Area

This study is focused on irrigation districts located in the Phoenix and Pinal AMAs, with twelve districts included in this works' empirical analysis. These represent the largest and most significant districts in each AMA, both in terms of planted area and water deliveries. Figure 1 provides an illustration of the differences in scale and spatial distribution of the irrigation districts in the study area.

Irrigation districts are responsible for delivering water to subscribers within their boundaries. However, these districts are each unique in their approach to fulfilling this responsibility. Figure 1 clearly shows large differences in scale between irrigation districts, but they are also dissimilar in their policy structure, water sourcing, and in the composition of their end users (i.e., agricultural, municipal, industrial, et cetera). This being said, there are a few important common characteristics present in the districts chosen for this analysis. Firstly, as shown in Figure 1, each district in this analysis is located within an AMA. This means each of these districts must annually report water sourcing and deliveries to ADWR, which in turn makes these reports available to the public. This reporting requirement is (happily) what makes this analysis possible, and (unhappily) what disallows the inclusion of any irrigation districts outside these AMA boundaries. Secondly, each of these districts are constrained in their groundwater use according to water rights fixed at the time of the 1980 Groundwater Management Act. Even when groundwater is the cheapest source of water available to districts, this constraining factor means that districts' water supply portfolios are often composed of water from multiple sources. Finally, each of these districts has access to Colorado River water delivered by CAP. This allows districts latitude in their water sourcing decisions.

The Pinal AMA is home to only four irrigation districts: Central Arizona Irrigation and Drainage District, Hohokam Irrigation District, Maricopa-Stanfield Irrigation and Drainage District, and San Carlos Irrigation and Drainage District. These districts manage 87% of water deliveries in the county, with water use in excess of 800,000 acre-feet per year (Lahmers et al., 2018).

The San Carlos Irrigation and Drainage District structurally differs from the other districts within the Pinal AMA, as its primary function is to deliver Gila River water stored in San Carlos Lake. Water storage in the reservoir has declined drastically in recent years, leaving many irrigators in the district faced with a choice between attempting to grow crops in the face of uncertain water availability, or accepting Catostrophic Risk Protection (CAT) payments due to failure of irrigation supply before ever planting crops (United States Geological Survey, National Water Dashboard, 2021; Bickel, 2021). CAT coverage is available to growers at subsidised rates through the Federal Crop Insurance Reform Act of 1994 (Sall and Tronstad, 2021). While groundwater and CAP deliveries are available to the district, growers receiving water from the San Carlos Irrigation and Drainage District may have their planting decisions affected by this limited water supply and opportunity to receive CAT payments (Bickel, 2021). The theoretical and empirical models presented in Chapters 3 and 5 of this work assume constraints on water deliveries are not binding. This assumption can be held within the arid study area due to CAP's ability to deliver Colorado River water to Central Arizona irrigation districts, as well as the system of flex credits and generous groundwater rights established by the 1980 Groundwater Management Act. Due to CAT coverage allowing farmers to opt out of planting in the face of declining reservoir storage, the San Carlos Irrigation and Drainage District is in essence constrained by water supply, even though alternate sources of irrigation water are available. As such, it is omitted from this work. The other three Pinal County districts are included in this analysis, although this is not without some complication which is discussed further in Chapter 4. Even with the San Carlos Irrigation and Drainage District omitted, the remaining districts account for over 87% of the total planted area within the Pinal AMA.

In sharp contrast to the Pinal AMA, the Phoenix AMA contains thirty-nine irrigation districts, ranging in size from small owner-operated cooperatives to the massive Salt River Project (Arizona Department of Water Resources, 2020). This analysis focuses on the nine largest irrigation districts by average annual water deliveries. Taken together, these nine districts include over 95% of the total planted area in all irrigation districts within the Phoenix AMA. In order of size, these are: Salt River Project, Roosevelt Water Conservation District, Roosevelt Irrigation District, Maricopa Water District, New Magma Irrigation and Drainage District, Buckeye Water Conservation and Drainage District, Queen Creek Irrigation District, Arlington Canal Company, and the Tonopah Irrigation District. In 2019, these nine districts delivered almost 925,000 acre-feet of water to end users, with the Salt River Project alone having delivered over 550,000 acre-feet. It must be noted that much of this water goes to serving municipal needs as opposed to agricultural uses, as many districts established early in the 20$^{\text{th}}$ century on what had then been rural land today find themselves subsumed by the Phoenix metro area.

Planting in Central Arizona is largely centered around two crops: alfalfa hay and cotton. Figure 2 presents annual planted acreage in different common crops in Maricopa and Pinal Counties from 2008 through 2019. It is easy to see at a glance that alfalfa and cotton dominate Central Arizona cropping. Alfalfa is consistently the most copiously planted crop throughout the two county area, and only in 2011 does cotton acreage come close to matching alfalfa acreage. Alfalfa is considered a fairly water intensive crop for the region, with consumptive water needs averaging around 6 acre-feet per acre per year (Erie et al., 1982). Owing to the fact the fact that some irrigation water applied will be lost to evaporation and evapotranspiration, this means that growers need to apply more than that average consumptive minimum to bring an alfalfa crop to harvest. How much more depends largely on weather and irrigation efficiency. It is worth mentioning that, unlike cotton, alfalfa crops are not planted annually. Instead, a stand of alfalfa typically lasts from 3 to 7 years, with growers able to take multiple cuttings in a season. This means that alfalfa stands must be irrigated year round to maintain the health of the crop. Despite this perennial character, alfalfa stands do not represent the same sort of structurally fixed investment as tree crops. If a grower had a strong enough reason to plant something different, an alfalfa stand could be removed without accruing significant financial losses.

Cotton is consistently the second most planted crop in the study area, with the exception of 2015, when durum wheat acreage exceeded cotton by about 13,000 acres. Cotton's consumptive water needs average around 3.5 acre-feet per acre per year, significantly less than alfalfa's (Erie et al., 1982). It should be mentioned that, unlike alfalfa, cotton is not grown year round, leaving the possibility for cotton to be rotated with other crops in the same growing season. Cotton is rotated most frequently with wheat and other small grains, with wheat crops in the region estimated to consume about 2 acre-feet of water per acre per year (Ottman, 2015; Erie et al., 1982). If these crops were planted on the same field in the same year, around 5.5 acre-feet of water per acre per year would need to be consumed for both harvests to be healthy. Assuming that many growers who decide to plant cotton will also plant wheat in the same field after their cotton crop is harvested, then this consumptive water use value may be compared to the year-round consumptive needs of alfalfa stands, meaning the two crops' annual water needs are fairly similar to one another.

Along with alfalfa, cotton, and durum wheat, corn and barley make up the five most common crops in the region. Figure 3, showing the percentage of overall planted area in each crop, reinforces the prevalence of alfalfa and cotton in Maricopa and Pinal Counties. In 2008, alfalfa and cotton taken together represent 66.5% of planted acreage in the study area. This is the only year these two crops together do not compose over two-thirds of planting in the region, and in seven out of twelve of the years presented these crops account for over 75% of planting. Another interesting trend notable in Figure 3 is cropping patterns moving away from the previously mentioned "top five" crops. All other crops make up more than 10% of planting in each year from 2016 onward.

One of these other crops worth mentioning briefly are pecan trees. Pecan orchards, and to a lesser extent pistachio orchards, have been planted with increasing frequency in recent years (Duval et al., 2019). These orchards are very different in character than annual crops, as they represent a substantial upfront investment to plant, followed by a period of many years before the trees will begin to bear fruit. Additionally, orchards require water to be applied year round, and some minimum amount of water to be applied to maintain the health of an orchard, even in a season where a grower might not intend to bring their trees to harvest. Mature pecans intended to bear fruit generally consume around 11.5 acre-feet of water per acre per year inches of water throughout the year, making them a very water intensive crop in an arid region (Sammis and Herrera, 1999). For these reasons, tree crops will be treated as a category of interest in this research.

Water Supply Considerations

As alluded to at the beginning of this section, water delivered to members by irrigation districts comes from a variety of sources. The districts discussed in this study all have some level of access to Colorado River water deliveries through CAP, as well as groundwater, the opportunity to participate in "in-lieu" water sourcing, and in some cases surface flows as well.

Prior to the passage of the 1980 Arizona Groundwater Management Act and the completion of the Central Arizona Project, groundwater was the most widely used water source in much of Central Arizona. While a few of the oldest districts in the state were established to deliver surface flows to their members, the relative ease with which groundwater could be pumped by the mid-20th Century meant that irrigated agriculture could expand to areas of the state far removed from rivers and streams (Lahmers et al., 2018; Silber-Coats et al., 2017). By the time the Groundwater Management Act was enacted, groundwater use in Central Arizona was at an all time high. Quantification of grandfathered water rights within the AMAs was based on pumping in the years between 1975 and 1980, meaning the amount of water allocated to each right is quite generous (McGreal et al., 2021). The later implementation of the "flex credits" system described in the previous section provided further flexibility in the use of these grandfathered water rights (Arizona Revised Statutes, 2016). Taken together, these factors lead to legal constraints on groundwater supplies not often binding groundwater use. However, physical access to groundwater may be constrained in some cases by deteriorated well infrastructure (Seasholes, 2021).

Given generous groundwater allotments which may be further supplemented through the purchase of flex credits, one might reasonably ask why well infrastructure in some cases might be deteriorated to the point of constraining groundwater use. The answer lies in the opportunity CAP provides growers to substitute groundwater pumping for surface water delivered from the Colorado River. With the introduction of CAP, irrigation districts (along well as other entities within the delivery) were given the option to purchase Colorado River water directly from the Bureau of Reclamation at cost (Central Arizona Project, 2016). In addition to purchasing water directly, grandfathered water rights holders are also able to enter into "in-lieu" contracts. Facilitated by ADWR as part of their on-going efforts to manage groundwater levels within the AMAS, "in-lieu" contracts allow water rights holders to receive some additional amount of Colorado River water delivered by CAP in exchange for reducing their permitted groundwater pumping by the same extent (McGreal et al., 2021). Well infrastructure has deteriorated in some cases due to rights holders taking advantage of these "in-lieu" contracts rather than physically withdrawing their groundwater allotments.

In addition to the water sourcing options described above, some districts, such as the Salt River Project, the Roosevelt Water Conservation Project, and the San Carlos Irrigation and Drainage District, are able to provide surface water flowers captured in reservoirs to their members throughout the season (Klawitter, 2021; Bickel, 2021; Fleck, 2013) It is important to note that most districts' water portfolios are composed of some combination of these sources, as opposed to meeting all of their members' water needs through a single water source (Fleck, 2013).

1.3 Methods and Data

This study involves data on water deliveries and key economic variables (such as crop prices and input costs) from 2008 through 2019. Water delivery data is sourced from the reports described above filed by irrigation districts with ADWR. These reports break down water deliveries by source type and category of end use. This distinction between end uses is what makes it possible for this analysis to specifically focus on water intended for irrigated agriculture. Crop price and input cost data are sourced from a number of different federal agencies, including the United States Department of Agriculture's National Agricultural Statistics Service (USDA NASS), the United States Department of the Interior's Bureau of Reclamation (USBR), as well as privately owned agencies such as the New York Cotton Exchange. Land cover data is sourced from USDA NASS's Cropland Data Layer (CDL), a satellite remote sensing product available to the public. The CDL contains satellite recorded information on land cover throughout the entire continental United States from 2008 onward. While CDL has data in some regions for years prior to 2008, Arizona was only added to the product in that year. As a result, the study period is restricted to years beginning in 2008. Climate data is sourced from the Global SPEI Database, a project based in Spain which provides Standardized Precipitation-Evapotranspiration Index data for any region in the world. The SPEI provides a measure of rainfall and temperature, with values ranging from -3 (hotter and drier than average) to 3 (cooler and wetter than average). These data, their sources, and steps taken to compose the data set used in this study will be discussed in much greater detail in Chapter 4.

This study assumes non-binding constraints on water availability and employs irrigation district level fixed effect regressions with robust standard errors to estimate the effects of various explanatory variables on water deliveries and irrigation intensity. The use of fixed effect regressions is necessitated by the heterogeneous nature of irrigation districts' water sourcing and policy structure. In order to preserve a high number of degrees of freedom, fixed effects are captured by "de-meaning" variables at the irrigation district level, rather than including a set of district level dummy variables. Exact model specifications and the decisions that lead to these specifications will be discussed at length in Chapter 5. Results will be evaluated for statistical significance beginning with those that can be said to be non-zero values with 90% confidence.

The econometric analysis presented in Chapter 5 finds that climate, cotton futures prices, CAP water costs and fallowed lands each have a statistically significant effect on water deliveries to Central Arizona agriculture. Meanwhile, irrigation intensity is significantly affected by climate, cotton futures prices, and CAP water costs. The marginal effects observed in both models for cotton futures prices and CAP water costs and for fallowed lands in the water deliveries model fall in line with economic intuition, while the marginal effects of climate runs contrary to expectations. These results are discussed extensively in Chapters 5 & 6.

2 Literature Review

This study is underwritten by a thorough review of recent economic literature regarding motivating factors behind irrigators' water withdrawals. Broadly speaking, the literature has informed this research in three ways. The first comes from background and insight into the specific history of irrigated agriculture and regulatory policy structure in Central Arizona. This background information has allowed for the sound application of broad economic principles (derived from studies performed in other parts of the world) to this specific study area. The second way the literature informs this research is in supporting the construction of a conceptual model of water demand. This model, described in detail in the next chapter, is built based on drivers of water demand observed in numerous instances throughout the literature. This grounding is important, as the model in Chapter 3 will be used to form the foundation of the econometric analysis presented in Chapters 5. The third and final way the literature contributes to this research is by providing examples of econometric analyses that further support the econometric analysis in this thesis.

This chapter discusses those studies which have directly contributed to this work. The literature is discussed in the order described in the previous paragraph: 1) those works providing background on irrigated agriculture in Central Arizona, 2) research which contributes to the conceptual model presented in Chapter 3, and 3) research which further informs the econometric analysis presented in Chapters 5. Often the same article falls into more than one of these groupings. When this is the case, the paper is discussed in parts, meaning the article is mentioned more than once in this chapter.

2.1 Understanding Irrigation in Central Arizona

Literature which has studied irrigation in Central Arizona is particularly valuable to this analysis. While some general economic principles are universally applicable, a sturdy understanding of the unique characteristics of the study area in question is a necessary component in any research. As briefly described in Chapter 1, Central Arizona's regulatory policy and water infrastructure are complex and elaborate. Both state and federal agencies are involved in water supply management. The 1980 Groundwater Management Act introduced a multi-tiered system of regulatory policies. The region also has policies allowing for water banking programs and the ability for in-lieu water exchanges, along with other idiosyncrasies not relevant to this research but present nonetheless. Prior academic research, both qualitative and quantitative, which is based in this study area has proven to be an invaluable asset.

Bickel et al. (2019) examine the economic impact reductions in water allocations would have on agricultural communities in Pinal County, Arizona, describing the water challenges currently facing the study area. These challenges include an increasing demand for water caused by population growth, a limited and declining supply caused by the initial 1922 over-allocation of Colorado River water in the southwest, declining water levels throughout the Colorado Basin driven by climate change, and resistance to the expansion of retention infrastructure due to concerns over its environmental impact. The study also details the structure of Pinal County's agricultural sector and distribution of water use, as well as describing the typical climate in the county. Bickel et al. (2019) also provide valuable insight into drivers of irrigators' profit-maximizing decision making, and so will be discussed further in the next section of this chapter.

Colby (2016) provides a broad overview of water banking programs throughout the United States and elsewhere in the world. While water banking is not specifically considered in any conceptual or econometric models presented in this study, background on the Arizona Water Banking Authority (AWBA), Arizona's public water banking agency, is useful in establishing a complete picture of Central Arizona's water management structure. The AWBA was created in order to fully utilize Arizona's annual Colorado River water allotment of 2.8 million acre-feet. Banking water in Arizona's aquifers assures supply reliability for

Indian Water Rights, helps to satisfy Arizona's groundwater use regulations, and even allows for interstate water banking. The AWBA purchases excess CAP water or treated effluent and transfers storage credits to ADWR or the Central Arizona Groundwater Replenishment District (CAGRD). These entities are able to use these credits to recover stored water in order to meet supply obligations or satisfy groundwater regulatory requirements at a later date. Colby also makes mention of the problem of spatial mismatch between groundwater storage sites and recovery wells, an issue which is likewise considered by Silber-Coats et al. (2017) and may be of particular interest to researchers concerned with localized shifts in groundwater elevation.

Fleck (2013) is a primary source of inspiration and information for this research. Fleck examines water use by many of the same irrigation districts as this analysis. Fleck's background on individual irrigation districts is extremely thorough and detailed, although additional irrigation districts are included in this study. While this means that the background Fleck provides on irrigation districts cannot be relied on exclusively, it is still a valuable source of information. Fleck's thesis is somewhat more limited in scale than this analysis, as his study contains data from ten irrigation districts between 1995 and 2011. Fleck also does not use remote sensed crop coverage data, which allows this work to expand on the framework he establishes, rather than just recapitulating his analysis with a few extra irrigation districts and more recent data. Like Bickel et al. (2019), Fleck's study includes conceptual and econometric elements which have also informed this study. These are discussed in greater detail later in this chapter.

Frisvold and Konyar (2012) compare multiple models' predictions of the effect of a reduction in irrigation water supplies in six southwestern US states. The study area includes California, Nevada, Arizona, Utah, New Mexico, and Colorado, states in which irrigation accounts for 82% of all water withdrawals and which collectively are facing declining flows in the Colorado River. The modeling approach involves a sophisticated nonlinear mathematical programming model which will not inform any part of this paper. However, this study does provide valuable background in terms of the large-scale water supply problems facing the southwestern US.

The University of Arizona's Water Resources Research Center publishes *The Arroyo*, an annual review of some particular facet of water management in Arizona. These publications are an extremely valuable source of background information regarding specific regulatory policy and management practices, both throughout the state and in Central Arizona. Three specific issues have informed this research. *Arroyo 2017* concerns water banking, groundwater recharge, and the recovery of water stored in Arizona's aquifers (Silber-Coats et al., 2017). *Arroyo 2018* takes a close look at irrigated agriculture in Central Arizona and elsewhere (Lahmers et al., 2018). Finally, *Arroyo 2021* discusses the history of the 1980 Arizona Groundwater Management Act and how state policymakers and water managers are preparing to meet Arizonans' water needs in the future (McGreal et al., 2021).

Arroyo 2017 provides an in-depth look at water banking in Arizona. This study supports and develops ideas mentioned in Colby (2016). The study touches on the 1980 Arizona Groundwater Management Act, and gives detailed descriptions of the CAGRD and the AWBA. Arizona's system of Active Management Areas and Irrigation Non-expansion Areas is also carefully explained, a topic of particular relevance to this study, since the effects of AMAs on irrigators' water use decisions are considered (Silber-Coats et al., 2017).

Arroyo 2018 examines irrigated agriculture throughout Arizona, with large sections devoted to the structure of agriculture in the central part of the state. The paper details the history of indigenous agriculture, the formation of the first irrigation districts (some of which even precede Arizona's statehood), and the economic importance of agriculture in Arizona today. The authors also describe how irrigated agriculture functions specifically in Central Arizona, where most irrigation districts are subject to restrictions imposed by the Groundwater Management Act. This provides useful information on irrigation districts' scales, water sourcing, and the scope of overall water use (Lahmers et al., 2018).

Arroyo 2021 describes the AMA regulatory structure mandated by the Groundwater Management Act. This study provides a detailed explanation of the requirements growers within the AMAs must adhere to, including the non-expansion of irrigated lands and the system of fixed water rights described in Chapter 1. *Arroyo 2021* is also useful in its description of the water challenges faced throughout Arizona today. Factors like the decline in Colorado River flows and shifting and unpredictable climate patterns are discussed in detail (McGreal et al., 2021). These challenges underscore the need for a better understanding of growers' water use decision making.

Megdal (2012) provides further explanation of the 1980 Groundwater Management Act, water banking in Central Arizona, and the important role that CAP plays in helping the AMAs meet their goals. The study pays particular attention to the operations of various entities within Arizona's water management structure, as the system has drawn considerable worldwide attention from those seeking to supply water in similarly arid areas. Megdal goes on to provide detailed information on groundwater storage and recovery in the Central Arizona AMAs. As in the cases of Colby (2016) and Silber-Coats et al. (2017), this level of detail regarding water banking does not pertain directly to this analysis, but could be potentially valuable for future researchers attempting to model how, when, and in what quantities stored water will be recovered.

Hanak et al. (2019) examine the water challenges facing not Arizona, but California. While the issues each state are faced with today differ slightly in their character, there also exist many similarities, indicative of the general degree of uncertainty surrounding water supply in the Western US. The study discusses various strategies and adaptations that could be implemented to address a future with diminished and less predictable flows in addition to strained aquifers. While this study does not seek to model the impact of various conceptual management schema on water use, future work on the topic in Arizona and other southwestern states would be well informed by Hanak et al.'s ideas and recommendations.

Condon and Maxwell (2019) further underscore the drastic dangers posed by allowing aquifers to decline unchecked. This hydrologic analysis provides important support for the primacy of groundwater use in irrigated agriculture globally. The authors also underscore the interdependance between groundwater levels, surface water flows, and climate, and assert that diminished groundwater levels can and will influence future climate equilibria. Condon and Maxwell make clear the fact that mismanagement of groundwater supplies could lead to drastic changes in the future availability of all water resources.

Ferris and Porter (2021) tie together many of the ideas described above. This sobering report examines the idea of "safe-yield" within the Arizona AMAs and finds it to be 1) likely unachievable given the current regulatory framework, and 2) likely inadequate in its goal of assuring water supply for future generations of Arizonans. In their analysis, Ferris and Porter point to the limits of conservation efforts' ability to outpace growing water demand, the existence of long-term rights to pump groundwater, and the hydrologic disconnect between recharge and storage. The authors see these factors as major hindrances to sustainable management of Arizona's water supplies. Ferris and Porter (2021) further support the need to understand irrigators' water use behavior, but also leaves the reader with the disquieting feeling that water managers throughout the state may already be a day late and a dollar short.

2.2 Models of Agricultural Water Demand

Chapter 3 of this thesis focuses on developing two conceptual models, one describing growers' profit maximizing behavior which pays special attention to water as an input, and the other describing climatic, economic, and policy factors that the literature has shown to drive demand for irrigation water. This section describes the economic literature that supports the construction of these models.

Being that Fleck (2013)'s analysis parallels this study's, it seems a fitting place to start this section. Fleck begins his conceptual modeling by constructing a simple one-input one-output profit maximizing model for a grower choosing how much water to apply to crops. This fundamental approach to first showing an individual grower's profit maximizing behavior before going on to describe factors that influence aggregate water demand has directly informed the analysis presented in Chapter 3, although the profit maximizing model presented in that chapter concerns two inputs and one output. The profit maximization model presented in Chapter 3 draws further inspiration from Griffin (2016). This text outlines a classic profit maximization model specifically concerned with water as a primary production input.

Fleck then goes on to define a function for water demand based on precedent found in the existing literature. His demand function includes crop acreage, the cost of irrigation water, and climate variables. He defines crop acreage as a function of the price of crops, a vector of other input prices, and once again the cost of water. This description of water demand as being motivated by choice of crop acreage reflects extensive/intensive decision making, where growers first choose to what *extent* they will plant crops before deciding how *intensely* they will irrigate. This approach will be seen in multiple studies reviewed here, as well as in the water demand model presented in Chapter 3. Fleck's demand model serves as a basis for the model in the next chapter, although this study's water demand model will include other water use drivers supported by the literature described below.

The extensive/intensive decision making model presented by Fleck is supported by Bickel et al. (2019). The study framework involves rationing models based on "the 'putty-clay' production function approach to modeling production relationships," an approach originally pioneered by Moffitt et al.. This approach allows for flexibility in terms of production relationships in producers' planning stages (malleable "putty"), before being constrained by their initial decisions in the next stage of production (hardened "clay"). This approach is very applicable to agricultural production, as growers have tremendous freedom to choose what to plant and to what extent before committing to some specific crop mix. After that point, farms' production becomes restricted by their earlier cropping decisions. The only choice that remains is how intensely to apply water and other inputs like pesticides or fertilizer (intensive choices). Incorporating this approach into the conceptual model presented in Chapter 3 is beneficial, being that it captures the constraining effect a grower's cropping decisions have once a commitment is made.

Once a water use model is established, Bickel et al. (2019) go on to quantify the effect of a 300,000 acre-foot reduction in water delivered to Pinal County agriculture based on data from 2017. The authors compare estimates from three different rationing models, each with varying degrees of complexity and therefore varying input requirements, although each considers expected profits per acre-foot of water applied for various crops. The study then quantifies the economic impacts that would result from the loss of revenue associated with such a reduction in cropped acreage. This study does not allow for growers to offset the initial 300,000 acre-foot reduction in water deliveries by pumping groundwater or through any other alternate source. While this is a likely outcome of any reduction in CAP deliveries in the county, a large-scale shift to groundwater use would be accompanied by its own set of costs, as many wells would have to be recommissioned or installed from scratch (Seasholes, 2021).

Whipple (2019) includes a profit maximizing model designed to determine a Pinal County cotton farmer's optimal profit and irrigation levels, based on water cost. This is especially relevant as Arizona is currently facing cutbacks to its allocations of Central Arizona Project water deliveries, which by design will impact growers in the central part of the state the hardest. Whipple's research is doubly valuable in this literature review, as Pinal county is within the study area observed in later chapters. Much of the background Whipple has gathered on groundwater levels, pumping costs, and the structure of the county's agriculture sector has been helpful in informing this analysis.

Whipple distinguishes between the costs of CAP water and groundwater pumping and optimizes his model based on both cost structures. He also includes a two-crop model which allows for farms' ability to split production between cotton and alfalfa. He is also able to optimize his choice variables under various water constraints. The study finds that, given a reduction in CAP deliveries, farmers will convert their water use to groundwater and irrigate at almost the exact same levels whenever possible. Whipple does note that some growers may not have the available pumping capacity (or resources to expand pumping capacity) necessary to replace all lost CAP deliveries with groundwater. Whipple's model also predicts an increased percentage of cotton acreage (relative to alfalfa) in any water shortage scenario.

Haacker et al. (2019) set out to determine the drivers of change in depth-to-water table measurements from wells on the Ogallala Aquifer. Specifically, the study is set on determining the effect of various comparable groundwater management schemes throughout the states which overly the aquifer. Smith et al. (2017) similarly examine the effect of policy on water use, which will be discussed in the next paragraph. It should be reiterated that Haacker et al.'s analysis is using change in groundwater levels to determine the impact of said policies, while this thesis involves agricultural water delivered as the variable to be explained, at both the conceptual and empirical levels. Because Haacker et al. (2019) are concerned with the effects of policy intervention, this supports the inclusion of policy programs as a driver of water use in a conceptual model.

Smith et al. (2017) analyzes the effects of a bottom-up attempt at groundwater management among growers in Colorado's San Luis Valley. In 2010, growers in the valley instituted a self-imposed groundwater pumping fee in an effort to forestall government regulation of groundwater resources. Again, this is an instance where the authors have been able to capture the effect of a policy intervention on groundwater pumping, in this case a direct tax on withdrawals. It should be noted that this study is made possible by all wells in the San Luis

Valley having had meters installed prior to the 2009 growing season.

Because the fee system was introduced gradually in different years throughout the valley, the authors are able to make use of a difference-in-difference framework to assess the impact of the intervention. The authors model groundwater pumping, both within and outside of the intervention area, at two enumeration levels: the well level and the parcel level. The parcel level analysis allows for the inclusion of parcel-specific explanatory variables, as one well may provide water to many parcels.

