Essays on Water Resource Economics

Daniel A. Brent

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2013

Reading Committee:

Hendrik G. Wolff, Chair

Neil Bruce

Joseph H. Cook

David Layton

Program Authorized to Offer Degree:
Department of Economics

University of Washington

Abstract

Essays on Water Resource Economics

Daniel A. Brent

Chair of the Supervisory Committee:
Assistant Professor Hendrik G. Wolff
Department of Economics

A canonical example in economics of the difference between marginal and total value is the diamond-water paradox. The high price of diamonds is derived from their rarity; whereas the price of water is low due to its abundance, even though it is essential to sustain human life. Scarcity, rather than abundance, better characterizes water availability for many people and this dissertation studies how applied economic principles can be utilized to manage water resources. The first chapter estimate the costs of water volatility in the agricultural sector through a hedonic analysis of heterogeneous water rights. Security for water rights is capitalized into the value of agricultural land, which informs the magnitude and distributional welfare effects of droughts. Tests for an endogenous changepoint fails find a time-varying price premium, indicating that the costs of increased water volatility due to climate change are not manifested in agricultural property markets. The second and third chapters focus on economic and behavioral incentives in urban municipal water demand. Chapter 2 presents a disaggregated model of water demand to separately estimate intensive and extensive margin demand elasticity. Identification is achieved through a novel method merging remotely sensed satellite data on vegetative cover with water metering records. The time series of vegetative cover captures changes in landscape over time and identifies the extensive margin elasticity - a parameter that has only been estimated implicitly through the difference in short run and long run demand. Households that maintain green lawns are less responsive to prices than households either change landscapes or have a mixed landscape. Higher water rates increase the probability of converting to low water-intensive landscapes, which in turn is a major driver of long-run demand. The extensive margin with respect to changing landscapes comprises 7%-48% of total elasticity for households with significant outdoor water use. The final chapter examines the impact of non-pecuniary incentives stemming from the behavioral economics literature on water demand. In a randomized field experiment social comparisons are found to significantly decrease water demand with substantial heterogeneity both across and within utilities. The utility with the highest average treatment effect saved three times as much water in percentage terms as the utility with the lowest average treatment effect. Higher users are more responsive to the program and there are important interactions between social norms and existing utility conservation programs. Water resources face stress due to population growth, rising incomes, and climate change and these stressors will only increase in the future. This dissertation addresses several key issues in agricultural and residential that aim to increase knowledge and aid public policy of managing water resources in times of scarcity.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ACKNOWLEDGMENTS

Graduate school produced the most difficult and enjoyable moments of my life. I learned that research is hard work and it is real challenge to finish a dissertation. There were times when I doubted whether I would finish my PhD, and the fact that I did is a testament to the help and support of a great number of people. First and foremost I would like to thank all of the academic and professional staff in the Economics Department at the University of Washington. My first two years of coursework trained me to think like an economist as well as arming me with the technical tools to pursue economic research. In particular I would like to thank all the members of my committee, especially my advisor Hendrik Wolff. Hendrik provided a consistent reminder to ensure that each research topic has a clear question and the importance of lucid and compelling writing. I also benefited from several experiences outside the economics department. Joe Cook has been an invaluable source of support and a mentor not only for water research, but also how to be an academic. Working on a grant under Sergey Rabotyagov funded a good portion of my graduate work and I learned how to work in an interdisciplinary setting. Sergey also proved a stellar employer, delicately balancing between allowing me to work autonomously and providing feedback when necessary. His guidance serves as a template for how to interact with my future research assistants.

I am incredibly grateful for the rewarding semester I spent as a predoctoral fellow[1] at Arizona State University through the Center for Environmental Economics and Sustainability Policy. The program granted unique access to some of the top established and emerging scholars in environmental economics. I honed my research skills at Arizona State and learned about the benefits of coordinating with industry and government partners outside

---

[1] See the website for more information on this unique program - http://legacy.wpcarey.asu.edu/CEESP/pre-doctoral-program/recipients.cfm.

of academia. I will reserve a special gratitude towards Kerry Smith, Michael Hanemann, Nick Kuminoff, and Josh Abbott for their guidance during my time at ASU as well as their continued support.

I am also indebted to many outside of academia for providing industry and policy expertise as well as incredible patience and accommodation with data requests. Doug Frost and Adam Miller at the City of Phoenix provided water metering records as well as essential advice that informed my analysis. Peter Yolles, Ora Chaiken, Chad Hayes, and Doug Flanzer of WaterSmart expended enormous time and effort providing me with data, as well as explaining the intricacies of their company and utility partners. If that were not enough they also opened their doors to allow me to work in their San Fransisco office for a week in each of my last two years of graduate school. I value all that I learned from those working in the water industry and hope I maintain a strong relationship those that can apply academic research in both the private and public sector.

Lastly, I would like to thank everyone in my personal life you helped me deal with the emotional highs and lows during graduate school. I leaned on the support of friends and family not only to cope with the stress of scholarly work, but to provide an outlet to refresh my body and mind. My classmates, particularly Skylar and Austin, turned into colleagues and coauthors who challenged and encouraged me academically while at the same time made work an enjoyable place to spend so much of my part of my life. My parents were more supportive and encouraging than any son could expect, and they will forever have my gratitude for helping my through this time. Of course I need to thank my dear wife Aliza, who could probably give my job market talk for all the times she has heard it. My lovely wife exhibited patience and understanding during lonely winter nights and gave me confidence that my research was worthwhile and interesting at time when my faith wavered.

It takes a village to raise a child, and, at least in my experience, it takes a community to produce a PhD.

DEDICATION

to my dear wife, Aliza

Chapter 1

THE VALUE OF HETEROGENEOUS PROPERTY RIGHTS: THE COSTS OF WATER VOLATILITY

1.1 Introduction

Research on the efficiency and distributional effects of water rights' institutions in the United States dates back to the inaugural issue of the American Economic Review [35] and as noted in [90] these issues are still relevant today. Water rights west of the 100th meridian in the United States are based on prior appropriation; a system where priority is defined by first in time first in right. Property owners that first establish rights in a water basin have senior rights that have priority in times of drought over junior rights that were instituted at a later date. This contrasts to the riparian regime in most of the eastern (and wetter) US that defines rights to reasonable use as tied to the land. [34] describes the the economic implications of differences between the two sets of institutions. The riparian regime includes legal uncertainty since new users along a water source may dilute the supply, while prior appropriation suffers primarily from physical uncertainty due to the water supply volatility. The complexity in the prior appropriation doctrine leads to costly adjudication, present in almost all western states, to resolve conflicts regarding both the quantity and priority of water rights.[1] As described in [23] in the absence of a competitive market for rights the appropriate system creates inefficiency due to the unequal sharing of risk between property owners with junior and senior rights. Risk in water rights, particularly in areas governed by the appropriative regime, is proportional to the water volatility in a watershed. Thus the institutional setting directly links water supply volatility to the welfare loss. This chapter models water volatility as the driver of a price premium for senior water rights in agricultural property markets and estimates the magnitude of this premium in a hedonic price model. In addition to estimating the value of senior water rights the hedonic model informs policymakers on the distributional effects of water scarcity, since the majority of costs associated with drought will fall on those that own junior water rights.

The obfuscation of property rights generates transaction costs impeding the efficient reallocation of water through markets to address threats of water scarcity associated with

---

[1] Data available at http://www.judges.org/dividingthewaters/dtw-links.html.

climate change [90]. The requirement to put water to a beneficial use, the "use it or lose it" clause for most areas that follow prior appropriation, adds additional uncertainty that make landowners reluctant to subject their water rights to the scrutiny necessary in a transaction for water rights. There are concerns climate change will exacerbate the inadequate risk sharing in agricultural water rights due to an increase in water volatility. Climate models predict that regions around the world will face more variable water supplies; due to changes in precipitation patterns and higher temperatures resulting in less water stored as snowpack. In particular, the Western United States is expected to experience more frequent and severe droughts in the summer the season of peak water demand [14]. Quantifying the heterogeneity in water rights is crucial for determining the distributional impact from climate change, as owners of low-priority rights will bear most of the costs of volatile water supplies. The value of priority in water rights, referred to as seniority, is difficult to directly estimate through water right transactions due the thinness of water markets. The goal of this paper is to first estimate the value of security in agricultural water rights, and second, to test if the premium paid for more secure rights increases over time in response to expectations about water volatility associated with climate change.

The idiosyncrasies of water institutions and the paucity of quality data on rights present challenges for a national or multi-state study on the economics of heterogeneous water rights. The Yakima River Basin in central Washington provides a suitable case study due to the dichotomous division of water rights in the basin and high quality data. High priority (also referred to as senior or non-proratable) rights were established before 1905 while all subsequently established rights are designated as junior (or proratable) and are subject to curtailment when water supply falls short of total entitlements.[2] There has never been an incidence of curtailment of senior rights; thus priority effectively insulates farmers from temporal shocks to the water supply. Downscaled climate models of the Pacific Northwest predict that the annual variance of the region's water supply will increase [135], resulting

---

[2] There are actually three levels of water rights, with tribal water rights having the highest level of priority. However, in practice there has never been any conflict between senior water rights and tribal rights.

in more years where the region experiences water shortages. Water shortages that have historically occurred in 14% of years will increase to 77% of years during the 2080s for the IPCCs A1B scenario [135]. Shortages do not stem from a decrease in total precipitation; rather climate change predominantly affects the water available during the irrigation season, from April to September. Intra-annual variation is predicted to be more extreme, with a higher percentage of rain falling during the winter. Lower volumes of snowpack will further reduce water available for irrigation. If farmers expect that climate change impacts their water resources, or will do so in the future, the value of land with senior rights will rise relative to land with junior rights.

To evaluate the theory first I employ the hedonic price model to estimate the premium associated with a senior water right. Next I test if the premium changes over time in accordance with farmers belief about increasing water variability, potentially due to climate change. Initial results indicate that the additional security associated with a senior water right adds 9-12% of the value of a farm. There is no a priori designation for when climate change begins to impact landowners expectations so I use a model with an endogenous change point to test for a time varying premium on senior water rights. Results show that the change point is at the end of the sample evidence that there really may not be a change point at all. While this approach imposes parametric restrictions on the form of the time-varying parameter it provides a starting point to test for behavioral response to climate change in the property market. Incorporating survey or polling data provides a more flexible definition of climate change and is left for further research. While this study focuses on the impacts in the Yakima River Basin of central Washington, the phenomenon of increasing volatility of water supply applies to many regions facing a changing climate, particularly those that rely on snowpack as a source of water supply in the summer. The rest of the paper is organized as follows. Section 1.2 presents a background of the literature and the study area, Section 1.3 introduces the economic model, Section 1.4 describes the data, estimation methodology empirical results, and a policy application, and Section 1.5 concludes.

1.2 Background

1.2.1 Existing Literature

The contribution of this paper is to value heterogeneous water rights and relate the relative value of secure rights to water volatility. The first application of the hedonic price model value to water rights was [37], who tests functional forms of the hedonic price function to determine characteristics of the water market. Later studies estimate heterogeneity in the value of water due to differences in the productivity of the land [51] and the ecological value of in-stream flow [104]. However, there are no studies that estimate the impact of variations in the right itself. Other research analyzing property rights with varying degrees of security focuses on land rights in the developing world [67]. [90] presents qualitative analysis on the appropriative rights system and its effects on the efficient allocation of water between and within sectors. The economic literature on estimating the costs of a variable water supply developed by [130] builds on the research of optimal groundwater extraction [24]. [130] coin the phrase Stabilization Value (SV) to explain the benefits from fixing a variable water supply at its mean. Research on the SV of water ranges from a static analysis outlining the benefits to buffering surface water with groundwater to a dynamic stochastic general equilibrium model [130, 42]. Production function approaches are appropriate in a setting where the production function is static; but are biased if farmers change crops, irrigation and fertilization technologies, or land use [97]. Alternatively, using property values to estimate the effect of water supply volatility incorporates the potential of landowner adaptation to changing economic and environmental conditions. [97] apply the Ricardian approach to estimate the impact of climate variables on the agricultural sector to avoid the bias in production function studies. The Ricardian approach utilizes the theory that land values should reflect the discounted value of expected profits, and therefore land rents are capitalized into farm values. National research on the economic value of water resources on agricultural land focuses on average precipitation see [120, 41] among others. [96] add surface water and a measure for water variance as independent variables in the Ricardian approach and find that surface water increases farm values while water variance depresses farm values. While these articles rely on county level data, [121] use farm-level data in California to show that water availability strongly capitalizes into farm prices. The Ricardian studies use spatial variation to identity climate variables. Since the results are derived for a state or the entire country they are more applicable in determining aggregate effects of climate change. Hedonic models, in contrast, often limit the sample to a small geographic area such as one particular county. This permits data with greater detail and inherently controls for factors that vary spatially such as precipitation, average temperature, and institutions. In fact [121] intended to use water rights to describe water access, but the system of water rights in California proved too tortuous to obtain water rights data of sufficient quality. Conversely, hedonic models explicitly value irrigation water or groundwater with micro-level data on water rights or permits for digging wells [26, 37, 51, 104, 109]. There is evidence that pooling irrigated and non-irrigated land is not appropriate in identifying the effect of climate on farmland values since precipitation and temperature have very different impacts when land is augmented by irrigation [120, 57]. An advantage of this research is that all land has access to irrigation, and thus circumvents the differential effects of climate on irrigated and dryland agriculture.

1.2.2 Agriculture in the Yakima River Basin

The Yakima River Basin in Central Washington State provides an excellent test case to examine the interaction of priority in water rights and water supply volatility because landowners with junior rights bear the preponderance of the costs due to drought. The Yakima River Basin is one of Washington's largest agricultural producer, contributing close to 20% of the state's $9.2 billion worth of agricultural output in 2011.[3] Much of the land east of the Cascade mountain range in Washington State is very dry and relies on irrigation for agriculture. The Yakima basin is therefore susceptible to severe economic losses from

---

[3] Data are available at http://agr.wa.gov/AgInWa/docs/126-CropProductionMap12-12.pdf - accessed 3/5/2013.

drought. The Yakima Basin Storage Alliance [133] estimates over \$130 million in economic losses from decreased agricultural production from the 2001 drought alone.[4] The vast majority of these losses fell on farmers with junior water rights, while farmers with senior rights still received their full water allotment, allowing them to proceed with normal farming operations. Increased frequency of severe drought years will diminish the relative value of farmland with junior water rights. Rational landowners will react to threats to water volatility, and this research tests whether they consider climate change as a real threat to their water supply. The next sections describe the features of the Yakima basin, and motivate the use of water rights to test for expectations of water supply volatility.

1.2.3 Water Supply in the Yakima Basin

The Yakima River basin contains parts of Kittitas, Yakima, and Benton County, though Benton receives much of its water from the Columbia River [132]. Most of the precipitation in the regions falls between October and March [131, 138], and this trend will increase in the future based on climate models by [135]. The major water use in the region is irrigated agriculture met predominantly by surface water. Five major reservoirs operated by the U.S. Bureau of Reclamation (USBR) with a combined total capacity of 1.07 million acre-feet (maf) serve six irrigation districts and a storage division that constitute the Yakima Project. Below Parker Gage, the major control point of the Yakima Project, the water supply is augmented by return flows from upstream use. The six irrigation districts served by the Yakima Project represent over 80% of the total water entitlements in the Yakima basin above Parker Gauge. This fraction increases when non-federally supplied irrigation districts are included, justifying the use of irrigation districts to analyze the impact on the region's agricultural sector.

The USBR operates reservoirs with the joint goals of flood control and the provision of irrigation water from April through September. Melting snowpack effectively acts as a

---

[4] Data are available at http://www.ybsa.org/agriculture.php - accessed 12/2/2012.

sixth reservoir typically allowing the USBR to wait until June to begin drawing down the reservoirs for irrigation [131]. Warmer temperatures cause earlier snowmelt, preventing the use of snowmelt during the irrigation season and reducing its substitutability with reservoir water. Therefore the quantity and timing of snowpack is crucial to the water supply system in the Yakima. Figure 1.1 illustrates historical deviations from mean withdrawals for each irrigation district in the Yakima project separated by the priority of water rights. There is a trend over time towards fewer withdrawals due to improvements in irrigation technology, conservation, and crop choice. Total annual diversions are relatively stable until around 1970 but since the 1990s the basin experiences violent dips in water use due to severe droughts that are particularly acute for the districts with a majority of junior rights. Kennewick Irrigation District (KID) only has junior rights but their position below Parker Gage allows some water to return in the form of recharge from upstream users as evidenced by smaller declines in withdrawals during droughts. The figure displays how senior water rights insulate landowners from water supply volatility, and motivates that the premium for this protection may be a function of climate change expectations.

1.2.4 Water Rights in the Yakima Basin

The institutions governing water rights in Yakima River basin simplifies estimating the costs of water volatility due to the dichotomous distinction of priority based on the date that beneficial use was established. All rights established prior to 1905 are classified as senior, or non-proratable, rights and all rights post-1905 are designated as junior rights, or proratable. The law requires that senior right holders receive their full water allotment before honoring any junior right. Therefore, when supply is insufficient to fulfill the total apportionment of water rights in the basin senior right holders receive their entire water commitment, and junior users divide the remaining water on a prorated basis. For example, consider 50 landowners with junior rights and 50 with senior rights where everyone has access to 1 ac-ft per year. If the water supply is 80 ac-ft in a specific year all the landowners with senior rights

Figure 1.1: Annual Deviations from Mean Diversions by District

Junior Rights Districts

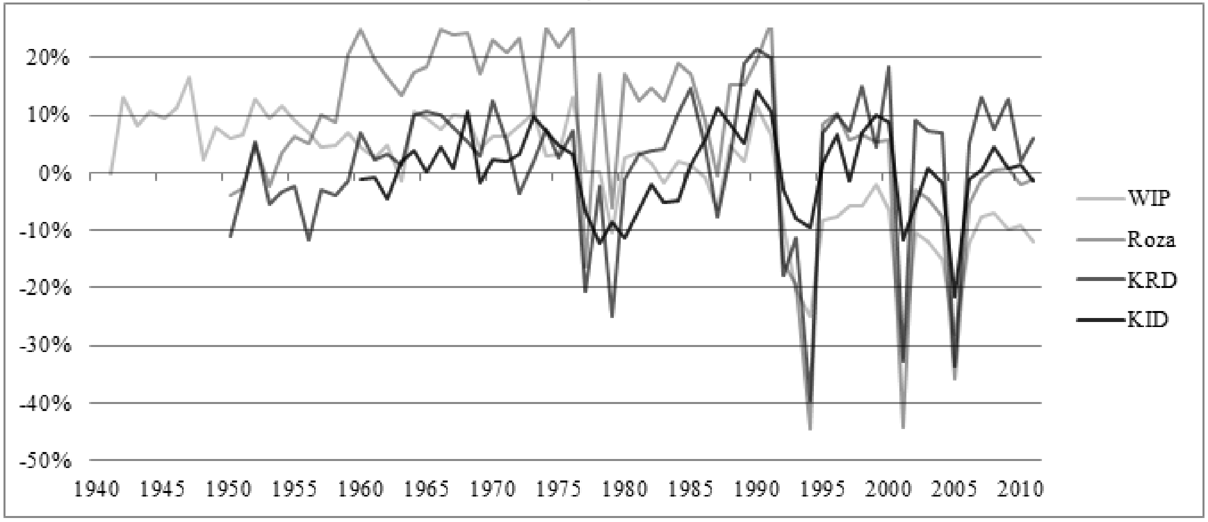

Senior Rights Districts

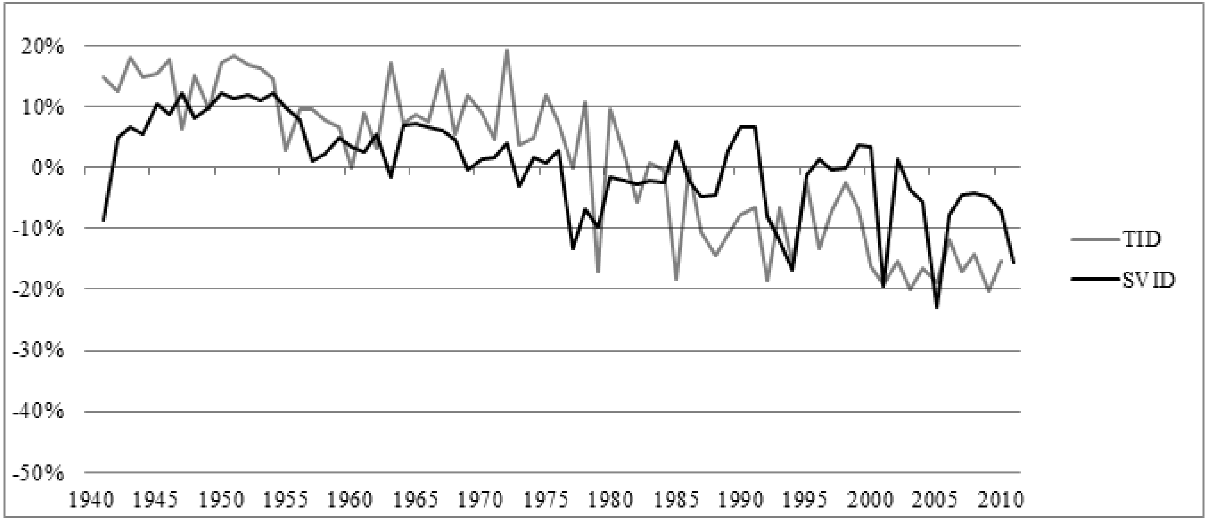

Notes: Annual deviations from the mean are shown in percentage terms by irrigation district. TID and SVID are identified as senior district while KRD, Roza and WIP are junior districts based on the Integrated Plan (USBR 2012). Even though KID owns predominantly junior rights it receives recharge water from withdrawals upstream and is therefore less susceptible to droughts. Data are from USBR via Chris Lynch.

get their full share (1 ac-ft each) while those with junior rights are prorated at 60% since the 50 junior landowners must split the remaining 30 ac-ft. The USBR determines the proration level at the beginning of the irrigation season based on forecasts of the Total Water Supply Available (TWSA), and adjusts the degree of prorating throughout the season in response to changing weather conditions. From 1970-2005 junior rights experienced prorating in 13 years whereas senior right holders have never been affected by prorating. Therefore junior water rights holders are more susceptible to seasonal and annual variation and will bear the majority of the costs due to climate change affects water volatility in the basin.

Approximately 55% of the surface water rights in the basin are proratable, leaving a significant portion of farmers without water during a drought. Several irrigation districts have all senior rights and some districts have a mix of both non-proratable and proratable rights. I distinguish the districts with both types of rights based on the two reports from the USBR [134, 132] that indicate Roza, KRD, and WIP all suffer severely from prorating during drought years. The highest proportion of senior rights in these districts is Wapato with 49% senior rights so I set this as the cutoff for a district that is defined as senior. This cutoff conforms with the literature [135] that prorating is particularly damaging below 70% and the fact that junior districts experience withdrawal reductions more than 30 below their historical average in Figure 1.1. Table 1.1 shows the properties in our sample by irrigation district with the percentage of non-proratable rights and a junior or senior designation for the district. The sample matches up closely with the population of water rights with 52% of properties having predominately junior rights.

In theory an active water market will alleviate some of the costs of water shortages by distributing water from low-value uses to activities with higher marginal value. While substantial gains to trade exist in years where junior water users suffer from severe prorating, water transactions developed slowly and there is still not a well-functioning water market in the region. Beginning in 2001 the Yakima basin initiated a water trading program during emergency drought conditions, as declared by the state. However, two complications prevent the operation of a competitive water market in the region. First, the necessary infrastructure to transfer water between all interested agents does not exist and second, legal obfuscations generate disinclinations to engage in trade. Water rights in Washington require the user to establish beneficial use, and if water remains idle for five consecutive years an owner relinquishes their right. Farmers are often hesitant to sell water because they need to prove they put the water to beneficial use if required to defend their right in court. Another concern is that the examination of the water right during the transaction may reveal that the right is not valid, or represents a smaller quantity of water than actively used by the farmer.

Table 1.1: Water Rights and Irrigation Districts

Name	# of Sales	% of Senior Rights	Senior Designation
Ahtanum	23	100	Yes
Buena	6	100	Yes
Cascade	39	100	Yes
Columbia	30	100	Yes
Ellensburg Water	35	100	Yes
Kittitas Reclamation	186	6.9	No
Kennewick	201	7.8	No
Moxee-Selah	16	85.6	Yes
Naches-Selah	32	91.1	Yes
Roza	471	0	No
Sunnyside Valley	0	72.5	Yes
Union Gap	0	79.1	Yes
Wenas	14	0	No
West Side	0	75.8	Yes
Yakima-Tieton	0	65	Yes
Yakima-Wapato	268	48.7	No
Total	2,166	-	1,026

Notes: The table shows observations from each district by row and the type of water rights in the columns. The sample has a slightly higher proportion (52%) of senior water rights than the population (55%) in the basin.

1.2.5 Climate Change in the Yakima Basin

Water curtailments occur relatively frequently for junior water users, though when prorating is above 70% of normal entitlements farmers can generally cope by changing variable inputs and the timing of irrigation [135]. So even though all prorating has costs, the most severe burden occurs in years where junior farmers receive less than 70% of their water right. According to downscaled climate models by [135] precipitation will increase in the cool months and decrease during irrigation season. Rising temperatures will decrease the snowpack available, exacerbating water shortages for the agricultural sector. Historically severe prorating occurred in 14% of years, but this is predicted to increase to 27-77% depending on the emissions scenario [135]. On the demand side rising temperatures will lead to higher evapotranspiration rates, increasing the water requirement of crops between 3% 9.8%, depending on the area and study methodology [134]. In summary, climate changes will exert pressure on water supply and demand through reduced precipitation during the irrigation season, earlier snowpack, and higher temperatures. Furthermore, rising water supply volatility will increase the years where prorating goes below 70%, predominantly impacting farmers with proratable water rights.

The Yakima River Basin Water Enhancement Program (YRBWEP) is evidence of the regions focus on addressing water scarcity. Beginning in 2009 the USBR and the Washington State Department of Ecology (ECY) began work on the YRBWEP with the goal of producing a Final Water Resources Integrated Management Plan (henceforth Integrated Plan). In addition to the two government agencies, members from the agricultural, environmental, legal, real estate, municipal and tribal communities participate as stakeholders in dealing with water scarcity in the region. If implemented, the Integrated Plan will cost between \$3.2-\$5.6 billion, with a base estimate of \$4 billion [133]. More than half of the expenditure will go towards enhancing the basin's storage capacity by constructing a new reservoir and upgrading existing storage facilities. A benefit cost study estimates that augmenting water resources through the Integrated Plan will increase irrigated agricultural production by \$400 million in net present value. This value comes solely from eliminating losses for farmers with junior water rights during droughts that cause less than 70% prorating under historical hydrologic conditions [133]. The cost estimates are biased downward because changes in water scarcity associated with climate change, and droughts resulting in prorating above 70% do not enter into the calculation. Conversely, the estimates do not account for adaptation such as crop switching or changes in irrigation technology, both of which ameliorate damages from droughts. Using property values to estimate the benefit of secure water availability will improve the methodology to quantify the benefits of the Integrated Plan

1.3 Economic Model

I use the hedonic price model to estimate the implicit value of a senior water right in the Yakima basin. [118] develops the hedonic price model in application to the residential housing market, and Palmquist and coauthors [107, 108] extend the model to land used for agricultural production. I derive the demand side of the market for agricultural land using per-acre variable profits gross of land payments, π_t^V

$$\pi_t^V = \mathbf{p_t} f_t(V_t, \mathbf{X}, W, \alpha) - c_t(V_t, \alpha) \tag{1.1}$$

where $\mathbf{p_t}$ is a vector of crop prices at time t, and f_t is the multiple output production function at time t that depends on $\mathbf{X}$, a vector of fixed attributes of the land, α, a farmer-specific unobserved skill parameter, W_t, the water availability on the land at time t, and $\mathbf{V_t}$, a vector of variable inputs. The cost function, c_t, depends on variable inputs and the idiosyncratic skill parameter. A farmer chooses $\mathbf{V_t}$ to maximize profits for any combination of $\mathbf{p_t}$, $f_{jt}()$, $\mathbf{X}$, W, and α, such that optimal profits can be expressed as,

$$\pi_t^{*V} = \pi_t^{*V}(\mathbf{p_t}, \mathbf{X}, W_t, \alpha) \tag{1.2}$$

The maximum bid that a farmer pays for a specific piece of land for use at time t is determined by the inputs of the profit function, as well as the desired net profits, π_t.

$$\theta_t(\mathbf{p_t}, \mathbf{X}, W_t, \alpha) = \pi_t^{*V}(\mathbf{p_t}, \mathbf{X}, W_t, \alpha) - \pi_t \tag{1.3}$$

By differentiating (3) it can be shown that $\frac{\partial \theta_t}{\partial X_i} = \frac{\partial \pi_t^{*V}}{\partial X_i}$ and $\frac{\partial \theta_t}{\partial W} = \frac{\partial \pi_t^{*V}}{\partial W}$. The derivative of the rental bid function is non-decreasing and concave in any desirable characteristic X_i and W_t, given typical assumptions of the variable profit function [43]. In equilibrium the marginal increase in variable profits must equal the marginal increase in the bid function, which in turn equals the rental price of land. The equilibrium rental schedule of land is an envelope of the bid functions. While equation (3) describes the decision for renting land for one-period, iterating the process into the future shows that the equilibrium sale price of land is equal to the expected discounted sum of future variable profits. In this context the increase in the market price, q_t, from a marginal increase in any attribute X, or W_t, will be the change in the discounted sum of expected current and future profits due to the extra amount of the attribute.

$$q_t(\mathbf{p_t}, \mathbf{X}, W_t, \alpha) = \sum_{h=t}^{\infty} E_h\left[\pi_h^{*V}(\mathbf{p_h}, \mathbf{X}, W_h, \alpha)\right] e^{-\beta h} \tag{1.4}$$

Analyzing the bid function for a permanent purchase of land as opposed to a one period rental iterates the process forward, where Θ_t is the bid for a permanent land purchase and $\bar{\pi}_t$ is the expectation of future net profits.

$$\Theta_t(\mathbf{p_t}, \mathbf{X}, \pi_t, W_t, \alpha) = \sum_{h=t}^{\infty} E_h\left[\pi_h^{*V}(\mathbf{p_h}, \mathbf{X}, W_h, \alpha)\right] e^{-\beta h} - \bar{\pi}_h \tag{1.5}$$

In this forward looking model $\frac{\partial \Theta_t}{\partial W_t} = \frac{\partial \sum_{h=t}^{\infty} E_h[\pi_h^{*V}]e^{-\beta h}}{\partial W_t}$; the marginal increase in the bid on for land with better water resources equals the increase in the expected sum of discounted profits due to the water. This setup models the farmers willingness to pay for secure water supply according to their expectations of the change in future profits. The literature on the stabilization value [42, 130] adds a theoretical background to the interpretation of a water right as an attribute of the hedonic price function. Let the premium on a senior water right S relative to a junior right J given all the characteristics of the property be defined as $E[P|X,S] - E[P|X,J] = \gamma$. Given that senior water rights are never prorated[5], the premium is equal to the revenue from a fixed quantity of water less the expected revenue from a variable water supply as seen in equation (6). Where the distribution of water $W \sim g(\mu, \sigma^2)$ can be described by its mean and variance , and $\tilde{\pi}()$ is the profit function optimized with respect

---

[5] This is likely a valid assumption considering senior rights have never been prorated.

to all other inputs conditional on W.[6]

$$\gamma = \sum_{h=t}^{\infty} \widetilde{\pi_h^{*V}}(\mu)e^{-\beta h} - \widetilde{\pi_h^{*V}}(W_h)e^{-\beta h} \tag{1.6}$$

Note that the expectation operator is only applied to the profit for a junior landholder since their water input depends on the random variable W while senior landowners profits depend on the constant μ. A Taylor series approximation of the junior landowner's expected profit, $E[\widetilde{\pi_t^{*V}}(W_t)]$, allows for the premium to be written as a function of the variance of the water supply.

$$\gamma = \gamma(\sigma^2) = -0.5\widetilde{\pi_t^{*V}}''(\mu)\sigma^2 \tag{1.7}$$

This value is positive if the production function is concave in the water input, implying a diminishing marginal value of water.[7] Whether (7) holds in practice likely depends on the setting, particularly the domain of $\tilde{\pi}$. In this setting it appears feasible due the public discourse on the costs of water scarcity. The analysis does not rely on this assumption, but rather is testing it directly by estimating γ as an attribute in the hedonic price with only the traditional assumptions in the hedonic model. The key point is that if $\gamma > 0$ then $\frac{\partial\gamma}{\partial\sigma^2} > 0$ is likely; and estimating a time-varying premium is an indication of changing expectations of water supply volatility.

1.4 Data & Estimation

1.4.1 Data

The primary data are sales of agricultural properties within an irrigation district located in the Yakima River Basin obtained from assessor offices for Kittitas, Yakima, and Benton

---

[6] This assumption is relatively mild because the important aspect is landowners perceptions of the distribution of the water supply which are unlikely to encompass anything beyond the first two moments.

[7] This assumption is difficult to assess because the profit function may not be continuous in water. There may be kinks where the water input causes the loss of a substantial portion of the crop or causes perennial crops such as fruit trees to die. Additionally certain regions of the support may reflect changes in crop choice or land use.

Counties in Washington State. The assessors office also provides sales, zoning, land use, market improvements, and irrigation district boundaries. The sales data and the irrigation district boundaries are both geo-referenced allowing each parcel to be placed within an irrigation district using Geographical Information Systems (GIS) software; dropping sales of parcels outside of irrigation districts. Using sales from irrigation districts alleviates the problem of tracking distinct water rights for individual parcels. While most water rights remain with a physical parcel of land it is possible for a landholder to sell all, or a fraction of, a water right; obfuscating the link between a water right and parcel. Irrigation districts hold rights and distribute water to their members, ensuring that a farmer within a district receives the water benefits associated with the rights of the district. Complete data water rights, including the priority date, for major irrigation districts in the region are publicly available through the documentation of the Acquavella adjudication [142].

I use sales from 1990-2011 to increase the likelihood of capturing changing expectations of water supply volatility. However, the long time horizon also poses challenges due to changing market conditions over time. I spatially match soil characteristics from the Unites States Department of Agriculture (USDA) SOILMART database to individual parcels using GIS. The Consumer Price Index from the Bureau of Economic Analysis (BEA) normalizes all monetary values to 2008 dollars. Distance to cities, major streams, and the Yakima River, as well as spatial data of supplemental water rights are obtained from the Washington State Department of Ecology (ECY) and generated through GIS. The additional water rights are spread evenly across all irrigation districts and between those that have junior and senior rights. These supplemental rights provide water to livestock and people on the farm, and may be used to supplement water from the irrigation district, but are generally not enough to sustain agriculture. Table reftab:datadescribe provides descriptions of variables used in the regression models and Table 1.3 displays summary statistics for continuous variables and sample percentages for binary variable.

Table 1.2: Data Description

Variable	Unit	Description
Price-per-acre	2008 USD	Dollar value agricultural sales divided by acres sold
Senior Water Right (Sr)	Binary	Dummy variable equal to one if a parcel lies within an irrigation district with more than 50% senior water rights
Residential	Binary	Dummy variable equal to one if there are any residential structures on the property
Groundwater Right	Binary	Dummy variable equal to one if a property has any supplemental water rights in addition to rights from the irrigation district
Rolling Avg	2008 USD	The rolling average of all sales from the previous twelve months normalized to 2008 USD
Acres	Acres	Total acres sold
Class 1-5	%	The percentage of soil in each of five soil classifcations where the classification determines the suitability of the land for agriculture; lower classes are more suited to agriculture, the best class being Class 1 and the worst being Class 5
Improvements-per-acre	2008 USD	This is the dollar value of infrastructure improvements to the parcel since the last sale
Distance to City	Miles	Distance of the parcel centroid to the nearest major city
Distance to Stream	Miles	Distance of the parcel centroid to the nearest major stream
Inverse Distance to UGA	1/Miles	One divided by the distance in miles of the parcel centroid to the nearest urban growth area
Inverse Distance to River	1/Miles	One divided by the distance in miles of the parcel centroid to the Yakima River
Kittitas	Binary	Dummy variable equal to one if the parcel is in Kittitas County
Benton	Binary	Dummy variable equal to one if the parcel is in Benton County

Notes: These are descriptions of all of the variables that were deemed important from the Bayesian Model Averaging and thus included in all regression models.

Table 1.3: Summary Statistics

Variable	Unit	Mean	Std.Dev.	Min	Max
Price-per-acre	2008 USD	7,054	5,971	516.6	29,647
Sr	Binary	0.474	0.499	0	1
Residential	Binary	0.295	0.456	0	1
Groundwater Right	Binary	0.0854	0.28	0	1
Rolling Avg	2008 USD	6,331	2,047	1,948	13,390
Acres	Acres	41.8	51.73	1	680.4
Class 1	%	0.381	0.383	0	1
Class 2	%	0.275	0.326	0	1
Class 3	%	0.186	0.283	0	1
Class 4	%	0.000502	0.0208	0	0.965
Class 5	%	0.156	0.276	0	1
Improvements-per-acre	2008 USD	3,593	4,903	0	29,756
Distance to City	Miles	36.53	21.67	2.884	92.36
Distance to Stream	Miles	1.981	1.613	0	8.123
Inverse Distance to UGA	1/Miles	54.19	316.5	0.0578	2029
Inverse Distance to River	1/Miles	0.881	5.691	0.0424	204
Kittitas	Binary	0.125	0.33	0	1
Benton	Binary	0.327	0.469	0	1

Notes: Similar to [51] observation are eliminated if they are less than 1 acre and greater than \$30,000 per acre. Class variables represent the percentage of each parcel that falls within that class. Slope is the average slope of the entire parcel. For binary variables the mean represents the proportion of observations for which the binary variable is equal to 1.

1.4.2 Econometric Model

I employ a Bayesian linear regression model with normal independent Gamma priors and a general covariance matrix as employed by [86]. The regression function is $y = X\beta + \epsilon$, where y is the real log sale price per acre, X is a matrix of covariates, β is a coefficient vector and ϵ is a heteroskedastic error term distributed $\epsilon \sim N(0, \sigma^2\Omega)$. A Box Cox test provides strong evidence for a log-linear model.[8] The notation for any parameter θ follows [86] where $\underline{\theta}$ represents the prior value that is chosen by the analyst and $\bar{\theta}$ is the posterior value as a function of the data and the prior. I use diffuse priors with zero mean and a

---

[8] [141] argue that the hedonic model may be mis-specified if there is potential for predicted values less than zero. The minimum predicted log per acre farm value is well above one, suggesting that there is not a cause for concern that the model will yield negative property values.

wide dispersion suggesting little prior information on the parameters the full description of the likelihood, priors, and joint posterior is available in the Appendix. This model produces a joint posterior distribution that is not of standard form. To estimate the model I draw directly from the conditional posterior distributions using the Gibbs sampler, a Markov Chain Monte Carlo (MCMC) method, to generate consistent estimates of the joint distribution. The Gibbs sampler sequentially draws from the full conditional posterior distributions of defined blocks, updating all the conditioning values in each run of the Gibbs sampler. The conditional posterior for β is the first block and is distributed multivariate normal, the second block is σ^2 with a gamma conditional posterior distribution, and Ω is estimated in the third block, with the distribution depending on the assumptions of the error term. The conditional posterior distributions for β and σ^2 are given by

$$p(\beta|y,\sigma^2,\Omega) \sim N(\bar{\beta},\bar{V}) \tag{1.8a}$$

$$p(\sigma^2|y,\beta,\Omega) \sim I\Gamma\left(\frac{\bar{\nu}}{2},\frac{\bar{\nu}\bar{s}^2}{2}\right) \tag{1.8b}$$

where $\bar{V} = (\underline{V}^{-1} + \sigma^{-2}X'\Omega^{-1}X)^{-1}, \bar{\beta} = \bar{V}(\underline{V}^{-1}\underline{\beta} + \sigma^{-2}X'\Omega^{-1}X\hat{\beta}(\Omega)),\ \bar{\nu} = n + \underline{\nu}$, and $\bar{s}^2 = \frac{(y-X\beta)'\Omega^{-1}(y-X\beta)+\underline{\nu}\underline{\mathrm{S}}^2)}{\bar{\nu}}$. I do not impose direct structure on the form of heteroskedasticity, but make parametric assumptions to aid in the computation. Specifically, I assume that Ω is a diagonal matrix with the precision distributed independent gamma. The intuition is that all error variances may be different, but they are drawn from the same distribution. The mean of the distribution is assumed to be zero, a trivial assumption, and the variance of the distribution is estimated within the model. This leads to two more parameters to estimate as additional blocks in the Gibbs sampler.

$$p(\lambda_i|y,\beta,\sigma^2,\nu_\lambda) = \Gamma\left(\frac{\nu_\lambda+1}{2}\right),\frac{2}{h\epsilon_i+\nu_\lambda} \tag{1.9a}$$

$$p(\nu_\lambda|y,\beta,\sigma^2,\lambda) \propto \left(\frac{\nu_\lambda}{2}\right)^{\frac{n\nu_\lambda}{2}}\Gamma\left(\frac{\nu_\lambda}{2}\right)^{-n}\exp\left(-\eta\nu_\lambda\right) \tag{1.9b}$$

where $\eta = \frac{1}{\nu_\lambda} + \frac{1}{2}\sum_{i=1}^{n}\left[\ln(\lambda_i^{-1}) + \lambda_i\right]$. The structure of the priors for Ω, shown in the Appendix, leads to the errors being distributed as a student-t with mean zero, variance σ^2 and degree of freedom ν_λ. The degree of heterogeneity depends onν_λv, the scale parameter in the distribution of λ, and is explicitly estimated within the model. Since the posterior for ν_λ is not of a standard form I use the Metropolis-Hastings algorithm to draw from a candidate generating function and then use an acceptance criteria to accept or reject a given draw. Figure A1 in the Appendix shows the histogram for the posterior estimates of ν_λ and the Metropolis-Hastings acceptance rate. The final model has four blocks in the Gibbs sampler that draws from the joint posterior of $p(\beta, \sigma^2, \lambda, \nu_\lambda|y)$. I employ Bayesian estimation techniques for two reasons. The first is the ease of adding additional elements to the model in the form of new Gibbs blocks, and the second is to alleviate omitted variable bias from misspecifying the empirical hedonic price function by using Bayesian Model Averaging (BMA). BMA accounts for the uncertainty inherent in model selection by weighting coefficients by the posterior model probabilities across all models. The posterior model probability for model i as shown in[86] is

$$p(M_i|y) = \frac{p(y|M_i)p(M_i)}{\sum_{m=1}^{M}\left[(y|M_m)p(M_m)\right]} \tag{1.10}$$

where y is the data, M is the total number of models, and $p(M_i)$ is the prior for model i that is set to $1/M$ for all models. There are 2^k potential linear models with k candidate regressors, making formal model selection computationally difficult as the number of candidate regressors increases. In this setting the 25 candidate regressors lead to over 33 million potential models and makes estimating and evaluating each unique model intractable. One form of BMA developed by [114] takes advantage of Markov Chain Monte Carlo Model Composition (MC3) that precludes estimating each separate model and converges to the region with the highest model posterior probabilities. The MC3 method selects new models by either adding or removing a variable from the current model M_i and then assigning an acceptance probability as a function of posterior probabilities that dictates whether the new model M_j will replace the current model M_i given by $p(\text{accept new model}) = \min\left[1, \frac{p(M_j y)}{p(M_i y)}\right]$.