Another instance of policy analysis, Zeff et al. (2019) seek to model a novel groundwater management policy being considered by growers in Diamond Valley, Nevada. Growers in the region have drastically overdrawn the local aquifer since the 1960s, leading to Diamond Valley being declared a "Critical Management Area" by the state. Similar to the situation described in Colorado by Smith et al. (2017), growers in Diamond Valley are attempting to prevent curtailments by instituting a homegrown groundwater management plan. The idea involves a gradual transition to universal water-rights cuts (rather than priority-based cuts), and/or transitioning to shares-based water allocations. This work provides further evidence of impact of anthropogenic choices on groundwater use, only in this instance rather than focus on irrigators' groundwater pumping, the study models the aquifer itself.

The study's modeling is hydroeconomic, as the authors are seeking to capture both the effect of the proposed interventions on the local economy and on Diamond Valley's aquifer. The model is run under five different water management scenarios and finds that many of the proposals would achieve the Valley's groundwater management goals while mitigating the economic decline associated with immediate priority-based curtailment. The findings here could be useful in making recommendations as to how Central Arizona growers (especially in Pinal County) might better manage their groundwater resources if CAP allocations one day become a thing of the past.

Another study from the Ogallala aquifer, Pfeiffer and Lin (2014) examine the effect of energy prices on groundwater use in Kansas. Like Haacker et al. (2019), this study is examining the impact of anthropogenic choices on water use, only in this case it is an economic driver that is the explanatory variable of interest. The study treats annual water pumped by individual farmers in a year as its outcome variable, which doesn't include any irrigation water from precipitation or surface sources, but allows for an incredibly granular statistical analysis.

The analysis once again consists of modeling designed to capture decisions made at the extensive and intensive margins. The first (extensive) model involves two simultaneous equations with crop choice and acres planted to each crop treated as dependent variables. The second (intensive) model treats groundwater extracted as the dependent variable, controlling for crop choice predicted by the first. In the empirical stage of their analysis, the authors find the total marginal effect of energy prices by summing the marginal effects from both models. Pfeiffer and Lin find that changes in energy prices result in a significant restructuring of crop planting patterns, with a small accompanying shift in acreage. They also find that change in energy prices results in significant decrease in groundwater pumping. This last result is robust across their principle model and two alternative specifications intended as robustness checks.

Kahil et al. (2015) take an approach somewhat similar to that presented in Bickel et al. (2019), but with a twist. The authors construct a simple but intricate hydroeconomic model of Spain's semiarid Jucar River Basin. The model is primarily concerned with the effect of droughts on the welfare of local economies and natural areas throughout the region, although crop choice, irrigation technology, land constraints, and the availability of other agricultural inputs are also considered. The authors are able to constrain their model by requiring that some baseline level of surface flows are put to environmental uses. This research illustrates the role that climate and water availability play in determining the fortunes of agricultural economies and natural spaces. Similarly to Bickel et al. (2019), the drought represented in

Kahil et al. (2015) could be taken as a proxy for a reduction in water allocated to Arizona growers from the Colorado River. Future researchers would certainly find this a useful foundation for any modeling that seeks to measure the impact of these reductions.

2.3 Econometric Approaches to Agricultural Water use

Chapter 5 of this thesis is devoted to constructing an econometric model of growers' water use based on the conceptual model of water demand outlined in Chapter 3. As such, much of the material described in the previous section has a direct impact on this empirical analysis by way of influencing this conceptual foundation. However, the conceptual model does not lend support to a particular econometric model. Instead, prior econometric studies provide inspiration for Chapter 5's modeling choices.

Once again, this discussion will begin by looking into Fleck (2013)'s econometric approach. Fleck's empirical modeling approach uses fixed-effects regressions, with total water used for irrigation treated as the dependent variable and a set of crop prices, water prices, weather measures, and land-use measures used as explanatory variables. Fleck includes irrigation district level fixed effects to account for heterogeneity across study subjects. Fleck uses a measure of retired irrigated acreage as reported to ADWR to track land use within each district.

Fleck finds that weather variables vary in significance across different functional forms, while crop prices and land use measures are almost always found to have significant effects on agricultural water use. The R^2 measure of fit is very high across many of the specifications in this analysis, which Fleck attributes mainly to the explanatory power of the fixed effect variables.

Deryugina and Konar (2017) seek to capture the effect of crop insurance on freshwater withdrawals throughout the United States (with some counties dismissed due to insufficient data). The study makes use of an instrumental variable approach which allows causal determination to be made. The authors treat water withdrawn for agriculture as an outcome variable subject to economic drivers, making it very relevant to this research.

The study first models water withdrawals as a function of insured acres using a standard ordinary least squares regression model, before making use of a two-stage least squares approach to establish a causal relationship. The instrumental approach makes use of "variation created by the 1994 Federal Crop Insurance Reform Act," which "had a large and immediate effect on insurance coverage." Treated as a variable, the policy change is found to have a high correlation with insured crop acreage and no significant correlation with agricultural water withdrawals.

Results vary widely between the two models, although both find that an increase in insured acreage results in an increase in water withdrawals. The OLS model estimates that a 1% increase in insured acreage is associated with a 0.051% increase in irrigation water use, while the 2SLS model estimates that a 1% increase in insured acreage *leads to* a 0.223% increase in water withdrawals. Using this instrumental approach to establish a causal relationship between crop insurance and agricultural water use can inform a researcher's initial modeling or provide a method for checking on a model's robustness. In addition to the masterful econometric approach, the establishment of crop insurance as a driver of water withdrawals supports the idea that economic factors can directly influence agricultural water use, which is vital to the conceptual model presented in Chapter 3.

As mentioned in the preceding section, Smith et al. (2017) uses Ordinary Least Squares regressions with a difference-in-difference framework to determine the effect of a new groundwater pumping tax on demand for water in Colorado's San Luis Valley. The study finds that the imposition of a pumping tax caused a significant decrease in groundwater pumping within the intervention area, consistent across both well-level and parcel-level models. The authors also make use of a Tobit model to assess the impact of the new policy on crop choice. The result of the crop choice modeling is less conclusive, in part because historic crop choices within the intervention area are not consistent with those elsewhere in the valley. The program in the San Luis Valley has led to a significant reduction in groundwater pumping and has thus far been successful in its goal of eliminating the need for state intervention. This straightforward and convincing study further underscores the need to include regulatory policy in both conceptual and empirical analyses of irrigators' water use.

Finally, while Larsen et al. (2015) do not directly influence the empirical analysis presented in Chapter 5, they do provide important support for the inclusion of land cover measurements drawn from USDA NASS's Cropland Data Layer. Specifically, Larsen et al. examine the accuracy of fine-scale spatial data with regard to the arrangement of crop types by comparing the CDL to data from the USDA Census of Agriculture. USDA NASS did not originally create the CDL to be used for precise statistical research, making work like Larsen et al. (2015) important support for any research depending on land cover measures taken from the CDL. The study finds that accuracy for individual crops is high in regions where that crop is predominant. In areas with many crops, large discrepancies can be commonplace between the CDL and the Census of Agriculture. This result is encouraging for the purposes of this research, as Central Arizona agriculture is largely dominated by two major crops: cotton and alfalfa. It is important to note that, given the timing of this study, the authors were only able to compare the CDL to the Census of Agriculture on a national scale for the year 2012, as the CDL was not expanded to the entire continental US until 2008 and the Census of Agriculture is only taken in years ending in 2 or 7.

2.4 Contributions

The body of work presented above represents many accomplished economic researchers in the field of agricultural water use and its implications. While the work presented throughout the rest of this thesis is only possible due to the foundations established by these and other authors, this analysis provides several contributions to the existing literature.

Of the studies mentioned above, only Fleck (2013) is concerned specifically with modeling demand for irrigation water in Central Arizona. As such, Fleck's work is an important foundation for this research. However, Fleck is not able to incorporate the land cover variables included in this analysis, as the CDL has only been available in Arizona from 2008 and his study period ends in 2011. The inclusion of land cover variables provided by the CDL builds on Fleck's existing econometric analysis, and the analysis of more up-to-date data will further develop understanding of agricultural water use choices. As precipitation and temperature patterns in the Colorado Basin continue to change, and CAP cuts loom, there is every reason to continue to develop our understanding of what factors motivate agricultural water use.

Remote sensed land cover data has not been commonly used in research involving agricultural water use. Over the past decade, these data products have become more readily available and increasingly accurate. The work presented here may encourage future researchers to incorporate remote sensed data products into their agricultural economic analyses.

Finally, many (although by no means all) of the studies mentioned above are primarily interested in the effect of a reduction in available water on agricultural economies. This work is essentially interested in the inverse: how does anthropogenic activity (including climate change) influence the water used to maintain and maximize farm profits? This line of questioning is particularly critical with CAP cuts on the horizon, as growers in Central Arizona may choose to replace a large part of these reduced deliveries with groundwater. As Condon and Maxwell (2019) points out, a significant reduction in aquifer storage would contribute to some new environmental equilibrium in the region, which could in turn make the maintenance of agricultural economies and water supply that much more difficult. Like Smith et al. (2017) and Deryugina and Konar (2017), this work seeks to better understand the effect of human activity on nature, an increasingly necessary area of research.

The content of the coming chapters owes everything to the works presented above. Ideally, this research will be a reflection of the consideration and ingenuity on display in these studies.

3 Theoretical Models

This study's conceptual modeling is presented in two parts. The first section describes a simplified version of an individual grower's water use decision making with the goal of maximizing profits. While Chapter 5's empirical analysis will be enumerated at the irrigation district level, it is important to bear in mind that irrigation districts' water delivery decisions are based upon said district's constituent members' water needs, along with any constraints (either legal or physical) on water availability. The second section describes a conceptual model of demand for water at the irrigation district level in a given year, based on existing literature. This conceptual model will provide foundational support for choices made in the empirical model specifications described later in this thesis.

3.1 Individual Profit-Maximization Model

The profit-maximizing model presented describes a circumstance where a grower concerned with a single crop is faced with the choice between applying water and and other inputs in order to bring crop yields to profit maximizing levels. This model then assumes that cropped acreage is already established. It may seem to some that this simple model is skipping a crucial step, as choice of crop mix and planted acreage is a vital part of growers' profit maximizing behavior. However, consider an alfalfa grower, whose stands are already planted and established at the beginning of a season. This agent would then only have the choice of what inputs to employ in order to bring yields to maximize profit. Thus, the first model presented is a one-output, two-input model.

In this model, W will represent water applications and X will represent some other choice input (labor, fertilizer, et cetera). In practice, X could be denoted as $\underline{X}$, a vector of multiple agricultural inputs, but for simplicity's sake this model will be restricted to two inputs. The model assumes yield to be a function of water applied and the other choice input, $Y = f(W, X)$, and that a grower may choose to apply varying quantities of either input. This crop production function is assumed to be "well behaved", meaning that at some point, as W and X increase, $f(W, X)$ will begin to increase at a decreasing rate and eventually begin to diminish. It is important to note that there may be some degree of substitutability and complementarity between water and other inputs, depending on the nature of the production function $f(W, X)$. Finally, the model assumes that no water is applied to or removed from the crop beyond the grower's chosen level of irrigation. Consequently, variations in rainfall effects (water applied to crops) and in temperature effects (water removed from crops via evapotranspiration) are not considered. Alternatively, one could make the assumption that typical variance in rainfall and temperature is already incorporated into the production function. Later, in the empirical work, rainfall and temperature effects are considered through the Standardized Precipitation-Evapotranspiration Index.

Define a producer's profits (π) as total value of a producer's output (TVP) less total operating costs (TC).

$$\pi = TVP - TC \tag{1}$$

Applying this very simple framework to a one-crop grower: quantity of output is given as a function of water applied and the other choice input, $f(W, X)$. The expected output price is given as P, and the total value product therefore is found by multiplying output by expected price. The cost of water, $c(W)$, is here specified as a function of water used. This could be due to a tiered cost structure for water. For example, a grower might only have access to some limited quantity of water from a low cost source. If this is exhausted, the grower would have to turn to a more high-cost alternative. If the cost of water is assumed to be a static unit price, than $c(W)$ would be replaced by some fixed value (r) multiplied by water applied. Finally, the cost of the other choice input is here represented as C_X. Total costs then are merely the sum of the cost of water applied and the cost of the other choice input.

$$
\begin{aligned}
TVP &= Pf(W, X) \\
TC &= c(W) + C_X X
\end{aligned}
\tag{2}
$$

Having established these relationships, the next step is to substitute prices and quantities into the initial profit function, yielding the following:

$$\pi = Pf(W, X) - c(W) - C_X X \tag{3}$$

Being that this profit function is now described fully in terms of only two variables, it can now be maximized by choosing optimal levels of W and X.

$$\max_{W,X} \pi = Pf(W, X) - c(W) - C_X X \tag{4}$$

At this point, constraints can be introduced if applicable. For example, a grower might be limited by the amount of water available. This constraint o available water could be due to either regulatory policy or physical availability, or some combination of the two. If this were the case, then $W \leq W_{max}$, with W_{max} representing the maximum water available.

In order to maximize this profit function, first order necessary conditions must be found by taking the partial derivatives of profits with respect to each input.

$$
\begin{aligned}
\frac{\partial \pi}{\partial W} &= \frac{\partial [Pf(W, X) - c(W) - C_X X]}{\partial W} \\
\frac{\partial \pi}{\partial X} &= \frac{\partial [Pf(W, X) - c(W) - C_X X]}{\partial X}
\end{aligned}
\tag{5}
$$

By setting these partial derivatives equal to zero, the profit function is forced to some extrema. Because of the previous assertion that $f(W, X)$ is a "well behaved" function, it can be assumed that this extrema will be a maximum. Once these partial derivatives are set equal to zero, W and X are no longer variables, but some optimal fixed quantities. These will be represented with W^* and X^*. Taking the derivatives outlined above yields the following:

$$\begin{aligned} Pf_W - c'(W^*) &= 0 \\ Pf_X - C_X &= 0 \end{aligned} \tag{6}$$

These functions can now be rearranged with the output price (P) on the left hand side of each.

$$\begin{aligned} P &= \frac{c'(W^*)}{f_W} \\ P &= \frac{C_X}{f_X} \end{aligned} \tag{7}$$

At this point, the equations can be set equal to each other, as the ratio of marginal cost to marginal productivity of each input is equal to the output price under profit maximizing conditions.

$$\frac{c'(W^*)}{f_W} = \frac{C_X}{f_X} \tag{8}$$

Rearranging this equation reveals the following:

$$\frac{f_X}{f_W} = \frac{C_X}{c'(W^*)} \tag{9}$$

Here, it can be seen that the rate of technical substitution, the ratio of the marginal productivities of two units involved in the production of one output, is equal to the ratio of the marginal prices of these two inputs. Recall that C_X is some fixed value, and that if the cost of water were defined as a fixed value as well (say r), the ratio of these marginal input prices would merely be the ratio of the market prices of the two inputs ($\frac{C_X}{r}$). This relationship between the rate of technical substitution and the input price ratio is to be expected when profits are maximized.

It is also at this stage of profit maximization that, if $f(W, X)$ were defined, an expansion path could be solved for. The expansion path is a function of one input in terms of the other, which provides all optimal input combinations as the scale of production expands. Solving Equation 9 for one or the other choice variable would provide this expansion path, but this is only possible if $f(W, X)$ were defined.

Instead, Equation 9 can be rearranged as such:

$$f_X = \frac{C_X}{c'(W^*)} f_W \tag{10}$$

This can now be substituted into the first order condition described in Equation 6 (recall, $Pf_X - C_X = 0$).

$$P[\frac{C_X}{c'(W^*)} f_W] - C_X = 0 \tag{11}$$

This equation can be rewritten to isolate f_W on the left hand side, which, after some additional operations, is found to be equal to the ratio of marginal cost of water to output price.

$$f_W = \frac{C_X}{P} \frac{c'(W)}{C_X} = \frac{c'(W^*)}{P} \tag{12}$$

Rearranging this once again reveals the following:

$$Pf_W = c'(W^*) \tag{13}$$

Output price (P) multiplied by the marginal product of water (f_W) provides the marginal value of product of water (MVP_W), while $c'(W)$ is definitionally the marginal cost of water (MC_W). So, when optimal profit-maximizing choices are made, the marginal value product of water is equal to the marginal cost of water.

$$MVP_W = MC_W \tag{14}$$

This relationship between marginal value product and marginal cost can be expected to exist for all choice inputs.

3.2 Irrigation District Water Demand Function

In this section, a conceptual model of irrigation districts' water demand is presented, identifying drivers of water demand based on existing economic research. This approach informs the econometric modeling in Chapter 5 of this thesis.

Modelling begins by introducing the idea that total water applied in an irrigation district in a given year is a function of two factors: the number of acres of each crop planted in the district in that year, and the overall average intensity of water applied throughout the district throughout the year. Crucially, this modelling approach assumes that any constraints on water supply are non-binding. Moore and Dinar (1995) demonstrate that often water supply does act as a binding constraint in the Western US, and so this assumption must be applied carefully. The Central Arizona irrigation districts have historically had access to large amounts of Colorado River water (via the Central Arizona Project), limited but plentiful groundwater resources, and in some cases surface water from the Salt and Gila Rivers. If a fixed amount of water is available to a grower, the work presented in this Chapter and in Chapter 5 is potentially not applicable, depending on whether or not the constraint binds.

A water demand function based on planted area and intensity of water application is described mathematically below. W_{it} represents the total water delivered by an irrigation district (i) to member growers in a given year (t). This is a function of acres planted in each crop (c), represented by A_{cit}, and intensity of irrigation (acre-feet applied per acre), represented by I_{it}.

$$W_{it} = f(A_{cit}, I_{it}) \tag{15}$$

Much of the agricultural economics literature describes growers' planting decisions as a two stage process. Before the growing season, growers decide what crops to plant and how many acres to devote to each. This is sometimes referred to as decision making on the extensive margin, (as in "To what *extent* should I plant my crops this year?"). Once crops are planted, the principal remaining decision is how intensely water and other production inputs will be applied throughout the growing season. This is sometimes referred to as decision making on the intensive margin (as in "How *intensely* should I irrigate my crops?"). This two stage characterization of growers' decision making process can be seen in Deryugina and Konar's 2017 study of crop insurance's effect on irrigation water withdrawals, in Frisvold and Konyar's 2012 examination of reduced water availability in the Southwestern US, and in Zeff et al.'s 2019 modeling of various water trading scenarios in Nevada. These are only a few of many examples that apply this approach. The water demand function described above essentially describes water demand as a function of decision making along these margins.

This demand function is definitional: water demand will be driven by what is planted, how much of it is planted, and how intensely water is applied. The process of identifying drivers of water demand then becomes a question of identifying drivers of these contributing factors. The remainder of this chapter will be dedicated to addressing extensive and intensive water use decisions.

Extensive Decisions - What Crops to Plant and How Much?

Factors determining a grower's crop mix and planted acreage are much discussed in the literature, which supports the inclusion of key variables in any function designed to model these extensive decisions. These variables include (but may not be limited to) crop prices, climate, input costs, regulatory policy, and irrigation and pumping technology. Each is discussed in greater detail below.

Expectations about crop prices are a foundational driver of planting decisions. It is intuitive to assume that a grower's expectations about the profitability of various crops will inform their decisions of what to plant in what quantities. Profitability depends on both revenues (crop prices, crop yields) and on production costs. Haacker et al. (2019) identifies crop prices as a primary driver of irrigation demand, and includes them in a Classification and Regression Tree analysis of drastic changes in groundwater levels in the High Plains Aquifer. Pfeiffer and Lin's 2014 study of the effect of energy price on groundwater withdrawals in the same region likewise includes expected crop prices in its econometric analysis. Smith et al. (2017) includes crop prices in their analysis of groundwater use in Colorado's San Luis Valley. Additionally, Deryugina and Konar (2017), Frisvold and Konyar (2012), and Zeff et al. (2019), all briefly described in the preceding section, each incorporate crop prices in their modeling. It is worth noting that pre-season expectations of revenues earned for a particular crop are often influenced by crop insurance or crop-specific federal commodity programs (Deryugina and Konar (2017), Yu et al. (2018)).

Expectation of the growing season's climate is another factor that research indicates can affect extensive decision making. Climate is included as a driver of regional irrigators' water demand in Kahil et al.'s 2015 hydroeconomic modeling of water use in Spain's Jucar Basin. Haacker et al. (2019) incorporate a drought index in their econometric analysis. Qiao (2018) describes climate as an important driver of farm profits, which themselves are a direct result of crop mix and planted acreage. Smith et al. (2017) includes surface water availability in their analysis of groundwater use, which is partially the result of temperature and precipitation in the region. While this could be seen as somewhat of a proxy for climate, it is also important to note that regulatory policy and physical water delivery infrastructure would also have some role to play in the availability of surface water.

Input prices are also considered in econometric modeling of crop mix and water demand. This study specifically focuses on the cost of water as an input of interest, and the cost of energy, a necessary component of surface water delivery and groundwater pumping. Haacker et al. (2019), Kahil et al. (2015), and Zeff et al. (2019) each include some measure of the cost of water in their analyses. Zeff et al. (2019) includes groundwater lift costs in the study's hydroeconomic optimization model, a function of energy prices and depth-to-groundwater. This provides a price for water that is a function of the price of energy. Smith et al. (2017) likewise incorporates pumping costs in it's empirical analysis. As briefly mentioned above, Pfeiffer and Lin (2014) is focused specifically on the effect of energy prices on groundwater use, and so energy prices are of course incorporated into their empirical modeling. Taking a different approach in their 2012 study, Frisvold and Konyar consider not the direct cost of water or energy, but the elasticity of substitution between various agricultural inputs, including irrigation water and energy. This is a valuable alternative to including direct measures of input costs in an explanatory model if data regarding input prices is difficult to come by.

Finally, the literature suggests that any effort to model water demand at the irrigation district level should include a set of district-specific characteristics, such as regulatory policy, and irrigation and pumping technology. Federal, state, and local policies can have a major effect on water demand. Haacker et al. (2019), Smith et al. (2017), and Zeff et al. (2019) each involve examining water policy's effect on water use in one way or another. While policy isn't necessarily a district-specific characteristic, it can potentially differ between irrigation districts. For example, the Central Arizona Irrigation and Drainage District and the Maricopa Stanfield Irrigation and Drainage District are each located in the Pinal AMA and thus subject to the same regulatory policies, while the Central Arizona Irrigation and Drainage District and the Buckeye Water Conservation and Drainage District are located in different

AMAs and thus subject to differing water regulations. In addition to considering policy, Haacker et al. (2019) also includes available technology as a driver of crop mix and acreage. Frisvold and Deva's 2012 study of the adoption of new irrigation technologies in the Western US tells us that "improving irrigation efficiency is seen as key to reducing water pollution and *easing competition for scarce water in the west*" (italics added for emphasis). Any of the aforementioned studies which include some measure of lift costs, such as Smith et al. (2017) and Zeff et al. (2019), must know or assume what pumping technologies are present in their study areas. In fact, due to not knowing the prevalence of specific pumping technologies in their study area, Pfeiffer and Lin's 2014 study includes multiple model specifications in order to account for a range of possible technologies being present.

With all this in mind, a function that captures choice of crop mix and planted acreage would include the following:

$$A_{cit} = f(P_t, C_{it}^{exp}, r_{it}, e_{it}, D_{it}) \tag{16}$$

Here (as above) A_{cit} represents acreage planted in each crop c in irrigation district i in year t. P_t is a vector of all crop prices (including the price of crop c and the price of crops not planted) in year t. C_{it}^{exp} is some measure of the expected climate in district i in year t, as extensive decisions occur before the start of the growing season. As in the profit maximization model, r_{it} represents the cost of water for district i in year t, while e_{it} represents the cost of energy for district i in year t. Finally, D_{it} is a vector of characteristics specific to district i in year t, including regulatory policy and the prevalence of different irrigation and pumping technologies.

Having specified this extensive function, the next step in assessing a function of overall water demand is to examine intensive decision making.

Intensive Decisions - How Much Water to Apply?

Recall that intensive decisions refer to the intensity of water applied to planted acres over the course of a growing season. Like the extensive decision making process, decision making regarding irrigation intensity is frequently discussed in agricultural economic literature. While there is not complete agreement on which factors are the most important drivers of these decisions, there are common threads throughout existing research. These include crop mix and acreage, crop prices, climate, and input costs, each of which is discussed in greater detail below.

Crop mix and acreage, i.e. a grower's extensive decisions once made, are the primary drivers of that same grower's intensive decision making. This relates to the putty and clay approach descried by Moffitt et al.. Once a crop is planted in the ground the malleable putty has set in to clay, meaning a course has been set for that growing season. The specific water needs of a growers' planted crops must now be met to a greater or lesser extent. These crop water needs are marginally driven by climate and a grower's desired yield. Deryugina and Konar (2017), Haacker et al. (2019), Kahil et al. (2015), and Smith et al. (2017) all highlight acres planted in each crop as a key driver of irrigation water demand.

Of particular interest to this analysis is the presence of fruit and nut trees in Central Arizona. Orchards represent a long term investment, as trees are expensive to plant and take some years before they begin to fruit and are expected to be productive for decades. As long term investments, it is import for growers to maintain the health of these trees. In any given year, orchards would require some minimum amount of water applied in order to remain healthy and able to produce crops in the future (Fereres et al., 2003). Remember that this modelling approach assumes that constraints on water availability for irrigation are non-binding. Were a binding constraint on water introduced, tree crops minimum water needs would act to further reduce the amount of water available to irrigate annual crops, as ceasing to irrigate tree crops entirely would typically lead to higher losses for a grower than would ceasing to irrigate annual crops. For this reason, tree crops will be treated separately from the annual cropping decisions represented in the extensive model.

As in the extensive model, crop prices again play a role in intensive decision making. This is shown in Frisvold and Konyar (2012), Haacker et al. (2019), Kahil et al. (2015), Pfeiffer and Lin (2014), and Zeff et al. (2019). As the growing season progresses, expectations of crop prices at the time of harvest may shift, leading growers to potentially alter their irrigation decisions. Also included once again are any changes in the costs of water and energy. In some cases, irrigation districts have established contracts for water and power which would not be subject to change throughout the growing season, but this cannot be taken for granted. Also, a shift in expected crop prices relative to these input costs may trigger some change in a grower's water application decisions. The inclusion of these variables in an intensive model is supported by Haacker et al. (2019), Kahil et al. (2015), Pfeiffer and Lin (2014), and Zeff et al. (2019).

Given all this, a model of an irrigation district level intensive decision making process can now be constructed. Formally, this function would resemble the following:

$$I_{it} = f(A_{cit}, P^o_{cit}, C_{it}, r_{it}, e_{it}, T_{it}) \tag{17}$$

Where I_{it} represents the intensity of irrigation in district i in year t. A_{cit} represents the outcome of district growers' extensive choices, including acreage planted in each crop present in district i in year t. P^o_{cit} represents a vector of only the prices for crops (c) planted in district i in year t, as extensive decisions have been made by this point and crops are now in the ground. The climate is included once again, although C_{it} now represents the true climate during the growing season, as opposed to grower's expectations. As in the extensive model, r_{it}, and e_{it} represent the costs of water and energy in district i in year t. Finally, T_{it} is included to indicate acreage of tree crops present in district i in year t.

Final Water Demand Model

Now that conceptual models have been established which include drivers of crop planting decisions and irrigation intensity, these can be substituted into the water demand model specified in Equation 15. The final form is shown below.

$$W_{it} = f(A_{cit}, I_{it}) = f(P_t, C_{it}^{exp}, r_{it}, e_{it}, D_{it}, C_{it}, T_{it}) \tag{18}$$

A_{cit} is omitted from this final specification as it can be generated by other variables contained within. P_{cit}^{o} is also omitted as any price information it contains would also be included in P_t.

This model accounts for all factors seen commonly throughout recent agricultural economic literature examining irrigation water use and informs the econometric analysis in Chapter 5.