1.4.3 Results

Table 1.4 presents the results from the BMA model. The coefficients are weighted by the posterior probabilities, and assigned a value of zero for models in which they do not appear. The last column displays the count of how many times a variable is selected in models that account for at least 1% of total posterior weight (a total of 15 models met this criterion). While the results from the BMA can be interpreted directly it is also useful to select a baseline model to see how it changes under certain scenarios that have economic implications. Additionally, BMA is computationally intensive and becomes intractable for more complicated models such as adding a Gibbs bloc for an endogenous changepoint. To be conservative when selecting a base model I include all variables that show up at least once in any model comprising at least 1% of the posterior mass in the BMA. I also include an amelioration set of variables suggested by the prior literature. We can compare variants of the base model to the base model itself and the BMA results. Since the Gibbs sampler uses a Markov Chain process to draw from the joint posterior distribution it is important to ensure that the effect of the initial values has disappeared. I perform several MCMC diagnostics to test whether the Gibbs sampler has converged to the true joint posterior. First, as seen in Table **??** in the Appendix looking at the autocorrelation in the draws suggests that by 5 lags the autocorrelation has mostly disappeared. Additionally, I perform the Geweke [62] chi-square test for the equality of means in two separate intervals of the Gibbs draws. For all parameters the p-values are greater than 10% failing to reject the hypothesis of different means as displayed in Table A.2 in the Appendix.

The posterior distribution on the senior right variable reveals whether land with senior water rights sells at a premium. A dummy variable identifies parcels in a district with access to sufficient senior water rights to insulate them from water supply shocks. Table 1.5 shows the estimates for the base model and Figure 1.2 displays the posterior distribution for the senior water right with the dashed lines designating the 95% highest posterior density interval the Bayesian analog to 95% confidence intervals. Variables that appear consistently in all

Table 1.4: Bayesian Model Averaging

Variable	Posterior Mean	t-statistic	t-probability	Count in top 15 models
Constant	7.166056	79.159821	0	NA
Sr	0.095496	2.81305	0.004951	NA
Time	0.013288	11.741287	0	15
Time2	-0.000019	-0.586345	0.557705	4
Acres	-0.002493	-7.632254	0	15
Improve	0.000041	11.905792	0	15
Slope	0.000013	0.005014	0.996	0
Rolling Avg	0.044642	0.657378	0.511007	4
Yield	0.000269	0.088185	0.929738	0
Class 1	0.000908	0.020276	0.983825	0
Class 2	0.000533	0.010417	0.991689	0
Class 3	-0.000033	-0.000556	0.999556	0
Class 4	0.003041	0.003906	0.996884	0
Class 5	-0.056869	-0.960185	0.337068	6
Dist City	0.000927	0.634151	0.526049	15
Dist UGA	-0.000304	-0.04071	0.967531	0
Dist River	-0.000191	-0.047907	0.961795	0
Dist Stream	0.00141	0.131781	0.89517	1
Inv Dist City	0.029193	0.042794	0.96587	0
Inv Dist UGA	0.000006	0.123553	0.901681	1
Inv Dist River	-0.001001	-0.351012	0.725613	2
Inv Dist Stream	0	-0.00204	0.998372	0
Groundwater	0.024229	0.412568	0.679964	2
Residentail	0.15089	3.797551	0.00015	14
Kittitas	0.597486	8.045166	0	15
Benton	0.245797	3.538911	0.00041	15

Notes: Coefficients are weighted by the posterior odds probability and are zero when covariates do not appear in a model. The last column displays the number of times a variable is selected in one of the top 15 models that have at least 15 of the total probability mass. 110,000 initial draws were taken with 10,000 omitted resulting in 100,000 draws.

top 15 BMA models have similar coefficients while those that do not appear as often are drawn to zero in the BMA model.

In both the BMA model and the base model the mean of the posterior distribution for the senior water right coefficient is positive and significant from zero at the 1% significant level. The BMA model predicts a 9.5% increase in the price per acre of farmland, corresponding to

Figure 1.2: Base Posterior Distribution for Sr. Water Right Coefficient

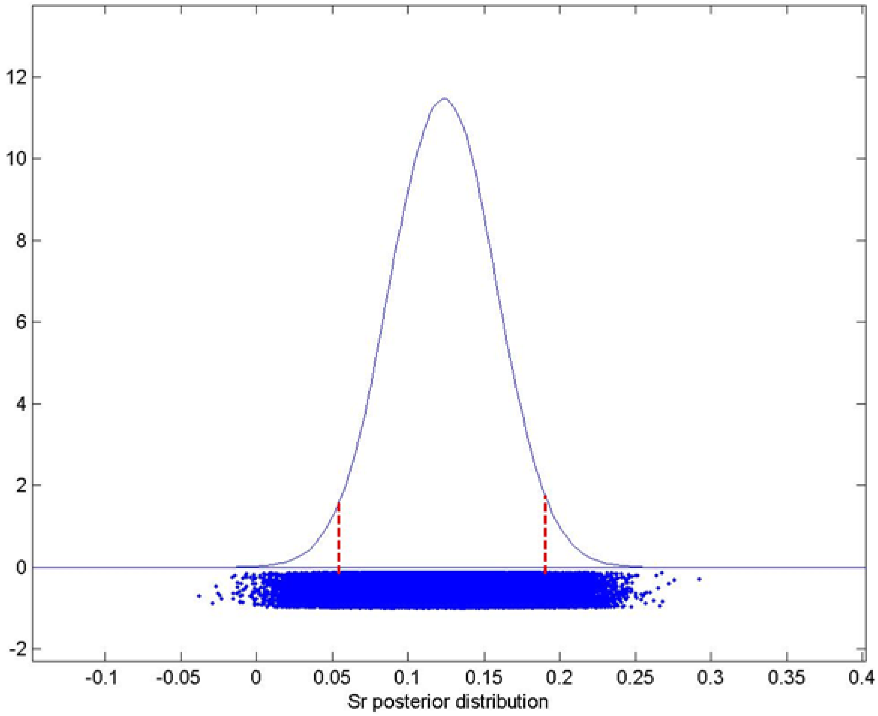

Notes: This is the posterior distribution for the senior water right coefficient in the base regression. Additional controls are shown in Table 3. 220,000 draws were taken in the Gibbs Sampler with the first 20,000 discarded. The points at the bottom are the raw draws and the dashed lines represent the 95% highest posterior density interval.

\$706 evaluated at the mean farmland value whereas the base model predicts values of 12.2% and \$861 respectively.[9] The annualized value per acre-foot of water is approximately \$10.09 in BMA and \$13.08 and the base model, which fits in the context of [51] that finds that an acre-foot of irrigation water ranges from \$9-44. A key distinction is that in this study I find that \$10.09-13.08 is the range of values for more secure irrigation water suggesting it is critical to account for heterogeneity in water rights. Additionally, the notion that irrigated and non-irrigated land may respond differently to climate change is noted in [120], and makes this study attractive by having a sample comprised exclusively of irrigated land. The other parameters have intuitive results. The real per-acre value of land in the Yakima basin is

---

[9] The interpretation of the marginal effect of the coefficient on the dummy for senior water rights on the per-acre sale price equals $100[exp(\beta_D - 1)]$ as shown by [70].

Table 1.5: Base Regression

Variable	Posterior Mean	Std Deviation	t-statistic
Senior	0.122015**	0.034335	3.55361
Time	0.019463**	0.003204	6.07511
Time$\hat{2}$	-0.0001**	0.000033	-3.0828
Acres	-0.00263**	0.000345	-7.62564
Improvement	0.000043**	0.000004	11.95061
Rolling Avg	0.193034**	0.068329	2.82506
Class 1	0.248335	0.700075	0.35472
Class 2	0.2519	0.699176	0.36028
Class 3	0.240286	0.699997	0.34326
Class 4	0.74422	1.043965	0.71287
Class 5	0.037268	0.701178	0.05315
Distance City	0.001623	0.001472	1.10269
Distance Stream	0.022726*	0.010716	2.12068
Inv. Distance UGA	0.00011*	0.000051	2.1694
Inv. Distance River	-0.00666*	0.002996	-2.22312
Groundwater	0.155692**	0.057715	2.6976
Residential	0.154672**	0.038647	4.00215
Kittitas	0.614138**	0.076476	8.03043
Benton	0.267071**	0.070776	3.77346
Intercept	5.529015**	0.896351	6.16836

Notes: The dependent variable is the natural log of the per acre sale price of a parcel. Class variables represent the percentage of each parcel that falls within that class. Slope is the average slope of the entire parcel. **, and * designate significance at the 1% and 5% level respectively. Posterior distributions are based on a heteroskedastic error term with the degree of heteroskedasticity estimates within the model. 220,000 initial draws were taken with 20,000 omitted resulting in 200,000 draws.

increasing over time, where time is defined as quarters from the first observed sale. The dummy for residential structures is positive and significant as is the coefficient for market improvements on the land. I add the percentage of land in each soil class based on the suggestion of [51] even though they do not show up in the top BMA models. As expected based on the BMA results, the means of these variables posterior distributions are not significantly different from zero.

1.4.4 Robustness

An interesting result is the significance and magnitude of the mean of the distribution for supplemental water rights in the base model the mean value of a supplemental right at 15.6% is higher than the premium for senior rights. As explained earlier these rights are mostly likely not the primary source of irrigation but offer some extra water in the form of a stream, groundwater wells, or additional Yakima River surface water rights. According to the theory the value of a senior right stems from insulating the landowner from water supply shocks so the benefits of supplemental rights should be greater for those with junior primary rights. Figure 1.3 shows the posterior distributions for the coefficient on supplemental water rights for the sample divided by the priority of primary rights. The mean of the posteriors suggest that supplemental rights add 0.8% and 25.5% to the farm values for senior and junior rights respectively; however there is significant shared probability mass between the two distributions. The finding that supplemental water predominantly benefits those with junior primary rights supports the initial claim that water rights are heterogeneous and priority insulates landowners from the effects of drought. It is important to consider the economic significance of these results. There is a relatively equal proportion of land with junior versus senior rights in the basin; and water supplied through irrigation districts is the primary source of irrigation water. Meanwhile, only 8.5% of properties have supplemental rights.[10] The low prevalence of supplemental rights along with the rarity of successful new water right applications in the regions suggest that supplemental rights play a limited role in dealing with water scarcity in the region.

I also run regressions for separate counties in the sample and Table 1.6 displays the posterior means along with the 95% highest posterior density intervals. One element that stands out is that the posterior for senior rights has a mean close to zero for Benton County. Recall that the biggest irrigation district in Benton County with predominantly junior rights,

---

[10] The proportion of supplemental rights on land with junior primary rights is 56%, slightly higher than the 52% of properties with junior primary rights in the sample.

Kennewick Irrigation District, has a more stable water supply than other junior districts due to recharge from upstream users. So the county regressions further support the hypothesis that the value of a senior water right is connected to the water supply volatility.

Figure 1.3: Posterior Distributions for Supplemental Water Right Coefficients

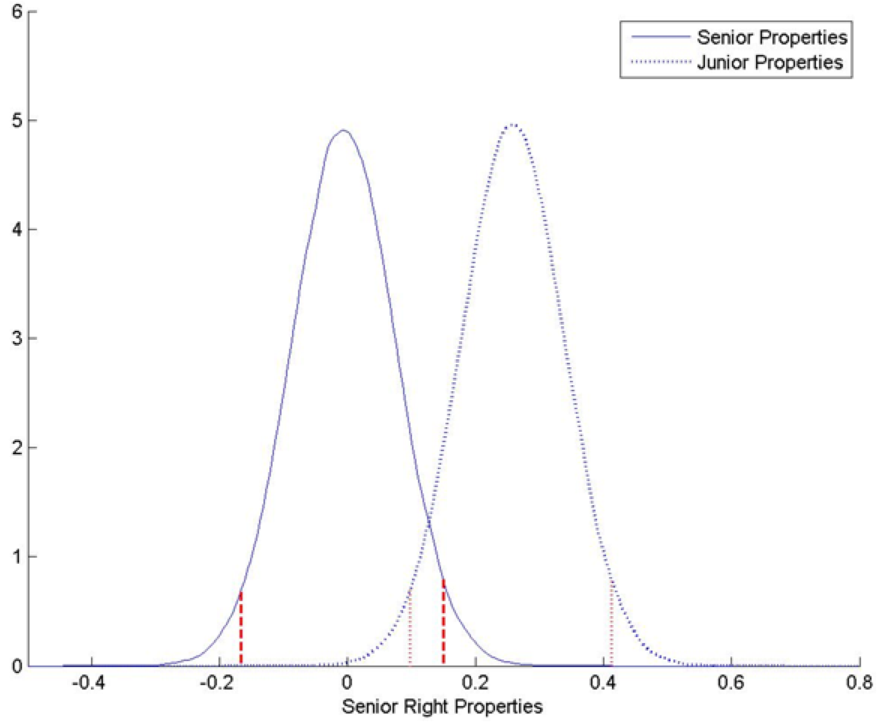

Notes: This is the posterior distribution for the supplemental right coefficient with all controls in the Base regression except for senior, since the sample is pooled by senior vs. junior. Posterior distributions are based on 220,000 draws in the Gibbs sampler with the first 20,000 omitted. The distribution with the solid is the based on properties with senior primary rights, and the distribution with the dotted line is based on properties with junior primary rights. Thick and thin vertical dashed represent the 95% highest posterior density interval for the distributions with senior and junior primary rights respectively.

1.4.5 Policy Scenario

The Integrated Plan is a strategy to address water scarcity in the region as farmers face droughts and the dearth of new water rights constrains developments. For agricultural producers, enhancing storage capacity will decrease the volatility of water deliveries to junior districts, making them more similar to senior districts. Using the estimates of the relative premium for farmland with senior rights in calculating the benefits to the agricultural sector

Table 1.6: County Regressions

Kittitas	Yakima	Benton
0.1135	0.2267	-0.0032
[-0.090-0.3173]	[0.1404-0.3133]	[-0.138-0.1325]

Notes: The first row is the posterior mean for the coefficient on a senior water right dummy, and the second row is the 95% highest posterior density interval for that parameter. All controls in the base result are included in the regressions except for the county dummies, which are perfectly multicollinear.

from storage enhancement in the Integrated Plan provides an alternative to the production function approach used by the USBR. I calculate the gains to agricultural production by multiplying the per-acre premium for land with senior water rights by the irrigable acres of land with junior rights. I only use land served by the Yakima Project since data are readily available, making the results an effective lower bound on the benefits for the whole basin. I believe this approach is justified since these districts represent a significant proportion of total agricultural land and the estimates can be directly compared to the results in the Integrated Plan. Using the hedonic approach the benefits from more secure water rights range between $136 and $234 million depending on the using the BMA or base results and whether Kennewick Irrigation District is included. These results are significantly lower than the $400 million estimate in the Integrated Plan suggesting an upward bias in the production approach that does not account for landowner adaptation. This fits into the general discourse that economic research does not fully permeate through to water policy. Research suggests the relative benefits of more fluid water markets compared to large government funded infrastructure projects, but unfortunately policy has not caught up with the economics (Olmstead, 2010). These calculations rest on the assumption that the premium for land in senior districts is constant throughout the sample period. This assumption is not valid if climate change alters landowners expectations about future water supply volatility which we address in the next below.

1.4.6 Non-Stationary Costs of Water Volatility

Establishing a climate change scenario is the first step in testing for changing expectations about water volatility. The first approach assumes that severe droughts serve as an information shock that, coupled with news and research about climate change, changes landowners' expectations. There have been two severe droughts that reduced prorating to below 70% since the year 2000: in 2001 and 2005. I pool the data into two periods: pre-2005 and post-2005 and run regressions, and then repeat the process for pre and post-2001. Figure 1.4 shows the posterior distributions for the senior water right coefficient when pooling the data before and after the two most recent major droughts. Using either drought to partition the data indicates that the premium on a senior right may not be stationary as the central tendency of the distributions shifts to the right in more recent years. However, there is a significant shared probability mass between the two distributions suggesting that an ad hoc approach to testing for climate change is not sufficient.

Since there is no established climate change treatment period I look to the data to find evidence for dividing of the sample. This is akin to testing for parameter instability on the coefficient on a senior water right, and follows the logic of models with a changepoint commonly used with time series data. I augment the normal independent Gamma model by adding an additional Gibbs block to estimate the full distribution for a changepoint parameter. The methodology is used to test for a structural break in the time trend of U.S. temperature (Li and Tobias 2011). First I order the data by sale date - it should be noted that this is cross sectional data so there are multiple observations per time period. The augmented model partitions the data dependent on θ. For our purposes $X_2(\theta)$ will contain all the same covariates as $X_1(\theta)$ as well as an interaction term of the senior dummy with a time trend, allowing for parameter instability in the coefficient on a senior water right.

$$y_{it}|\beta, \tilde{\beta}, \sigma^2, \lambda, \nu_\lambda, X_{1(\theta)}, X_{2(\theta)} \sim \begin{cases} N(X_1\beta, \sigma_i^2) & \text{if } t \leq \theta \\ N(X_2\tilde{\beta}, \sigma_i^2) & \text{if } t > \theta \end{cases}$$

Figure 1.4: Posterior Distribution for Sr. Water Right Coefficient Pooled by Date

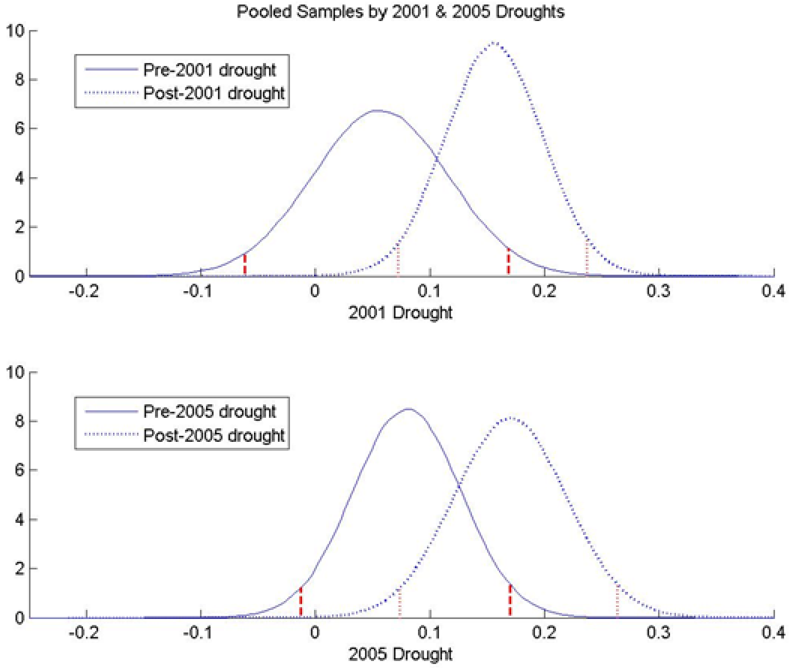

Notes: This is the posterior distribution for the senior right coefficient with all controls in the Base regression. Posterior distributions are based on 220,000 draws in the Gibbs sampler with the first 20,000 omitted. The distributions with solid lines in each pane are the based on the sample before the 2001and 2005 drought respectively, and the dotted lines are based on the sample after the drought. Thick and thin vertical dashed represent the 95% highest posterior density interval for the pre and post drought sample respectively.

$X_{j(\theta)}$ for $j = 1, 2$ denotes the full set of regressors under each regime, where

$$X_{1(\theta)} = \begin{bmatrix} X_{1,1} \\ X_{2,2} \\ \vdots \\ X_{i,\theta} \end{bmatrix} \text{ and } X_{2(\theta)} = \begin{bmatrix} X_{i+1,\theta+1} \\ X_{i+2,\theta+2} \\ \vdots \\ X_{n,T} \end{bmatrix}$$

I use an uninformative uniform prior for θ, and since the θ is discrete-valued I can calculate the unnormalized ordinates for $\theta \in \{1, , T\}$. Normalizing the ordinates by dividing each ordinate by the sum of all unnormalized ordinates produces a discrete valued distribution from which I can draw values of θ. The prior for θ is simply $1/T$, all other parameters are as defined in the base model, and the posterior is given by

$$p(\theta|y) \propto p(\theta)|D_\theta|^{-\frac{1}{2}} \left[\frac{\underline{\nu s}^2}{2} + \frac{1}{2}\left(y - X_{(\theta)}\underline{\beta}\right)' D_{(\theta)}^{-1}\left(y - X_{(\theta)}\underline{\beta}\right)\right]^{-\frac{n+\underline{\nu}}{2}} \tag{1.11}$$

where $D_{(\theta)} \equiv I_n + X_{(\theta)}\underline{\mathrm{V}}X'_{(\theta)}$. I select quarters as the unit of time because the posterior is discrete and there are not many years, while using months requires significantly more computation time. The results are robust to different units of time at low levels of draws for the Gibbs sampler. The median of the change point parameter is 83 corresponding to a change point in the third quarter of 2010. This is very close to the end of the sample suggesting there the coefficient on the senior water is stable over time. Figure 1.5 shows graph of the full set of frequencies of posterior estimates for θ.

Using a linear interaction after a certain time may be an overly simplistic methodology to identify a climate change scenario; however it appears to be an appropriate starting point. There are many factors that contribute to belief in climate change. Creating a continuous or discrete index of climate variables may be more appropriate than a monotonic linear approach. This will allow farmers to expectations of climate change to ebb and flow as the conditions change. Discussions with employees at the USBR and irrigation districts suggest that farmers are aware of snowpack and TWSA estimates. Regressions run with lagged proration rates and TWSA estimates interacted with senior dummy do not produce significant results. Another approach is to employ continuous measures of climate perceptions such as surveys or opinion polls, and is left for future research.

Figure 1.5: Posterior Estimates for Changepoint Parameter

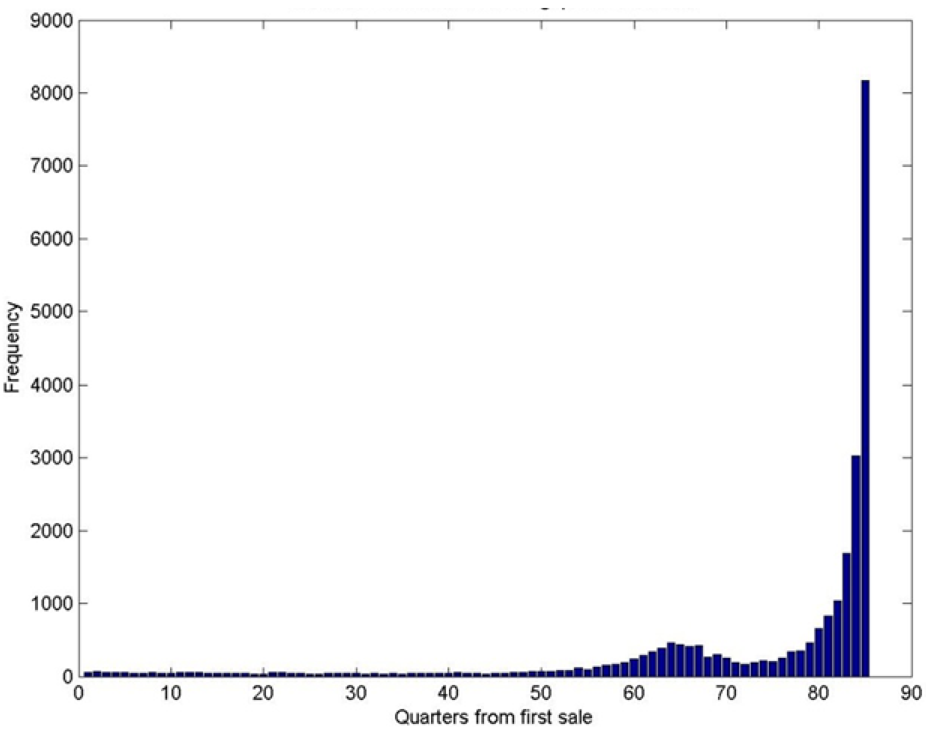

Notes: Frequency counts for the posterior of the changepoint parameter that determines the quarter where the senior right becomes time-varying. The horizontal axis represents the number of quarters from the first sale in the sample. The control variables are the same as in the base regression model. The first 2778 draws out of 27778 draws of the Gibbs sampler are omitted resulting in 25,000 draws.

1.5 Conclusion

Increasingly frequent demand and supply shocks from climate change are raising awareness of water scarcity for agricultural producers in the Western United States. The aggregate and distributional effects of water scarcity are intimately related to the institutions that govern water rights. This paper quantifies the value of priority for senior water rights as a mechanism to protect landowners against the effects of droughts. The central tendency of the posterior distribution for senior water rights is significant under all specifications, and comprises 9-12% of the per-acre value for the average farm in the Yakima Basin. Using the estimates from the posterior distribution of the senior water right coefficient in a benefit cost framework suggests that the methodology employed by the Integrated Plan overestimates the agricultural benefits associated with enhanced storage capacity in the basin. This finding depends on the assumption that the full sample of housing sales is representative of farmer's expectations of climate change induced water scarcity. If farmers expectations changed since 1990, the coefficient on senior water rights will also change. More specifically, if farmers respond to climate change by increasing the demand for water security, the premium paid for land in districts with senior rights will increase. Pooling the samples based on recent severe droughts suggests higher premiums for priority after the droughts. However, explicitly estimating an endogenous changepoint does not support the hypothesis of a time-varying premium. This research finds that while farmers in central Washington do pay for security in agricultural water rights, expectations about changing water supply volatility are not manifested in property markets. Further research can pursue the link between water volatility and the price of water security by exploiting cross sectional variation in water supply volatility. As is often the case with studying water rights gathering high quality data and finding regions with comparable institutions remains a challenge.

Chapter 2

ESTIMATING WATER DEMAND ELASTICITY AT THE INTENSIVE AND EXTENSIVE MARGIN: THE ROLE OF LANDSCAPE DYNAMICS

Abstract

Outdoor water use is a critical element of residential water demand with important implications for long run demand, vulnerabilities to extreme drought, and infrastructure planning. This research exploits a unique panel dataset of monthly water metering records and yearly remotely sensed landscape choices for nearly 187,000 households over 12 years. I utilize changes in landscape to empirically assess the role of outdoor water use, and estimate the price elasticity at the intensive and extensive margin. Intensive margin elasticity is estimated through conditional demand functions for households that maintain one type of landscape throughout the sample, while extensive margin elasticity is calculated through a discrete choice model of landscape conversion. The three major results are (1) elasticity for households with a commitment to maintaining green landscapes is significantly lower in magnitude than the general population, (2) lagged and current prices increase the probability of conversion from a green to dry landscape, and (3) landscape conversions reduce consumption between 18-24%. For consumers with green landscapes the extensive margin constitutes up to 48% of total elasticity, but there is a wide range depending modeling specification.

2.1 Introduction

Complementary goods that use water as an input critically shape residential water demand. These quasi-fixed goods play a similar role to sunk capital in the theory of the firm. Like a factory or a fleet of cargo planes enables a firm to produce a final good and earn a profit, complementary goods make up the capital stock through which consumers derive utility from water consumption. Given an increase in the price of water the literature assumes [2, 29, 103] that consumers first make behavioral adjustments before making large-scale investments to improve the efficiency of their goods that use water as an input. Therefore the adjustment of the water capital stock represents changes along the extensive margin and distinguishes the short and long run effects of changing water prices. While explicit treatment of the capital stock has been studied in the energy and gasoline markets there is little research devoted to changes in the water capital stock, primarily due to data limitations.

In order to estimate changes along the extensive margin I augment a panel dataset of household water metering records with satellite data to measure changes in landscaping, one of the most critical complementary goods in residential water demand [12, 137]. The resulting dataset allows this paper to analyze the role of landscape as a driver of demand parameters and quantify the impact of water rates on landscape conversion decisions. Observing the capital stock concurrently with consumption separates intensive and extensive margin elasticity; parameters often measured implicitly through the difference between long and short run demand elasticity. Additionally, using satellite data to measure landscape choices over time establishes a promising methodology for estimating outdoor water demand. Many utilities' conservation plans concentrate on outdoor use due to the discretionary nature of outdoor water consumption, its role in shaping peak demand, and the exhaustion of many indoor efficiency gains.[1] The decision to maintain a lush green landscape, or switch to drought-resistant vegetation, alters consumer responses to water rates and weather conditions causing structural changes for demand in arid regions. Research suggests that long-run price elasticity is greater in magnitude than short-run elasticity [2, 15, 39, 103] implying that there are meaningful reductions in water consumption along the extensive margin by replacing

---

[1] One example of the penetration of indoor water efficiency is regulation at the federal and state level that requires low-flow toilets to be installed in new construction.

complementary goods such as turf lawn.

I estimate the impact of landscape dynamics by merging a time series of satellite data on vegetative cover from Landsat TM5 at the parcel level to monthly water metering records and structural characteristics of the house. The result is a novel panel dataset of water and landscape for nearly 187,000 households over 12 years in Phoenix, AZ. Observing landscape choices over time identifies consumers who maintain, or convert, their landscape. Estimating conditional demand functions for those with no variation in landscape represents the intensive margin, treating a house-landscape combination as the primary unit of analysis. Price elasticity is up to 26% and 51% lower in absolute value for households that maintain a wet and dry landscape respectively compared to the population with mixed or changing landscape. The wet group's imperviousness to price stems from the fixed commitment to a lush landscape. It is notable that the wet group is less elastic than the general sample even though outdoor demand is relatively more price sensitive as evidenced by [94] among others. In contrast, consistently dry households primarily use water indoors, and indoor use is known to be less elastic than outdoor use [13, 95, 94]. The all-dry households also have one less margin for adjustment compared to those with significant outdoor use. Furthermore, while all groups consume more water during drought conditions, the response is significantly greater in the wet group relative to the dry and mixed groups.

In a discrete choice model I estimate factors that impact the decision to convert from a wet to dry landscape. Increases in current and lagged water rates significantly increase the probability of conversion. Coefficients on lagged prices are larger in magnitude than current prices, suggesting that the price mechanism is predominantly manifested through consumers' past experiences of high water bills. Neighbors' landscape choices are also positive and significant determinants in the conversion decision. This result reinforces anecdotal evidence of the importance of neighborhood effects that lead to clusters of households transitioning to drought-resistant landscapes as social norms evolve. The results demonstrate that price mechanisms, as opposed to commonly employed mandatory watering restrictions, can effectively curtail outdoor water use; an outcome that economists have shown improves social welfare [69, 94]. In fact, the impact of price on landscape conversions, the extensive margin in this analysis, represents 7-48% of total price elasticity for households with significant outdoor water use. The model also reveals the temporal dimension of water rates on landscapes, and subsequently water consumption.

Key demand parameters such as the responsiveness to price and weather variables evolve as the composition of residential landscapes changes. Since climate change is anticipated to increase the probability of severe sustained droughts, it is critical to measure the distribution of green landscapes and understand the heterogeneity of consumer response to policy interventions. Utilizing satellite data adds depth to water demand estimation and informs policy makers about the long run effects of rate increases and the potential impacts from climate change.

This paper is organized as follows. The next section reviews the relevant literature. Then, Section 2.3 discusses the background and describes the data. Section 2.4 overviews the estimation strategy the empirical results. Concluding remarks are offered in Section 2.5.

2.2 Existing Literature

The conventional treatment of the extensive margin in the water demand literature relies on the assumption that households cannot fully respond to prices in the short term due to capital investments in complementary goods. Therefore the long run equilibrium, where full adjustment takes place, encompasses changes along the extensive margin in addition to the intensive margin. Early research employs various partial adjustment models, such as the flow-adjustment model that includes lagged consumption in the demand equation [18, 29, 40]. The long run elasticity is calculated by dividing the price coefficient by one minus the autoregressive, or flow adjustment, coefficient. [110] uses a maximum likelihood approach to overcome the simultaneity bias in increasing block rate tariffs, and continues to apply the flow-adjustment coefficient to distinguish between long and short time horizons. A structural model based on the water provider's optimization problem by [103] results in a dynamic panel data model that the authors estimate using GMM to address the individual and time-varying heterogeneity. [15] employ an error correction model to estimate the contribution of long-run adjustment in a dynamic short-term annual model using aggregate data. In past studies the autoregressive coefficient is generally positive and less than unity, producing values for long-run elasticity that exceed the short-run elasticity in absolute value. The insight gleaned from the long-run water demand literature is that since long-run elasticity is larger in magnitude than short-run elasticity important consumer responses to price must come at the extensive margin. The partial adjustment models are appealing because they simply estimate both short-run and long-run elasticities, but it also imposes structure on the adjustment process. [27] suggest that partial adjustment models do not appropriately model large shocks, and [82] find that a simultaneous equation model outperforms a flow adjustment model to model electricity demand. The strategy here is to first control for important consumer durables, and then explicitly model the demand for these durables.

As the study of allocating scarce resources, economics focuses a great deal of attention to demand management in times of drought, see among others [69, 84, 94, 102]. There are multiple studies that measure the efficacy of demand side management policies that come into effect during water shortages. One of the most common command and control policies is limiting outdoor water for lawns and gardens during times of drought. Estimates of the welfare gain from using prices as opposed to mandatory restrictions range from \$96 to \$152 per household for a given season [69, 94]. Water managers face the challenge of quickly reducing demand under the constraint that baseline water consumption is a function of the capital stock of complementary goods. The most important consumer durable for both individual and aggregate demand during times of drought is the type of landscape outside a home; evidenced by the multitude of command and control policies that target outdoor water use [30, 69, 94, 116, 115]. In cities with an arid climate, such as the southwestern United States, outdoor water use represents 50% or more of aggregate demand, of which up to 90% is for landscape [12, 40]. Despite the importance of demand for watering landscape during drought conditions, there is no research that directly incorporates landscape into water demand or models the landscape conversion decision.

[15] highlight the importance of seasonality by showing that models estimated for each calendar month are significantly different than the pooled model, and weather interactions have an important effect on elasticity measures. Their results, as well as demand estimates from other research [49, 39] suggests that pooling summer and winter periods together is inappropriate due to the importance of seasonality. Additionally, outdoor water use is greatest during the summer peak demand period, and consequently greatly influences infrastructure investments. Summer demand, and more specifically outdoor water use, draws attention to the counter-cyclical nature of water demand with negative supply shocks such as droughts correlated to positive demand shocks such

as increased water requirements for lawns and pools. For these reasons, and the improvements in indoor water efficiency, water managers are increasingly focusing conservation efforts on outdoor water demand. This research exclusively models summer peak demand and develops an approach to incorporate landscape as a proxy for outdoor use - the primary driver of seasonality in water demand.

As climate change is expected to increase the frequency of drought it is crucial to recover demand parameters for households with landscapes of varying water intensity. Even without any impacts from climate change utilities need to develop resiliency plans for dealing with multi-year droughts. This research adds to the literature by estimating the direct impact of price on the water-capital stock. Apart from several studies using a pulse-flow dataset that distinguishes different end uses of water [94, 106], outdoor water demand is not explicitly estimated. Prior studies use proxies such as monthly dummies, or the difference between summer and winter consumption, to account for outdoor use. This research shows that the key driver of outdoor use, and peak summer demand, is the type of landscape in single-family homes. Data on water and landscape over time facilitates a new methodology for explicitly estimating demand for households with different outdoor water capital.

2.3 Background & Data

2.3.1 Water and landscape in Phoenix

Phoenix lies in a desiccant arid climate and its history is inextricably tied to importing water from external sources. The Hohokam occupied what is now present day Phoenix for almost 2,000 years and built complex irrigation systems to smooth out short-run volatility in precipitation. What happened to this early society remains a mystery, but some scholars believe they perished from a long and severe drought a prescient tale for the region's current inhabitants [11]. The immense water infrastructure projects conducted by the Bureau of Reclamation in the early 20th century enabled the development of a strong agricultural sector by securing water rights from the Colorado via the Central Arizona Project, with additional water sourced from the Verde and Salt Rivers via the Salt River Project. The experience of engineering solutions for water scarcity by transporting and storing a vast volume of water has been replicated in many Western regions, and has allowed the rapid population growth to the southwestern United States. As of the decennial census in 2010 Phoenix was the sixth largest city in the United States, with nearly 1.5 million inhabitants and a metro area that includes 6 other municipalities with over 200,000 people. During Phoenix's transition into a major metropolis in the second half of the 20th century the water rights freed up from converted agricultural land allowed residential developments to establish lush green landscapes with immense water requirements. As water rates rise and the environmental issues related to water scarcity become more prominent, households are beginning to convert their green landscapes to drought-resistant native vegetation, known as xeriscape.

I focus on summer demand, defined as June through September, because it is the peak demand period and the link between landscape and water is strongest during the hot dry months. Landscape conversion normally takes place before May in order to avoid the extreme heat of the summer months. Therefore, observations of the landscape at any point between May and September are good indicators of the landscape for the summer season. Additionally, while summer landscape can be viewed as an irreversible decision due to the high fixed costs of conversion; winter landscape does not have the same features. Bermuda turf, the most common grass in the Phoenix area, lies dormant in the winter and consumers need to reseed each season in order to have a green lawn in the winter.[2] Therefore some households that have a green lawn in the summer will appear dry in the winter even if they have turf grass. The reverse situation is not true since dry summer landscapes generally require ripping out turf grass, making summer landscape an appropriate measure of the water capital stock.

Water rates in Phoenix consist of a monthly allotment subject to a small unit environmental charge combined with a volumetric charge for all units above the allotment. Both the allotment and price change within each fiscal year during the high, medium, and low demand periods. Consumption, as well as price, generally increases from the low to medium and medium to high periods due to both observed and unobserved factors. Grouping all periods together may lead to a positive bias on elasticity parameters since higher prices are correlated with unobserved factors that increase demand. The focus on the summer high demand exploits annual price variation, abstracting from

---

[2] Average yard size is approximately 7000 square feet (lot size less square footage of house). If half of that the yard is turf and conversion costs range between \$1.5-\$2.5 per square foot, then conversion costs range from \$5,250-\$8,750; from http://www.mesaaz.gov/conservation/convert.aspx.

the intra-annual variation that is confounded with demand shocks. Additionally, there is evidence summer water demand is inherently different and lumping winter and summer demand together is not appropriate [39, 49, 15].

2.3.2 Data

Data limitations constrain empirical studies of water demand. Many studies estimating long-run demand use time series data aggregated for entire municipalities [15, 29, 100, 103]. Household level datasets often have small sample sizes and are reused for different research applications [2, 18, 3]. The lack of data prevents researchers from observing important complementary goods simultaneously with water consumption; a goal established early in the literature [3]. There are five primary datasets that I merge together: water metering records, satellite data on vegetative cover, weather data, structural housing characteristics, and census data. Monthly metering records for 186,813 single-family homes in the City of Phoenix Water Department's service area from 1998-2009 measure water consumption. This rich dataset is a balanced panel containing over 26 million observations. Since this period corresponds with the collapse of the housing market, using data from *active* accounts ensures that landscape conversions are not merely the neglected lawns of foreclosed homes.

The second primary dataset is a time series of satellite data obtained from the National Aeronautics and Space Administration's (NASA) Landsat 5 Thematic Mapper series, henceforth referred to as Landsat. Landsat data, publicly available for download from the USGS Glovis system,[3] is a valuable data resource used extensively in the natural sciences. However, Landsat is less common in economic research and to my knowledge has never been merged with household-level water data. The Landsat satellite captures reflected and emitted energy in six bands of the electromagnetic spectrum as well as one thermal band. Landsat TM5 takes an image of the same location every 16 days with 30-meter resolution, but is often obstructed by clouds rendering the image unusable. The Normalized Difference Vegetation Index (NDVI), one of the most frequent measures of vegetative cover, serves as a proxy for landscape choices [1, 127, 128]. In the remote sensing literature there is a meticulous debate over the best index to measure vegetative cover see among others [64, 78, 136].

---

[3] The web portal is available at http://glovis.usgs.gov/.

NDVI is widely used and is appropriate as an introduction of remote sensing data into water demand, thought the optimal index depends on the environment, the natural phenomena that is being captured, and the collection instrument and is beyond the scope of this research.

Since Phoenix is an arid environment with few cloudy days I can acquire high-quality images for each year in summer and winter. Each image represents an observation at one point in time and there are several steps to process the data to ensure comparability over time and space, described in detail in Section A.1 of the Appendix. The final landscape dataset is a panel where the cross-sectional unit is geographical location and the time series is the year of an individual image. I spatially merge the NDVI data to each parcel in the water consumption dataset using Geographic Information Software (GIS)[4]. The size of each Landsat pixel is $900m^2$, which is slightly larger than the average lot size of $861m^2$. Since the pixels do not perfectly match up with the parcels I downscale the NDVI to $100m^2$ and take a spatially weighted average of the downscaled pixels. The merging process introduces noise into the landscape variable and is representative of the tradeoff between classification accuracy and scale when processing remote sensing data. While the remote sensing literature focuses on precise and accurate classification, the purpose in this research is to develop a proxy for landscape changes over time for a large geographical area.