4 Data

Data used in the upcoming empirical analysis comes from multiple sources and has been compiled with guidance taken from existing literature and consultations with experienced researchers. Data is chiefly sourced from state and federal agencies, with additional data coming from one foreign public agency and one private entity (discussed in greater detail below). The end product is a set of panel data with each observation representing a Central Arizona irrigation district in a given year. As previously stated, this analysis includes data from twelve Central Arizona irrigation districts, all located within either the Phoenix or Pinal AMA. Annual observations occur between 2008 and 2019. 2008 is necessarily the lower threshold for observations due to the fact that the source for cropland data, USDA NASS's Cropland Data Layer (CDL), does not have information on agriculture in Arizona prior to 2008. Similarly, 2019 is the upper threshold since 2020 water delivery data for irrigation districts has been published by the Arizona Department of Water Resources (ADWR) at the time of this writing. The result is a total of 144 observations, annual observations from twelve districts over the course of twelve years.

The rest of this chapter will discuss data sourcing and procedures used in compiling the data set, as well as the characteristics of the data itself.

4.1 Water Use Variables

Data on water use by irrigation districts in a given year are taken from reports filed annually by irrigation districts with ADWR each year. Under the 1980 Groundwater Management Act, irrigation districts located within an Active Management Area are compelled to file reports listing water use by water right, which are in turn published by ADWR. This provides information on both end use (agricultural, municipal, industrial, etc.), and water source (groundwater, surface water, CAP, etc.), with water accounted for in acre-feet. Most districts summarize water usage by water right into these grouped categories, and these summary data serve as the primary source for district water use data.

Regressions in this analysis consider two dependent variables: water deliveries to agricultural uses, and the intensity of water applied to crops. Summary statistics describing these dependent variables are provided in Table 1. This table also describes overall water deliveries (including but not limited to those intended for agricultural use), the percentage of districts' water use intended for agriculture, and the natural log of deliveries to agriculture.

Figure 4 presents agricultural water deliveries to irrigation districts included in this study, with overall average annual deliveries, as well as average annual deliveries to districts in each AMA plotted. The actual values plotted are the natural log of water deliveries, in order to make the higher volume Pinal AMA districts and the lower volume Maricopa AMA districts comparable. It is plain to see that water deliveries to Central Arizona agriculture have generally been in decline since 2008. This is likely a result of the GMA's regulations involving the non-expansion of irrigated area (within the AMAs, existing agricultural lands can fall permanently out of production but new lands can not be brought under irrigation). Another explanation for this general decline could be the expansion of developed area in and around Phoenix diverting more water away from agricultural uses. However, the Pinal AMA's average annual agricultural water use would seem to dispute this last idea, as the general trend of reduced agricultural water use is similar in both the Phoenix AMA and the much less developed but faster growing Pinal AMA (World Population Review, 2021). Table 2 clearly illustrates the large differences in agricultural water deliveries across districts.

Because of these substantial differences in delivery volumes, the regression models presented in Chapter 5 are specified with the natural log of agricultural deliveries as the dependent variable (as opposed to the raw value). Water deliveries to agricultural uses are reported in acre-feet directly, meaning the only data "cleaning" that is necessary is to take the natural log of those values reported by the irrigation districts themselves. The natural log is chosen due to the wide range of values and high variance observed in deliveries to irrigation districts. This high variance is due to the differences in scale between irrigation districts. In 2011, the Central Arizona Irrigation and Drainage District delivered over 338,000 acre-feet of water to agricultural end users, while in 2010, the Tonopah Irigation District delivered just under 13,800. These two values differ by orders of magnitude and are in no way outliers: of 144 agricultural water delivery values observed, 24 are over 200,000 acre-feet while 13 are under 20,000. Taking the natural log of these values "flattens" these scalar differences, allowing regression model results to be more easily interpreted.

Figure 5 shows the natural log of water deliveries to agriculture in each year in each irrigation district. A number of trends are apparent. Firstly, large districts such as the Central Arizona Irrigation and Drainage District routinely deliver large quantities of water, while small districts such as the Arlington Canal Company deliver less. This is a product of the amount of cropland available in these districts. Some districts show substantial variation from year to year, while others are relatively stable in their water deliveries. Certain districts also seem to be gradually reducing their water deliveries over the course of the study period, such as the Hohokam Irrigation District.

Irrigation intensity is a measure of the average volume of water (acre-feet) applied to every 900m^2 planted area in an irrigation district. To find this value, it is necessary to divide agricultural water deliveries by planted area in each district in each year. Planted area is found by summing the area of all crops planted in an irrigation district in a given year, and is listed in Table 2. These cropped areas are found using CDL data, described in much greater detail later in this chapter. CDL's pixel resolution is 30 meters by 30 meters, or 900m^2, which is why this unit of measure is chosen for land area. However, this could easily be scaled up (to an acre or hectare) without altering the fundamental information the variable conveys.

Figure 6 describes irrigation intensity by irrigation district across the study period. When examining Figure 6, some notable trends are apparent. Due to irrigation intensity being the ratio of agricultural water deliveries and planted area, this measure is not bound by the amount of planted area available in a district. District with small planted area can exhibit high irrigation intensity and vice versa. Therefore, intensity can be more easily compared across smaller and larger irrigation districts than raw delivery volumes. Generally, most irrigation districts studied exhibit some decline in irrigation intensity over the course of the study period. In some districts, such as the Hohokam Irrigation District, this decline is very pronounced.

Finally, it is important to touch on the relationship between water deliveries and planted area. Agricultural water deliveries and CDL planted area measures are highly positively correlated, returning a correlation coefficient of 0.9549. This is to be expected, as the amount of planted area heavily governs water deliveries in arid regions where all planted acreage is irrigated (an exception to this rule being if crops are abandoned mid-season). As mentioned above, his high correlation is not observed with planted area and irrigation intensity. The correlation coefficient is only 0.0153 between the two variables, indicating almost no correlation whatsoever. Notice in Table 2 and Figure 6 that districts with relatively small planted areas sometimes exhibit high irrigation intensity. Likewise, the very large Central Arizona Irrigation and Drainage District exhibits relatively low irrigation intensity. Growers' extensive decision making in terms of planted acreage largely determines agricultural water deliveries (as demonstrated in Chapter 5). The high correlation of 0.9549 still leaves some room for other factors to influence overall water use, such as crop mix and economic variables. On the other hand, it seems as though intensive decisions regarding irrigation intensity are not influenced by planted area (per Figure 6). Exogenous factors (such as crop prices or climate) may be playing a large role in determining this behavior. These relationships are further explored in Appendix A.2.

4.2 Land Cover Measures

Land cover measures are taken from the Cropland Data Layer (CDL), which is published annually by USDA NASS. The CDL contains remote sensed land cover data covering the contiguous United States at a 30m resolution (with roughly 4.5 of these 900 square meter pixels equaling an acre). While the CDL's primary purpose is measuring crop cover, it returns some value for every 900m^2 pixel within the area recorded. This means that developed area, fallowed lands, open water, grassland, forest, and other non-cropped land covers are all classified and recorded, some of which are used later in this analysis.

CDL raster data can be integrated with shapefiles from ADWR which define the boundaries of Arizona's irrigation districts. In this way, one can determine the acreage of individual crops within any given irrigation district, as well as fallowed area, developed area, and the area of other land covers. The variables detailed in Table 3 are defined in this way, with a measure taken for each irrigation district in the study area for each year in the study period. These variables are used in this work in two ways: first, to examine the extent to which irrigation districts' water deliveries depend on crop mix and planted acreage (Appendix A.2), and second, as explanatory variables in the empirical analysis.

Table 2 breaks down average planting behaviors by irrigation district, as well as including average agricultural water deliveries and some fixed characteristics, such as district size and which AMA a district falls within. Sometimes there is a large disparity between overall district area and average planted area. For example, the Maricopa-Stanfield Irrigation and Drainage District on average plants just over 50% of the total district area available. The Salt River Project, in the heavily developed Phoenix metro area, regularly sees crops planted on just over 6% of the district's area. For this reason, when calculating any variables dependent on area (such as irrigation intensity described above), planted area is used rather than total area within the district. Fallowed or idle cropland will also be included in the empirical analysis presented in Chapter 5, but will not be included in measures of planted area.

Looking further at planted area reveals one of the major differences between the Phoenix and Pinal AMAs. Although there are only three irrigation districts within the Pinal AMA, taken together they make up 59% of the average planted area observed in this study. Among irrigation districts in the Phoenix AMA, only the Roosevelt Irrigation District averages a larger planted area than the smallest Pinal district. This speaks to a difference between the two AMAs and their approaches to agricultural water management. Agriculture is a key element of Pinal County's economy, while the economy of the Phoenix metro area is as largely diversified as one might expect a major American city to be (Lahmers et al., 2018). This lends further weight to the inclusion of differences in AMA water regulatory policies in econometric modeling of water use in Central Arizona, achieved through the inclusion of an AMA indicator variable.

Table 2 also illustrates both the extent to which alfalfa and cotton dominate Central Arizona's agricultural landscape, and the degree to which irrigation districts differ in their approach to the two crops. While alfalfa is the predominant crop in all but one irrigation district, some districts split their planting of the two crops almost evenly, while others plant almost no cotton at all. Districts within the Phoenix AMA heavily favor alfalfa in their crop mix, while Pinal AMA districts' cropping patterns are more diversified. Also notable is the fact that the only instance of a district regularly planting more cotton than alfalfa comes from the truly massive Central Arizona Irrigation and Drainage District. The average area of cotton planted in this district is greater than the average area of alfalfa planted in any district in the study area.

Moving away from alfalfa and cotton, it is evident from Table 2 that some districts favor planting cereal grains over cotton, and in the case of the Maricopa Water District, over alfalfa as well. The grains category is made up of five distinct cereal crops, not one of which is nearly as common as alfalfa or cotton (recall Figure 2). Tree crops and irrigated pasture are not particularly prevalent throughout the study area, with neither category regularly representing more than 3% of the planted area of any irrigation district. In terms of raw area, tree crops are almost twice as common as irrigated pasture in each district observed. This study is particularly interested in the effect of tree crops on water use, as they are structurally unlike any other crop group defined in Table 3. Tree crops are necessarily long term investments, and require substantial up-front costs to plant. While specific characteristics vary depending on the species, tree crops generally take several years after being planted to bear fruit, and require some minimum amount of water to survive from year to year (Fereres et al., 2003). As a result, there exists something close to a fixed lower bound for water needed to maintain an orchard's health. Finally, the "other crops" category makes up a significant portion of a few districts' planting, although it represents a small proportion of overall planting in the study area. Being that this grouping includes over a dozen crops, it may be hard to assess the implications of any significant effect related to this grouping observed in an econometric model.

In addition to the crop mix variables described above, the CDL provides data on fallowed/idle croplands, described in the final column of Table 2. Whether fallowed area has a significant effect on either water deliveries or irrigation intensity is explored in econometric analyses presented in Chapter 5. One hypothesis is that the greater the extent of fallowed area within a district, the less irrigation water need be delivered. However, fallowed area may also play a role in driving growers' irrigation intensity decision making. Under some types of water entitlements, more fallowed area in a district might mean there is more water available to irrigate those crops which have been planted. If the fallowing is due to water leasing by growers to non-agricultural water uses, then fallowing generally does not create unused water that can be applied to other crops on growers' land. Table 2 also shows that, like agricultural water delivery volumes, fallowed area differs substantially between districts, so once again the natural log of fallowed area will be included in Chapter 5's econometric analysis.

USDA NASS reports metadata for the CDL, including accuracy estimates for land use variables reported by state for every year data is published. Figure 7 reports accuracy estimates for cotton, alfalfa, and fallowed area measures in Arizona during the study period. It is notable that accuracy or measurement for fallowed lands is consistently lower than that of cotton and alfalfa. This is likely due to how the CDL trains its algorithm to process remote-sensed data. USDA NASS uses self reported information from growers on what crops are being planted in individual plots. In the case of fallowed area, growers report lands that are not in production currently but have been recently and which they intend to plant again in the future. This self reported information is then used to train processing algorithms to interpret remote sensed satellite imagery. It is very likely that, in arid state like Arizona, algorithmic processing may have a harder time distinguishing fallowed crop land from other lands than it would distinguishing lands planted in crops like cotton or alfalfa (Willis, 2021).

4.3 Economic and Climatic Variables

In addition to water use data taken from ADWR reports and land cover data compiled from the CDL, this analysis will consider a number of economic and climatic variables' effect on irrigation decisions. These include prices for alfalfa and cotton, the costs of Central Arizona Project (CAP) water and electricity, and the Standardized Precipitation Evapotranspirtation Index, a climate measure that incorporates both temperature and rainfall. Basic summary statistics relating to each of these variables are presented in Table 4. The sources and procedures used in compiling these data, as well as more detailed discussion of the variables' behaviors, will be discussed in detail in this section.

Being that alfalfa planting is predominant throughout the Phoenix and Pinal AMAs, the prices growers expect to receive for alfalfa are of particular interest. As there aren't many major federal commodity programs supporting alfalfa production, market prices for alfalfa are generally seen as a fair measure of growers' price expectations (Frisvold, 2021).

Alflafa market prices used in this analysis are sourced from USDA NASS's Quickstats service, which provides annual prices paid for alfalfa in Arizona in dollars per ton. These prices are then adjusted for inflation using the US Bureau of Labor Statistics Consumer Price Index (CPI) for the year 2019 to generate real prices that are comparable across the study period, henceforth referred to as 2019$. This same index will be used to adjust all other price variables in this study.

Figure 8 illustrates change in real alfalfa prices in Arizona over the study period. There is no immediately recognizable pattern or trend at work in the pricing pattern. The relatively high standard deviation presented in Table 4 speaks to the high variability in price on display.

Like alfalfa, cotton is an extremely commonly grown commodity crop throughout Central Arizona. Unlike alfalfa, cotton is subject to many federal commodity programs. This results in a scenario where prices that cotton growers' receive may differ considerably from the price the cotton is ultimately sold for in the marketplace. Because of this, simply considering market price of cotton as a motivator of growers' cropping and irrigation decisions is likely inadequate. Instead, a fair proxy for a grower's expectation of price received must be found. One such proxy variable commonly seen throughout the literature is the New York Cotton Exchange December Futures (2021) price just prior to planting. The December Futures price from the last Friday in February is seen most informative for grower's planting decisions, as this is provides some expectation of price as late in the year as possible before planting decisions are made (Tronstad, 2021). While some research incorporates the end-of-February December Futures price into larger functions which generate an expected price of cotton (Sall, 2019), other work has indicated that these data alone serve as a strong proxy for growers' expectations of price (Frisvold, 2021). Along with other price variables, these December Futures prices have been adjusted to 2019$.

Figure 9 illustrates change in the real price of December cotton futures over the course of the study period. Generally, these prices seem to be on the decline since 2011, although this perceived trend could simply be due to the small sample of years pictured. December cotton futures display substantial variability, although the standard deviation is not nearly as high as the statistic reported for alfalfa prices.

CAP water prices are fortunately much more straightforward than cotton prices when it comes to compiling useful data. The CAP, administered by the US Bureau of Reclamation (BoR), provides low cost Colorado River water to Central Arizona end users. Water prices for growers are available from from CAP, which publishes annual fee schedules that include fees for the current year as well as projected prices in future years (Central Arizona Project, 2021b). These fee schedules are listed in dollars per acre-foot of water delivered and are readily available through BoR. Once compiled, the typical adjustment to 2019$ makes these prices comparable to each other as well as other economic variables.

Figure 10 illustrates change in the real price of CAP water deliveries over the course of the study period. Unlike the behavior observed in Figures 8 and 9, CAP prices trend consistently up over the eight year period between 2008 and 2015 (notwithstanding a small blip in 2012). This trend has vanished in the last four years, as water costs first stagnated before dropping substantially between 2017 and 2018. Variability is considerably lower than that observed for alfalfa or cotton futures prices, which is consistent with Figure 10's graphical representation.

Another economic factor considered but ultimately not included in this study's empirical analysis are energy costs. This may seem like a glaring omission, and indeed, existing economic literature supports the inclusion of some measure of the cost of energy to account for costs associated with groundwater pumping. Recall from Chapter 2 and 3 that Pfeiffer and Lin (2014), Smith et al. (2017), and Zeff et al. (2019) each consider energy costs in some form in their empirical analyses. So then, why are they omitted from the econometric models specified in Chapter 5 of this study? Irrigation districts in Central Arizona often have their power supplied through a state organization called the Arizona Power Authority (APA). The APA contracts with the federal Western Area Power Administration to purchase energy generated at Hoover Dam, Parker Dam, and other BoR hydroelectric projects located along the Colorado River. This energy is provided at extremely low cost, with rates being provided to agricultural uses for as little as $0.06 per kilowatt hour. In addition, APA contracts with irrigation districts and their subsidiary electrical districts are often very long term, meaning these low rates for energy can be fixed for decades (Arizona Power Authority, 2021). Of the twelve irrigation districts included in this study, eleven can be conclusively shown to have their power provided through the APA (George Cairo Engineering, Inc., 2021; Arizona Power Authority, 2021; Western Area Power Administration, 2021). Only the Arlington Canal Company lacks available documentation linking them to the organization. While this may mean that their power is supplied from other sources, it could also be the case that this district is indeed an APA customer, and that documents indicating this are not forthcoming. Therefore, at least eleven of twelve districts in this analysis are paying very similar if not identical rates for energy which are also exceptionally low and do not vary from year to year. These low costs likely have little impact on districts' pumping decisions. Because energy costs are low, nearly ubiquitous (with the possible exception of the one district mentioned above), and unchanging, they have been omitted from the analysis presented in the next chapter.

Finally, climate behavior in this analysis is captured through the Standardized Precipitation-Evapotranspiration Index (SPEI). The SPEI is a publicly available data product developed by Santiago Begueria, Borja Latorre, Fergus Reig, and Sergio Vicente-Serrano under the auspices of Spain's Climatology and Climate Services Research Center (Global SPEI Database, 2021). The SPEI is a statistically robust multi-scalar drought index which includes measures of precipitation and temperature in the form of potential evapotranspiration (McEvoy et al., 2012). The index is calculated at different time scales and standardized, returning values between -3 and 3, corresponding to hotter, drier periods and cooler, wetter periods respectively. Short time scales (between 3 and 6 months) are most efficient for assessing drought conditions related to soil water content, and so this analysis will use SPEI data standardized over a moving three month period (Global SPEI Database, 2021). These data are further compiled into variables representing own-year average, prior-year average, winter average, growing season average, and prior rainy season average. The exploration of these various weather measurements is detailed in Appendix A.3. Spatially, SPEI data is available through the Global SPEI Database (maintained by the group listed above), at a pixel scale of 0.5° squared. This is fortunate for the purposes of this analysis, as the Global SPEI Database's spatial scale roughly corresponds to the boundaries of Maricopa and Pinal counties (Global SPEI Database, 2021). Model specifications presented in the next chapter will consider SPEI in year $t-1$ when predicting water deliveries in year t, as a grower would have no way of knowing climate patterns in year t when making cropping decisions at the outset of the growing season.

Figure 11 illustrates SPEI measures for Maricopa and Pinal Counties throughout the study period. 2007 is included as a lagged measure will be employed in the following chapter. The measures generally follow each other throughout the study period, but with significant differences in some years. Pinal County tends to be cooler than Maricopa, with the exception of 2007-08 and 2015. This could be attributable to monsoon patterns, changes in elevation, and/or the growing urban heat island effect around the Phoenix metro area. In many years, Maricopa County is considerably hotter and drier than its neighbor to the South. There are a handful of observations with positive values (indicating a year that is cooler and wetter than normal), but the study area is typically hotter and drier than average in these years. Notable in Figure 11 is a semi-oscillating pattern of hotter, drier years followed by cooler, wetter years in each county. This oscillation does not occur 100% of the time, but is seen in 18 of the 22 year-to-year changes observed in the study period. Figures 12 & 13 illustrate annual precipitation and temperature in Maricopa and Pinal Counties, measured in cumulative inches of rainfall and number of days with a high temperature at or above

100 degrees Fahrenheit respectively. An inspection of Figures 12 & 13 shows a recurring pattern of oscillating values in many years. This lends credibility to the behavior observed in Figure 11. It is possible that, over a small sample of years, this oscillating behavior might confound coefficient estimates obtained through regression analysis. Results obtained in the next chapter relating to the effects of prior-year SPEI on outcome variables will therefore be treated with additional scrutiny.

5 Models and Results

Initial statistical analysis is conducted using simple Ordinary Least Squares (OLS) regressions, which include much of the data described in Chapter 4. These OLS regressions provide an important foundation for this research, but standard robustness checks indicate that the variances of error terms returned by these regressions are potentially heteroskedastic. This finding is not surprising, being that different irrigation districts vary in policy structure, physical infrastructure, and constituent composition (as described in Chapter 2). Descriptive analysis of Cropland Data Layer (CDL) variables also leads to some concern regarding systemic measurement error in the remote sensed data. In order to compensate for these issues, the final statistical models presented include variables demeaned at the district level to account for district level fixed effects, as well as robust standard errors to compensate for possible systemic measurement error occuring in the CDL data.

This chapter explains the approach outlined above in fine detail. After describing the basic OLS model specifications initially used, the process of implementing robustness checks is discussed, before going on to specify the final fixed effect regression models. Finally, the last section of this chapter describes and interprets the results returned by these fixed effects models.

5.1 OLS Models

Initial statistical analysis involves the simple OLS regressions described below. As mentioned in Chapter 4, regressions are run considering two dependent variables: water delivered to agriculture and irrigation intensity. The deliveries model is specified as such:

$$
\begin{aligned}
ln(D_{it}) = \alpha_D + \beta_1 W_{at-1} + \beta_2 Alf_t + \beta_3 Cot_t + \beta_4 CAP_t \\
+\beta_5 TrA_{it} + \beta_6 ln(FaA_{it}) + \beta_7 PlA_{it} + \beta_8 Pin_i + e_{it}
\end{aligned}
\tag{19}
$$

Here, D_{it} represents to deliveries to agriculture in district i in year t. Recall from Chapter 4 that the natural log operator is applied to water deliveries due to large differences in scale and high variance of deliveries to irrigation districts. W_{at-1} represents the annual average Standardized Precipitation-Evapotranspiration Index (SPEI) metric in AMA a in year $t-1$. As mentioned in Chapter 4, the prior year SPEI is the preferred explanatory variable for weather, as a grower would have no way of observing weather patterns in year t when choosing what crops to plant and in what quantity at the beginning of a growing season. The process of evaluating alternative weather variables is described in detail in Appendix A.3. Alf_t, Cot_t, and CAP_t respectively represent real alfalfa prices, December Futures prices for cotton, and Central Arizona Project (CAP) water costs in year t. Again recall from Chapter 4, all price variables are adjusted for inflation and presented in 2019 dollars. TrA_{it} and FaA_{it} represent area planted in tree crops and fallowed area in irrigation district i in year t as reported by the CDL, with the natural log of fallowed area employed to account for large differences in the scale of irrigation districts. PlA_{it} represents all planted area in irrigation district i in year t. This is included as a control variable to account for extensive decisions made at the beginning of the growing season. As Appendix A.2 shows, extensive decision making regarding crop mix and acres planted drives most (but, crucially, not *all*) water delivery decision making, and so this variable is included to compensate for this effect. Pin_i is a simple binary variable describing whether irrigation district i belongs to the Pinal AMA or not. Recall that there are only two AMAs included in the study area, Pinal and Phoenix. Therefore, Pin_i is equal to zero for those districts located in the Phoenix AMA and equal to one for those located in the Pinal AMA. Finally, an error term (e_{it}) is included.

The irrigation intensity model presented below is similar to the water deliveries model in many aspects, but includes a handful of crucial differences.

$$I_{it} = \alpha_I + \delta_1 W_{at-1} + \delta_2 Alf_t + \delta_3 Cot_t + \delta_4 CAP_t + \delta 5 TrP_{it} + \delta_6 FtP_{it} + \delta_7 Pin_i + e_{it} \quad (20)$$

The most obvious difference is the dependent variable (I_{it}) representing irrigation intensity in irrigation district i in year t. As discussed in Chapter 4, irrigation intensity is the measure of water delivered to district i per planted area in district i in year t. The weather measure (W_{at-1}) and the three price variables (Alf_t, Cot_t, and CAP_t) remain unchanged from the water deliveries model. The Pinal AMA binary variable (Pin_i) likewise remains unchanged. Another major difference between the two models is in the inclusion of measures of planted area. While the water deliveries model includes overall area measures for tree crops and fallowed lands in order to account for the effect these areas have in districts of varying sizes, the irrigation intensity models include ratios of these areas compared to total planted area. TrP_{it} represents the percentage of total planted area in district i in year t planted in tree crops, while FtP_{it} represents the ratio of fallowed lands to planted lands in district i in year t. One characteristic of note is that, while area planted in tree crops can never exceed total planted area, no such restriction exists regarding area left fallow, meaning that while TrP_{it} is bound between zero and one, FtP_{it} is under no such constraint. Because planted area is used in calculating the dependent variable, it is omitted from this specification in order to avoid any issues with endogeneity. Finally, e_{it} is once again included as an error term.

The results of these OLS models are presented in Table 5. In both models, alfalfa prices and cotton futures are seen to be having a significant impact on water use. The negative coefficient estimates interacting with alfalfa prices in both models are somewhat confounding, as the law of demand would generally dictate that a higher output price should lead to increased input use. This may be due to the nature of alfalfa planting described in Chapter 1, in which alfalfa stands have a productive life of three to seven years. Alfalfa planting may be less subject to current prices than purely annual crops, as stands can typically exist for up to seven years. Cotton futures' impact on both water deliveries and irrigation intensity fall in line with standard economic expectations: as prices rise, so too does demand for an essential production input. CAP water costs are interacted with a significant and negative coefficient estimate in both models, which also falls in line with standard economic expectations.

Lagged SPEI is seen to have a significant effect on irrigation intensity, but not on overall water deliveries to agriculture. The lack of significance in the water deliveries specification could be the result of weather having little effect on growers' planting decisions, as farm operations in Pinal and Maricopa Counties likely expect lots of heat and little rainfall for most of the growing season, and plan accordingly. Another possibility is that the Pinal AMA variable may be picking up some of the effect of weather, as SPEI is recorded at roughly the county level. This possibility is explored further when robustness checks are applied. The coefficient estimate associated with lagged SPEI in the irrigation intensity model is significant, but, similar to the alfalfa price variable, is of a sign that runs contrary to expectations. Recall that the SPEI index returns negative values for time periods that are hotter and drier than normal. The statistically significant and positive coefficient tells the story that growers in the study area apply irrigation water less intensely following a hot, dry, year. This confounding observed effect could be the result of the oscillating behavior of the SPEI variable over the short study period, as discussed in Chapter 4.

Area planted in tree crops is seen to have a significant impact on both water deliveries and irrigation intensity. This is perhaps due to tree crops' water needs differing from field crops, including irrigation water being required year round and some base level of water required in order to maintain healthy orchards even when growers don't intend to produce a crop. The significant effects observed may also be due to the tree crop variables "picking up" the effects of other irrigation district characteristics. Tree crop planting may change more slowly than other land cover variables modeled in these specifications, as orchards are expensive to plant and remove. If the area planted in tree crops in some irrigation district remains relatively static throughout the study period, than this variable has the potential to capture the impact of other district-level characteristics. Once again, this possibility will be explored further once robustness checks are applied.

Fallowed area is seen to have a significant impact on both water deliveries and irrigation intensity, although the observed effects are once again are somewhat confounding to interpret. A large fallowed area significantly increases irrigation water deliveries, while a high fallow-to-planted ratio significantly reduces irrigation intensity. These effects are vexing, as it does not seem unreasonable to expect a district with more fallowed land would require less water to be delivered, or that more lands left fallow relative to planted area would provide the opportunity to intensify irrigation. Planted area is seen to have a highly significant and positive effect on water deliveries. Like cotton futures prices, this is very much in line with intuitive expectation.