Although there is a continuum of landscaping options in Phoenix the two overarching categories are drought-resistant native plants, known as xeric, and lush green vegetation, usually comprising turf lawn, defined as mesic. The key distinction for this research is that xeric is much less water-intensive than mesic. Exact classification is not the primary concern; rather I develop a variable that captures the general water requirements for landscaping at the parcel level. Since NDVI captures the greenness of a given area, and water is required to maintain almost all vibrant green vegetation in Phoenix, the index is appropriate for this coarse classification. For clarification, xeric and dry will be used interchangeably to define low-NDVI landscapes as will mesic, green, and wet for high-NDVI landscapes. In order to obtain accurate landscape classifications I compare quantiles of NDVI, which ranges from -1 to 1, with the data from an widely-cited[5] existing remote sensing study [128]. NDVI performs well in classifying landscape at the tails of the distribution, and is less

---

[4] According to a confidentiality agreement with the City of Phoenix I merge the NDVI values at the parcel level and then City officials attached it to an anonymous identifier representing a water account.

[5] As of 10/15/2015 Stefanov et al. (2001) had 282 citations in Google Scholar.

accurate in the middle of the distributions. Table **??** displays the comparison of NDVI quantiles with the classification of [128].

A limitation of validating the use of NDVI for landscape classification in this study is that data in [128] are only available for one year, and the purpose of this research is to observe both water and landscape over time. Comparing quantiles of NDVI over time is problematic because NDVI is correlated with time-varying weather conditions; Figure **??** in the Appendix visualizes the problem of using raw NDVI quantiles across years. In order to compare NDVI across years I regress NDVI on weather variables to parse out the annual variation due to weather. Weather data are collected from three sources: the National Oceanic and Atmospheric Administration's National Climatic Data Center,[6] Oregon State Universitys PRISM Climate Group,[7] and the University of Arizonas AZMET Weather Data.[8] There are advantages to each of the weather data sources. NOAA's data are one of the most commonly used weather data and are a natural starting point. The PRISM data are delineated on a 4-kilometer grid and thus have much finer resolution than data from a single, or several weather stations. The shortcoming of PRISM data is that only average monthly minimum, maximum, and precipitation are available and vegetative growth is known to respond non-linearly to weather. The AZMET data calculates evapotranspiration for the Phoenix area, which measures a combination of the consumptive use of turf grass and evaporation. This has been used in agronomy [10, 77] and water demand research [106, 102], and accounts for variables such as wind speed and moisture deficit that specifically impact plant growth.

In order to improve comparability of NDVI values over time I normalize the NDVI to control for fluctuations in weather condition. Section A.2 of the Appendix describes the normalization procedure and presents results from the weather normalization regressions. Using the residuals from the weather normalization regressions I create quantiles over the full distribution of normalized NDVI. Each parcel is the merged to the normalized yearly NDVI quantiles, monthly water metering records, structural housing characteristics, selected census demographic variables, and weather variables. The resulting dataset is a panel with two sources of time-varying data. Water

---

[6] The data are available online at http://www.ncdc.noaa.gov/.

[7] The data are available online at http://www.prism.oregonstate.edu/.

[8] The data are available online at http://ag.arizona.edu/azmet/azdata.htm/.

consumption, water rates, and weather all vary on a monthly level, whereas NDVI varies annually. Static data include structural features of the house derived from assessor data and census data at the tract and block level. The structural characteristics of the house are recorded at the time of the sale and thus can vary over time, but the vast majority of the structural features of the house remain constant during the sample[9].

I form three groups based on the time series of satellite data: Wet, Dry, and Mixed. The Wet and Dry groups contains households that, for every year in the sample, have a normalized NDVI value above the 80th percentile or below the 30th percentile respectively[10]. The Mixed group makes up the remainder of the sample and consists of those that converted from wet to dry landscapes, converted from dry to wet, or have at least one normalized NDVI observation that lies between the 30th and 80th percentiles. To clarify the notation, wet/dry are general descriptors or refer to an observation at one point in time, whereas the capitalized versions Wet/Dry correspond to the formal groups in the sample that are consistently wet or dry. Table 2.1 below provides summary statistics for each of the three landscape groups.

In addition to the classification diagnostics with respect to the [128] study, merging the landscape data with water consumption records confirms that the satellite data performs well as a proxy for landscape. Figure 2.1 shows the average historical consumption for the three groups. While all groups have a cyclical dimension to consumption, seasonality is much stronger for the Wet group, and in fact these households are the primary driver of peak summer demand. In the winter months all three groups converge relative to the extreme differences in summer usage. To put these consumption figures in perspective according to the EPA the average household consumption is 400 gallons per day. So while the Dry group uses about the average amount in the summer the Wet group uses approximately twice the amount of water as the average U.S. household.

---

[9] For example only 0.8% of households in the sample either added or removed a pool.

[10] The results from the classification diagnostics in Section A.3 of the Appendix reveal that performance of NDVI for dry extends out from the tails to a greater degree than for wet. For this reason, and in order to obtain a relatively balanced group for wet and dry, I have non-symmetrical cutoffs for landscape classification.

Table 2.1: Summary Statistics

		Landscape Group		
Variable	Unit	Mix	Dry	Wet
Yard	% (of lot)	0.797	0.825	0.82
Year Built	Year	1973	1975	1960
Pool	%	0.343	0.298	0.428
Rooms	#	6.065	5.638	6.454
Bathrooms	#	2.162	1.981	2.304
House size	Square feet	1,631	1,432	1,970
Lot size	Square feet	8,934	9,961	13,325
Total Sales	#	1.618	1.634	1.541
Sale Value	2000 USD	110,207	93,312	189,658
Same House as 1995	%	0.493	0.47	0.542
Household Size	#	3.071	3.091	2.565
Renters	%	0.145	0.187	0.142
College	%	0.22	0.2	0.35
Median Income	2000 USD	52,236	49,284	61,428
Water	Gallons per day	557	428	758
Observations	-	6,937,962	780,455	844,308
Households	-	143,112	16,079	17,598

2.3.3 Water rate structure

Before estimating water demand I describe the water rate structure in Phoenix and some of the issues it raises in demand estimation. A graphical representation of the Phoenix's water rate structure is shown in Figure B.1 in the appendix. The total bill a consumer faces is wp_1 if the consumption is less than or equal to $\bar{w}$, and $[\bar{w}p_1 + (w - \bar{w})p_2]$ if consumption is greater than $\bar{w}$. The marginal price a consumer faces depends on how much water she chooses to consume, creating a simultaneity problem when estimating demand functions [71, 76, 105]. Since consumers only pay p_2 for the units above $\bar{w}$ their total bill is less than the marginal price applied across all units. This additional income relative to a linear price structure is referred to as virtual income, and equal to $\bar{w}(p_2 - p_1)$. Virtual income and marginal price are endogenous since they are simultaneously determined with the quantity of water consumed. A common approach to deal with the endogeneity is to use the full rate structure as instruments since the variables are correlated with actual marginal price and exogenous to the household [105, 106]. In this case this is the high block rate, the low block rate, and the fixed charge. Phoenix has a zero profit constraint the fixed cost must decrease when the marginal cost increases, provided that costs stay constant. This generates a negative

Figure 2.1: Water Consumption by Landscape Group

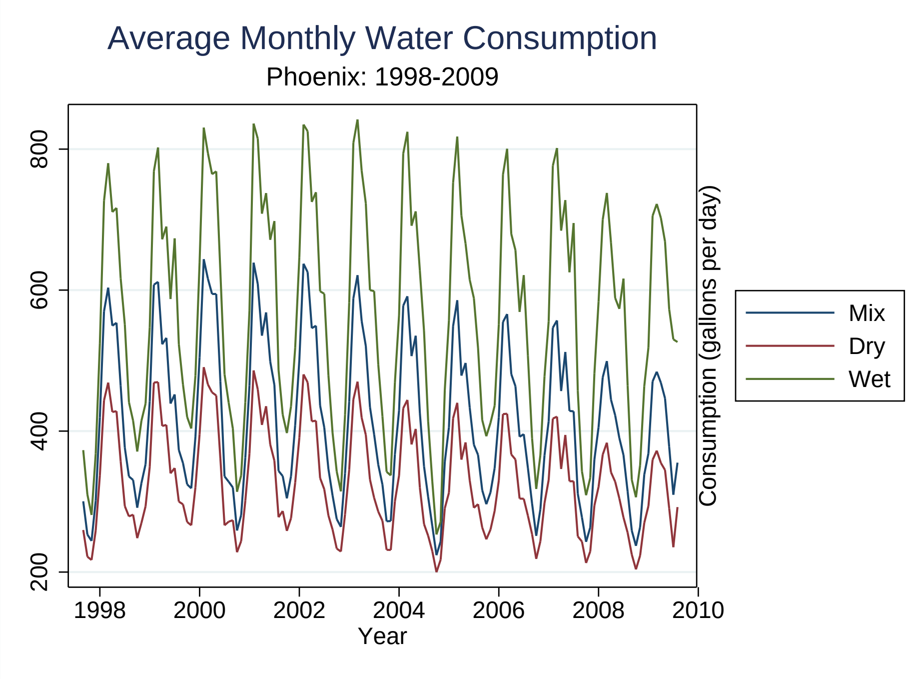

Notes: Wet and Dry groups are determined by those that were continuously above the 80th percentile and below the 30th percentile of NDVI respectively. The Mix group is the remainder of the sample.

correlation between the fixed cost and the marginal price.

Phoenix is in a unique situation where the lower block is almost free, except for a small environmental charge, and the majority of the volumetric charge comes from of the volumetric charge in the second block. It is an open question whether consumers respond to the average or marginal price; and the true measure of consumer response likely is heterogeneous and depends on the rate structure [79, 102, 125]. In Phoenix the price signal primarily stems from the volumetric charge in the second block, and it is possible that even consumers in the first block are responding to the price in the second block. The first block is set at 10 ccf[11] in the summer and the price in this block ranges from $0.09-0.39 over the sample; yielding a maximal variable cost in the first block of less than $4.00 per month. The structure of Phoenix's water rate indicates that the appropriate economic interpretation of elasticity might be to measure behavioral response to the marginal price

---

[11] One ccf is one hundred cubic feet and is approximately 748 gallons.

at the second block, regardless of the actual marginal price they face. Looking at the data, almost 80% of monthly observations are in the high block and fewer than .6% of households never visit the high block.

One limitation in this study is that the primary source of variation is time series. Figure 2.2 displays the variation in both the fixed charge and the high block over the sample period. Many studies use cross sectional variation [94, 106] and quasi-experimental approaches [80, 79, 84, 102] to recover elasticity parameters because there is often more price variation across utilities than within a utility over time. There are advantages to using restricting analysis to one utility that justify the analysis for this setting. [106] bring up the concern of endogeneity for rate schedules as water scarcity or environmental concerns relate to water use and also factor into rate structure. Changes in consumption across different utilities may be the result of differences in rate structures as opposed to the marginal price. Comparing satellite data from multiple geographic areas is difficult due to differences in climate and characteristics of native vegetation. Many utilities offer various incentives and rebates that obscure the pure price effect from responses to other policies. In Phoenix there were no landscape rebates during the sample, which allows for a clean identification of the pure water rate effect on landscape conversion. For these reasons I believe that using time series variation is worthwhile, though cross-utility analysis that incorporates changes on the extensive margin is a promising area for future research.

Figure 2.2: Variation in Phoenix Water Rates

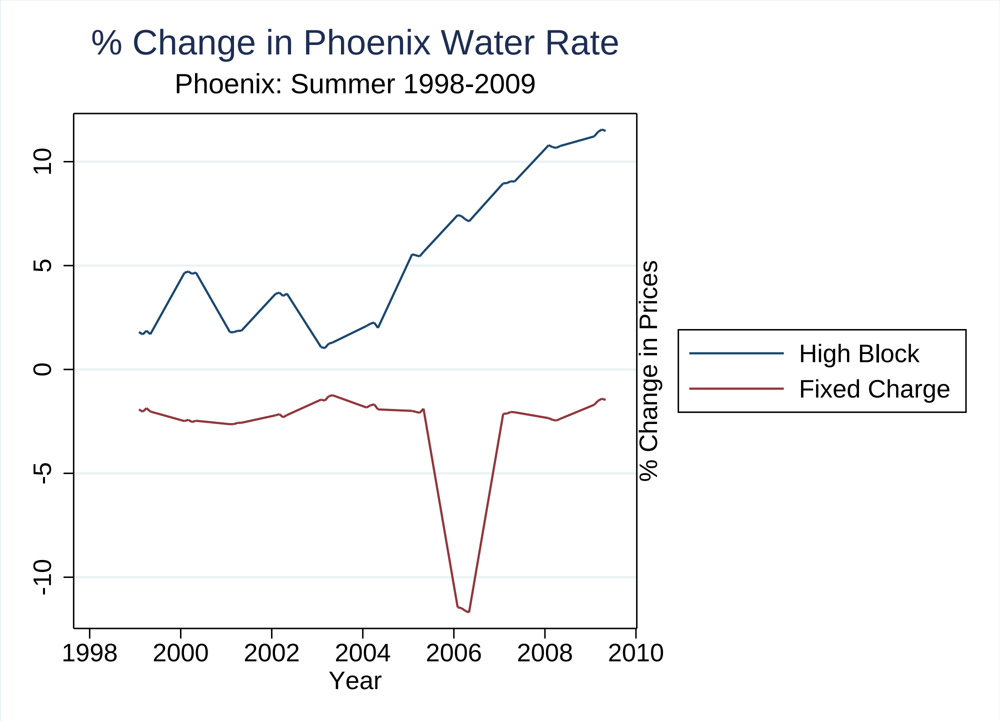

Notes: Real percentage changes are based on prices in constant 2009 USD

2.4 Estimation Strategy

The goal of this paper is to estimate intensive and extensive margin elasticity. The ideal dataset includes observable variation in all complementary goods, water consumption, and prices. I only observe changes in landscape and therefore treat landscape as the water capital for each home. Since landscape plays a major role in urban water demand this treatment is a useful foundation to estimate extensive margin elasticity. To reiterate, changes in water consumption holding landscape constant represents the intensive margin, while the change in consumption associated with converting landscape is the extensive margin. For notational simplicity I assume there are only two types of landscape: wet and dry.[12] Similar to [47] I disaggregate water consumption into conditional probability weighted averages based on the landscape. In this setting average water consumption is designated as $E[w] = P_{dry}E[w|dry] + P_{wet}E[w|wet]$, where $E[w|dry]$ and $E[w|wet]$ are the expec-

---

[12] Note that these are not the same definitions as Wet and Dry, rather these are colloquial terms that designate two general landscaping regimes.

tations of water consumption given a dry and wet landscape; and P_{dry}, P_{wet} are the probabilities of dry and wet landscapes. Equation (1) therefore divides aggregate elasticity into the probability weighted averages of conditional elasticity and the landscape conversion term.

$$\begin{aligned}\epsilon(E[w], p_w) = {} & \epsilon(E[w|dry], p_w) P_{dry} \left(\frac{E[w|dry]}{E[w]} \right) \\ & + \epsilon(E[w|wet], p_w) P_{wet} \left(\frac{E[w|wet]}{E[w]} \right) \\ & + \epsilon(P_{dry}, p_w) P_{wet} \left(\frac{E[w|dry] - E[w|wet])}{E[w]} \right)\end{aligned} \tag{2.1}$$

The elasticity of x with respect to y is $\epsilon(x, y)$, w is the quantity of water p_w is the price of water. The first two terms are the probability weighted averages of the conditional elasticities for dry and wet households respectively, and the last term captures the impact of price on landscape conversions. There are two key insights in equation (1). First, heterogeneity exists in the intensive margin elasticity based on the type of landscape, displayed as $\epsilon(E[w|dry], p_w)$ and $\epsilon(E[w|wet], p_w)$. Second, the extensive margin elasticity measures the impact of price on the proportion of households with dry landscapes, $\epsilon(P_{dry}, p_w)$, scaled by the change in consumption associated with converting from wet to dry, $P_{wet}\frac{E[w|dry]-E[w|wet]}{E[w]}$. Estimating the separate elements of equation (1) requires a time series for landscape to identify changes in landscape as well as households that preserve a fixed landscape. Using the Wet and Dry groups I isolate changes in consumption for households that maintain one type of landscape throughout the sample, thus defining the intensive margin. Likewise the Mixed group consists of changes along the extensive margin that will be estimated in a discrete choice model of landscape conversion.

Before estimating the conditional demand functions I perform model specification based on the full sample pooled over all landscapes by estimating water demand presented in equation (2).

$$\ln(w_{it}) = \alpha + \gamma ln(p_{it}) + \beta X'_{it} + \nu_{it} \tag{2.2}$$

Here w_{it} is water consumption for household i at time t, $p_{w,it}$ is the price of water, X_{it} is vector of controls, and ν_{it} is an idiosyncratic error term, where $\nu_{it} = \mu_i + \xi_{it}$. Table 2.2 compares the demand estimation for random effects and fixed effects models either instrumenting for the observed marginal price or using the marginal price in the second block. The dependent variable is the log of monthly water consumption, with panel cluster-robust standard errors as defined by [139].

In columns (1) and (2) the price variable is the log of the marginal price for the second block and in columns (3) and (4) price is the log of observed marginal price instrumented with the full rate structure to address the simultaneity of consumption as suggested by [105]. Net evapotranspiration is a function of weather variables and quantifies the consumptive use of turf grass in the Phoenix metro area net of recorded precipitation measured in millimeters per square foot. Temperature is the average monthly maximum temperature, and PHDI is the Palmer Hydrological Drought Index that serves as a proxy for drought conditions with lower values signifying more severe droughts. Household level characteristics include the percentage of the lot that can be landscaped, the year the house was built, a dummy for a pool, the number of bathrooms, and square footage of the house.[13] Household size and whether the same occupant lived in the house in 1995 are obtained at the census block level from the 2009 5-year American Community Survey.

Comparing the results across specifications shows that there is not much variation in price elasticity, or other parameter coefficients. Price elasticity is slightly higher in absolute value in the IV model, as are the coefficients on temperature and net evapotranspiration. While random effects is efficient, consistency rests on the assumption of no correlation between the random effect μ_i and the static covariates. A cluster-robust version of the Hausman test [99, 32, 139] rejects the null of no correlation, requiring the fixed effects model. In order to examine the differences between modeling demand with the marginal price and the high block price I perform a cluster-robust version of the Hausman test [139] to test for differences between these regression equations. While the fixed effects equations in column (2) and (4) overall are significantly different from each other, the price coefficient is not significantly different at the 10% level. It is important to consider the interpretation of the two designations of marginal price. The elasticity using the observed marginal price represents the change in consumption due to a change in price that a consumer actually faces. Alternatively, employing the high block price the elasticity simply represents the response to a change in the high block rate. It is intuitive that using observed price produces elasticities of greater magnitude because it represents prices that affect consumers who consume below the

---

[13] The yard variable is created by subtracting the footprint of the house from the total square footage.

Table 2.2: Specification for Water Demand

VARIABLES	(1) RE No-IV	(2) FE No-IV	(3) RE IV	(4) FE IV
ln Price	-0.233***	-0.233***	-0.243***	-0.242***
	(0.00667)	(0.00667)	(0.00842)	(0.00841)
Time	-0.0121***	-0.0121***	-0.0128***	-0.0129***
	(0.000364)	(0.000364)	(0.000419)	(0.000418)
Net Evapotransporation	0.0209***	0.0209***	0.0266***	0.0266***
	(0.000130)	(0.000130)	(0.000268)	(0.000268)
Temperature	0.0237***	0.0237***	0.0275***	0.0273***
	(0.000100)	(0.000100)	(0.000189)	(0.000188)
PHDI	-0.00502***	-0.00502***	-0.00504***	-0.00504***
	(7.74e-05)	(7.74e-05)	(9.97e-05)	(9.96e-05)
Yard	1.896***		2.088***	
	(0.0199)		(0.0255)	
Year Built	-0.00353***		-0.00436***	
	(9.90e-05)		(0.000129)	
Pool	0.212***		0.280***	
	(0.00266)		(0.00399)	
Bathrooms	0.0396***		0.0417***	
	(0.00299)		(0.00370)	
House sq. ft.	0.000429***		0.000497***	
	(3.49e-06)		(4.84e-06)	
Same House in 1995	0.192***		0.229***	
	(0.00794)		(0.0102)	
Household Size	0.0843***		0.120***	
	(0.00158)		(0.00235)	
Houshold Fixed Effects	No	Yes	No	Yes
Observations	8,564,134	8,564,134	8,564,964	8,564,961
Households	186,813	186,813	186,834	186,831

Notes: Dependent variable is the natural log of monthly household water consumption. For columns (1) and (2) the price variable is the log of the marginal price for the second block and in columns (3) and (4) price is the log of observed marginal price instrumented with the full rate structure to deal with simultaneity with consumption. Household fixed effects are used in and robust standard errors clustered at the household level are given in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

threshold. Since it is unclear which specification is preferable I report results from regressions with both marginal price variables.

2.4.1 Conditional demand functions

In order to capture the intensive margin elasticity and explore heterogeneity in demand parameters I estimate conditional demand functions defined by

$$\ln(w_{it}^l) = \alpha_i^l + \gamma^l ln(p_{it}) + \beta^l X_{it}' + \xi_{it}^l \tag{2.3}$$

Equation (3) is similar to equation (2) except that demand is conditional on landscape group $l = \{M, D, W\}$, and I focus on the fixed effects model. The other terms include household fixed effects α_i, a vector of controls X_{it}, and an idiosyncratic error term ξ_{it}. The parameters of interest are the values of γ^l for each of the landscape, interpreted as elasticities in the log-log specification. Table 2.3 presents the results from the conditional demand functions, using both price specifications, for the three landscape groups: Mixed, Dry, and Wet.

Table 2.3: Conditional Demand Functions

	(1) Mixed FE	(2) Dry FE	(3) Wet FE	(4) Mixed FE-IV	(5) Dry FE-IV	(6) Wet FE-IV
ln Price	-0.248***	-0.127***	-0.213***	-0.266***	-0.122***	-0.184***
	(0.00791)	(0.0239)	(0.0236)	(0.0106)	(0.0240)	(0.0247)
Time	-0.0129***	-0.0162***	-0.00154	-0.0140***	-0.0161***	-0.00169
	(0.000464)	(0.00131)	(0.00128)	(0.000555)	(0.00138)	(0.00152)
Net ET	0.0208***	0.0194***	0.0240***	0.0268***	0.0226***	0.0294***
	(0.000353)	(0.000521)	(0.000920)	(0.000491)	(0.000896)	(0.00136)
Temp	0.0237***	0.0195***	0.0256***	0.0277***	0.0216***	0.0281***
	(0.000156)	(0.000398)	(0.000447)	(0.000251)	(0.000656)	(0.000616)
PHDI	-0.00530***	-0.00431***	-0.00338***	-0.00545***	-0.00410***	-0.00296***
	(9.16e-05)	(0.000273)	(0.000262)	(0.000122)	(0.000307)	(0.000311)
Obs	6,937,925	780,455	844,308	6,937,891	780,450	844,305
HHs	151,247	17,236	18,327	151,247	17,236	18,327

Notes: Dependent variable is the natural log of monthly water consumption, and conditional demand functions for subsets of the sample determined by NDVI over time. For equations (1)-(3) price is specified as the marginal price in the top block and columns (4)-(6) use the instrument for observed marginal price. Household fixed effects are used in and robust standard errors clustered at the household level are given in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Since landscape remains constant for the Dry and Wet groups, households in these groups can only change indoor water use and the intensity of watering so the coefficients on price represent intensive margin elasticity. Note in Table 2.3 that the Dry and Wet groups have lower point estimates for elasticities in both specification relative to the Mixed group. Wald tests reveal that these differences are significant at the 95% level for all combinations except the Wet-Mixed groups in the non-IV specification (the price coefficients in columns (2) and (3)). Households in the Wet group have a fixed commitment to maintain a green landscape that dampens their responsiveness to rising water rates. The theoretical basis for this result can be traced back to [111] in his theory of conditional demand and quasi-fixed goods. This is particularly noteworthy due to the extensive literature that finds outdoor water use to be more price elastic than indoor use; and the Wet group surely consumes a large amount of water outside in order for the satellite to detect substantial vegetation. The Dry group's relative inelasticity stems from removing one margin of adjustment, namely outdoor water use; and because indoor use is less elastic than outdoor use. These results reveal the degree of heterogeneity in water demand parameters due to landscape.

Several other comparisons between the conditional demand functions are noteworthy. The coefficient on the yearly time trend shows that the Mixed and Dry group's consumption is steadily decreasing by approximately 1.5% per year, while the Wet group usage stays relatively flat over time controlling for other factors. Across the country indoor water use has declined due to improved efficiency in appliances and regulations that require low-flow toilets. These savings are manifested in both the Mixed and Dry group, but outdoor water use dominates any indoor efficiency gains for the Wet group. Similarly, the Wet group is significantly more responsive to both temperature and net evapotranspiration than the Dry and Mixed group because more water is necessary to sustain green plant life in hot dry conditions. It is interesting that the coefficient on PHDI, which measures longer-term weather conditions, is larger for the Mixed group than either Wet or Dry. Long-run drought conditions may prompt the Mixed group to change their landscape or at least stop maintaining it a margin of adjustment not available to the Wet group. Issues of sample selection arise in the conditional demand functions since their may be underlying differences, irrespective of landscape, between the three groups. In the random effects specification I run a selection model incorporating terms to correct for sample selection [74, 88, 122, 38]. These models produced similar results to the simple random effects model, but when the sample selection terms drop out in the fixed effects model because they do not vary over time. Landscape in Phoenix is trending towards more xeriscaping; probably due to changing environmental attitudes and rising water rates. As seen in Table 2.3, as the proportion of households with dry landscapes increases overall water demand will be less responsive to price and weather variables. In the next section I estimate a model of landscape conversion to determine the role of price in the landscape decision, and the subsequent impact on water demand.

2.4.2 Landscape conversion

The decision to convert a landscape from turf grass to native desert plants is a major investment for a household. The benefits are a string of savings that come from lower water consumption as well as reduced labor costs if the xeriscape requires less maintenance, as is often the case. Costs of conversion consist primarily of the upfront fixed cost of the conversion; for example one set of estimates from Las Vegas range from \$1.37-1.93 per square foot [126][14]. Landscape conversion is treated as an irreversible investment since it is unlikely that a household would re-install grass after an expensive xeriscape investment due to the high fixed costs. This is similar to the decision to invest in residential energy efficiency [73, 117], development and land use [25, 119], technology adoption [50], and plant exit decision [19].

Landscape is a highly visible feature of the house and landscape is capitalized at both the household and neighborhood level. [83] find that living in a neighborhood with predominantly green landscapes adds more value to a property than if the property itself has a green landscape. Part of the value rests on the fact that green neighborhoods have lower nighttime temperatures a valuable attribute in southwestern cities. Households in predominantly mesic neighborhoods that convert their lawns to xeric landscapes free-ride on neighborhood-level amenities by avoiding the expense of maintaining their own yard. I model the timing of landscape conversion as the product of a household optimization problem such that a household chooses the time of conversion T to minimize:

$$V = \int_0^T [p_t\bar{W} + (m_t - b_t)]e^{-\rho t}dt + \int_T^\infty (1-\theta)p_t\bar{W}e^{-\rho t}dt + K_t e^{-\rho t} \tag{2.4}$$

---

[14] The [126] estimates are from 2001 and match up with those for the Phoenix area after accounting for inflation. In 2010 dollars the conversion costs are \$1.80-2.54, within the range of \$1.50-2.50 reported for the Phoenix area.

where p_t is the price of water at time t, $\bar{W}$ is the water requirement for a green landscape, m_t is maintenance cost (outside of water costs) of a green landscape relative to a dry landscape, and b_t is the dollar value of the relative benefits of a mesic landscape compared to xeric There is a one-time cost of K_t to convert a landscape, taken to be the numeraire, and I assume a conversion achieves a proportional reduction in consumption by a factor θ. The discount rate is ρ and is less than one. The first order condition to this optimization problem that dictates whether a household will convert is

$$\theta p_T \bar{W} + (m_T - b_T) - \rho K_T \geq 0 \tag{2.5}$$

If the water savings and non-water relative net costs exceed the discounted capital cost then a household will convert. The term $(m_t - b_t)$ captures the non-water component of the landscape decision with b_t representing the visual appeal and recreational value of a grass lawn[15], whereas m_t consists of labor and material costs associated with landscape maintenance. This term is likely to be negative for households that have not yet converted at the start of the sample. The non-market benefits, b_t, distinguish the landscape conversion decision from a conventional model of residential investment in efficiency that exclusively focuses on minimizing expenditure. For example, outside of the difference in energy costs consumers are likely indifferent between weather proofed windows and old drafty windows. Another factor that in landscape conversions that does not necessarily exist in other contexts is the role of social norms and attitudes. Neighbors' decisions may alter the decision to convert landscape since there may be pressure to either maintain a predominantly green neighborhood, or to prove environmental credentials by adopting native desert vegetation. As social norms evolve having a xeric landscape may become more socially acceptable and may even become a desirable signal of environmental stewardship. Regardless of the direction of the effect neighbors' decisions likely play a role in the decision to convert from a wet to dry landscape, and are therefore included in the model.

Landscape conversion is coded through observing a series of NDVI values over time. After normalizing for weather as described in section A.2 of the Appendix, I generate quantiles of NDVI to define either wet or dry landscape years. Since there is noise in the NDVI data I am conservative

---

[15] [83] estimate how green landscapes capitalize into housing prices and find that living in a green neighborhood increases the monthly rental rate by over $100 for a green neighborhood and $17 for a green lawn.

in defining a conversion as a sequence of at least two wet years followed by at least two dry years. Through this sequence I can not only observe whether a household converted from wet to dry, but also the year of conversion. The dependent variable, C_{it}, is equal to one if household i converts at time t and is otherwise equal to zero. The price of water is observed as are the landscape choices for each household and their neighbors. The fixed amount of water and the efficiency factor are not directly observed but are likely a function of observable household characteristic such as the size of the yard and household size. Similarly the elements that make up $(m_t - b_t)$ are not observed, but I do have data that can be used as proxies such as whether the house is owned or rented, neighbors' landscape decisions, and the value of the home. Household level data for the fixed capital costs of landscape conversion and variable landscape maintenance costs for both wet and dry landscapes are not available. Therefore identification requires the assumption that water rates are not correlated with these unobserved costs over time.

To estimate the model I separate the decision to convert into a deterministic and a random component:

$$C^*_{it} = x_{it}\beta + \epsilon_{it} \tag{2.6}$$

such that household i converts at time t if:

$$C_{it} = \begin{cases} 1 & \text{if } C^*_{it} > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{2.7}$$

I parameterize ϵ_{it} in equation (5) as type I extreme value leading to a panel data logistic regression. Since there are likely important unobserved variation at the neighborhood level, for example homeowner association's restrictions and social pressures, I use a multilevel model [61] that allows for random neighborhood-level intercepts. The probability of conversion for household i at time t is defined as:

$$P(C_{it} = 1) = \frac{e^{\alpha_N + \beta' x_{it}}}{1 + e^{\alpha_N + \beta' x_{it}}} \tag{2.8}$$

where α_N is a random neighborhood-level intercept for household i. I also run a version of the model with both individual and neighborhood level random effects, and the main results do not change substantively. The random effects logit model does not have a closed form solution and is therefore estimated through numerical quadrature simulation methods. Table 2.4 provides the parameter estimates for the random effects logit model for landscape conversion. Since the predicted probabilities depend on the random components I present the raw coefficients. These coefficients represent the marginal change in the log odds ratio, or logit, for a unit change in x_{it}. The interpretation behind the coefficient's magnitude depends on the probability of conversion, which is very low in our sample. Since the logit is a monotonic transformation positive coefficients increase the probability of conversion while negative coefficients have the converse effect.

Table 2.4: Landscape Conversion Random Effects Logistic Regression

VARIABLES	(1)	(2)	(3)	(3)	(5)
Price	5.484***	2.164**	4.826***		
Price (t-1)		4.276***	-4.793***	5.538***	-1.453*
Price (t-2)			19.58***		18.71***
# Dry Neighbors	0.130***	0.128***	0.112***	0.126***	0.109***
House sold w/i 1 year	0.234***	0.242***	0.263***	0.243***	0.265***
Yard size	1.24e-05***	1.24e-05***	1.25e-05***	1.24e-05***	1.25e-05***
Rooms	-0.0803***	-0.0804***	-0.0811***	-0.0805***	-0.0814***
Year Built	-0.0121***	-0.0119***	-0.0105***	-0.0117***	-0.0102***
Pool	0.0334	0.0337	0.0346	0.0338	0.0349
House Price	-2.08e-06***	-2.09e-06***	-2.23e-06***	-2.11e-06***	-2.25e-06***
% Renter Occupied	-0.135	-0.130	-0.100	-0.127	-0.0928
% College Educated	-0.0221***	-0.0223***	-0.0238***	-0.0224***	-0.0242***
Same House as 1995	0.336	0.334	0.330	0.334	0.333
Time	2.332***	2.105***	1.254***	1.906***	0.840***
Time$\hat{2}$	-0.283***	-0.257***	-0.245***	-0.224***	-0.173***
Constant	22.37***	20.23***	-11.96**	19.88***	-9.055**
log-sd	-0.508***	-0.513***	-0.557***	-0.518***	-0.565***
SE of log-sd	(0.0509)	(0.0511)	(0.0523)	(0.0511)	(0.0525)
Observations	876,665	876,665	876,665	876,665	876,665
Chi-Sq	584.1	576.7	520.3	572.6	511.8
Chi-Sq p-value	0	0	0	0	0

Notes: Observations after a conversion are dropped, since they are not able to convert twice in the sample. Random intercepts are implicitly estimated at the neighborhood level; their log of their standard deviation is reported. Chi-square statistic is the for the hypothesis test of ignoring the random intercepts and running a pooled logit. Bootstrapped standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Examining the results in Table 2.4 shows that higher water prices increase the probability of landscape conversion. I use the marginal price in the high block because this is the relevant price signal to those with significant water consumption for landscape. The various columns designate different temporal dimensions of price variables. If a consumer bases their landscaping decisions on their expected summer water costs (water rate increases are publicly listed on the City of Phoenix website) then the current price is the appropriate price to use. However, it is plausible that high water bills in the previous season drive landscape decisions; motivating lagged water rate as the applicable price. In scenarios such as columns (2), (3), and (5) where more than one price variable is included, one derives the total impact of price by the cumulative effect of the price coefficients. In all scenarios the cumulative effect is positive and significant indicating that consumers do respond to rising water rates by changing their landscapes.

Other variables also have interesting, and usually intuitive results. Having more neighbors with a dry landscape significantly increases the probability of converting to from wet to dry.[16] This reinforces anecdotal evidence from water managers at the City of Phoenix of clusters of conversions at the neighborhood level. The positive coefficient on neighbors' conversions suggests that alleviation of social pressure from being an isolated xeric landscape, in a green neighborhood, and/or evolution of social norms to favor water-efficient landscape, dominates the desire to free ride on a neighborhood-level amenity. While the results associated with neighbors landscapes are intriguing I caution full causal interpretation of these results as peer effects. The difficulty of identifying peer effects is noted in [92, 93], and future research will explore whether neighbors decisions have a causal effect the probability of conversion or if the results are driven observed or unobserved correlation across space. The spatial pattern of conversion also has implications for dealing with the urban heat island effect since there is a non-linear relationship between landscape and nighttime temperatures.[17] Houses that have been sold within one year also are more likely to convert; probably because the new owner may have different landscape preferences than the seller. Census blocks with a higher proportion of renters are less likely to convert, which is consistent with the theory since they will recoup less of the costs of fixed investment if they plan to move in the future. Interestingly, the size of the yard has a small and insignificant effect on conversion. Similar

---

[16] The twenty closest non-contiguous neighbors are used in order to avoid contamination in NDVI data with direct neighbors.

[17] [66] shows that the change in nighttime temperature associated with converting one landscape from mesic to xeric depends on the current composition of landscapes.

to neighbors' landscapes there are two opposing forces driving this effect; first a larger yard means a higher fixed cost of conversion but at the same time will translate into higher future water savings. Large expensive houses with pools may be less likely to convert since water bills likely make up a smaller portion of their income. Chi-square tests for the random effects model over a simple pooled logit model consistently reject the null that the random effects logit is appropriate suggesting that the neighborhood level random intercept is an important criterion in the landscape decision.

A shortcoming to the random effects logit model is that it does not account for individual level heterogeneity [31, 75]. To deal with any bias due to the idiosyncratic error term I also estimate a fixed effects, or conditional, logit model. This model conditions the likelihood on the sum of household decisions, $\sum_{t=0}^{T} C_{it}$, so all households that do not have any variation over time drop out. This restricts the sample to households that converted, and only parameters for time-varying covariates can be estimated. The results from the fixed effects logit are displayed in Table 2.5. The main results are the same for the time varying parameters, though, since there is a relatively small proportion of conversions, direct comparisons across the random effects and conditional models are likely inappropriate. The challenge of model comparison is exacerbated for logit models because coefficients are only estimated up to a normalization that depends on the observed variation in the independent variables. Still, it is promising that controlling for individual level heterogeneity primarily changes the magnitude of the results, but not the sign.

In addition to the two sets of results in Tables 2.4 and 2.5, I estimate models based on a lower threshold for conversions and without the weather normalization. The base estimation identifies the string of conversions by classifying a wet and dry observation of NDVI above the 70th percentile and below the 40th percentile respectively. I relax the classification to allow wet observations to be above the 60th and dry to be below the 50th percentile, therefore identifying more conversions[18]. Additionally, I define conversions without the weather normalization by classifying landscape observations based on yearly quantiles of NDVI. In this specification I focus on the placement of a household within the distribution of NDVI for a given year. The advantage to this approach is it eliminates the risk of mis-specification in the weather normalization stage, but it does not account

---

[18] Discussions with the City of Phoenix suggest that the initial classification is likely underestimating the number of conversions, warranting an approach that increases the number of conversions.

for the fact that the distribution of NDVI is likely changing over time due to a move towards xeric landscapes. The same pattern of price and other key variables is exhibited across all the modeling specifications and landscape classification procedures. The results are provided in Table A.3 of the Appendix.

Table 2.5: Landscape Conversion Fixed Effects Logistic Regression

VARIABLES	(1)	(2)	(3)	(3)	(5)
Price	12.34***	12.50***	13.28***		
	(0.774)	(1.074)	(0.942)		
Price (t-1)		-0.244	-5.759***	7.439***	3.453***
		(0.749)	(0.790)	(0.525)	(0.695)
Price (t-2)			10.84***		9.428***
			(0.969)		(1.006)
# Dry Neighbors	0.303***	0.303***	0.279***	0.275***	0.250***
	(0.00583)	(0.00735)	(0.00776)	(0.00699)	(0.00666)
House sold w/i 1 year	0.128**	0.127*	0.136**	0.132**	0.140**
	(0.0598)	(0.0672)	(0.0573)	(0.0618)	(0.0657)
Time	2.540***	2.551***	2.028***	1.441***	0.919***
	(0.108)	(0.118)	(0.131)	(0.0415)	(0.0645)
Time^2	-0.384***	-0.385***	-0.360***	-0.200***	-0.171***
	(0.0180)	(0.0189)	(0.0196)	(0.00634)	(0.00710)
Observations	32,760	32,760	32,760	32,760	32,760
Number of key	4,109	4,109	4,109	4,109	4,109
Pseudo-R2	0.270	0.270	0.276	0.261	0.267

Notes: Observations after a conversion are dropped, as they are not able to convert twice in the sample. Random intercepts are implicitly estimated at the neighborhood level. Bootstrapped standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

2.4.3 Landscape conversions and water demand

The purpose of estimating a model for landscape conversion is ultimately to recover the extensive margin elasticity and determine the impact of conversions on water demand. A simplistic approach is to multiply the marginal change in the conversion probability by the difference of in average consumption between the Wet and Dry group. A problem with this technique is that the Wet and Dry groups may have fundamental differences because, by definition, they do not convert during the sample. In order to estimate the impact of conversions on water demand I create a variable that designates whether household i experienced a conversion at time t defined by

$$LC_{it} = \begin{cases} 1 & \text{if converted prior to time } t \\ 0 & \text{otherwise.} \end{cases}$$

Augmenting the water demand model with LC_{it} estimates the impact of conversion on consumption. Incorporating conversions into water demand also provides a validation test for the conversion classification by linking it back to water data. Since the NDVI data is relatively coarse there is a concern that the landscape conversion model is actually picking up landscape decisions of neighboring parcels instead of the parcel itself[19]. To test for this problem I also create a variable for how many neighbors have converted their landscape. Neighbors' conversions is simply defined by the sum of conversions for non-contiguous neighbors of household i such that $NLC_{it} = \sum_n LC_{nt}$, where n is a neighbor to household i I remove contiguous households from the sample of neighbors because these observations may suffer from measurement error since the NDVI pixels can span multiple parcels. I test this across all the specifications for classifying a conversion as discussed above. The new water demand model is defined by:

$$\ln(w_{it}) = \alpha_i + \gamma p_{it} + \beta X_{it} + LC_{it} + NLC_{it} + \xi_{it} \tag{2.9}$$

Table 2.6 presents the estimates of conversions and neighbors' conversions on water demand. Conversions have a negative impact on water consumption across all specifications and the magnitude is not only statistically, but also economically, significant ranging from 17.9-23.8% of monthly water demand. Similarly, neighbors' conversions cause a statistically significant reduction in demand, but the magnitude is roughly one-tenth of that for an actual household conversion. This is likely due to the fact that the probability of converting increases when a household's neighbors convert; therefore neighbors conversions may act as a proxy for partial conversions that are not picked up by the NDVI data.

These results, along with the parameter estimates from the landscape conversion model, permits

---

[19] This is unlikely because the immediate contiguous neighbors are removed when creating variables for neighbors landscapes precisely to address this concern.