Finally, districts in the Pinal AMA are shown to have significantly lower baseline expectations for both water deliveries and irrigation intensity. The lower intensity may be due to the slightly cooler climate of Pinal County, and the generally larger scale of irrigation districts within the Pinal AMA. However, it is strange that these large-scale irrigation districts would generally require fewer deliveries overall. As mentioned above, it is possible that the AMA binary variable is picking up some effect of weather, and this will be investigated further in the coming sections.

Some of the results described in the preceding paragraphs are a bit confounding. Due to this, as well as the varied nature of Arizona irrigation districts, robustness checks must be applied.

5.2 Robustness Checks

As mentioned in the opening paragraphs of this chapter, the irrigation districts within the study area structurally differ in many ways. It is for this reason that the OLS regressions outlined above are likely to generate heteroskedastic error terms. In this section, checks for heteroskedasticity are reported. Year-to-year change in the tree crops variable will also be evaluated in an effort to determine whether tree crop area may be acting as an irrigation district fixed effect in the agricultural deliveries model.

The first step in assessing heteroskedasticity is a visual inspection of the error terms returned by the OLS regression models plotted against the fitted values of the dependent variables. Figures 14 & 15 show these results for deliveries to agriculture and irrigation intensity respectively. A visual inspection of Figure 14 seems to show clear signs of clustering, with groupings occurring when fitted values are between 11 and 12, and greater than approximately 12.5. On the other hand, Figure 15 does not immediately show evidence of heteroskedasticity. It could well be the case that the structural differences between irrigation districts do not affect irrigation intensity to the same degree that they do water deliveries.

The next step involves performing statistical tests on the regression outputs themselves. Bruesch-Pagan tests are run on both models to determine the likelihood of heteroskedastic error terms. The Breusch-Pagan test for heteroskedasticity compares fitted values to residual errors in an effort to validate a null hypothesis of constant variance. If this null hypothesis can not be validated, it is likely that residual error terms are heteroskedastic. The results of these checks are presented in Table 6. What was apparently evident in Figures 14 & 15 is now statistically quantified by these results. The agricultural deliveries model exhibits a relatively high χ^2 value of 2.66, corresponding with a p-value of 0.1028. This means the null hypothesis can be rejected with some confidence, but 0.1036 is close to the typical threshold of 0.05 and extremely close to a slightly less stringent threshold of 0.1. For this reason, although the test would suggest controlling for heteroskedasticity could be foregone, a fixed-effects regression will be run for deliveries to agriculture. On the other hand, the small χ^2 value associated with the irrigation intensity model (0.02), and its corresponding p-value of 0.8918, indicate that the null hypothesis of constant variance can be rejected with confidence. This is once again consistent with a visual inspection of Figure 15. For this reason, one could accept the OLS intensity model results as sound, although for the sake of completeness, a fixed-effects regression will be run for irrigation intensity as well.

Rather than include a binary variable for each irrigation district in the study, fixed effects models will instead demean variables at the district level in order to preserve as many degrees of freedom as possible. This process involves subtracting the average value of a given variable in a district from each individual observation recorded for said district. This demeaning process works mathematically in the same way as including a set of binary variables, but allows for increased degrees of freedom and avoids potential "overfitting" due to dummy variables. The Pinal AMA binary variable will be dismissed in these specifications. This serves two purposes. First, demeaning the Pinal AMA variable at the district level will result in zero values being returned for all observations, as this is a fixed, district-level characteristic. Secondly, dropping the Pinal AMA variable will address the potential issue of the AMA binary capturing some effect of weather. SPEI will no longer be directly tied to an existing binary variable.

The next robustness check to perform is an evaluation of the tree crop variable, specifically focused on whether planting of trees in each irrigation district varies widely from year to year. Figure 16 charts area planted in tree crops for each irrigation district throughout the study period. Two notable trends seem to stand out. Firstly, for many of the irrigation districts with relatively little area devoted to tree crops, the area measure seems to change very little from year to year. Figure 17 highlights this further by omitting the two irrigation districts with the most tree crops planted (Central Arizona Irrigation and Drainage District and Maricopa-Stanfield Irrigation and Drainage District). Because some districts show so little variation in area planted in trees, this variable will be excluded from any fixed-effect models, due to the potential for collinear relationships affecting results. As a result, tree crops will be considered a fixed characteristic of irrigation districts for the purposes of this study. The true effect of tree crops on agricultural water use in Central Arizona is therefore a subject ripe for further research.

The second discernible trend is more concerning. Visual inspections of Figures 16 & 17 clearly show systemic peaks and troughs in the remote-sensed land cover data. Most notable is the large spike across almost all districts in 2010, but also potentially evident are a peak in 2014 and a trough in 2016. While tree crops do tend to be planted in batches, it seems unlikely that many districts would see a huge spike in planted acreage in 2010 followed by a huge reduction the following year. Instead, these peaks and troughs could be due to ongoing refinement of the remote-sensed data collection process, either at the level of photographic collection of data itself, or the algorithmic processing of said data (Willis, 2021). The first question that must be answered is whether these systemic patterns are specific to the tree crop variable, or if they are seen across all crops planted in the study area. Figure 18 shows all planted area by irrigation district throughout the study period. When charted, the planted areas appear much smoother than the areas planted in tree crops, but evidence of the same trends described above can still be seen in 2010, 2014, and 2016. Because irrigation intensity is determined by dividing water deliveries by planted area, these CDL remote-sensed variables cannot be dismissed completely without greatly reducing the scope of this research. Instead, all fixed-effects regressions will employ robust standard errors in order to account for the potential of systemic perturbations in the remote-sensed data.

5.3 Fixed-Effects Models

With these robustness checks in mind, presented below are the fixed effects models which will be analyzed in this research. Beginning with water deliveries to agriculture, all variables described in the OLS model specifications are included, with the exception of area planted in trees and the Pinal AMA binary variable.

$$ln(D_{it}) = \alpha_D + \beta_1 W_{at-1} + \beta_2 Alf_t + \beta_3 Cot_t + \beta_4 CAP_t + \beta_5 ln(FaA_{it}) + \beta_6 PlA_{it} + e_{it} \quad (21)$$

Here, shorthand for any variable specified in the OLS model remains unchanged. As described in the previous section, two additions to this model are irrigation district level fixed effects through demeaning and the inclusion of robust standard errors.

The model of irrigation intensity is much the same as the OLS specification, only with percentage of planted area represented by trees and the Pinal AMA binary variable having been omitted. Here again irrigation district fixed effects area accounted for through demeaning and robust standard errors are employed. The model is formally specified below.

$$I_{it} = \alpha_I + \delta_1 W_{at-1} + \delta_2 Alf_t + \delta_3 Cot_t + \delta_4 CAP_t + \delta_5 FtP_{it} + e_{it} \quad (22)$$

These fixed effect models R^2 measures will be reported in two ways: R^2 Within and R^2 Between. R^2 Within reports how much variation in the dependent variable within an irrigation district over time is accounted for. R^2 Between reports how much variation between irrigation districts is captured by the model. R^2 Within is typically of greater interest, as the application of fixed effects is due to skepticism over a model's ability to otherwise explain variation in the dependent variable between districts. With these fixed-effects regression models fully specified, and having accounted for heteroskedasticity and systemic variation in CDL remote-sensed data, the coefficient estimates produced will now be discussed.

5.4 Results

With the effects of individual irrigation districts accounted for and additional complications arising from the land cover variables' systemic error addressed, the results of the fixed effects model can now be interpreted. Table 7 presents coefficient estimates from both the water deliveries and irrigation intensity models. Compared to the results from the OLS regressions, the fixed effect models' report less confounding significant coefficient estimates in many cases. The fixed effect models' R^2 measures have decreased compared to their OLS counterparts. This is likely due to the omission of the tree crops and Pinal AMA indicator variables, which almost certainly capture the effect of other district-level fixed characteristics. Therefore, it can be assumed that these fixed effect models' R^2 measures are more precise. The water deliveries model returns an R^2 Within of 0.37 and an R^2 Between of 0.748, meaning the model captures the effect of the explanatory variables on variation in water deliveries between districts more efficiently than it does the effect of the explanatory variables on variation in water deliveries within the individual districts over time. The irrigation intensity model flips this dynamic, with an R^2 Within of 0.4287 and a vanishingly small R^2 Between of 0.0034, meaning this model has almost no ability to capture the effect of the explanatory variables on variation in irrigation intensity between districts. Taken in their entirety, the fixed-effect models seem to have improved on the OLS modelling approach. The rest of this section will be spent discussing the specific coefficient estimates returned, and the implications of these results.

The lagged SPEI weather metric is found to have a highly significant effect in both fixed-effect models. This finding differs from the results returned by the OLS modelling, and is likely due to the Pinal AMA binary variable having been omitted. The positive sign on both coefficient estimates is somewhat unexpected. Recall from Chapter 4 that the SPEI measure ranges from -3 to 3, with negative values meaning a hotter, drier year. This means that, when interacted with the highly significant coefficient estimate in either model, a hot and dry year $t-1$ indicates lower water deliveries in year t. Visual inspection of Figure 11 seems to show hotter, drier years and cooler, wetter years alternating throughout the study period. This alternating pattern is likely driving the positive sign of the coefficient estimate, and is not necessarily a reliable trend. For this reason it would be useful to investigate the effect of SPEI over a longer study period, once more remote-sensed data become available.

The cotton futures price variable returns statistically significant values for both models and CAP water costs are seen to have a significant impact on irrigation intensity, with the signs of all significant coefficient estimates conforming to standard economic expectations. Higher prices of cotton futures are seen to increase water deliveries to agriculture and irrigation intensity, while higher prices for CAP water are seen to decrease irrigation intensity. It is no surprise that planted area's impact on water deliveries to agriculture likewise conforms to expectations. The fixed effect model returns a highly significant positive coefficient estimate for this variable. Recall that planted area is included as a definitional control variable and any other coefficient behavior would certainly be alarming.

Just as was seen in the OLS model results, alfalfa prices return negative coefficient estimates, which run contrary to economic intuition. Happily, this confounding behavior is not found to be significant in the fixed effect models, meaning these negative values may be dismissed out of hand. This lack of significance could be explained by alfalfa's semi-annual/semi-perennial nature. Because stands may be productive for many years, water use decisions with respect to this crop may be less responsive to changes in alfalfa price in any one year. Another possible explanation for this coefficient estimate's confounding behavior include could be related to trends in the dairy industry. If alfalfa purchasers are concerned not only with the prices of alfalfa as an input, but also the prices of milk and other inputs necessary to produce milk, the effect of alfalfa prices on water use decisions could be diluted. Being that alfalfa is the predominant crop in the region, further research on the factors that drive alfalfa planting would likely be very valuable, as year t prices alone perhaps do not tell the entire story.

Coefficient estimates for fallowed land variables remain significant for water deliveries to agriculture in the fixed effect models. As in the OLS models, the fallow-to-planted ratio's effect on irrigation intensity is negative, but is no longer statistically significant and may be disregarded. The effect of the natural log of fallowed area on agricultural water deliveries has gone from positive in the OLS specification to negative in the fixed effect modelling. This negative effect of fallowed area on water deliveries seems to fall more in line with expectations, as it seems intuitive that more fallowed area would lead to reduced water deliveries. The lack of an effect of the fallow-to-planted ratio on irrigation intensity may be due to more efficient irrigation technology and other best practices being implemented. Recall from Figure 6 that irrigation intensity has declined in almost all districts studied over the course of the study period. If this is due to growers being able to use less water to harvest their crops, then fallowing lands would not be undertaken to make more water available to irrigate lands in production. Instead, holders of grandfathered water rights could choose to fallow lands in order to reduce overall irrigation supply and thereby accrue “flex credits”, which may be sold or saved for a not-so-rainy day.

6 Conclusion and Policy Implications

Arizona and the desert southwest are currently facing water challenges with no precedent in recorded history. Climate change in recent decades has led to hotter temperatures, reduced snowpack which melts much earlier in the year, and sporadic and undependable rainfall throughout the region. As a result, flows in the Colorado River are greatly diminished relative to recent decades, and Lake Mead's surface elevation sits at an all time low. In 2021, this depletion of Lake Mead's water reserves led, for the first time, to the federal government declaring a Tier 1 water shortage in the Colorado's Lower Basin. This shortage condition will result in reduced river water deliveries to Arizona, Nevada, and Mexico, with cutbacks in Arizona mainly affecting irrigated agriculture in the central part of the state. Colorado River water has not always been available in Central Arizona, with growers' in the region being able to meet their water needs in the past by utilizing groundwater reserves. This practice was curtailed to a large extent when the 1980 Arizona Groundwater Management Act established Active Management Areas (AMAs) in those parts of the state where groundwater reserves were most imperiled, including much of Maricopa, Pinal, and Pima counties. The restrictions imposed by the Groundwater Management Act have greatly reduced the depletion of Central Arizona's aquifers, but this outcome was made possible in part by the availability of Colorado River water to Central Arizona growers due to the completion of the Central Arizona Project (CAP) canal in the late 1980s. With the amount of water available through CAP now reduced under the Colorado's Tier 1 shortage condition, uncertainty exists as to whether or not growers will once again turn to groundwater in order to meet their irrigation needs.

Many of Arizona's policymakers and water managers have anticipated this current challenge, and some action has been taken in order to mitigate the fall out resulting from reduced Colorado River water deliveries. Water banking has been a component of water management in the state since the 1986 passage of the Underground Water Storage and Recovery Act, allowing excess CAP water to be stored in underground aquifers for some future not-so-rainy day. These banked reserves will likely play some significant part in preserving water supplies in Central Arizona if shortages in the Colorado's Lower Basin become a fixture of the 21st century. Additionally, the Groundwater Management Act has reduced the number of acres in Central Arizona available to irrigated agriculture, and the state has provided incentives for growers to follow best management practices in terms of irrigation water use (Arizona Department of Water Resources, 2021). Beyond these policies, many municipalities throughout the state have begun using treated waste water for non-potable uses. There are state level discussions regarding alternative means of augmenting Arizona's water supply, including desalination of groundwater and further developing water treatment methods to allow for potable reuse (McGreal et al., 2021).

While effectively managing and seeking to augment existing water supplies are often the primary focus of Arizona water managers, a firm understanding of the demand for water throughout the state's various economics sectors is indispensable. The research presented in the preceding chapters is intended to contribute to this understanding, specifically in terms of irrigators' demand for water in the Phoenix and Pinal AMAs. As mentioned in Chapter 1, agricultural water use accounts for 68% of total water use in Arizona, meaning that understanding the needs and motivations of the agricultural sector is a necessary part of understanding water demand in the state overall. This research chiefly contributes to developing understanding of agricultural water demand in three ways.

First, Chapter 3 presents a conceptual model for water demand enumerated at the irrigation district level. This model is constructed based on findings from the agricultural economics literature, a review of which is summarized in Chapter 2. The only regional characteristics necessary for this conceptual model to be relevant is significant dependence on irrigation in local agriculture and a non-binding constraint on water supply. The model considers the drivers of irrigation water demand in two categories: extensive and intensive decisions. Extensive decisions are what crops to plant in what quantities, and intensive decisions determine how intensely water is applied once crops are planted. Future research can continue to build upon this model in order to provide a framework for empirical analysis in various regions.

The econometric models presented in Chapter 5 explore how these extensive and intensive decisions highlighted in the conceptual model are driven by crop prices, growers' expectations regarding climate, the cost of water, and irrigation district characteristics (including regulatory policy, irrigation technology, and the extent of tree crops present). Recall that the districts included in this study account for 95% of irrigated lands in the Phoenix AMA, and 87% or irrigated lands in the Pinal AMA. This means the sample observed is extremely close to the overall population of agricultural water users in the study area over the course of the study period. The intention of Chapter 5's econometric analysis is to treat this population of Central Arizona water users as a representative sample of Central Arizona water users in future years. Interpreted this way, the results obtained can be used to predict agricultural water use in coming years using only a small set of explanatory variables. Prediction is not easy, and sampling from only the first twelve years of a time period extending into the future can certainly not be said to be random sampling. Follow up research undertaken after some time has passed would be a useful way to evaluate the results presented in Chapter 5.

This empirical work contributes to a growing literature on the incorporation of remote sensed land cover data into models of water demand. Remote sensed data products are today widely available, and commonly used in agricultural economics to provide climate data (Dell et al., 2014; Donaldson and Storeygard, 2016). However, remote sensed land cover data, such as those provided by USDA NASS's Cropland Data Layer (CDL), have been utilized less frequently. Chapter 5's empirical analysis makes use of CDL data in three primary ways: planted area in an irrigation district is included as a control variable in the water deliveries model, the dependent variable irrigation intensity is found by dividing water deliveries by planted area in the second model, and fallowed area within an irrigation district is included as an explanatory variable in both specifications. Additionally, area planted in tree crops had been included in the OLS models presented initially, but was dropped from the fixed effects models due to statistical issues with the fixed effect specifications. Incorporating the CDL's land cover data into these models allowed for greater specificity than would otherwise have been possible. As the time periods for which CDL data is available accumulate in various regions of the US, the use of such data will soon be more commonly seen in agricultural economic modelling.

This work also contributes toward understanding agricultural water demand in Central Arizona through the empirical analysis discussed in Chapter 5. This research is specific to irrigation districts in the Phoenix and Pinal AMAs, though findings may inform future research and water management efforts elsewhere in the world. Given a hotter, drier Colorado basin, curtailments to Arizona's allotment of Colorado River water are likely to occur more frequently in the coming decades. As such, awareness of the key factors which drive demand for irrigation water is valuable to anticipate Central Arizona's agricultural water demand. The analysis presented in Chapter 5 finds that December Futures prices for cotton have a significant, direct effect on total water deliveries to agriculture and irrigation intensity, while CAP water costs have a significant inverse effect on these outcome variables. Fallowed area is seen to have a negative impact on overall water deliveries to agriculture, and no significant effect in terms of irrigation intensity. This lack of a significant impact on irrigation intensity could be due to improved irrigation technology and best practices reducing the need for additional application of water for irrigation. The Standardized Precipitation-Evapotranspiration Index (SPEI) is shown to have a significant, direct effect on the outcome variables as well. This is a somewhat confounding finding, as SPEI values are negative in hot, dry years and positive in cool, wet years. As discussed in Chapter 5, this direct effect could be the result of the pattern of oscillating cool and warm years seen in Figure 11. This trend is most likely random, and a study over a longer time period could corroborate or falsify the effect of SPEI returned by the empirical analysis in this work. Recall that the time frame of this study is limited by the number of years CDL data is available. Prices for alfalfa, the region's most widely grown crop, are not seen to have any statistically significant impact on either outcome variable once district-level fixed effects are accounted for. As mentioned in Chapter 5, this could be due to alfalfa stands being less responsive to year-to-year price fluctuations due to their semi-perennial nature, and certainly warrants further research.

In addition to alfalfa's role in driving water use decisions, several other aspects of Chapter 5's empirical analysis provide areas for future research. Firstly, as discussed above, the confounding signs associated with the SPEI index merit further scrutiny. Next, the question of tree crops true effect on water deliveries and irrigation intensity remains unsolved, as collinearity issues meant that the area of an irrigation district planted in tree crops had to be omitted from the final fixed effect specifications. As mentioned in Chapter 1, pecan orchards have much higher water requirements than the crops considered in this study's empirical models (Sammis and Herrera, 1999). Given the prevalence of pecan trees in Arizona, with over 17,000 acres planted in 2017, the effect of orchards on agricultural water use is very much of interest and will require further study (Lahmers et al., 2018).

Another question this work does not address is whether or not the effects observed in Chapter 5's empirical analysis would be consistent in irrigation districts outside of Central Arizona's AMAs. Throughout the state, many irrigation districts exist outside the AMA system, with notable clusters in and around Yuma and Safford (Arizona Department of Water Resources GIS Data, 2021). Being outside of any AMA, these districts are not required to report their water use to the Arizona Department of Water Resources (ADWR), nor are they constrained in their groundwater withdrawals. This lack of water use reports presents a major obstacle for anyone interested in comparing the effects seen in this work to those districts outside the AMAs. While challenging, this problem is not insurmountable, and research focused on water use decisions outside the AMAs would certainly be valuable.

When parsing the implications of this study's findings, one must bear in mind the looming cuts to Central Arizona agriculture's allotment of Colorado River water. Tables 5 & 6 show that, while irrigation intensity generally declined between 2008 and 2019, overall water deliveries did not. How will growers respond when these CAP deliveries are significantly curtailed? Will growers who hold "flex credits" supplement their water needs through additional groundwater pumping, and if so, will large investments be necessary to improve well infrastructure? Water managers at ADWR, CAP, and elsewhere will likely seek to project demand for irrigation water under these shortages, to either assess economic losses due foregone planting, or to assess groundwater depletion due to growers turning to pumping to offset reduced CAP deliveries. We have seen that cotton futures prices have a significant direct effect on water use outcomes and that fallowed area has a significant inverse effect on overall water deliveries, meaning both might be good barometers of growers' water demand at the outset of a season. Conversely, we have seen that CAP water prices act as a disincentive to increased irrigation water use. This responsiveness to cost could be critical if water managers must decide on a mechanism to equitably distribute reduced CAP allocations.

The American West has a water problem. Arizona has a water problem. The state's development over the last hundred years was made possible largely through development of non-renewable groundwater resources. More recently, the CAP allowed for the Colorado River to ease some of the strain on Central Arizona's aquifers and provide water, not only to agriculture, but to the growing population centers located along the I-10 corridor. For growers in Central Arizona, this Colorado River water may not be available for much longer, as drought conditions throughout the already arid region are intensified as a result of anthropogenic climate change. For water managers and policymakers seeking to curate and augment existing water supplies, understanding the drivers of water demand is now more crucial than ever. The hope for this research is that it can contribute in some small way to that goal.