Table 2.6: Water Demand and Landscape Conversions

VARIABLES	(1)	(2)	(3)	(3)	(5)	(6)	(7)	(8)
Price	FE-IV	FE-IV	FE-IV	FE-IV	FE	FE	FE	FE
Normalization	Yes	Yes	No	No	Yes	Yes	no	no
Classification	40/70	50/60	40/70	50/60	40/70	50/60	40/70	50/60
Conversion	-0.223**	-0.229**	-0.238**	-0.231**	-0.175**	-0.180**	-0.187**	-0.179**
	(0.010)	(0.009)	(0.011)	(0.008)	(0.008)	(0.007)	(0.008)	(0.006)
Neighbors Conversions	-0.0257**	-0.0264**	-0.0267**	-0.0233**	-0.0211**	-0.0208**	-0.0218**	-0.0182**
	(0.00171)	(0.00157)	(0.00181)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	8,101,921	8,101,921	8,101,921	8,101,921	8,101,934	8,101,934	8,101,934	8,101,934
Households	174,216	174,216	174,216	174,216	174,229	174,229	174,229	174,229

Notes: Price treatment for IV instruments for marginal price, and GLS uses the high marginal price. Normalization refers to whether the landscape classifications were based on weather normalized NDVI (Yes) or raw NDVI (No). Classification designates the threshold for landscape being defined as either dry or wet in a given year. Cluster robust standard errors are reported in parentheses. ** $p<0.01$, * $p<0.05$

calculation of extensive margin elasticity. I estimate extensive margin elasticity by simulating the effect of a price increase on the probability of conversion and multiplying this by the change in consumption associated with conversions. This is then compared to the intensive margin elasticity for the Wet group from the conditional demand functions. I present a simple scenario for a 10% increase in the price of water to generate the changes in consumption across both the intensive and extensive margins. The predicted probabilities cannot be directly used for simulation because they rely on the idiosyncratic term.[20] Instead of predicted probabilities I use the coefficients to determine the change in the log odds ratio similar to [73]. Since the initial probability is known I can generate a new probability of conversion based on the price increase. The logit is a nonlinear model and therefore the new conversion probability depends on the initial probabilities and values of the dependent variable. For the initial values in the elasticity calculations I use the probability of conversion excluding all households that are dry at the beginning of the sample, and the median high block price.

Total elasticity is assumed to be the sum of intensive margin elasticity from the Wet group estimated from equation (3) and the price-induced change in consumption from landscape conver-

---

[20] In the conditional logit this is not directly estimated, and in random effects model the random component can be backed out but is estimated simultaneously with the coefficients and therefore is not appropriate for use in simulations that change the independent variables.

sion. Using the random effects specification produces an extensive margin elasticity that is 7-48% of the total elasticity, depending on the number of lags used in the model. The model with price lags of two years generates very high new probabilities resulting in the upper end of the range. Parameters from the conditional logit generate extensive margin elasticity that comprises 16-43% of total elasticity. It is important to remember that this definition of the extensive margin refers exclusively to changes in landscape, and the intensive margin includes behavioral changes and investments in water efficient appliances inside the home. The wide range should provide cause for caution in interpreting the precise magnitude of the effect, and is an important concern in using logit parameters to identify the size of causal effects [98]. The important finding is that water rates do increase the probability of landscape conversion and that landscape conversions have a large impact on water use. The durability of price effects is a critical element of extensive margin changes, since once a landscape is converted water savings will likely persist into the future.

These results suggest promise for incorporating satellite data capturing landscape choices in modeling water demand. In order to put the landscape conversions in Phoenix into context I overlay the spatial distribution of long-run changes in consumption with the prevalence of landscape conversions. Figure 2.3 displays the quantiles for block-level changes in consumption, defined as average consumption in 2008-2009 divided by average consumption in 1998-1999. Lower quantiles represent area where water consumption dropped the most, whereas higher quantiles show regions that did not experience much change in demand over the course of the sample. Consumption is averaged over all houses at the Census block group level to preserve anonymity[21]. The map visualizes the spatial heterogeneity of changes in consumption over time. The hatched regions of Figure 2.3 represent census blocks where at least 5% of the houses in the sample converted their landscape. The intense effect of landscape conversions on changes in long-run demand is striking in Figure 2.3. Central and northeastern Phoenix had very few conversions and consumption in those areas is relatively constant, or even slightly higher than 10 years ago. Conversely, western Phoenix experienced many conversions and those areas reduced consumption by 20-30% over the course of the sample. Figure B.6 in the Appendix shows the same data in percentage changes as opposed to quantiles of percentage change. Observing the distribution of homes with water-intensive landscape

---

[21] For confidentiality concerns I drop all census block groups with fewer than 20 houses from the map.

also provides a baseline for potential water reductions in the future. If most households have already converted further reductions in outdoor use may not be feasible, suggesting that policies to address water scarcity should focus on indoor use or augmenting supply.

Figure 2.3: Landscape Conversions & Consumption over Time

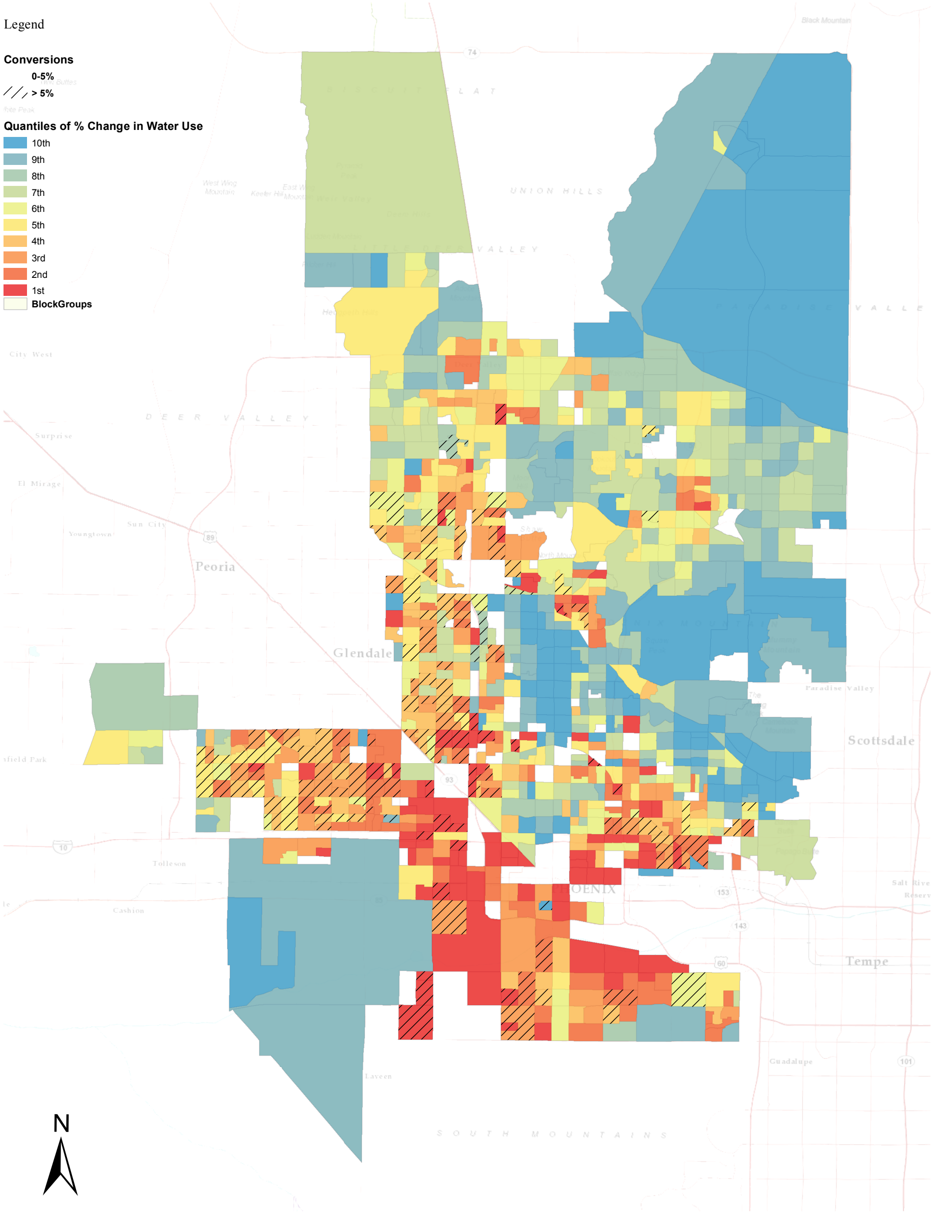

Notes: The shaded coloring are based on the average change in water consumption during the course of the sample. Specifically it measures the average of consumption in 2008-2009 divided by average consumption in 1998-1999. Thus the lower quantiles have reduced consumption by more than the higher quantiles. Census blocks with less than 20 houses in the sample are removed for confidentiality concerns. The hatched blocks are where more that 5% of the households in the sample experienced landscape conversions.

2.5 Conclusion

This paper examines the role of landscape dynamic in water demand by merging a time series of satellite data with monthly water metering records. Jointly observing water demand and changes in complementary goods over time enables estimation of demand elasticity on the intensive and extensive margin. Prior research treats the extensive margin implicitly by examining the difference in short and long run demand with the assumption that changes in the extensive margin can only occur in the long run. As concerns of water scarcity increase it is critical to develop a deeper understanding of demand as seen in energy economics [63, 68]. I examine the heterogeneity in water demand due to an important complementary good by conditioning on landscape decisions over time. Key differences exist between households across the landscape spectrum. Maintaining either a wet or dry landscape decreases price elasticity, and households that maintain a green lawn are more responsive to weather conditions.

In a discrete choice model I find that water rates increase the probability of converting from water intensive vegetation to a xeric landscape. Landscape conversions have a significant impact on demand; reducing consumption by up to 23%. Changes along the extensive margin may constitute up to 48% of aggregate demand elasticity. There appears to be a strong social component to the landscape choice, as the number of neighbors with a dry landscape is a strong predictor of landscape conversions. As social norms evolve, and consumers sort to live close to like-minded people or collocate along correlated unobserved heterogeneity, the pattern of conversion clusters is likely to continue. There are important implications for clusters of conversions in a desert city as xeric landscapes exacerbate the urban heat island effects in a highly nonlinear way. Establishing xeric landscapes may lead to a tradeoff between water and energy conservation as hotter neighborhoods use more energy for cooling. This highlights the notion that a green landscape provides direct benefits and indirect benefits in arid regions.

Analyzing long-run changes in demand shows that those who maintain a green landscape have a relatively constant trend in consumption over time, while areas that convert appear to be driving large scale reductions in demand. If, as expected, the transition of landscaping practices in Phoenix lead to monotonic increases in xeric landscape there will be important implications for water demand. There will be smaller peaks in the summer and demand will be less responsive to changes in water rates and fluctuations in the weather. Conversions smooth water consumption, but also reduce the potential savings from outdoor water restrictions during droughts. Understanding the stock of existing landscape, and speed of conversion, has enormous consequences for long and medium run planning by quantifying the potential savings from outdoor use.

Introducing satellite data as a proxy for landscape is a promising development to understanding water demand, and there are many rooms for advancement. While the data I use have a consistent time series back to 1984 and are publicly available there are much higher resolution data that will improve the accuracy of landscape classification, and can potentially identify specific types of vegetation. Introducing multiple geographical locations introduces challenges when using remote sensing data, but also adds cross sectional variation in water rates and utility policies, and thus remains a propitious area of future research.

Chapter 3

HETEROGENEOUS RESPONSES TO SOCIAL NORMS FOR WATER CONSERVATION

3.1 Introduction

Water utilities, particularly those in the Western United states, face pressure to meet rising water demand. California, for example, has required that utilities reduce demand by 20% by 2020[1] and Texas recently adopted a $2 billion water infrastructure fund.[2] To achieve these conservation goals, water utilities are exploring every strategy in their arsenal. Incorporating social comparisons into information campaigns is an attractive policy instrument aimed at changing consumer behavior when the price mechanism is not feasible due to market structure or political concerns. Social comparisons are particularly well-suited for markets where consumers may have difficulty optimizing their level of use due to a lack of salience over personal consumption quantities, the costs of water or energy saving investments, or the commodity's price signal itself.

Relative information has the additional advantage of leveraging motivations that have been extensively studied by psychologists and behavioral economists. This fits into a broader context of applying behavioral economics in public policy by making relatively small, and seemingly innocuous, changes to a decision framework such as changing the default option. Numerous examples of behavioral economics in policy initiatives span food choice, retirement allocation, organ donation, and utility demand. [46, 81, 16, 52, 45]. Opower, an early pioneer for social comparisons in utility demand, sends households information comparing their energy consumption to a group of their peers. Opower's proven success, with almost fifty utility clients spanning six countries and saving over 2.5 billion kilowatt hours, is a shining example of the new application of behavioral economics to achieve environmental goals.[3] Research on social norms for water consumption assesses how this relatively new tool fits into the array of conservation options at the disposal of water managers seeking to influence demand with a very different profile than energy.

We harness a rich dataset from WaterSmart Software, a clean technology company establishing itself as the Opower of water, to analyze the effects of social norms as a water conservation tool. WaterSmart is currently working with twelve utilities, in different stages of development. Three began as randomized pilot programs and are mature enough for analysis. Randomized field

---

[1] See http://www.swrcb.ca.gov/water_issues/hot_topics/20x2020/ for more information.

[2] See http://openstates.org/tx/bills/83/SJR1/ for the full text of the bill.

[3] Information on Opower can be found at their website - http://opower.com/customers.

experiments are an exciting area of economic research that addresses issues of endogeneity due to selection into treatment in the data generating process; see among others [91, 72, 65, 87, 89]. The randomization process allows for simple estimation of average treatment effects with easily defensible assumptions. This contrasts to using observational data that requires more stringent assumptions and complicated econometric techniques to identify causal effects. Eliminating endogeneity allows us to devote energy to understanding the mechanisms through which consumers reduce consumption by analyzing the heterogeneity in treatment response.

Exploring heterogeneity uncovers the mechanisms through which the program reduces consumption as well as identifies ways to improve the program's performance and cost effectiveness. Our results show that there is significant heterogeneity in the response to the program across initial water use within pilots and stark differences in the average treatment effect (ATE) between pilots. The pilot with the largest ATE, at 5.2%, saved three times as much water as the pilot with the smallest. We relate the pattern in ATE magnitudes to general pilot differences in Section 3.5. Within-pilot heterogeneity shows that baseline water consumption, defined as average water use prior to the treatment, has the most significant and consistent effect on treatment response. Based on quantile regressions and regressions that interact the treatment effect with baseline consumption, households that use more water have a greater response to treatment in all three pilots. Households that use high levels of water have a greater financial incentive to reduce consumption and may be more able to significantly reduce consumption through low marginal cost mechanisms such as fixing leaks and changing irrigation practices. Other conservation programs using social norms found a similar pattern of results for both energy [5] and water [53].

The intra-pilot heterogeneity of baseline consumption is not reflected in the results for ATEs across pilots. Though there are only three pilots to analyze, the magnitude of the ATE is inversely related to the average pilot-level consumption. The incongruence between the inter- and intra-pilot effects of baseline water use is not surprising and lies in the nature of the social norm campaign. The social norm message compares a household's consumption to a group of peers with similar determinants of water demand in the same utility. For households with similar, or lower, consumption than their local peers the program may cause little, or negative, water savings due to the favorable normative message they receive.[4] By design, the social norm message does not influence consumer behavior through the relative water use between pilots. To explain inter-pilot heterogeneity, differences in specific utility policies and the demographics of the customer base are more relevant and will be discussed later.

While baseline water use influences the intra-pilot heterogeneity in treatment response in a consistent way across pilots, this is not the case for ideology and housing values, which we use as a proxy for income. The significant results for conditional average treatment effects (CATEs) from interactions with ideology quintiles display a different role for within-pilot variation in ideology in the separate pilots. Interactions with housing value quintiles only produce significant results for the CATE in one of the three pilots.

We examine the mechanisms through which consumers reduce water use in response to the social norm by examining the time pattern of savings and interactions with existing conservation programs. Since monthly savings do not increase in the summer months the behavioral response likely focuses on indoor water use. Different patterns across utilities imply that in some pilots consumers make investments in consumer durables or establish permanent behavioral changes, whereas in one pilot the savings appear to be transient even while messaging persists. Interacting the social norm with existing conservation programs finds that social norms act as a complement to existing programs. The social norm message contributes to higher participation rates in other utility programs, which then contributes up to 20% of the average treatment effect assigned to the social norm message.

This paper is organized as follows. The next section briefly discusses the motivations for the use of information campaigns and social norm messaging over pecuniary mechanisms. The following section presents the existing literature estimating the impact of social norm campaigns in the context of utility use. Then, Section 3.4 discusses the background of WaterSmart, describes the data, and provides details on the experimental design. Section 3.5 overviews the estimation strategy and presents the empirical results. Section 3.6 provides plans for future research as we receive more data. Concluding remarks are offered in Section 3.7.

---

[4] In other studies [5, 53] the categorical specification of the normative message did not cause low-use household to increase consumption. In future research we hope to analyze the effect of the specific messages through a regression discontinuity design similar to [5].

3.2 Theoretical Background

Using social norms in public policy is a relatively recent development. Although a water utility's tariff structure is an important tool in managing demand, there are a number of reasons why the study of non-pecuniary approaches is appealing. First, although some utilities incorporate scarcity pricing into water rates the primary objective in designing rate structures remains cost recovery. Those utilities that do account for scarcity may not fully price the complex risks to water supplies associated with climate change. Second, even under the assumption that utilities appropriately incorporate environmental externalities and increasing scarcity into price schedules, many utilities face zero profit constraints that limit their ability to raise rates.[5] In addition to zero profit constraints, the public utilities that serve most major population centers may find it difficult to raise rates due to concerns over equity and political backlash by constituents of local governments. While the price mechanism is the most efficient way to reduce consumption, other instruments must be considered when raising prices or introducing Pigouvian taxes is not feasible. Command and control interventions, such as imposing mandatory restrictions on outdoor water use, are a common alternative to raising water rates but are shown to reduce consumer welfare [69, 94]. In this policy climate water managers must seek alternative mechanisms to reduce demand .

Incorporating social comparisons into information campaigns is an attractive policy instrument aimed at changing consumer behavior when the price mechanism is not feasible due to market structure or political concerns. These benefits are augmented in settings where it is difficult for consumers to optimize their level of use due to a lack of salience over personal consumption quantities, the costs of water or energy saving investments, or the commodity's price signal itself. There are reasons to suspect that consumers do not, or are unable to, choose their monthly water use based on a careful maximization process. Given that the outlays on water generally comprise less than 2% of household income, and that a proportion of water use will remain constant to fill basic human needs, many consumers may not devote much attention to their water bill. The consumers that do scrutinize their bills not only need to know the water intensity of various household ac-

---

[5] While utilities could restructure water rates to raise the marginal price while leaving the average price unchanged, there is a debate whether consumers actually respond to the average or marginal price of water [80, 102, 60, 125].

tivities, but also how to translate complex rate structures into marginal costs. The complexity in water rates may cause consumers to respond to average price instead of marginal price as suggested by [80, 79]; leading to increased water use since the average price can be substantially lower than marginal price in utilities with increasing block rates. [124] shows that automatic energy bill payments increase energy consumption by decreasing the salience of the price signal and [60] finds that water utilities that do not present the marginal price on their bills have lower price elasticity. If households are consuming at the optimal level prior to receiving a social comparison, the information should have no effect on their water use. These studies, and the results of this paper, suggest that salience matters, and information campaigns in water may work in part by bringing water use to the forefront of consumer decision.

Relative information has the additional advantage of leveraging motivations that have been extensively studied by psychologists and behavioral economists. [48] attempts to reconcile the divide between theories in economics, where rational agents are ”pulled” by incentives, and sociology, where agents are ”pushed” by their surrounding environment. [129] describe a third route between these two extremes whereby agents still make rational decisions, but can be ”nudged” by changing seemingly inconsequential elements in the choice environment. Comparing household use to a group of peers establishes a baseline level of conventional consumption. Knowledge of this convention may generate civic pressure to refrain from resource profligacy. Alternatively, discovering that personal use levels are far above the convention may signal the availability of low cost investments to reduce use. Recent research applies these ideas from behavioral economics to a range of public policy objectives spanning food [46], organ donation [81], retirement savings [16], and several examples within environmental economics [5, 52, 45]. Research on social norms for water consumption assesses how this relatively new tool fits into the array of conservation options at the disposal of water managers.

3.3 Existing Literature

A series of papers by Hunt Allcott and coauthors [7, 4, 5, 9, 8, 6] on various dimensions of energy economics help establish a role and set of techniques for applied environmental economists to investigate motivations explored in behavioral economics. Allcott’s work analyzing data from

Opower [5, 8, 9] parallels our research on water conservation. [5] finds that Opower's program of providing social comparisons along with technical information reduced energy demand by an average of 2%. There is considerable heterogeneity in the treatment effect across the consumption distribution, with the right tail experiencing the largest savings.