Figures and Tables

Figure 1: Irrigation districts in the Phoenix and Pinal AMAs

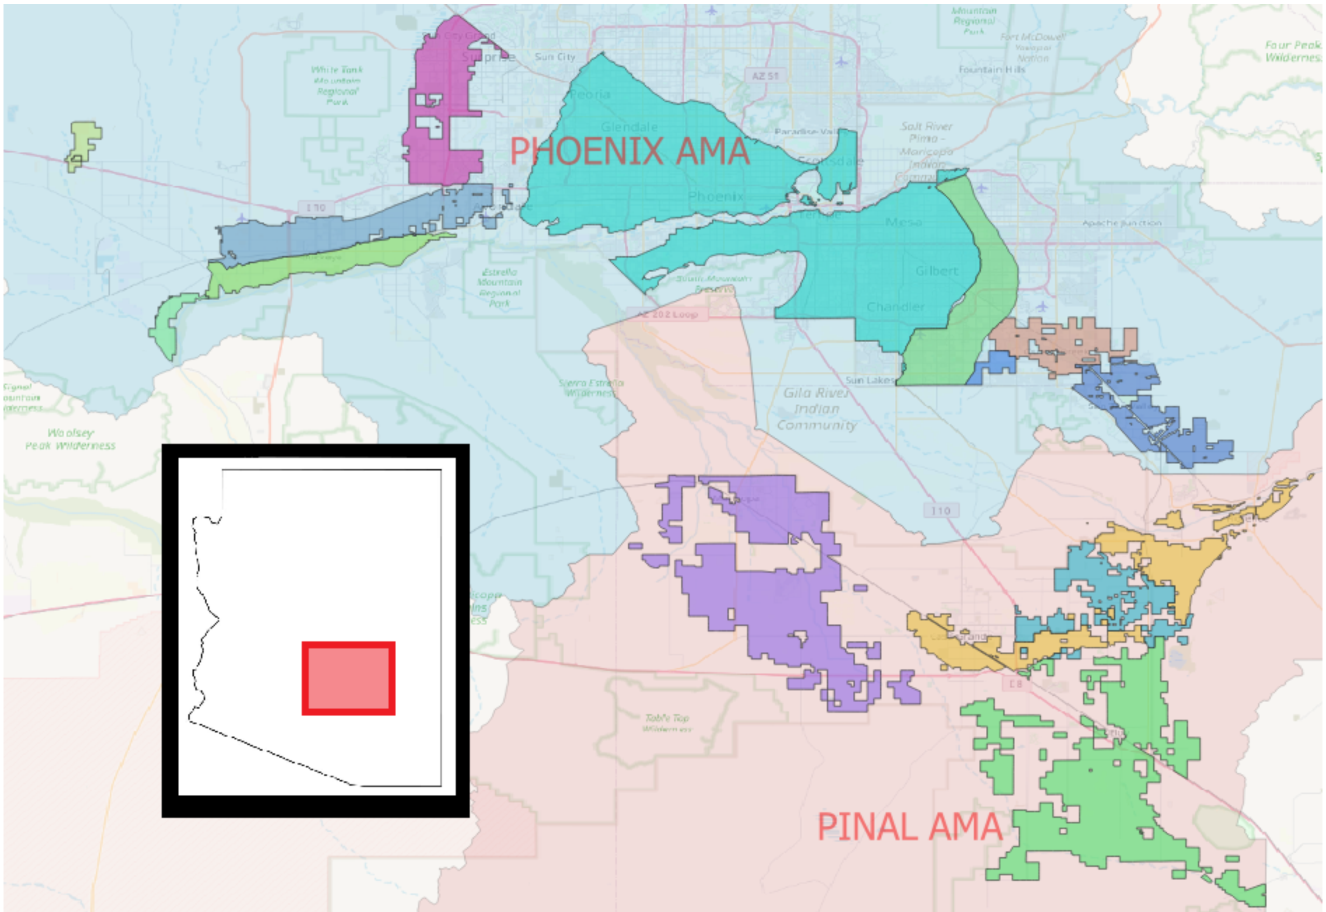

Data Source: ADWR

Figure 2: Crops by Planted Acreage in Maricopa and Pinal Counties

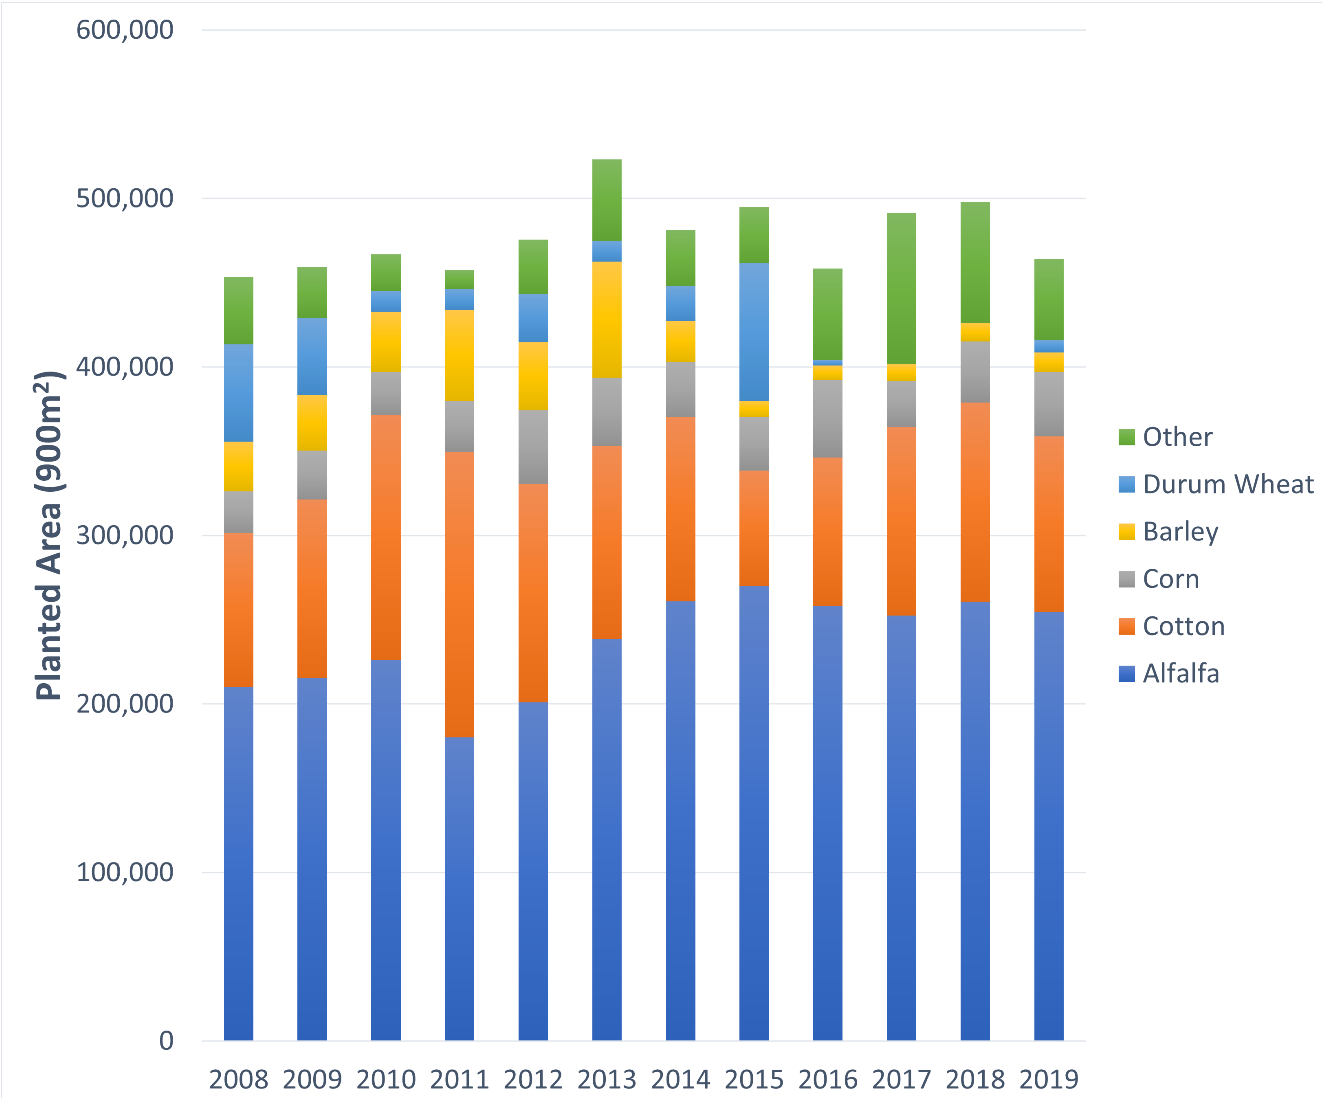

Data Source: USDA NASS Cropland Data Layer

Figure 3: Crops by Percent of Planted Acreage in Maricopa and Pinal Counties

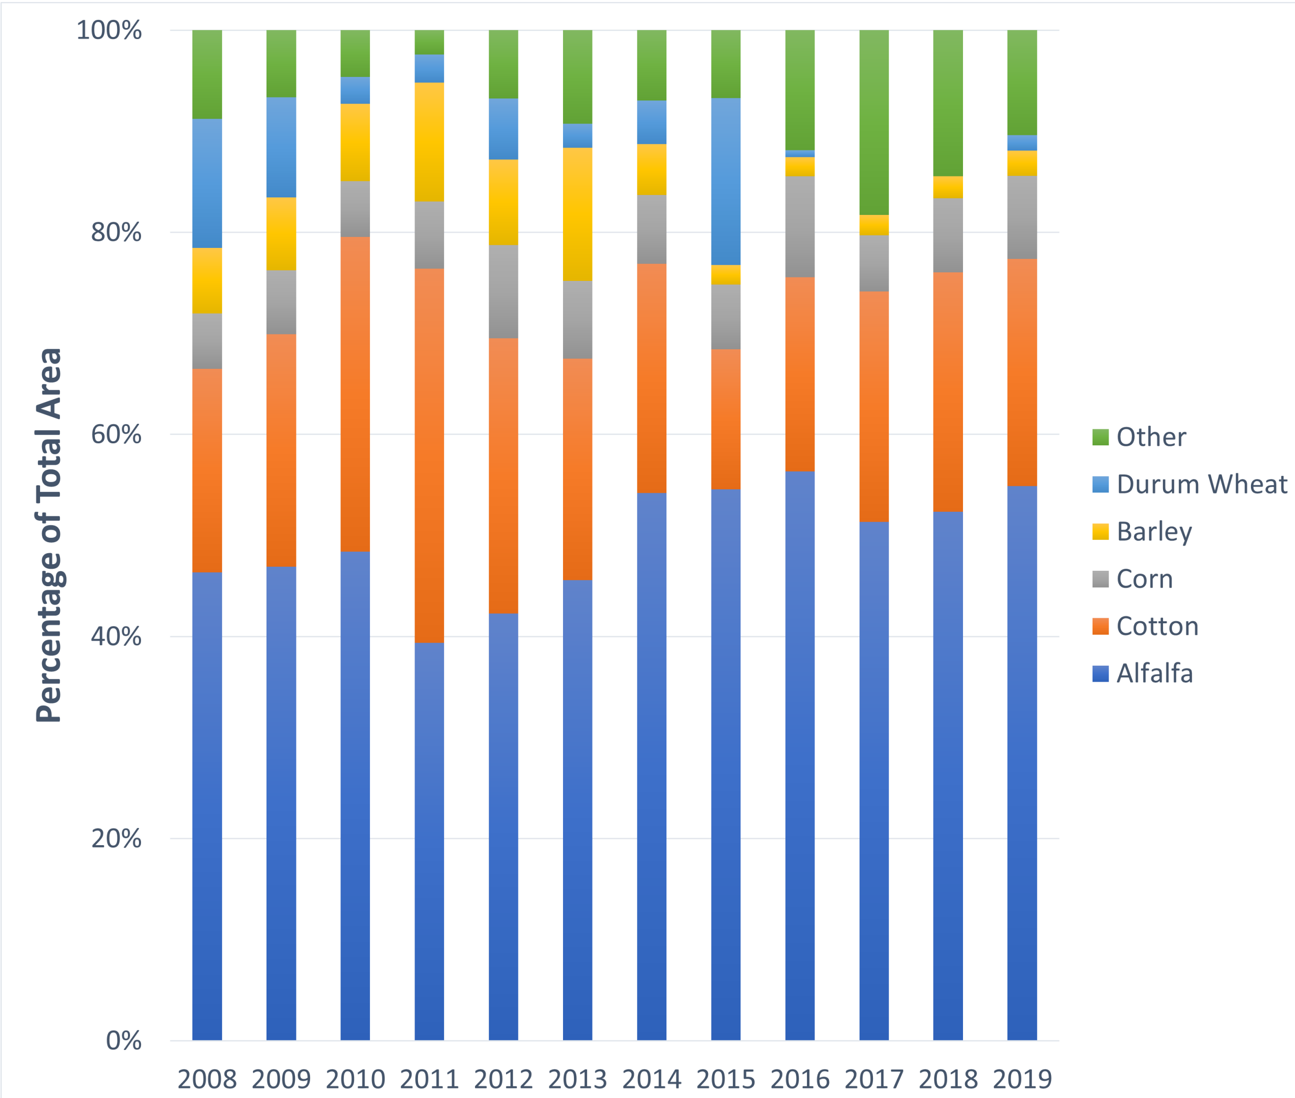

Data Source: USDA NASS Cropland Data Layer

Figure 4: Annual Average Water Deliveries to Agriculture

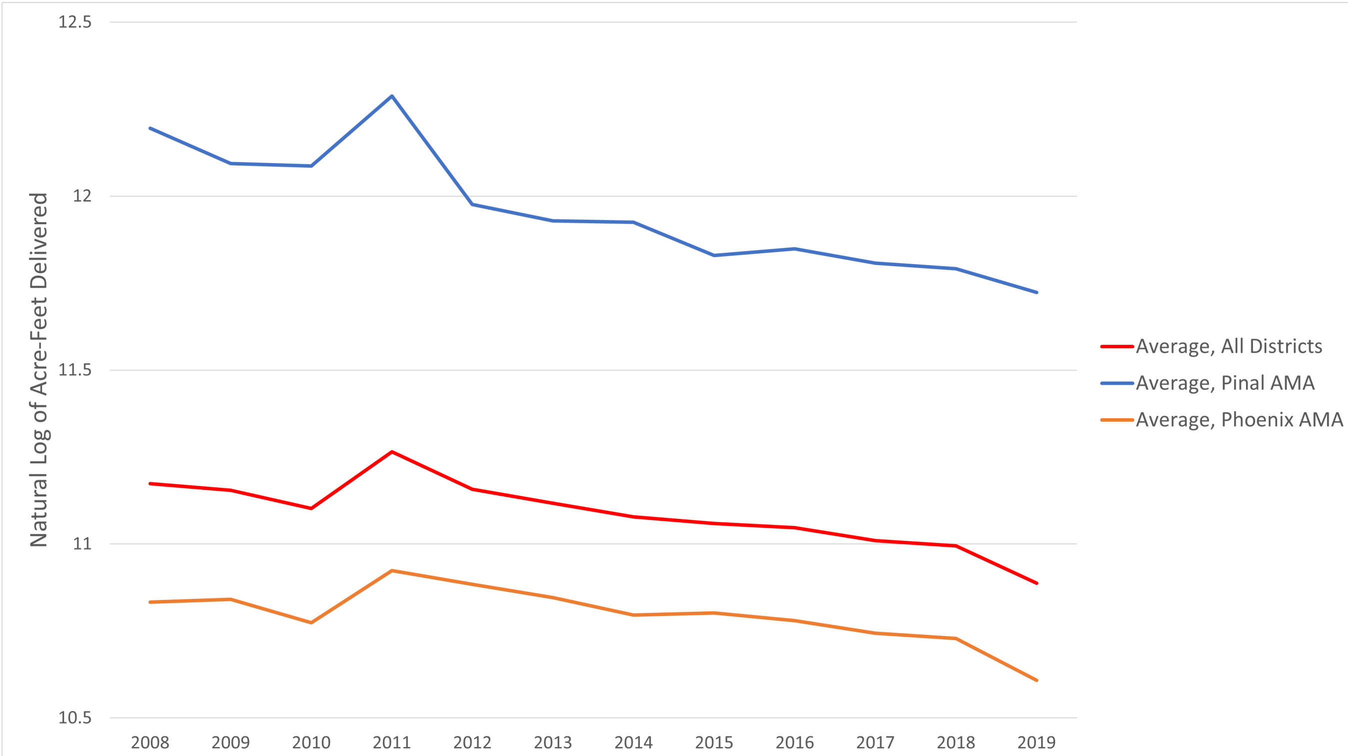

Data Source: ADWR

Figure 5: Natural Log of Water Deliveries to Agriculture by District

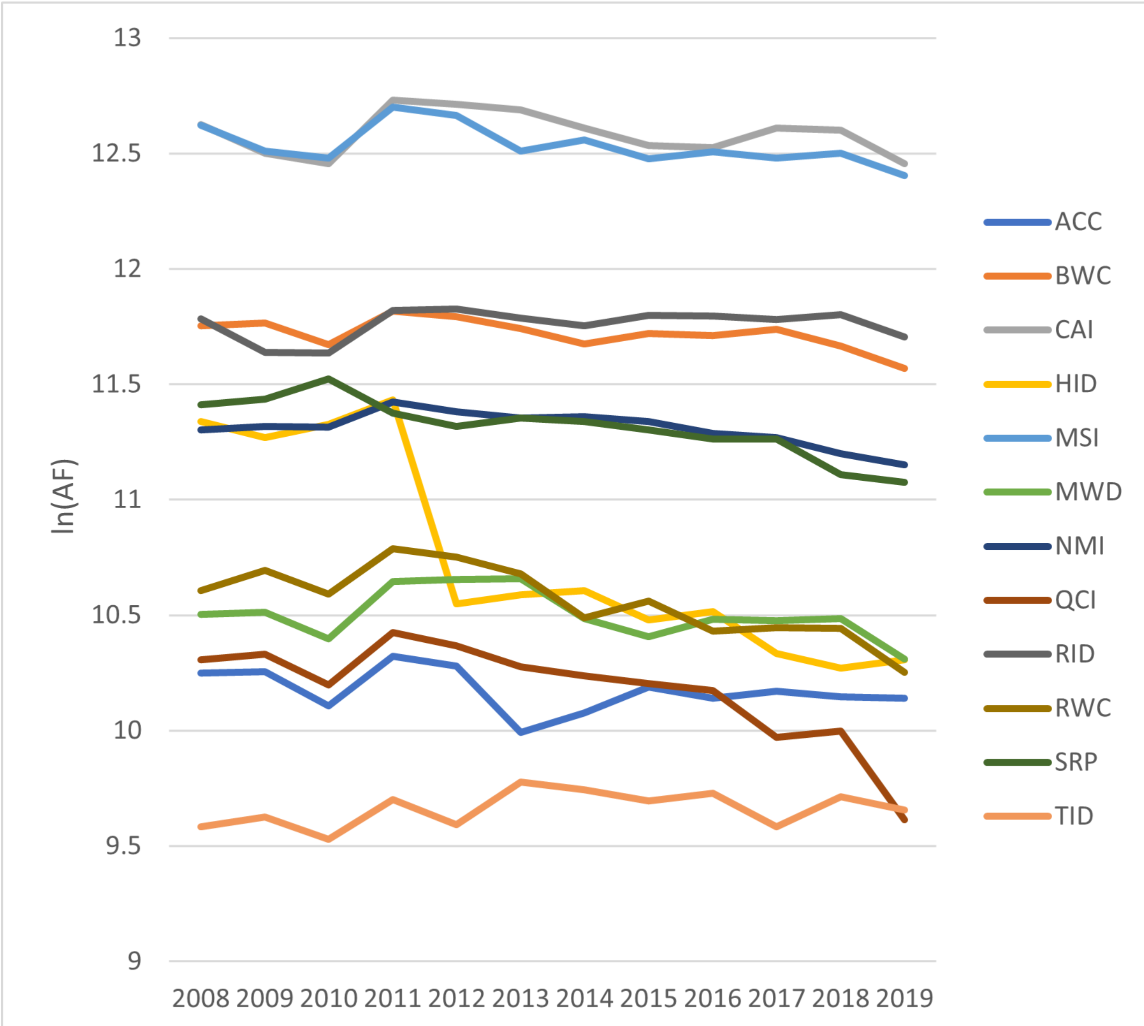

Data Source: ADWR

Figure 6: Irrigation Intensity by Irrigation District

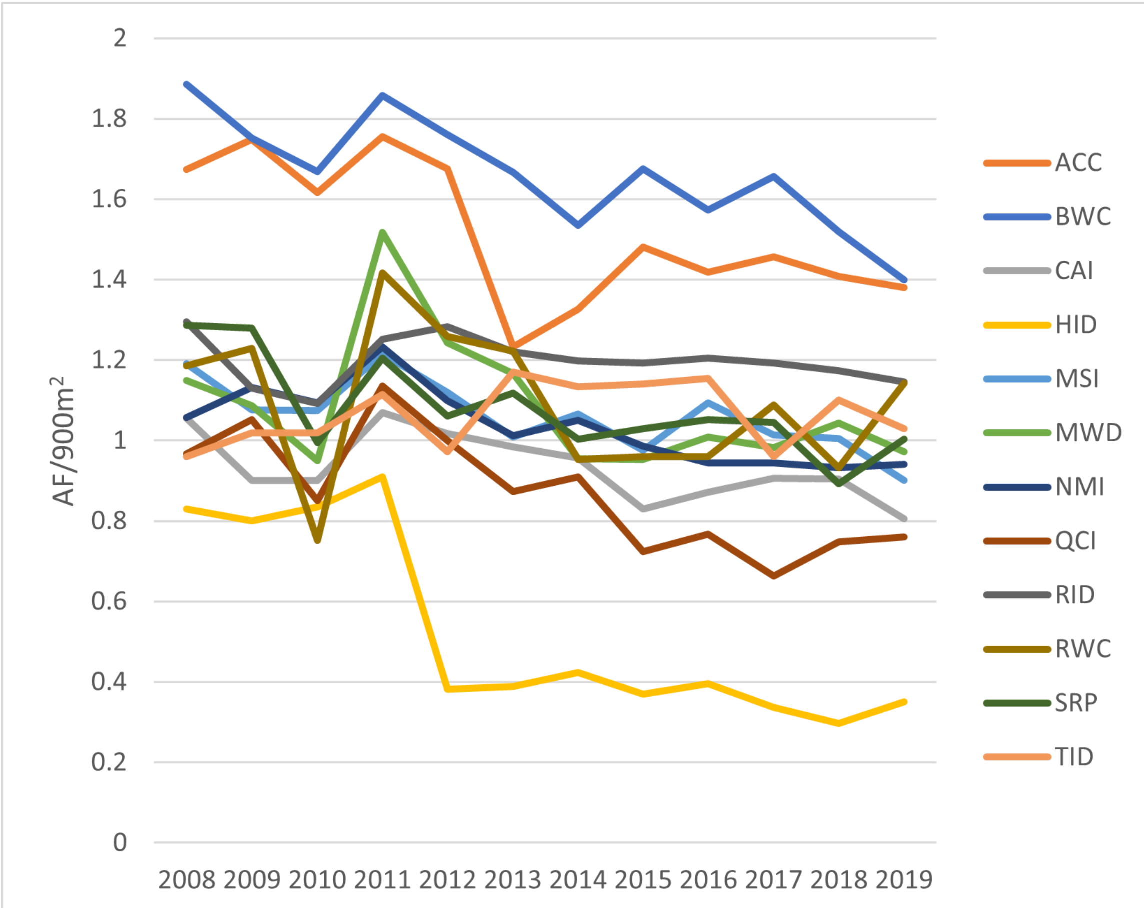

Data Source: ADWR, USDA NASS Cropland Data Layer

Figure 7: Cropland Data Layer Accuracy - Arizona

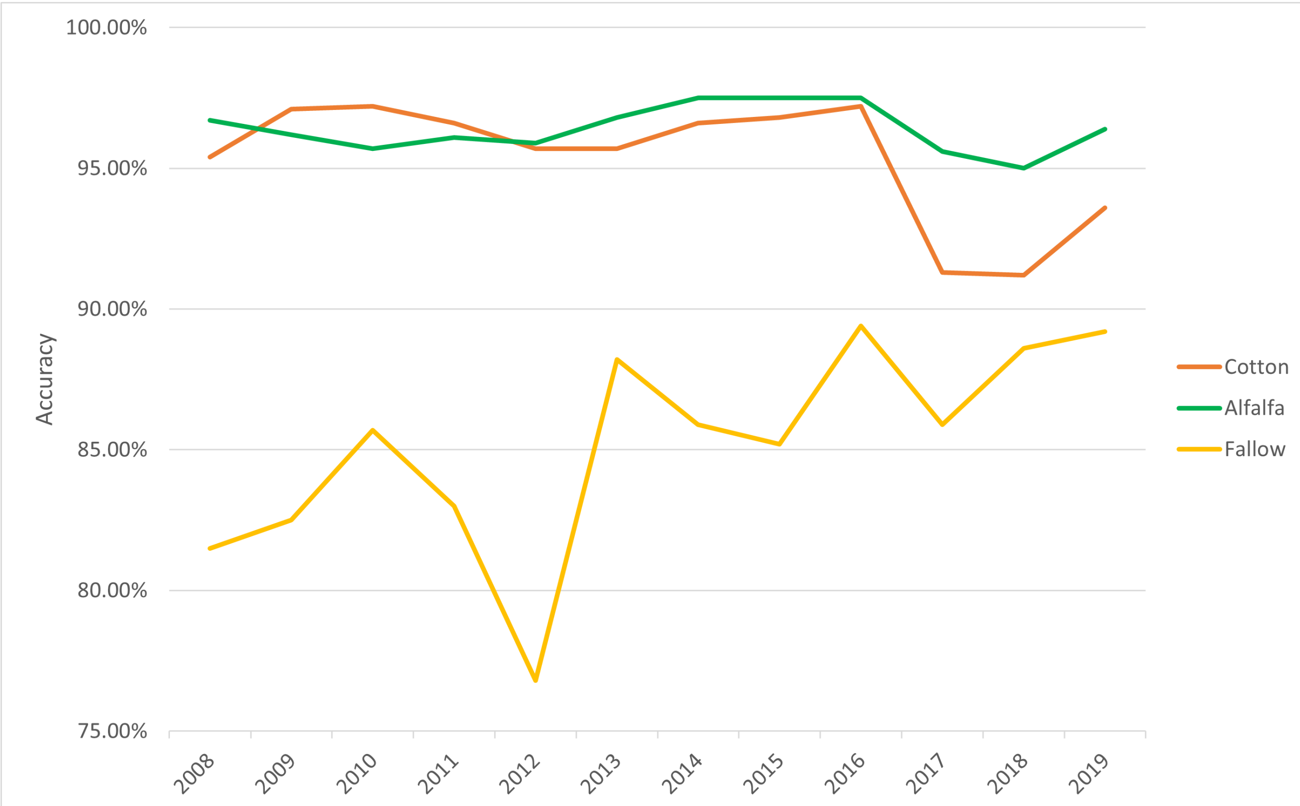

Data Source: USDA NASS Cropland Data Layer

Figure 8: Real Alfalfa Prices (AZ)

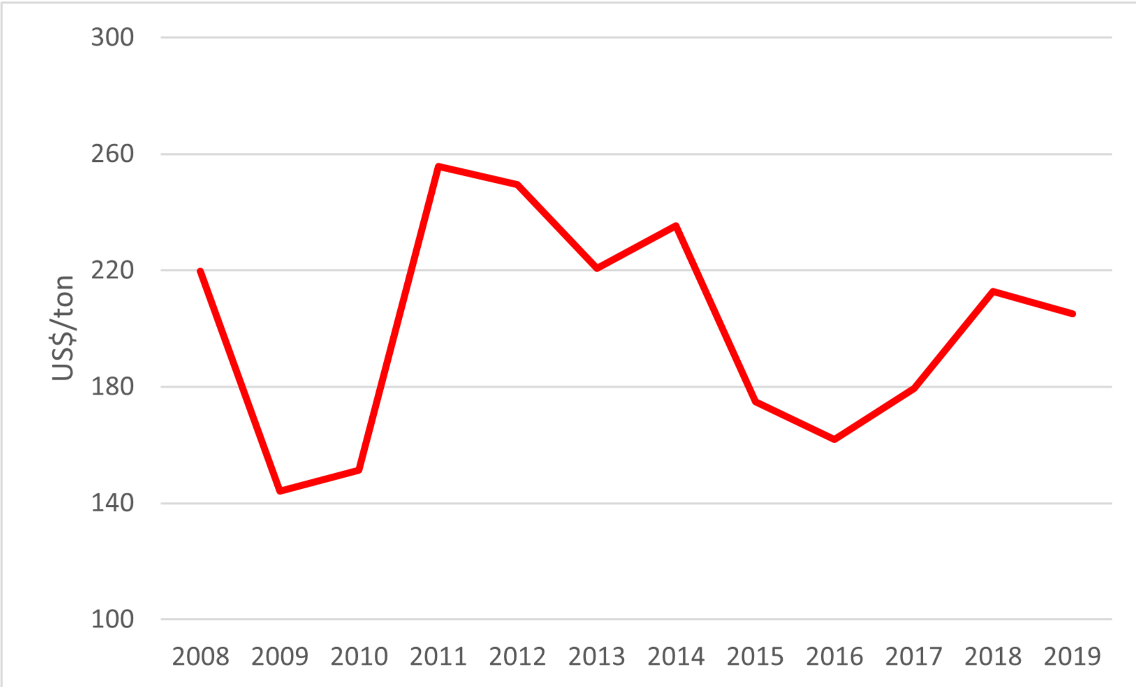

Data Source: USDA NASS

Figure 9: Real December Cotton Futures Prices

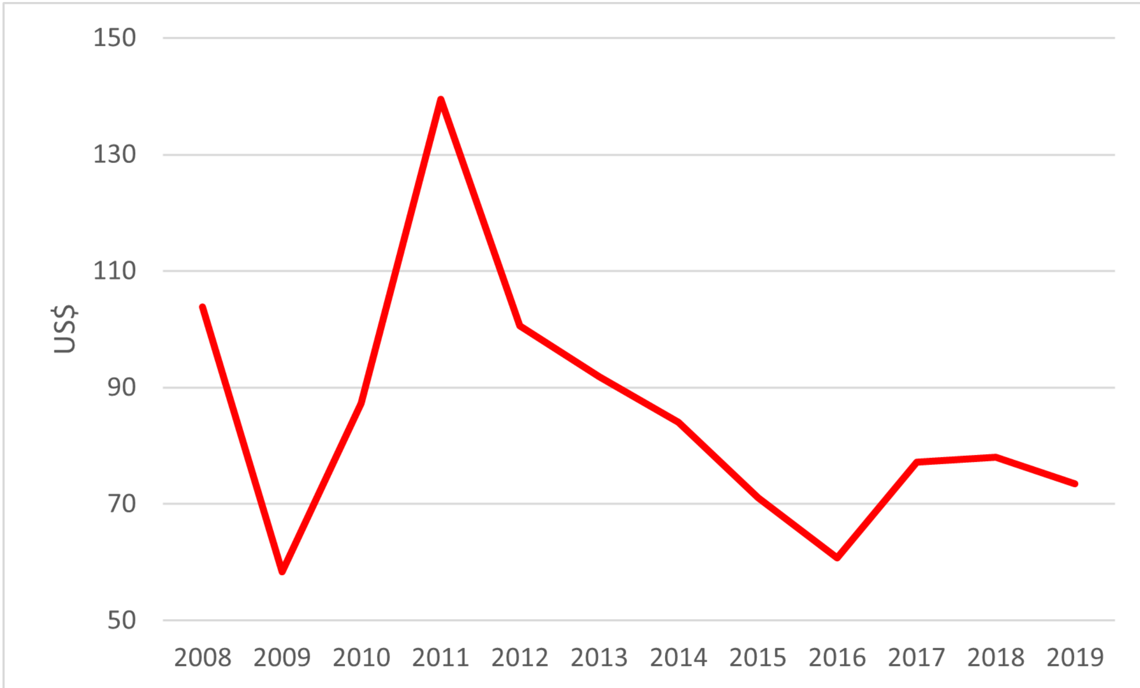

Data Source: New York Cotton Exchange

Figure 10: Real CAP Water Prices

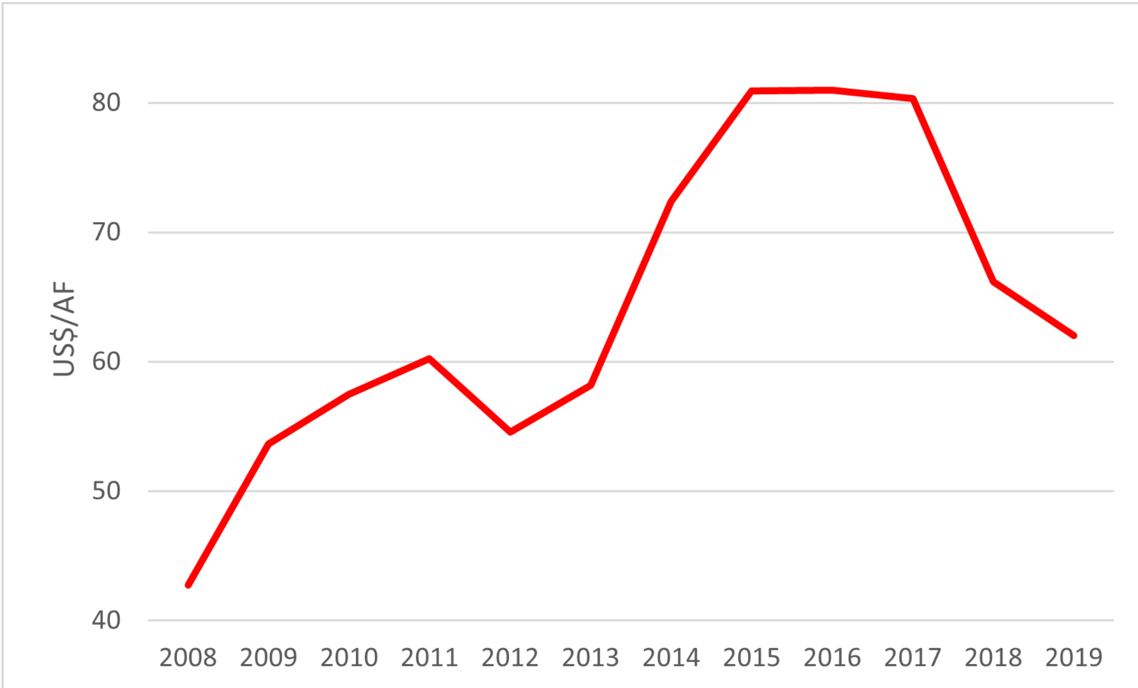

Data Source: US Bureau of Reclamation

Figure 11: Central Arizona SPEI

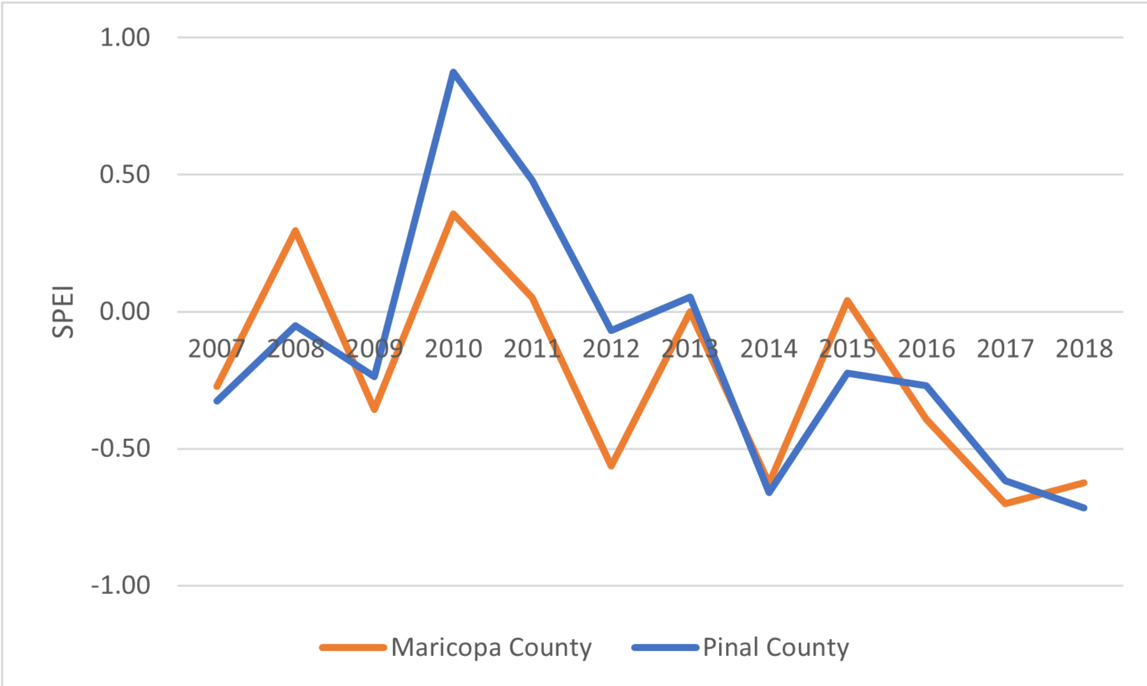

Data Source: Global SPEI Database

Figure 12: Central Arizona Precipitation

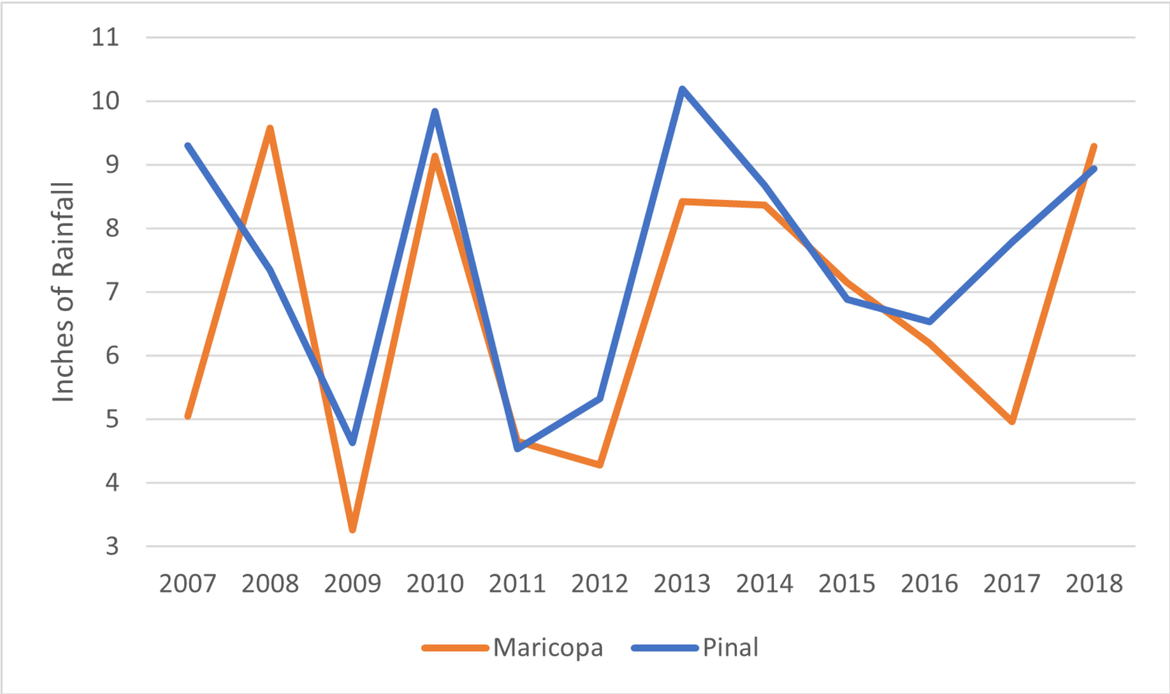

Data Source: NOAA National Weather Service

Figure 13: Central Arizona 100+ Degree Days

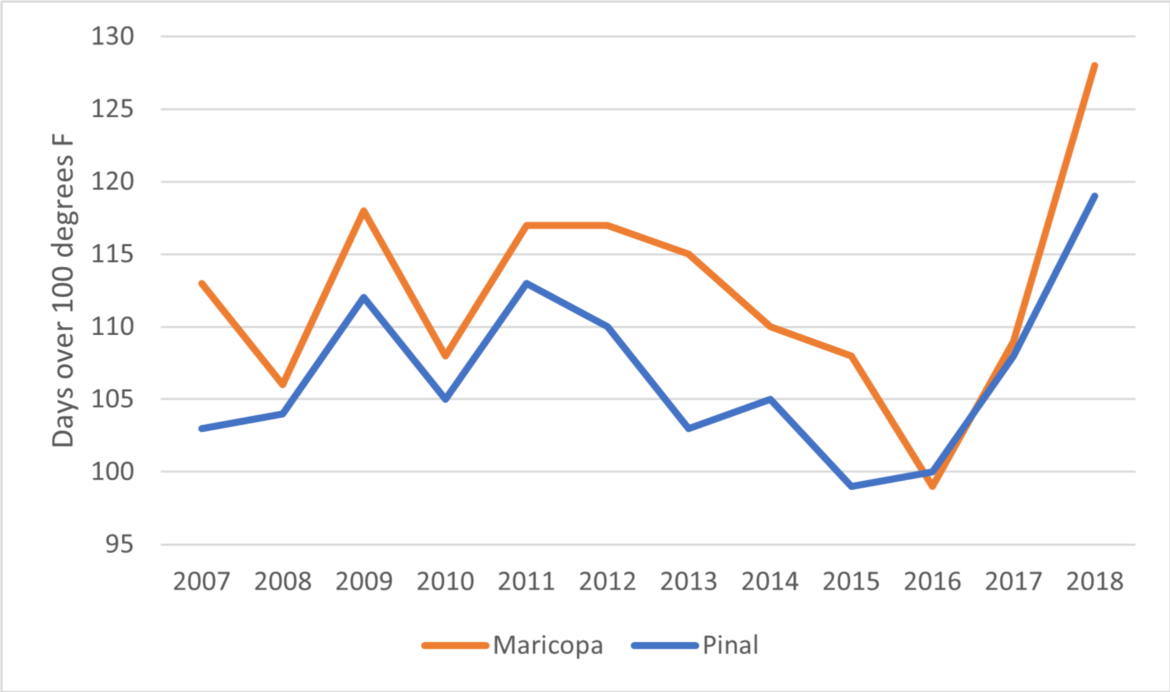

Data Source: NOAA National Weather Service

Figure 14: Residuals - Deliveries to Agriculture

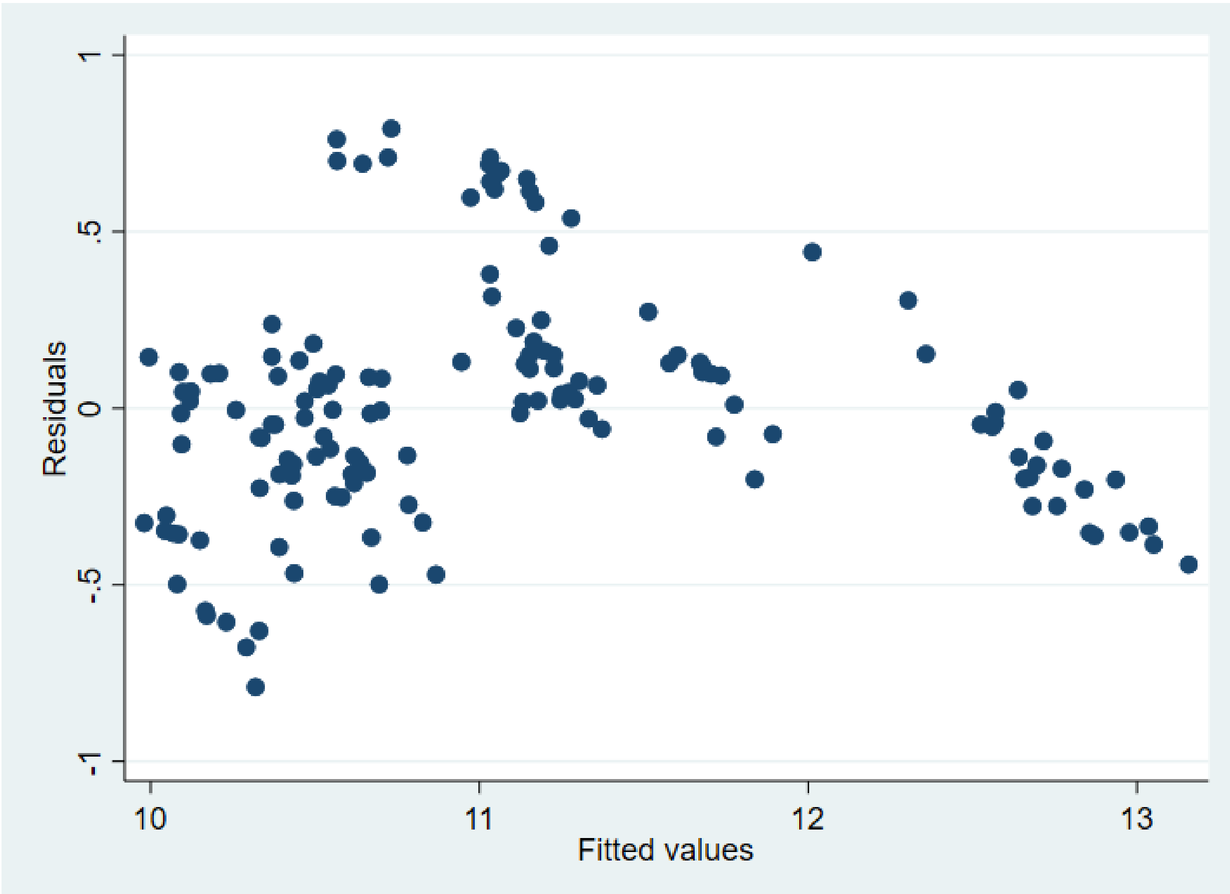

n = 144

Figure 15: Residuals - Irrigation Intensity

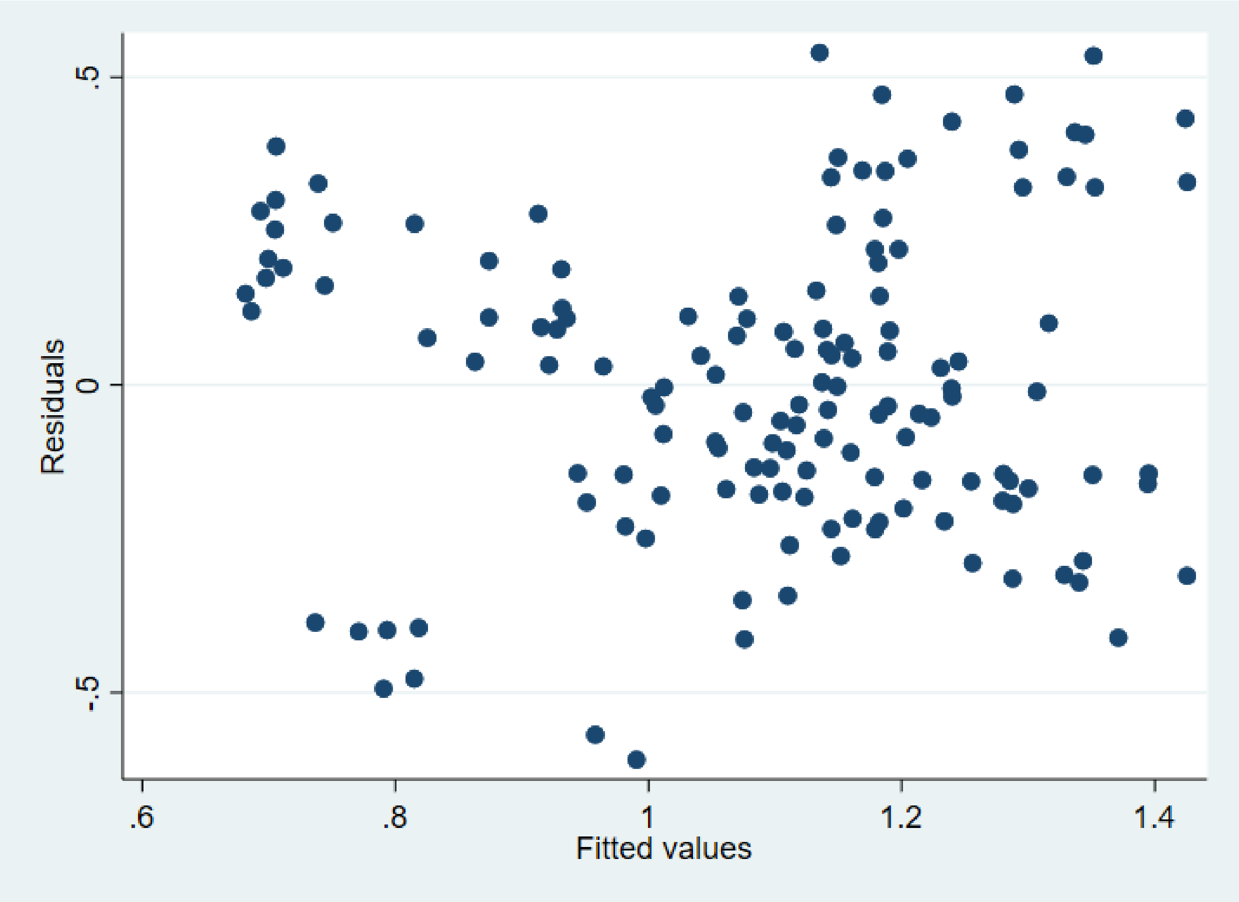

n = 144

Figure 16: Tree Crop Area by Irrigation district

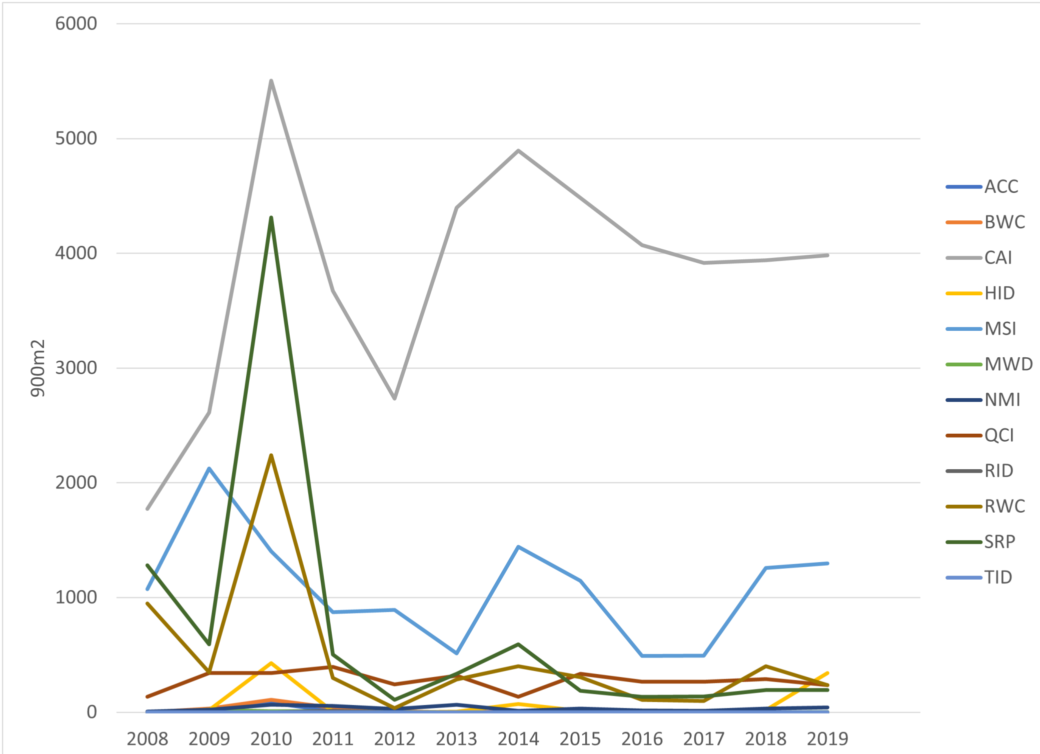

Data Source: USDA NASS Cropland Data Layer

Figure 17: Tree Crop Area by Irrigation district, detail

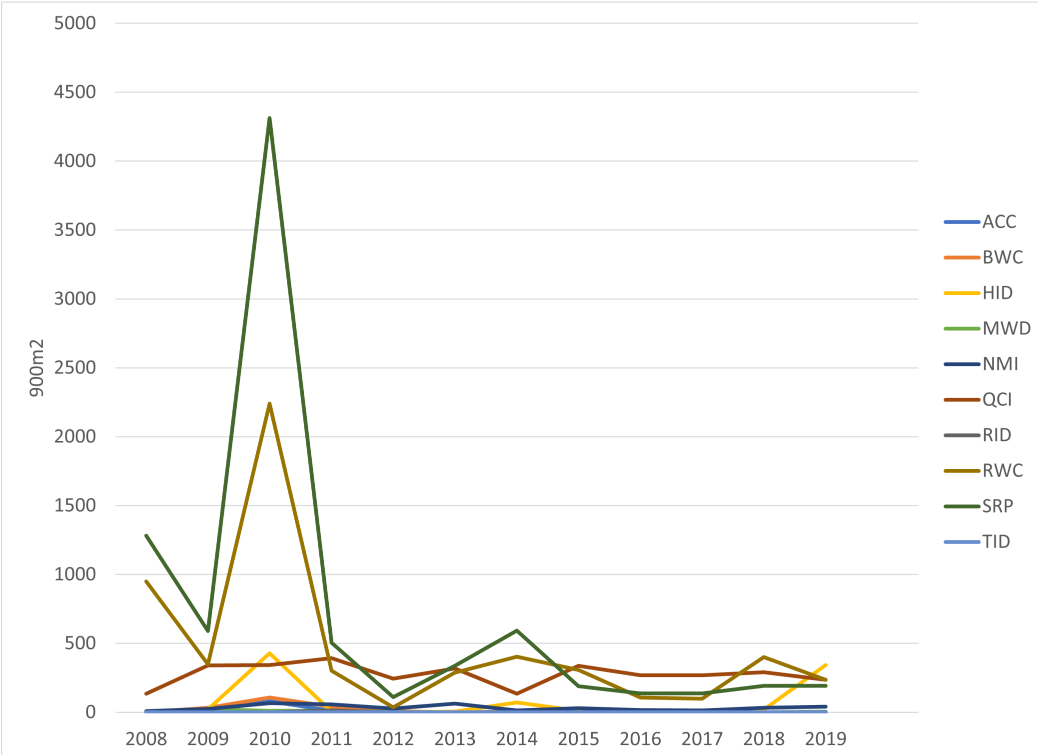

Data Source: USDA NASS Cropland Data Layer,
omitting CAIDD and MSIDD

Figure 18: Planted Area by Irrigation district

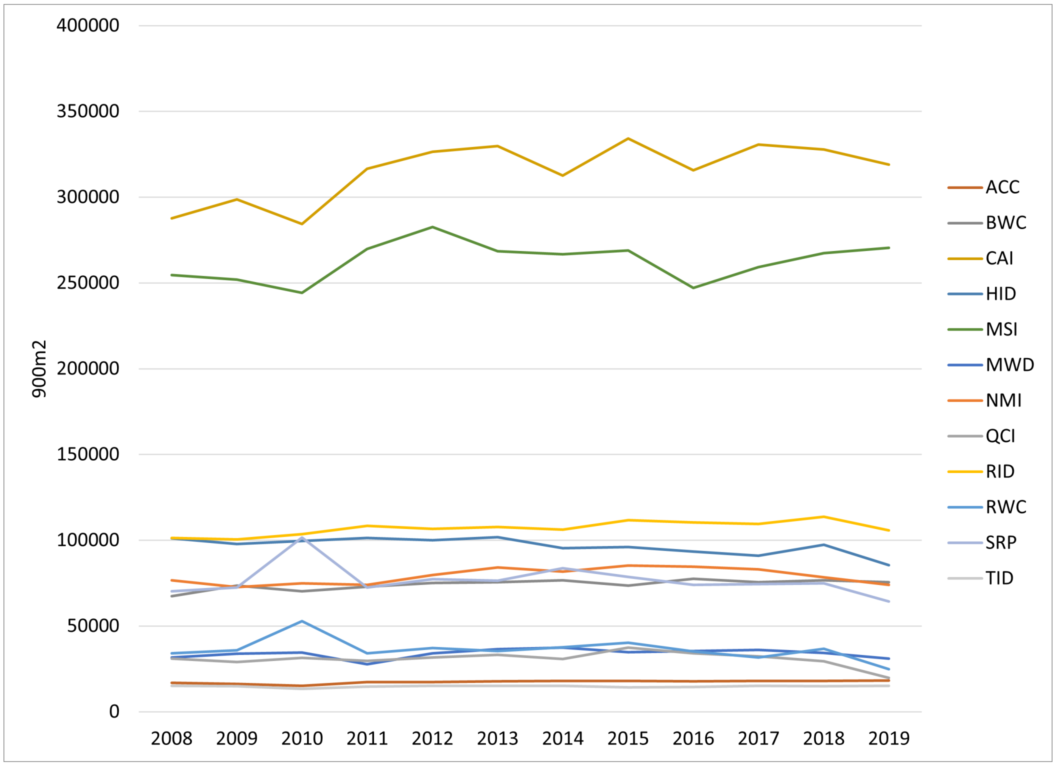

Data Source: USDA NASS Cropland Data Layer

Table 1: Summary Statistics - Water Deliveries

	Water Deliveries	**Agricultural Water Deliveries**	**Natural Log Ag Deliveries**	**Percent Agricultural**	**Irrigation Intensity**
Unit of Measure	acre-feet	acre-feet	ln(acre-feet)	percentage	af/900m^2
Mean	144,182	98,710	11.09	84.2%	1.09
Min	13,883	13,766	9.53	11.6%	0.3
Max	593,828	338,502	12.73	100%	1.76
Standard Deviation	155,017	92,408	0.9	25.7%	0.31

n = 144

Table 2: Summary Statistics - Irrigation Districts

District	**AMA**	**Average Ag Deliveries**	**Total Area**	**Average Planted Area**	**Average Alfalfa Area**	**Average Cotton Area**	**Average Grains Area**	**Average Trees Area**	**Average Pasture Area**	**Average Other Crops Area**	**Average Fallowed Area**
Arlington Canal Company	Phoenix	26,271	24,645	17,407	13,566 (78%)	687 (4%)	1,375 (8%)	9 (<1%)	12 (<1%)	1,757 (10%)	1,410
Buckeye Water Conservation Dist.	Phoenix	123,114	98,497	74,210	57,394 (77%)	3,404 (5%)	8,255 (11%)	17 (<1%)	18 (<1%)	5,122 (7%)	5,727
Central Arizona Irrigation and Drainage Dist.	Pinal	294,129	489,731	315,340	63,219 (20%)	183,711 (58%)	54,174 (17%)	3,832 (1%)	926 (<1%)	9,479 (3%)	89,188
Hohokam Irrigation Dist.	Pinal	51,480	127,571	96,696	40,377 (42%)	36,212 (38%)	17,845 (19%)	77 (<1%)	29 (<1%)	2,157 (2%)	19,583
Maricopa-Stanfield Irrigation and Drainage Dist.	Pinal	278,757	465,940	262,647	109,151 (42%)	48,975 (19%)	83,465 (32%)	1,084 (<1%)	1,065 (<1%)	18,907 (7%)	115,443
Maricopa Water Dist.	Phoenix	36,561	163,449	33,925	10,174 (30%)	965 (3%)	10,762 (32%)	4 (<1%)	312 (1%)	11,708 (35%)	28,769
New Magma Irrigation and Drainage Dist.	Phoenix	81,676	122,301	79,088	47,492 (60%)	14,433 (18%)	12,626 (16%)	33 (<1%)	38 (<1%)	4,466 (6%)	15,220
Queen Creek Irrigation Dist.	Phoenix	26,778	91,182	30,811	14,393 (47%)	6,170 (20%)	8,212 (27%)	276 (1%)	146 (<1%)	1,613 (5%)	11,862
Roosevelt Irrigation Dist.	Phoenix	128,394	176,294	107,117	56,731 (53%)	19,755 (18%)	24,756 (23%)	2 (<1%)	54 (<1%)	5,820 (5%)	25,569
Roosevelt Water Conservation Dist.	Phoenix	39,011	186,831	36,328	25,274 (70%)	894 (3%)	6,998 (19%)	476 (1%)	697 (2%)	1,988 (6%)	15,179
Salt River Project	Phoenix	82,613.81	1,158,801	76,752	50,955 (66%)	9,031 (12%)	11,126 (15%)	715 (1%)	541 (1%)	4,385 (6%)	27,782
Tonopah Irrigation Dist.	Phoenix	15,736	18,623	14,784	7,389 (50%)	2,156 (15%)	4,512 (31%)	0 (0%)	36 (<1%)	691 (5%)	1,395

n = 12 per district, ag deliveries in acre-feet, all area units in $900m^2$
(percentage of planted area in parentheses)

Table 3: Cropland Data Layer Variables

Variable	**CDL Category**
Alfalfa Area	Alfalfa
Cotton Area	Cotton
Grains Area	Corn, Sorghum, Barley, Durum Wheat, Winter Wheat
Trees Area	Citrus, Pecans, Pears, Pistachios, Olives, Oranges, Grapes
Pasture Area	Grassland/Pasture
Other Crop Area	Cantaloupes, Watermelon, Rye, Spring Wheat, Lettuce, Other Hay (not Alfalfa), Oats, Dry Beans, Potatoes, Carrots, Chick Peas, Millet, Broccoli, Cabbage, Honeydew, Double Croppings
Fallowed Area	Fallow/Idle Cropland
Developed Area	Developed Open Space, Low Development, Medium Development, High Development

Table 4: Summary Statistics - Explanatory Variables

	Alfalfa Price	**Cotton Price (Dec. Futures)**	**CAP Water Cost**	**SPEI** (lag)
Unit of Measure	2019$/ton	Futures Price (2019$)	2019$/af	Index
Mean	200.84	85.49	64.12	-0.211
Min	144.19	58.34	42.75	-0.716
Max	255.73	139.50	80.96	0.874
Standard Deviation	36.4	21.19	11.79	0.382

n = 144

Table 5: OLS Model Results

	Water Deliveries	**Intensity**
R^2	0.8743	0.3708
n	144	144
Lagged SPEI	0.0581 (0.082)	0.1075* (0.061)
Alfalfa Price	-0.0021* (0.001)	-0.0017* (0.001)
Cotton Futures	0.0042* (0.002)	0.0037** (0.002)
CAP Water Cost	-0.0045* (0.003)	-0.0035* (0.002)
Trees (area/pct)	-0.0003*** (<0.001)	-6.1341* (3.315)
Fallow (ln(area)/ratio)	0.0917*** (0.029)	-0.24*** (0.088)
Planted Area	<0.0001*** (<0.001)	n/a
Pinal AMA	-1.0472*** (0.112)	-0.336*** (0.049)
Constant	9.642	1.5493

*** = $p < 0.01$, ** = $p < 0.05$, * = $p < 0.1$
Standard Errors in Parentheses

Table 6: Breusch-Pagan Tests for Heteroskedasticity

Model	χ^2	p-value
Deliveries to Agriculture	2.66	0.1028
Irrigation Intensity	0.02	0.8918

Table 7: Fixed-Effect Model Results

	Water Deliveries	**Intensity**
R^2 - Within	0.37	0.4287
R^2 - Between	0.7484	0.0034
n	144	144
Lagged SPEI	0.1282*** (0.026)	0.1162*** (0.02)
Alfalfa Price	-0.0015 (0.001)	-0.0009 (0.001)
Cotton Futures	0.0032** (0.001)	0.0027** (0.001)
CAP Water Cost	-0.0025 (0.001)	-0.0029** (0.001)
Fallow (ln(area)/ratio)	-0.0405* (0.022)	-0.0519 (0.064)
Planted Area	<0.0001** (<0.001)	n/a
Constant	10.9883	1.2585

*** = p < 0.01, ** = p < 0.05, * = p < 0.1

Robust Standard Errors in Parentheses

A Appendices

A.1 Data Collection Procedures

Where possible, publicly available data were collected for this study directly online, without the need for transcription or transformation, in order to minimize any risk of human error affecting data quality. In practice, the only raw data directly applicable to this study's empirical analysis are average annual alfalfa prices obtained from the USDA's National Agricultural Statistics Service, as these data are available in the form of annual averages at the county level. All other data used had to be transcribed or transformed in some way. Some of these procedures were simple transcriptions from PDF files or averaging monthly observations to find annual means, while other were more involved. Each will be described below, with greater detail provided for the more complex data collection procedures.

Simple Procedures

The simplest procedures performed in the data collection process apply to the variables for expected cotton price, CAP water cost, and the Standardized Precipitation-Evapotranspiration Index (SPEI) climate measure.

Recall from Chapter 4 that the New York Cotton Exchange December Futures prices taken from the last Friday in February of a given year are seen as a strong proxy for growers' price expectations. Data from the New York Cotton Exchange are publicly available through a wide variety of third party online resources which document market patterns. Expected cotton price data are found by manually transcribing the price listed on the last Friday in February of year t. Data used in this study were transcribed from futures.tradingcharts.com, which reports historical data from the New York Cotton Exchange for December Futures prices for cotton as far back as 1971.

CAP water cost data likewise require transcription but are easier to source than cotton futures data, as they are made directly available by the US Bureau of Reclamation. Each year, the bureau publishes rate schedules for CAP water, the cost of which is dependent on end use. For example, municipal and industrial water users are charged a higher rate per acre-foot of Colorado River water delivered through the CAP than agricultural water users. These rate schedules are available through the Bureau of Reclamation's website as far back as 2007. All that is necessary for this study is to transcribe the agricultural water cost for year t from each schedule.

Finally, SPEI data from the Global SPEI database are available monthly at a moving three-month average. Recall from Chapter 4 that, while these data are not available at the level of individual counties, their pixel scale of 0.5° squared closely matches the boundaries of Maricopa and Pinal counties. Because these data are available monthly, it is necessary to average the values by year to find an annual measure. This average annual value is then used in the empirical analysis presented in Chapter 5.

Irrigation Water Delivery Variables

Beginning with water delivery variables, data collection procedures begin to become a bit more complex. Water delivery variables are taken from annual reports filed with the Arizona Department of Water Resources (ADWR) by irrigation districts. With the exception of one district, every report contains a schedule of water deliveries made in that year enumerated by water right number, called Schedule D. The type of right is also included for each line item listed in Schedule D. Many of these reports also contain a summary page, called Schedule D1S. Schedule D1S contains a breakdown of water deliveries by type (non-exempt irrigation, exempt irrigation, municipal providers, etc.), and by source (groundwater, CAP, in-lieu, etc.).

As Schedule D1S is concise and well organized, it is the first-best option for obtaining water deliveries data. Data from Schedule D1S is used whenever this summary page is included in a report. Using data from Schedule D1S is considered a first-best procedure for two reasons: 1) summing up line items from Schedule D is tedious, and 2) summing up line items from Schedule D is more prone to human error. When Schedule D1S is available, it is a simple task to record the sum of deliveries to non-exempt irrigation and exempt irrigation uses.

Where Schedule D1S is unavailable, Schedule D becomes the next best source for water deliveries data. Line items from Schedule D are included whenever the water right number begins with '58-' (denoting farm-owned irrigation grandfathered rights), '57-' (denoting district-owned irrigation grandfathered rights), or '88-' (denoting irrigation rights on a farm registered as having "Best Management Practices"), as well as any lite items denoted as exempt water deliveries (Tyler Fitzgerald, AZ Dept. of Water Resources, personal communication, June 25, 2021). When pulling data from Schedule D, deliveries are first summed by individual water rights type, before summing these water rights type groupings to find total deliveries to irrigation. This allows for arithmetic checks to occur in multiple places, in order to diminish the possibility of human error.

The exception mentioned above relates to the San Carlos Irrigation and Drainage District (SCIDD). Rather than filing water use reports including Schedule D1S or Schedule D, reports filed by SCIDD simply include a front page describing overall water use by the district. As is evident in other districts' reports, not all water used in a district necessarily is put to agricultural uses, and so the lack of any indication of agricultural use in the SCIDD's reporting was initially distressing. Table 2 shows the SCIDD to be one of the largest in the study area in terms of water deliveries, and omitting it from this study due to uncertainty over the nature of its' water use is an outcome to be avoided if possible. Thankfully, the Arizona Department of Water Resources were able to provide more detailed breakdowns of water use in the Pinal AMA for years ranging from 2016 to 2019. These detailed breakdowns allowed for the comparison of specific agricultural water use within the SCIDD to the overall water use listed on the front page in the district's reports. In the four years for which these comparisons were possible, deliveries to non-exempt irrigation in the SCIDD closely resembled those values reported in the district's reports as "Total Water Delivered To Lands". Due to these favorable comparisons, the assumption can be made that SCIDD's reported "Total Water Delivered To Lands" variable accounts for water delivered to agricultural uses in years prior to 2016. It is unfortunate that it is necessary to make an assumption regarding these data, but preferable to having to omit the SCIDD from the study entirely.

Land Cover Variables

Finally, irrigation district land cover variables are generated with QGIS software using irrigation district shapefiles from the Arizona Department of Water Resources, and raster data from the USDA National Agricultural Statistics Service's Cropland Data Layer. USDA NASS publishes the Cropland Data Layer annually, going back as far as 2008 in Arizona. It is for this reason that the study period in this analysis only goes back as far as 2008, as land cover variables are restricted to observations from 2008 forward. Using QGIS's raster calculator tool, binary raster layers are generated from the Cropland Data Layer to reflect specific land cover categories. These categories include alfalfa, cotton, major grains, tree crops, pasture, other crops, fallowed/idled lands, and developed area. The pixels reflecting the land cover of interest in each of these raster layers are then summed by irrigation district using QGIS's zonal statistics tool. From here, it is simple to divide the pixel area of a land cover grouping within an irrigation district by the pixel area of the entire district to find the percent land devoted to a particular land cover group in a given year. For crop groups, the pixel area devoted to a particular cropping category within an irrigation district can also be divided by the overall planted area within the district, which is simply the sum of the alfalfa, cotton, major grains, tree crops, pasture, and other crops categories. This procedure yields the percentage of planted area within a district planted in a certain crop or crop group.

A.2 Effects of Irrigation Districts, Crop Mix, and Prior-Year Planted Area

Purpose

This section describes measures taken to ensure that the modeling approach presented in the main body of this work is valid in it's assertions that water deliveries and irrigation intensity are partially driven by exogenous variables such as price, climate, and policy. A concern regarding these assertions is that water delivered and application rates could be essentially predetermined by a growers' choice of acreage planted in a particular set of crops. If deliveries and irrigation intensity are completely defined by planted area, the econometric analyses reported in Chapter 5 are unnecessary. A another concern is that crop area may be relatively invariant over time, with growers simply planting the same acreage and crop mix as they had selected the prior year. A final concern is that irrigation districts' fixed characteristics (i.e. scale, location, water rights, etc.) may predetermine water deliveries and irrigation intensity. This section shows that, while current and prior year cropped area carries much explanatory power for both water deliveries and irrigation intensity (water delivered per planted area), the opportunity remains for additional factors to influence water delivered and irrigation intensity. Likewise, while irrigation districts' fixed characteristics carry significant explanatory power, they do not comprehensively describe water use decisions.

Models Focused on Planted Area

The first stage of this process considers two Ordinary Least Squares (OLS) regressions. The first model examines deliveries as a function of area within an irrigation district planted in different crops and crop groups. The crops and crop groups included are alfalfa, cotton, major grains (including corn, sorghum, barley, and wheat), tree crops (including pecans, pistachios, olives, citrus, and others), irrigated pasture, and a catch-all grouping for all other crops (including but not limited to cantaloupes, oats, potatoes, millet, lettuce, beans, and carrots). Recall, Table 2 describes the average size of area planted in each of these crop groups by district. In addition to planted area, this first model also considers the effect of fallowed lands on an irrigation district's water deliveries.

The second model describes intensity as a function of the percentage of total district planted area represented by different crops and crop groups. Percentage of planted area is chosen due to irrigation intensity being a per-acre measure. Fallowed area is not included in this second specification. Instead, the ratio of fallowed area to planted area is considered. This measure ranges from a minimum observed value of 0.01 to a maximum value of 3.27, with an average of 0.43 describing just under half as much fallowed area as planted area. Because irrigation intensity decision making is modeled as a function of the percentage of a districts' total planted area represented by various crops, an issue arises with collinearity. By definition, these percentage variables sum to 1, as they are represented in decimal form. This means that each observation contains six explanatory variables that add up to 1, which happens to be the value assigned to the constant term in an OLS regression's X matrix. For this reason, any model specifications mentioned in this section which include all six crop category percentage variables will not include a constant term.

Formally, these models are specified as follows:

$$W_{it} = \alpha + \beta_a A_{ait} + \beta_c A_{cit} + \beta_g A_{git} + \beta_t A_{tit} + \beta_p A_{pit} + \beta_o A_{oit} + \beta_f F_{it}$$

$$I_{it} = \delta_a P_{ait} + \delta_c P_{cit} + \delta_g P_{git} + \delta_t P_{tit} + \delta_p P_{pit} + \delta_o P_{oit} + \beta_f FR_{it}$$

Here, W_{it} represents water delivered and I_{it} represents intensity of irrigation (total water applied divided by planted area) in district i in year t. A_{xit} represents area planted in crop x in district t in year i. In the second model, P_{xit} represents the percentage of overall planted area planted in crop x in district i in year t. The six crop groups described above are represented by subscripts a (alfalfa), c (cotton), g (grains), t (trees), p (pasture), and o

(other crops). Fallowed area in district i in year t is given as F_{it}, and the ratio of fallowed area to planted area in district i in year t is given as FR_{it}.

The results of these models are presented in Table A1. Of greatest interest are the models' measures of fit (R^2). An R^2 value approaching 1 means that the explanatory variables' collective effect on the outcome variable is almost completely definitional. In this case, little-to-no explanatory power is unaccounted for and inclusion of other potential variables is not fruitful. The identity specifications presented in Table A1 return R^2 values of 0.9368 for water deliveries to agriculture, and 0.9285 for irrigation intensity.

The relatively high measure of fit returned by both models seems to indicate that current year water deliveries and irrigation intensity can largely be explained by current year crop mix and acreage planted. This is intuitive. Once planting decisions are made, there is some minimum amount of water needed to maintain the health of a grower's crop. There could be instances where external factors lead to a decision to abandon a crop that's already been planted, but these events are likely rare. While an R^2 value of 0.9368 is high, it is not 1. This means that, while planted area accounts for much of growers' water use, the inclusion of these explanatory variables alone does not explain deliveries completely. Some other factors influence growers' water use decisions.

The irrigation intensity model, with its R^2 value of 0.9285, indicates that crop mix and planted area also have a substantial effect on water delivered per planted area, but that there also exist other influential factors to be accounted for. This is also intuitive, as one might easily assume that growers must apply a predictable amount of water per acre, based on what has been planted. However, the literature has shown that growers will also make intensive choices throughout the growing season (Foster et al., 2014). These may be based on a number of exogenous variables, such as a change in the cost of inputs or a shift in the expected price of crops at harvest time. The modelling presented in Table A2 underscores this point, and will be discussed at greater length in the coming paragraphs.

Inclusion of Interaction Terms

In an effort to decompose the effect of various exogenous climatic and economic factors from the effect of crop mix and area planted, regressions are run which include interaction variables on the right hand side. These interaction variables are generated by taking the product of some external factor (such as SPEI value in a given year) and each land cover variable contained in the specifications above. A formally stated example follows, with water deliveries specified as the outcome variable.

$$W_{it} = \alpha_W + \beta_x A_{xit} + \gamma_x A_{xit} SPEI_{it}$$

In this example, A_{xit} represents a vector of areas planted in different crops in irrigation district i in year t. These planted area variables are included once in their original form, and included again having been interacted with SPEI in irrigation district i in year t. Of specific interest here is the vector of coefficients for these interaction terms, γ_x. Statistical significance in these coefficients indicates that the external factor included in the model specification is having some impact beyond that explained by planted area variables.

The interaction models are run using real alfalfa prices, real expected cotton prices, real CAP water costs, and lagged SPEI values (prior year) as interaction variables in separate models. Results are presented in Tables A2 and A3. Most specifications return either one or two significant coefficients which correspond with interaction terms. This is consistent with the measures of fit seen in the planted area models above, as much explanatory power is already captured by the planted area variables. It must be noted that CAP costs can be seen to have no significant effect on water deliveries when planted area is accounted for. Likewise, cotton prices appear to have no significant effect on irrigation intensity when percentage of planted area is accounted for, a somewhat surprising finding. Once again, R^2 values are generally high for for all interaction models, as much explanatory power is captured by the planted area variables. While these R^2 values are high, but not so close to 1 as to preclude the addition of other explanatory variables. This, along with the statistically significant interaction terms, reinforces the notion that there are other influential factors driving agricultural water deliveries and irrigation intensity that have yet to be accounted for.

Irrigation Intensity Over Time

Next, the irrigation intensity model specified at the top of this section is run once again, with the observations split into subgroups by year. Recall, the model is specified as follows:

$$I_{it} = \delta_a P_{ait} + \delta_c P_{cit} + \delta_g P_{git} + \delta_t P_{tit} + \delta_p P_{pit} + \delta_o P_{oit} + \beta_f FR_{it}$$

Observations are grouped into four three-year sets: 2008-10, 2011-13, 2014-16, and 2017-19. The intention here is to observe how irrigation intensity is affected by the extent to which a given crop is planted in a district from year to year. Observations are grouped because, with only 14 irrigation districts, running this model at the level of each individual year would include so few observations that results may be unreliable.

Table A4 outlines the results of this modelling approach. There is some significant variety between the coefficients returned, meaning that some factors not accounted for in this modeling are at work driving growers' intensive decision making. The percentage of planted area represented by alfalfa and percentage of planted area represented by major cereal grains aboth significant across all four groups. The coefficient estimated for the alfalfa variable is much higher for the first grouping (2008-10) than for the other three. The coefficient estimated for the grains variable floats close to a value of 1 with no discernible pattern in its fluctuations. The "other crops" category changes wildly from year to year, but this may be due to a change in the mix of the crops making up this catchall category, and so should be taken with a grain of salt. The fallow-to-planted ratio is seen to have a consistently significant negative effect on irrigation intensity, but once again the coefficient values vary considerably. In general, the frequent shifts in coefficient values observed from group to group allude to the presence of exogenous variables not included in the above specification affecting the intensity of irrigation water applied.

Effects of Prior-Year Planted Area

Next, checks are performed to examine to what extent a crop's planted area can be predicted based on plantings of the same crop in the previous year. This process begins with looking at simple regressions that include only the natural log of a single crop's area in a given year (t) as the outcome variable and the natural log of the same crop's area the year prior ($t-1$) as the sole explanatory variable. Formally, the models is specified as follows:

$$ln(A_{it}) = \alpha + \beta ln(A_{i(t-1)})$$

This simple framework will then be built upon with the inclusion of alfalfa and cotton prices in year t, described below.

$$ln(A_{it}) = \alpha + \beta ln(A_{i(t-1)}) + \delta_a P_{at} + \delta_c P_{ct}$$

It is important to mention that, because these models involve a lagged land cover variable, and because the Cropland Data Layer only publishes measurements from Arizona as far back as 2008, these regressions involves fourteen fewer observations than others in this analysis, as it is impossible to construct a lag variable for an observation where $t = 2008$. Results from these models are presented in Table A5.

At a glance, two details stand out in these results: the high significance of the lagged variables, and the relatively high measures of fit. These go hand in hand. An irrigation district heavily planted in a particular crop or crop group in year $t-1$ is likely to be planted with a similar crop mix and acreage in year t. The R^2 values for the alfalfa models (0.9603 and 0.9618) are very high. This can be explained by the fact that alfalfa is not planted annually. Recall, alfalfa is a semi-perennial crop and so, unless external factors drive a grower to remove their stands of alfalfa in order to plant something different, it is likely that much of the alfalfa planted in Central Arizona in one year will remain there into the next. R^2 values are much lower for cotton, major grains, and other crops, indicating greater flexibility in cropping decisions for these annual crops. One unexpected result observed is the lower measure of fit seen in the tree crop models. The idea of area planted in trees shifting greatly from year to year is non-intuitive. However, tree crops have a finite life. Planting may occur in "pulses" and removal of some portion of planted area would likely be happening in any given year. This low measure of fit may also be explained by USDA NASS's Cropland Data Layer generally being more accurate in its identification of cash crops such as cotton and alfalfa than tree crops (US Department of Agriculture, National Agricultural Statistics Service, 2021).

In those specifications which include alfalfa and cotton prices, these price variables are frequently significant. These crop prices play an important role in determining planting decisions from year to year. It is not surprising that the prices of these two crops would affect each other's cropped areas, but the prices of alfalfa and cotton also have a major impact on acreage planted in the catchall "other crops" category. These prices do not affect area planted in tree crops, as orchards are unable to shift cropping patterns on an annual basis. The significance of these crop prices attests to the fact that planting decisions are not solely predetermined by plantings in the prior year.

Effects of Irrigation Districts' Fixed Characteristics

Finally, a fixed effect regression is specified for each dependant variable (the natural log of water deliveries to agriculture and irrigation intensity). These regressions contain no explanatory data aside from district level fixed effect variables. These regressions are run in order to assess the extent to which this study's outcome variables can be explained solely due to fixed district-level characteristics.

The results of these fixed effects regressions are presented in Table A6. In each model, the Arlington Canal Company is not assigned a fixed effect variable in order to avoid issues with collinearity. Instead, the Arlington Canal Company serves as a baseline, whose water deliveries and irrigation intensity ratio are represented by the constant in each specification. In an interesting statistical quirk, standard errors reported are uniform across all fixed effect variables. This is due to the standard errors of the coefficients of the district fixed effect variables depending on the root mean square error of the regression as a whole. Because each fixed effect variable has the same number of annual observations, the same standard error is returned for each coefficient estimate. This is to be expected, and it is because of this that standard error is only reported once for each model.

The statistics of greatest interest reported in Table A6 are the R^2 values for each model specification. The water deliveries model exhibits a high R^2 value of 0.9681. This indicates that much of districts' water need can be explained by fixed characteristics, such as scale, location, composition of water rights, and which AMA they belong to. Similar to the cropped area models discussed earlier in this appendix, the R^2 value, while high, is not so high that it could reasonably be called vanishingly close to a value of 1. There is still room for additional explanatory data to contribute to a statistical understanding of districts' overall water deliveries.

The irrigation intensity model reports a relatively modest R^2 measure of 0.8073. Once again, this behavior is similar to that of the cropped area models. Much of irrigation water application rates may be attributed to district-level fixed characteristics, which could include soil composition, elevation, and irrigation technology. However, just as irrigation intensity is not able to be explained solely by crop and acreage choice, neither can it be explained only through fixed district level characteristics. An R^2 value of 0.8073 means there is still much to be explained by factors other than fixed characteristics.

Finally, both models return almost entirely significant coefficient estimates. Every fixed variable included in each specification returns a significant coefficient estimate, aside from the dummy value associated with the Queen Creek Irrigation Districts annual water deliveries to agriculture. This evidence of the high significance of district-level fixed characteristics on the outcomes examined in this study further supports the inclusion of district level fixed effects in Chapter 5's empirical analysis.

Table A1: Cropping Identity Models

	Water Deliveries	**Irrigation Intensity**
R^2	0.9368	0.9285
Alfalfa	1.367*** (0.271)	1.3581*** (0.179)
Cotton	0.6914*** (0.196)	0.5688 (0.42)
Grains	1.2662*** (0.205)	1.2857*** (0.278)
Trees	11.9739 (8.633)	-12.486 (8.503)
Pasture	-3.6183** (1.474)	3.414** (1.381)
Other Crops	1.1443 (0.759)	1.4035*** (0.34)
Fallow	-0.3933** (0.141)	-0.3432*** (0.082)
Constant	-3737.64	n/a

*** = p < 0.01, ** = p < 0.05, * = p < 0.1
(Standard Errors in Parentheses)

Table A2: Decomposing Drivers of Irrigation Intensity

	Alfalfa Price	Cotton Price	CAP Water Cost	Lagged SPEI
R^2	0.9316	0.9353	0.9376	0.9341
Percentage Alfalfa (Interacted)	-0.0026* (0.001)	-0.0004 (0.004)	-0.0059 (0.004)	0.1415 (0.179)
Percentage Cotton (Interacted)	0.0012 (0.002)	0.0042 (0.005)	0.0019 (0.012)	-0.1237 (0.25)
Percentage Grains (Interacted)	0.0062* (0.003)	0.0075 (0.007)	-0.0032 (0.008)	-0.0149 (0.39)
Percentage Trees (Interacted)	-0.114 (0.074)	0.1656 (0.128)	-0.1293 (0.284)	6.2886 (8.552)
Percentage Pasture (Interacted)	0.038 (0.061)	0.049 (0.262)	-0.6409 (0.736)	26.6578*** (7.725)
Percentage Other Crops (Interacted)	-0.0013 (0.01)	0.0139 (0.018)	0.0053 (0.016)	0.1517 (0.248)
Fallow-to-Planted (Interacted)	0.0014 (0.001)	$<$0.0001 (0.001)	-0.0069* (0.004)	-0.2182*** (0.054)
Percentage Alfalfa	1.8819*** (0.239)	1.3995*** (0.221)	1.7886*** (0.398)	1.5013*** (0.246)
Percentage Cotton	0.3162 (0.427)	0.1751 (0.783)	0.5224 (0.751)	0.4929 (0.4)
Percentage Grains	0.0365 (0.622)	0.5817 (0.502)	1.3307* (0.706)	1.138*** (0.295)
Percentage Trees	10.8613 (15.366)	-26.2503 to (13.448)	-5.0086 (23.706)	-5.8873 (13.537)
Percentage Pasture	-4.4962 (10.955)	-1.4434 (23.324)	39.4191 (43.033)	31.464*** (8.886)
Percentage Other Crops	1.7142 (1.822)	0.3819 (1.415)	1.2889 (1.252)	1.6477*** (0.352)
Fallow-to-Planted	-0.6238*** (0.201)	-0.34** (0.126)	0.0562 (0.221)	-0.5448*** (0.088)

*** = $p < 0.01$, ** = $p < 0.05$, * = $p < 0.1$
(Standard Errors in Parentheses)

Table A3: Decomposing Drivers of Water Deliveries

	Alfalfa Price	Cotton Price	CAP Water Cost	Lagged SPEI
R^2	0.9435	0.9499	0.9419	0.939
Alfalfa Area (Interacted)	-0.0023 (0.002)	-0.0017 (0.004)	-0.0006 (0.007)	-0.0683 (0.304)
Cotton Area (Interacted)	0.0041 (0.003)	0.0088*** (0.002)	-0.0108 (0.019)	0.009 (0.281)
Grains Area (Interacted)	0.0098*** (0.003)	0.0184* (0.009)	-0.0073 (0.014)	0.4762 (0.486)
Trees Area (Interacted)	-0.0199 (0.169)	-0.1622 (0.127)	0.2202 (0.8)	3.1527 (14.376)
Pasture Area (Interacted)	0.1148 (0.072)	-0.1067 (0.25)	-0.4893 (0.6123)	5.378 (5.241)
Other Crops Area (Interacted)	-0.0035 (0.007)	0.0083 (0.023)	-0.0043 (0.041)	-0.871 (0.584)
Fallow Area (Interacted)	-0.0044*** (0.001)	-0.0062 (0.006)	-0.0027 (0.008)	-0.2822* (0.149)
Alfalfa Area	1.8873*** (0.354)	1.5731*** (0.279)	1.4666** (0.628)	1.2857** (0.494)
Cotton Area	-0.2717 (0.541)	-0.2536 (0.37)	1.3478 (1.003)	0.6733 (0.417)
Grains Area	-0.8692 (0.722)	-0.5258 (0.928)	1.4177* (0.687)	1.6812** (0.634)
Trees Area	19.8313 (30.048)	32.1394** (14.359)	-0.5888 (44.51)	14.9925 (20.889)
Pasture Area	-26.4385 (14.913)	3.9463 (18.297)	24.2592 (38.434)	-0.27 (2.553)
Other Crops Area	1.8237 (1.294)	0.9536 (1.487)	2.1494 (3.496)	0.5598 (0.331)
Fallow Area	0.5749* (0.275)	0.2105 (0.379)	-0.147 (0.582)	-0.5601** (0.193)
Constant	-4277.74	-5708.09	-5950.22	-4124.79

*** = p < 0.01, ** = p < 0.05, * = p < 0.1
Standard Errors in Parentheses

Table A4: Irrigation Intensity over Time

	2008-10	2011-13	2014-16	2017-19
R^2	0.9746	0.9152	0.9541	0.9554
Percentage Alfalfa	1.7054***	1.3529***	1.3357***	1.336***
	(0.149)	(0.285)	(0.143)	(0.123)
Percentage Cotton	0.4734*	0.872	0.713	0.5774
	(0.266)	(0.502)	(0.685)	(0.502)
Percentage Grains	0.9479***	1.143**	1.0331***	0.9506**
	(0.222)	(0.465)	(0.263)	(0.313)
Percentage Trees	4.9668	-22.4887**	-12.679*	-14.1173*
	(4.732)	(8.678)	(6.112)	(7.371)
Percentage Pasture	-4.3074***	8.4326	0.8778	2.947
	(1.118)	(8.197)	(18.449)	(2.526)
Percentage Other Crops	2.4167***	1.5189**	1.9666***	1.6261***
	(0.541)	(0.675)	(0.337)	(0.274)
Fallow-to-Planted	-0.4348***	-0.2339**	-0.5748***	-0.4311***
	(0.042)	(0.084)	(0.107)	(0.123)

*** = $p < 0.01$, ** = $p < 0.05$, * = $p < 0.1$
(Standard Errors in Parentheses)

Table A5: Checking for Temporal Correlation

Crop	Alfalfa		Cotton		Grains		Trees		Other	
R^2	0.9603	0.9618	0.8807	0.8849	0.8271	0.8284	0.6616	0.6837	0.5707	0.5991
Crop Area (Lagged)	0.975*** (0.016)	0.9745*** (0.016)	0.9169*** (0.028)	0.9176*** (0.027)	0.9207*** (0.034)	0.9216*** (0.035)	0.8238*** (0.057)	0.8365*** (0.056)	0.7651*** (0.055)	0.746*** (0.054)
Alfalfa Price		0.0012* (0.001)		-0.006** (0.003)		$<$ 0.0001 (0.003)		-0.0059 (0.004)		0.0103*** (0.004)
Cotton Price		-0.003** (0.001)		0.01** (0.005)		0.003 (0.004)		-0.0043 (0.007)		-0.0199*** (0.006)
Constant	0.26	0.276	0.719	1.047	0.624	0.344	0.82	2.31	1.84	1.6

*** = $p < 0.01$, ** = $p < 0.05$, * = $p < 0.1$
(Standard Errors in Parentheses)

Table A6: Irrigation District Identity Models

	Water Deliveries	**Irrigation Intensity**
R^2	0.9681	0.8073
Standard Error	0.0697	0.0575
Buckeye Water Conservation and Drainage District	1.5466***	0.148**
Central Arizona Irrigation and Drainage District	2.4154***	-0.5809***
Hohokam Irrigation District	0.579***	-0.9879***
Maricopa-Stanfield Irrigation and Drainage District	2.3624***	-0.4525***
Maricopa Water District	0.3291***	-0.4288***
New Magma Irrigation and Drainage District	1.1357***	-0.479***
Queen Creek Irrigation District	0.0025	-0.6435***
Roosevelt Irrigation District	1.5886***	-0.316***
Roosevelt Water Conservation District	0.3884***	-0.4228***
Salt River Project	1.1424***	-0.4337***
Tonopah Irrigation District	-0.5113***	-0.4497***
Constant	10.1723	1.5145

*** = $p < 0.01$, ** = $p < 0.05$, * = $p < 0.1$

A.3 Alternate Variables Considered

Alternate Land Cover Variables

Data incorporated into this study includes remote sensed land cover data provided publicly by the US Department of Agriculture's Cropland Data layer (CDL). The CDL provides annual data for land cover at a pixel resolution of 30m. This data product is invaluable to this research, as it allows variables such as planted area and fallowed area to be incorporated into the empirical analysis presented in Chapter 5. Additionally, without the CDL, the irrigation intensity dependent variable could not have been constructed, as irrigation intensity is defined in this study as water deliveries to agriculture divided by planted area. Appendix A.2 reports further details of CDL data by breaking planted area down into subgroups (cotton, alfalfa, grains, etc.) in order to investigate the overall effect of extensive cropping choices on water deliveries. None of these subgroups are included in Chapter 5's model specifications. However, there were extensive considerations regarding the use of two specific subgroups of land cover variables. After preliminary regressions, these were ultimately excluded from this work's empirical analysis: the area within an irrigation district planted in major grain crops and the area within a district classified by the CDL as "developed". Each is discussed in some detail below.

Recall from Figures 2 & 3 that barley, corn and durum wheat make up some of the most commonly planted crops in Central Arizona after alfalfa and cotton. The major grains category of land cover used in Appendix A.2 incorporates these three crops, as well sorghum and winter wheat. Taken as a group, this category represents a substantial portion of Central Arizona's planted crops, with this combination of major grains exceeding cotton in terms of planted area on occasion (once again, note Figure 2). It is for this reason that the idea of including this grains group in Chapter 5's empirical analysis was strongly considered. Ultimately, grains were omitted from the econometric models presented in that chapter for the following reasons. While grain crops taken as a group occasionally exceed cotton in terms of planted area, in only one year out of the twelve observed in this study does a single grain crop, durum wheat, outpace cotton planting. While together grain crops make up a significant part of Central Arizona's agricultural portfolio, these crops individually are still generally being planted on a much smaller scale than alfalfa and cotton. Finally, as discussed in Chapter 1, cotton and wheat crops are frequently rotated throughout the year, with wheat being planted in cotton fields as winter crop after harvesting. Because of this, including variables related cotton and grains in the same econometric model could potentially confound regression results.

The decision not to include developed area in Chapter 5's empirical analysis likewise required careful consideration. Early iterations of this work would have included all water deliveries within a district, as opposed to water deliveries specifically made to agriculture. Overall water deliveries data is more readily available through the Arizona Department of Water Resources and so an initial idea was to include developed area as a control for municipal/industrial water use in any total water deliveries model. This approach was ultimately dropped for the following reasons. First, the CDL's metadata seeks to evaluate the veracity of remote sensed observations through ground truthing. These evaluations are available for all crop categories (as well as fallowed/idle cropland) in each state in each year. These evaluations are *not* made for the CDL's developed area categories, leaving no way of knowing how accurate these measurements are (US Department of Agriculture, National Agricultural Statistics Service, 2021). The intention of the CDL is to provide information on cropping, not all forms of land cover, meaning that in the absence of ground truthed evaluations, data on developed area are less than reliable (Patrick Willis, personal communication). Happily, irrigation districts water use reports were eventually found to contain more detailed breakdowns of water use by sector (see Appendix A.1 for more detail). These breakdowns allowed for water delivered specifically to agriculture to be included as an outcome variable, eliminating the need to control for developed area.

Alternate SPEI Periods

The effect of climate is an area of interest in this study, and so, choosing an observational period for climate relevant to growers' water use decision making in year t is an important consideration. As such, four different time periods over which the climate could be measured were evaluated during the process of specifying empirical models of water demanded and irrigation intensity. These are: 1) the average climate in year t, 2) the average climate in the prior year (year $t-1$), 3) the average climate in the prior rainy season, defined as July to December of year $t-1$, and 4) the average climate in the prior winter, defined as November and December of year $t-1$ and January of year t. The water deliveries model outlined in Chapter 5 is specified using each of the three prior year SPEI measures. Appendix A.2 shows that extensive cropping choices made at the start of the growing season account for the majority of water deliveries to agriculture. Because of this, only climate measures from time periods prior to the start of the growing season are considered for the water deliveries model. On the other hand, intensive decisions are made throughout the growing season, and Appendix A.2 shows that, while extensive decisions have some effect on irrigation intensity, other factors play a much larger role in driving these choices. Therefore, the irrigation intensity model is specified using all four time period measures, as it seems likely the climate during the growing season may play a role in intensive decision making. Coefficient estimates from these models will be compared in order to assess the suitability of the various time period measures for use in the main body of this work.

Recall from Chapter 4 that the climate measure of choice in this study is the Standardized Precipitation-Evapotranspiration Index (SPEI). The index includes measures of precipitation and drought in the form of potential evapotranspiration (McEvoy et al., 2012). These measures are standardized at various time scales, with values returned ranging between -3 (meaning a period is hotter and drier than average), to 3 (meaning a period is cooler and wetter than average). This study makes use of SPEI values standardized over a moving three month time period, a time frame considered most efficient for assessing drought conditions related to soil water content (Global SPEI Database, 2021). These three month standardized values are then averaged over the periods described in the last paragraph to create four SPEI measures.

The first step in comparing the four potential variable measures is observing the correlation coefficients between every pair. A correlation matrix is presented in Table A7. Because each of the three prior year measures contain some amount of the same information, it is not surprising that the correlation coefficients are much higher between any combination of these prior year measures than any prior year variable's correlation with the own year measure. Because of the high correlation between these prior year measures, it is less likely to observe large differences in regression coefficient estimates' signs and significance levels. On the other hand, the low correlation between the own year measure and each of the prior year measures indicates that large differences between the sign and significance of regression coefficient estimates are quite possible.

The water deliveries model presented in Chapter 5 (Equation 21) is restated here for convenience.

$$ln(D_{it}) = \alpha_D + \beta_1 X + \beta_2 Alf_t + \beta_3 Cot_t + \beta_4 CAP_t + \beta_5 FaA_{it} + \beta_6 PlA_{it} + \overline{\Gamma_D} * \overline{FE_i} + e_{it}$$

where errors are clustered at t

The only difference between the above model and Equation 21 is X standing in for the various SPEI time period measures to be evaluated. Therefore, β_1 will be the coefficient of interest in this evaluation. Table A8 presents coefficient estimates obtained from regressions run using all three prior year time period measures: the full year, the rainy season (July to December), and the winter (November to January). Results are consistent across all three measures, with each returning a positive and highly significant coefficient estimate. As mentioned in the previous paragraph, this consistency is likely due to the high correlation between these measures, as each contains some amount of the same information. There is no clear measure of choice, nor a clear measure to be rejected. The prior year time period is included in Chapter 5 because, being a twelve month average, it contains more information than the other two.

Once again, the irrigation intensity model presented in Chapter 5 (Equation 22) is restated here for convenience, and once again X stands in for the various time period measures.

$$I_{it} = \alpha_I + \delta_1 W_{at-1} + \delta_2 Alf_t + \delta_3 Cot_t + \delta_4 CAP_t + \delta_5 FtP_{it} + \overline{\Gamma_I} * \overline{FE_i} + e_{it}$$

where errors are clustered at t

Table A9 presents coefficient estimates returned by running this regression model using each of the three prior year time period measures, as well as the own year average. The prior year measures again return results that are consistent with each other in terms of sign and significance. The coefficient estimates for these prior year measures are positive and highly significant. On the other hand, the own year average SPEI measure returns a negative coefficient estimate of no statistical significance. Recall that a negative SPEI value indicates a hotter, drier year, so the negative value for the own year coefficient estimate is more in line with an intuitive explanation of climate's effect on irrigation intensity. However, the lack of statistical significance means confidence cannot be placed in this result. It is possible that in Central Arizona, growers generally expect hot and dry conditions and plan their irrigation schedules accordingly, meaning the year-of climate would not impact irrigation intensity. Adversely, the somewhat unexpected positive coefficient estimates returned by the prior year measures are highly significant, and included in Chapter 5's empirical analysis. As above, this measure is chosen over the prior rainy season or the prior winter due to a greater amount of information being captured by a twelve month average. Additionally, choosing the same SPEI measure for both empirical models creates consistency in terms of the data required to specify each. It should be noted that, based on the results presented in this section, were one to choose to run the same models specified in Chapter 5 using one of the alternate prior year SPEI measures described above, coefficient estimates and overall measure of fit would likely change very little.

Table A7: Correlation Coefficients of SPEI Measures

	Full Own Year	**Full Prior Year**	**Prior Year Rainy Season (July - Dec)**	**Prior Year Winter (Nov - Jan)**
Full Own Year	- -	0.2072	0.0739	0.0365
Full Prior Year	0.2072	- -	0.7347	0.7771
Prior Year Rainy Season	0.0739	0.7347	- -	0.8791
Prior Year Winter	0.0365	0.7771	0.8791	- -

*** = $p < 0.01$, ** = $p < 0.05$, * = $p < 0.1$
(Standard Errors in Parentheses)

Table A8: Lagged SPEI Measures: Deliveries Models

	Full Prior Year	**Prior Year Rainy Season (July - Dec)**	**Prior Year Winter (Nov - Jan)**
R^2	0.9755	0.9768	0.9752
Coefficient Estimate	0.1104*** (0.035)	0.1301*** (0.021)	0.0753*** (0.015)

*** = p < 0.01, ** = p < 0.05, * = p < 0.1
(Standard Errors in Parentheses)

Table A9: Lagged SPEI Measures: Intensity Models

	Full Own Year	**Full Prior Year**	**Prior Year Rainy Season (July - Dec)**	**Prior Year Winter (Nov - Jan)**
R^2	0.8957	0.9052	0.9047	0.9005
Coefficient Estimate	-0.0415 (0.048)	0.1009*** (0.027)	0.0879*** (0.023)	0.0558*** (0.016)

*** = p < 0.01, ** = p < 0.05, * = p < 0.1
(Standard Errors in Parentheses)

A.4 Meetings and Contacts

Matthew Ford and Danielle Tadych, Research Assistants
University of Arizona, Department of Hydrology and Atmospheric Sciences
mford4@email.arizona.edu, dtadych@email.arizona.edu
November 19, 2020 - Zoom call
Topics: groundwater/pumping data, GIS data integration, tribal lands, academic literature

George Frisvold, Professor and Extension Specialist
University of Arizona, Department of Agricultural and Resource Economics
frisvold@ag.arizona.edu
January 7, 2021 - Zoom call
Topics: crop rotations, crop price data, irrigable area vs. planted area, pumping costs, irrigation technologies, federal commodities programs

Dave DeWalt, Arizona State Statistician
United States Department of Agriculture - National Agricultural Statistics Service
dave.dewalt@usda.gov
January 15, 2021 - Email
Topics: irrigation technology, Cropscape

Russell Tronstad, Professor and Extension Specialist
University of Arizona, Department of Agricultural and Resource Economics
tronstad@ag.arizona.edu
January 19, 2021 - Zoom call
Topics: effective prices for crops, expected revenue for farms, Central Arizona cropping patterns

Thomas Whipple, Hydrologist
Arizona Department of Water Resources
twhipple@azwater.gov
January 20, 2021 - Zoom call
Topics: water costs, energy providers, Cropscape, irrigation district water use

Tyler Fitzgerald, Agricultural Water Resources Specialist
Arizona Department of Water Resources
tfitzgerald@azwater.gov
February 5, 2021 - Email
Topics: irrigation district water use, public records requests, water rights/program codes

Ken Seasholes, Manager - Resource Planning & Analysis
Central Arizona Project
kseasholes@cap-az.com
February 10, 2021 - Zoom call
Topics: CAP and Central Arizona water supply, urbanization, ADWR data, tribal water data, pumping capacity, flex credits

Tyler Fitzgerald, Agricultural Water Resources Specialist
Arizona Department of Water Resources
tfitzgerald@azwater.gov
March 1, 2021 - Email
Topics: further clarification on irrigation districts' water use reporting

Hannah Hansen, Graduate Research Assistant
University of Arizona, Department of Agricultural and Resource Economics
hannahhansen@email.arizona.edu
March 6, 2021 - Zoom call
Topics: checking data compilation

Russell Tronstad, Professor and Extension Specialist
University of Arizona, Department of Agricultural and Resource Economics
tronstad@ag.arizona.edu
March 8, 2021 - Email
Topics: Expected prices for cotton

Laura Condon, Assistant Professor
University of Arizona, Department of Hydrology and Atmospheric Sciences
lecondon@arizona.edu
March 18, 2021 - Zoom call
Topics: GIS processing of Cropscape data

Craig Wissler, Associate Professor
University of Arizona, Department of Natural Resources and the Environment
cwissler@arizona.edu
March 23, 2021 - Zoom call
Topics: GIS processing of Cropscape data

Ron Klawitter, Principal - Water System Projects
Salt River Project
Ronald.Klawitter@srpnet.com
April 6, 2021 - Zoom call
Topics: District water reporting (water classifications, lands classifications, etc.), water delivery forecasting, Salt River Project characteristics

Dan Scheitrum, Assistant Professor
University of Arizona, Department of Agricultural and Resource Economics
dpscheitrum@arizona.edu
April 8, 2021 - Zoom call
Topics: Econometric modeling, robustness checks

Danielle Tadych, Graduate Research Assistant
University of Arizona, Department of Hydrology and Atmospheric Sciences
dtadych@email.arizona.edu
April 22, 2021 - Zoom call
Topics: GIS data processing

George Frisvold, Professor and Extension Specialist
University of Arizona Department of Agricultural and Resource Economics
frisvold@ag.arizona.edu
May 27, 2021 - Zoom call
Topics: econometric modeling, robustness checks, cotton price variables

Patrick Willis, Contractor
United States Department of Agriculture
patrick.willis@usda.gov
June 22, 2021 - Zoom
Topics: Cropscape, fallow/idle cropland assessment

Tyler Fitzgerald, Agricultural Water Resources Specialist
Arizona Department of Water Resources
tfitzgerald@azwater.gov
June 25, 2021 - Google Meet
Topics: Water rights classifications, San Carlos Irrigation and Drainage District, statewide irrigation district water deliveries data

Dan Scheitrum, Assistant Professor
University of Arizona, Department of Agricultural and Resource Economics
dpscheitrum@arizona.edu
September 2, 2021 - Zoom call
Topics: Econometric modeling, robustness checks

Tyler Fitzgerald, Agricultural Water Resources Specialist
Arizona Department of Water Resources
tfitzgerald@azwater.gov
October 16, 2021 - email
Topics: Deliveries out of district, deliveries to groundwater storage facilities

Ashley Kerna Bickel, Economic impact Analyst
University of Arizona, Department of Agricultural and Resource Economics
ashley.bickel@arizona.edu
November 17, 2021 - zoom
Topics: San Carlos Irrigation and Drainage District, gross crop revenues per acre

References

Alexander, T. G. (2002). Irrigating the Mormon Heartland: The Operation of the Irrigation Companies in Wasatch Oasis Communities, 1847-1880. *Agricultural History 76*(2), 172–187.

AMS/USDA (2021). U.S. Department of Agriculture, Agricultural Marketing Service, Cotton Market News, various issues. Online. Available at http://www.marketnews.usda.gov/portal/cn.

Arizona Department of Water Resources (2020). Irrigation districts within the active management areas.

Arizona Department of Water Resources (2021). Conservation.

Arizona Department of Water Resources GIS Data (2021). Irrigation district.

Arizona Power Authority (2021). About us.

Arizona Revised Statutes (2016). Arizona revised statutes 45-467.

Bayman, J. M. (2001). The Hohokam of Southwest North America. *Journal of World Prehistory*, 55.

Bickel, A. K. (2021). Private Communication, see Appendix A.4.

Bickel, A. K., D. Duval, and G. B. Frisvold (2019). Simple Approaches to Examine Economic Impacts of Water Reallocations from Agriculture. *Journal of Contemporary Water Research & Education 168*(1), 29–48.

Central Arizona Project (2016). Agriculture and the central arizona project. *Central Arizona Project*.

Central Arizona Project (2021a). Arizona heads into tier 1 colorade river shortage for 2022.

Central Arizona Project (2021b). Water rates.

Colby, B. (2016). Developing system conservation programs: A guide for the bureau of reclmation and partner organizations. *United States Bureau of Reclamation*, 1–123.

Condon, L. E. and R. M. Maxwell (2019). Simulating the sensitivity of evapotranspiration and streamflow to large-scale groundwater depletion. *Science Advances 5*(6), eaav4574.

Dell, M., B. F. Jones, and B. A. Olken (2014, September). What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature 52*(3), 740–798.

Deryugina, T. and M. Konar (2017). Impacts of crop insurance on water withdrawals for irrigation. *Advances in Water Resources 110*, 437–444.

Donaldson, D. and A. Storeygard (2016, November). The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives 30*(4), 171–198.

Duval, D., A. Bickel, G. Frisvold, and S. Perez (2019). Arizona's tree nut industry and its contributions to the state economy. *University of Arizona Cooperative Extension*, 49.

Erie, L., O. French, D. Bucks, and K. Harris (1982). Consumptive use of water by major crops in the southwestern united states. *US Departments of Agriculture Economic Research Service Conservation Research Report Number 29*, 42.

Fereres, E., D. A. Goldhamer, and L. R. Parsons (2003). Irrigation water management of horticultural crops. *HortScience 38*, 7.

Ferris, K. and S. Porter (2021). The myth of safe-yield: Pursuing the goal of safe-yield isn't saving our groundwater. *Kyl Center for Water Policy at Morrison Institute - Arizona State University*.

Fleck, B. E. (2013). Factors affecting agricultural water use and sourcing in irrigation districts of central arizona. *Universtiy of Arizona Department of Agricultural and Resource Economics - Master's Thesis*, 1–187.

Foster, T., N. Brozović, and A. P. Butler (2014). Modeling irrigation behavior in groundwater systems: Research Article. *Water Resources Research 50*(8), 6370–6389.

Frisvold, G. (2021). Private Communication, see Appendix A.4.

Frisvold, G. B. and S. Deva (2012). Farm size, irrigation practices, and conservation program participation in the us southwest. *Irrigation and Drainage 61*(5), 569–582.

Frisvold, G. B. and K. Konyar (2012). Less water: How will agriculture in Southern Mountain states adapt? *Water Resources Research 48*(5), 1–15.

George Cairo Engineering, Inc. (2021). Irrigation districts.

Global SPEI Database (2021). Global spei database.

Griffin, R. C. (2016). Water resource economics: The analysis of scarcity, policies, and projects. *MIT Press Books*.

Haacker, E. M., K. A. Cotterman, S. J. Smidt, A. D. Kendall, and D. W. Hyndman (2019). Effects of management areas, drought, and commodity prices on groundwater decline patterns across the High Plains Aquifer. *Agricultural Water Management 218*, 259–273.

Hanak, E., A. Escriva-Bou, B. Gray, S. Green, T. Harter, J. Jezdimirovic, J. Lund, J. Medellín-Azuara, P. Moyle, and N. Seavy (2019). Water and the Future of the San Joaquin Valley. pp. 100.

Kahil, M. T., A. Dinar, and J. Albiac (2015). Modeling water scarcity and droughts for policy adaptation to climate change in arid and semiarid regions. *Journal of Hydrology 522*, 95–109.

Klawitter, R. (2021). Private Communication, see Appendix A.4.

Lahmers, T., S. Eden, and J. Polle (2018). Arroyo 2018: Water and irrigated agriculture in arizona. *University of Arizona Water Resources Research Center*, 1–16.

Larsen, A. E., B. T. Hendrickson, N. Dedeic, and A. J. MacDonald (2015). Taken as a given: Evaluating the accuracy of remotely sensed crop data in the USA. *Agricultural Systems 141*, 121–125.

Libecap, G. D. (2007). Owens valley revisited, a reassessment of the west's first great water transfer. *Stanford University Press*.

MacDonald, G. M. (2010). Water, climate change, and sustainability in the southwest. *National Academy of Sciences*.

McEvoy, D. J., J. L. Huntington, J. T. Abatzoglou, and L. M. Edwards (2012). An evaluation of multiscalar drought indices in nevada and eastern california. *Earth Interactions 16*, 18.

McGreal, B., S. Eden, and J. Polle (2021). Arroyo 2021: Arizona groundwater management - past, present, and future. *University of Arizona Water Resources Research Center*, 1–16.

Megdal, S. B. (2012). Arizona groundwater management. *The Water Report 104*, 7.

Moffitt, J. L., D. Zilberman, and R. E. Just (1978). A "putty-clay" approach to aggregation of production/pollution possibilities: An application in dairy waste control. *American Journal of Agricultural Economics 60*, 452–459.

Moore, M. R. and A. Dinar (1995). Water and land as quantity-rationed inputs in california agriculture: empirical tests and water policy implications. *Land Economics 71*, 11.

Ottman, M. J. (2015). Planting dates for small grains in arizona. *University of Arizona Cooperative Extension*, 2.

Pfeiffer, L. and C. C. Lin (2014). The Effects of Energy Prices on Agricultural Groundwater Extraction from the High Plains Aquifer. *American Journal of Agricultural Economics 96*(5), 1349–1362.

Qiao, X. (2018). The water use and climate effects on farm profiability in colorado river basin. *University of Arizona Department of Agricultural and Resource Economics - Master's Thesis*, 1–51.

Reisner, M. (1986). Cadillac desert: The american west and its disappearing water. *Viking*.

Richter, B. D., S. Andrews, R. Dahlinghaus, G. Freckmann, S. Ganis, J. Green, I. Hardman, M. Palmer, and J. Shalvey (2019). Buy Me a River: Purchasing Water Rights to Restore River Flows in the Western USA. *Journal of the American Water Resources Association*, 1752–1688.12808.

Sall, I. (2019). Agricultural producers' decision-making and preferences in relation to economic incentives, climate and weather conditions. *University of Arizona Graduate College Arid Lands and Resource Sciences Program - PhD Dissertation*.

Sall, I. and R. Tronstad (2021). Simultaneous analysis of insurance participation and acreage response from subsidized crop insurance for cotton. *Journal of Risk and Financial Management 14*, 562.

Sammis, T. and E. Herrera (1999). Estimating water needs for pecan trees. *New Mexico State University Cooperative Extension*, 2.

Seasholes, K. (2021). Private Communication, see Appendix A.4.

Silber-Coats, N., S. Eden, and J. Polle (2017). Arroyo 2017: Arizona water banking, recharge, and recovery. *University of Arizona Water Resources Research Center*, 1–16.

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*(4), 985–1023.

Tronstad, R. (2021). Simultaneous analysis of insurance participation and acreage response from subsidized crop insurance for cotton. *Arizona Cooperative Extension - University of Arizona*, 1–51.

United States Geological Survey, National Water Dashboard (2021). San carlos reservoir at coolidge dam, az.

US Bureau of Reclamation (2021). Projects & facilities.

US Census Bureau (2021). Population.

US Department of Agriculture, National Agricultural Statistics Service (2021). Cropscape and cropland data layer - metadata.

Western Area Power Administration (2021). Desert southwest's customer list.

Whipple, T. (2019). Planning for depletion: Optimal irrigation in the pinal ama under changing water sources. *University of Arizona Department of Hydrology and Atmospheric Sciences - Master's Thesis*, 1–44.

Willis, P. (2021). Private Communication, see Appendix A.4.

World Population Review (2021). World population review - us counties.

York, A. M., H. Eakin, J. C. Bausch, S. Smith-Heisters, J. M. Anderies, R. Aggarwal, B. Leonard, and K. Wright (2020). Agricultural water governance in the desert: Shifting risks in central arizona. *Water Alternatives 13*(2), 418–445.

Yu, J., A. Smith, and D. Sumner (2018). Effects of crop insurance premium subsidies on crop acreage. *American Journal of Agricultural Economics*, 91–114.

Zeff, H., D. Kaczan, G. W. Characklis, M. Jeuland, B. Murray, and K. Locklier (2019). Potential Implications of Groundwater Trading and Reformed Water Rights in Diamond Valley, Nevada. *Journal of Water Resources Planning and Management 145*(6), 1 – 19.