Allcott's work on Opower serves as a template for the analysis since water and energy share many common features. However, there are key differences between water and energy that warrant studying the efficacy of social norms in the water sector. First, while many regions experience seasonal peak demand in the summer for both energy and water, the marginal value of activities driving the peak is likely different. Air conditioning, a major component of peak energy demand, has significant health benefits in combating heat stroke [123]. Conversely, research shows that summer demand due to discretionary outdoor use has higher demand elasticity than indoor use [94, 140, 39, 49] suggesting a lower marginal value for activities shaping peak demand. Additionally, while the challenge of peak energy demand is the cost of supplying daily maximum demand with the most expensive energy production, water utilities are primarily concerned with seasonal demand peaks that are correlated with stressed summer water supplies. Second, energy externalities are primarily global in nature associated with climate change risks from greenhouse gas emissions;[6] conversely, water consumption externalities such as groundwater exhaustion and reducing in-stream flows for ecosystem services have predominantly local impacts. Consumers may be more receptive to limiting externalities whose effects are concentrated within their region as opposed to providing the global public good of emission reductions. Due to both the consumption patterns and localism of environmental impacts, a priori, we may expect to see a different pattern of response to social norms in energy and water.

Ferraro, Price, and Miranda [52, 53, 54] analyze an information campaign using social norms in Georgia to reduce water consumption during a severe drought. [52] explicitly test the efficacy of different combinations of technical information, an appeal to support a public good, and a social comparison and find that a social comparison along with technical information is most effective in reducing consumption. Investigations of the heterogeneity for the same program find that wealthy

---

[6] There are local health benefits from energy conservation associated with air quality depending on the fuel source and location of power generation.

consumers with high water use are most responsive [53] and that the effect decreases over time [54]. A major difference between that program and our study area is that WaterSmart applies a continuous treatment over many months, whereas the Georgia program had only one mailing and one follow-up. The dynamic effects of the program in Georgia are analyzed after the program is over, while in our setting we investigate whether conservation gains from social norms are sustainable over many treatment periods. Additionally, that project only encompassed one service area; so while intra-program heterogeneity can be explored it is not possible to compare responses across various regions. Lastly, the program was in response to a severe drought during which public perception towards water consumption was likely heightened.

In assessing the efficacy of social norms as a water conservation tool it is essential to place this instrument in the context of multiple concurrent policies, generally defined as demand side management. Utilities provide home water audits to provide a personal recommendation for water savings as well as offer rebates to replace toilets and washing machines and convert lawns to low-water vegetation. Part of the attraction of social norms is that existing techniques are either not cost effective or do not produce the necessary reductions in demand. [17] shows that rebates for high-efficiency toilets do very little to incentivize new consumers; rather the policy merely rewards or speeds up purchases that would naturally occur in absence of the rebate. [94] and [69] find that outdoor watering restrictions implemented in arid regions or during extreme droughts cost households over $100 per irrigation season. Past research finds that non-pecuniary programs are less effective than price and mandatory restrictions [115] and that more research is needed to explore policy interactions [116]. Section 3.4.3 goes into details on the existing conservation programs in our study area and Section 3.5.4 conducts an empirical analysis of the interplay of social norms with these programs.

3.4 Background & Data

3.4.1 WaterSmart Software

WaterSmart Software[7] (henceforth WaterSmart) is a clean technology company that contracts with water utilities to reduce demand by providing consumers with Home Water Reports (HWR) containing a social comparison of water use and technical advice. Additionally, WaterSmart assists utility companies with the analysis and interpretation of water metering data. Utility companies desire reductions in demand to offset costly temporary water purchases, delay investments in new supplies, and meet regional or state objectives. Raising rates is often politically infeasible, and incentive programs such as toilet rebates are expensive and may not meet goals of additionality [17].

WaterSmart's business model is to send customized Home Water Reports (HWR) with three sections to customers. Some HWRs are mailed to households, while others are sent via email. The main component is a social comparison that contrasts a household's water consumption to that of its peers and then provides a three-tiered qualitative WaterScore based on their relative position. The second section is a list of three personalized recommendations for strategies to save water and rebates available from the utility. Lastly, there is a space that offers incentives for signing up for an online account for more detailed data on water use and additional water conservation tips. Figure 3.1 shows an example of the HWR for an above-average home[8]. A sample of the additional information available through the web portal is presented in Figure 3.2.

---

[7] Additional information on WaterSmart is available on their website http://www.watersmartsoftware.com/ .

[8] The other categories are "Good" and "Wise" and are very similar to this HWR and are available upon request.

Figure 3.1: Home Water Report

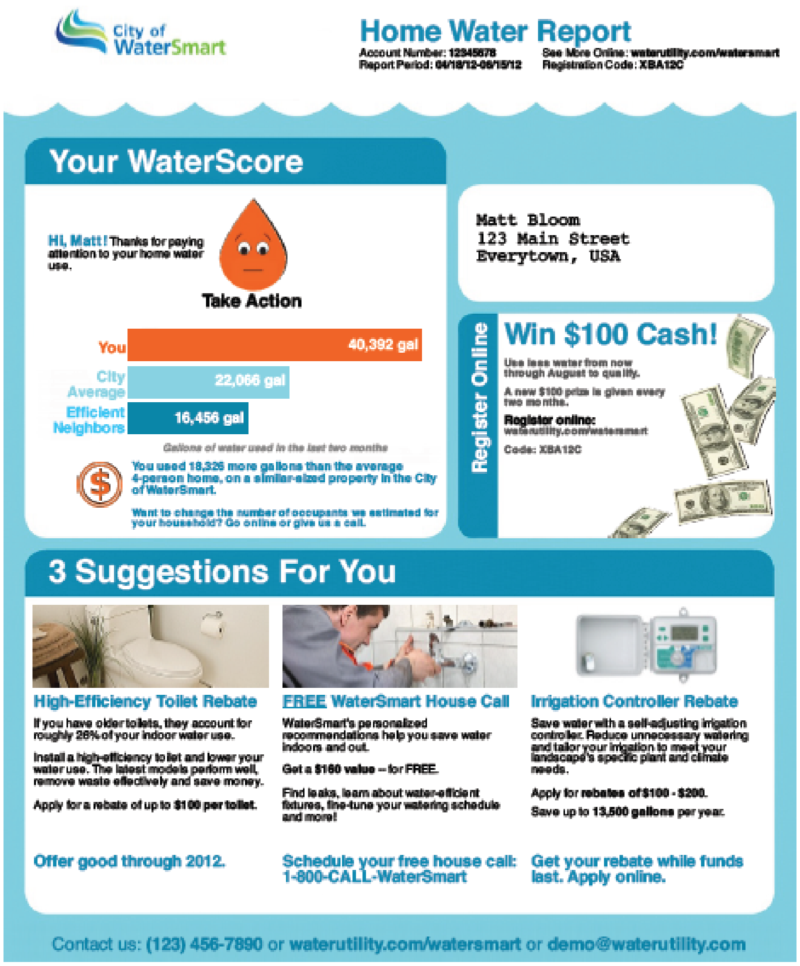

City of WaterSmart

Home Water Report

Account Number: 12345678
Report Period: 04/18/12-06/15/12

See More Online: waterutility.com/watersmart
Registration Code: XBA12C

Your WaterScore

Hi, Matt! Thanks for paying attention to your home water use.

Take Action

You	40,392 gal
City Average	22,066 gal
Efficient Neighbors	16,456 gal

Gallons of water used in the last two months

You used 18,326 more gallons than the average 4-person home, on a similar-sized property in the City of WaterSmart.

Want to change the number of occupants we estimated for your household? Go online or give us a call.

Matt Bloom
123 Main Street
Everytown, USA

Register Online

Win $100 Cash!

Use less water from now through August to qualify.

A new $100 prize is given every two months.

Register online:
waterutility.com/watersmart

Code: XBA12C

3 Suggestions For You

High-Efficiency Toilet Rebate

If you have older toilets, they account for roughly 26% of your indoor water use.

Install a high-efficiency toilet and lower your water use. The latest models perform well, remove waste effectively and save money.

Apply for a rebate of up to **$100 per toilet.**

Offer good through 2012.

FREE WaterSmart House Call

WaterSmart's personalized recommendations help you save water indoors and out.

Get a **$160 value** -- for FREE.

Find leaks, learn about water-efficient fixtures, fine-tune your watering schedule and more!

Schedule your free house call: 1-800-CALL-WaterSmart

Irrigation Controller Rebate

Save water with a self-adjusting irrigation controller. Reduce unnecessary watering and tailor your irrigation to meet your landscape's specific plant and climate needs.

Apply for **rebates of $100 - $200.**

Save up to **13,500 gallons** per year.

Get your rebate while funds last. Apply online.

Contact us: (123) 456-7890 or waterutility.com/watersmart or demo@waterutility.com

Notes: Home Water Reports have WaterScores of "Wise", "Good", or "Take Action". The bottom panels contain suggestions that are customized based on household data. The right "Win $100" offers different incentives or messages

Figure 3.2: Online Web Portal

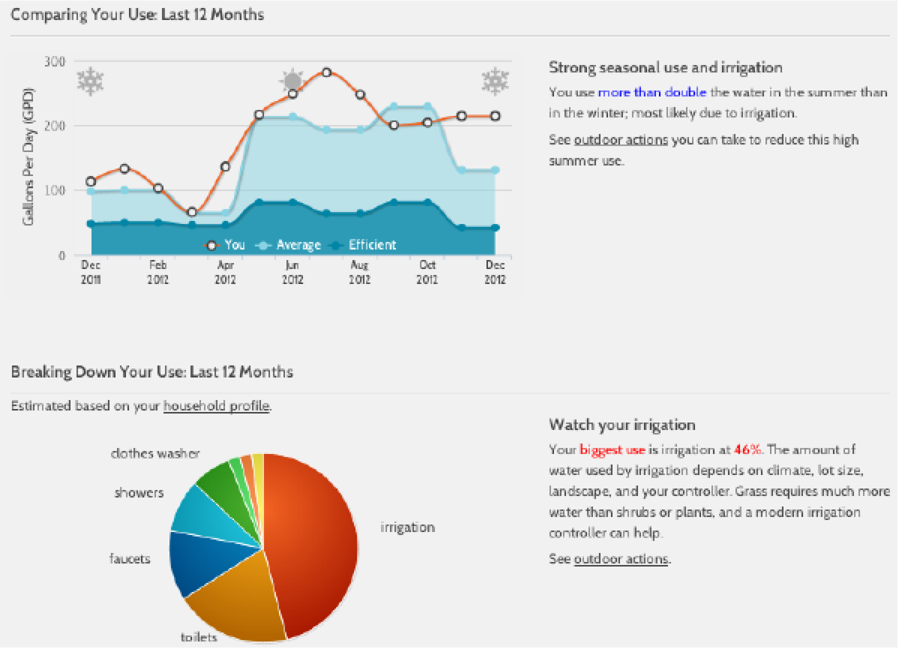

Comparing Your Use: Last 12 Months

Strong seasonal use and irrigation

You use more than double the water in the summer than in the winter; most likely due to irrigation.

See outdoor actions you can take to reduce this high summer use.

Breaking Down Your Use: Last 12 Months

Estimated based on your household profile.

Watch your irrigation

Your biggest use is irrigation at 46%. The amount of water used by irrigation depends on climate, lot size, landscape, and your controller. Grass requires much more water than shrubs or plants, and a modern irrigation controller can help.

See outdoor actions.

Notes: Online web portal is available to the treatment group and can be viewed as an additional "opt-in" treatment over the base treatment. It offers additional data on consumption, utility rebates, and technical information to help reduce water consumption.

3.4.2 Data

We combine several sources of data for the analysis. WaterSmart provides the water metering records at the household level, as well as structural features of the house and treatment status from three California utilities in pilot programs. Daily weather data is obtained from the National Oceanic and Atmospheric Administrations National Climatic Data Center.[9] We match households to the nearest weather station with available temperature and precipitation data and aggregate daily values within the irregular read periods of household water read observations.

Income data from the Census Bureau is not available at a fine enough geography to proxy for household level income. Instead, we use block level median housing values from Zillow, an on-line real estate company with a database of millions of homes and sales transactions across the US.[10] Lastly, we gather voting records from the Statewide Database hosted by the University of California, Berkeley to create ideological indices capturing the environmental sympathies of individual block groups.[11] The Data Appendix (coming soon) details the process of creating a working dataset through all the primary data sources, as well as quality assurance measures taken before conducting the analysis.

We expect that environmental attitudes influence the manner in which households interpret the social norm for water conservation. The Green Ideology Index (GII) is formed using six census block level measures from the 2008 and 2010 California elections: the percent who voted Democrat for U.S. Senator in 2008 and Governor in 2010, the percent who voted yes on "pro-environment" Propositions 7 (2008), 10 (2008), and 21 (2010), and the percent who voted no on Proposition 23 (2010), which would have suspended the state's 2006 Global Warming Solutions Act.[12] For

---

[9] We use the temperature and precipitation data from the Global Historical Climatology Network-Daily (GHCN-D), which is publicly available to download at http://www.ncdc.noaa.gov/data-access/quick-links#ghcn.

[10] http://www.zillowblog.com/research/2012/01/21/zillow-home-value-index-methodology/

[11] The data are publicly available at http://statewidedatabase.org/.

[12] Proposition 7 would have required California utilities to produce half their electricity from renewable resources by 2025. Proposition 10 would have allocated $5 billion as cash incentives for high fuel economy and alternative fuel vehicles and R&D for and education on renewable energy and alternative fuel technologies. Proposition 21 would have increased vehicle license fees in the state by $18 in order to raise $500 million a year dedicated to California State Parks. Proposition 23 would have suspended the Global Warming Solutions Act of 2006 until the state's unemployment rate decreased to 5.5% for four consecutive quarters, a condition not met for decades.

each measure, we assign a score of 100 to the top 0.1% of blocks with the greatest percentage of individuals voting along green-friendly lines and a score near zero to the least green-friendly census block. Averaging these ordinal scores together yields the GII and maintains the 0 to 100 scale across all of California (the greenest areas are consistently green across all votes). We match individual households to the GII via their census block as a proxy for their own ideology under the assumption that like-minded individuals tend to cluster around green-related amenities that cater to their preferences.

3.4.3 Utilities

The three pilots provide an opportunity to estimate the effect of social-norm messaging in various climates and within communities with different demographic characteristics. For confidentiality concerns, we will refer to the utilities by the pseudonyms: Watertown, Aquaville, and Hydroburg. footnoteWaterSmart partners with water utilities who prefer that their names are kept anonymous when presenting data on consumer water consumption. Table 3.1 illustrates basic differences between the three pilots. Watertown and Aquaville are located in more moderate climates than Hydroburg; therefore water demand exhibits a different seasonal component. Since seasonality in water demand is primarily manifested through outdoor use, such as maintaining landscapes and refilling pools, the mechanisms by which water conservation programs affect consumption may vary by climate. The pilots also vary by average income and environmental ideology - both of which have an effect on water consumption. Hydroburg, the warmest, richest, and least environmentally conscious pilot has the highest baseline consumption. Watertown, the wettest, poorest, and greenest pilot, consumes the least water on average. A priori, these differences across pilots lead us to expect varied responses to WaterSmart's social-norm messaging program. For example, [36] find that the response to social norms in energy is different for liberals and conservatives. In the case of income, more affluent consumers are more able to make larger water saving capital investments, however poorer households have a stronger financial incentive to reduce consumption because their water bill is a higher proportion of their disposable income.

The utilities also differ drastically in the characteristics of their service area. Aquaville has the largest customer base, serving over 1 million people. Due to its size there is greater variation in both

Table 3.1: Summary Statistics By Pilot

Pilot	Water Use	Temp (F)	Rain (in)	Income	Home Value	Green Index
Watertown	202	70	2.2	57,333	358,394	71
Aquaville	282	67	1.7	96,696	670,242	62
Hydroburg	348	72	0.8	120,943	815,231	34

Notes: The water use reported is baseline water consumption in gallons per day; Temperature is the average monthly maximum temperature in degrees Fahrenheit; Precipitation is the monthly average in inches, median income is averaged at the census block group, and the green index is averaged at the census block.

climatic zones and demographic characteristics. Hydroburg is also a relatively large utility with over 300,000 customers, while Watertown serves less than 10,000 people. In addition to differences in the customer base, as evidenced in Table 3.1, utilities of such different sizes do not have the same resources at their disposal. Luckily, they do have similar existing conservation programs. Table 3.2 shows the various conservation programs in the three pilots *in addition* to WaterSmart's program. All the utilities offer rebate programs, as well as other programs tailored to the community's needs.[13]

The three utilities have very different rate structures that may influence treatment response in different ways. Watertown has the simplest rate with a fixed cost and single volumetric charge for the duration of the sample.[14] Aquaville has a standard increasing block rate structure with three tiers. The rate structure is relatively flat, with a 51% increase from the lowest to the highest tier. Hydroburg has the most complicated and steepest rate structure. The rate is determined by the percentage of water a household uses relative to an allocation - determined by irrigable area and occupancy. The rate for the highest tier, reserved for consumption that is more than 200% above allocation, is over 600% higher than the base rate providing a strong financial incentive for conservation. Since the allocation acts as a signal for the appropriate level of use Hydroburg's water rate is structured to penalize households who use relatively more water than similar households. Therefore social norm messaging that highlights relative consumption resembles Hydroburg's rate structure, and may have a different impact compared to pilots where the rate structure does not

---

[13] Hydroburg offers programs to address outdoor water use, while Watertown focuses on community engagement.

[14] During the treatment period plans were in place to move to an increasing block rate tariff.

reflect relative use. Since water demand is a function of utility-level policies, such as water rates and conservation programs, as well as the demographics and preferences of the consumer pool, we focus on the identification of treatment heterogeneity within pilots. Differences across pilots will be left to qualitative assessments driven by the empirical results. We leave the analysis of interaction between demand elasticity and social norms for future research.

Table 3.2: Utility Conservation Programs

	Watertown	**Aquaville**	**Hydroburg**
Rebates			
Toilets	X	X	X
Clothes Washer	X	X	X
Lawn Conversion	X	X	X
Sprinklers			X
Irrigation Controller		X	X
Technical Advice			
Home Water Audits	X	X	
Community Classes	X		

Notes: Only programs that are available from the utility are included. Some programs are administered through regional bodies, and additional resources are available from state and regional agencies.

3.4.4 Experimental Design

This section serves to describe the experimental design, defend the randomization process, and motivate the ability to identify causal effects and estimate several measures of treatment heterogeneity. Randomization eliminates concerns over selection bias, but it is important to evaluate the differences between the treatment and control group to determine the effectiveness of the randomization. Sample sizes, start dates, and the number of households treated in these pilots can be found in Table 3.4. We did not have control over the randomization process and the sample sizes are small enough to warrant inspecting the balance in observable variables across the treatment and control groups. Table C.16 displays the difference in means and the associated p-values from t-tests for several important variables. There are some variables that do have significant differences that warrant further examination, the most important being baseline water use in Hydroburg. Ensuring balance over the observable variables that are expected to influence heterogeneity in the treatment effect is an important prerequisite for an investigation into heterogeneous treatment effects. To ensure balance we employ a matching algorithm, described in the Data Appendix. Household fixed effects in our preferred specification will absorb all of the difference in household-level variables, but it is still important to assess if the treatment and control group have the same trend over time. Figure C.8 shows the time series of water use in gallons per day for both the treatment and control groups. It is clear from the graphs that all the pilots have similar trends over time across treatment status. This alleviates concerns about the significant difference in mean baseline water use in the estimation of the average treatment effect for Hydroburg even without matching. (See the Appendix for water use time series for each pilot in terms of year-over-year percentage changes.)

Table 3.3: Differences in Means Between Control and Treatment Groups

(a) Variables used to explore heterogeneity

Pilot	Baseline Water	Green Index	Home Value	Income
Aquaville	-1.211 (0.861)	0.3 (0.48)	-12,107 (0.946)	-84 (0.77)
Hydroburg	13.89 (0)	2.7 (0.606)	109,701 (0.979)	3,337 (0.011)
Watertown	-6.541 (0.993)	0.5 (0.195)	-19,930 (0)	-530 (0.21)

(b) Structural Features of the House

Pilot	Sqft	Lot Size	Single Family Home	Year Built
Aquaville	-11 (0.753)	-178.4 (0.116)	0 (0.112)	1.7 (0.538)
Hydroburg	208 (0.292)	217 (0.738)	-0.015 (0)	8.5 (0.495)
Watertown	-32 (0.123)	-124.6 (0.854)	-0.029 (0.126)	1.2 (0.079)

Notes: Values reported are the control mean minus the treatment mean. P-values for the difference from zero are in parenthesis. Sqft is the square feet of the home, the Lot Size is measured in square feet, Single Family Home is a proportion of the sample in a detached house, and Year Built is the year the home was built. While we don't use income explicitly to explore heterogeneity, the use of home value is used as a proxy because of its finer geographical granularity.

One avenue for addressing heterogeneity that is gaining attention in both applied, and theoretical, literature is the estimation of quantile treatment effects (QTE) [55, 56, 5, 53, 22, 58, 59]. This is important for policymakers who want to know the type of households that are most responsive to an intervention and whether a program may have the perverse effect on certain segments

Figure 3.3: Baseline Water Use by Pilot

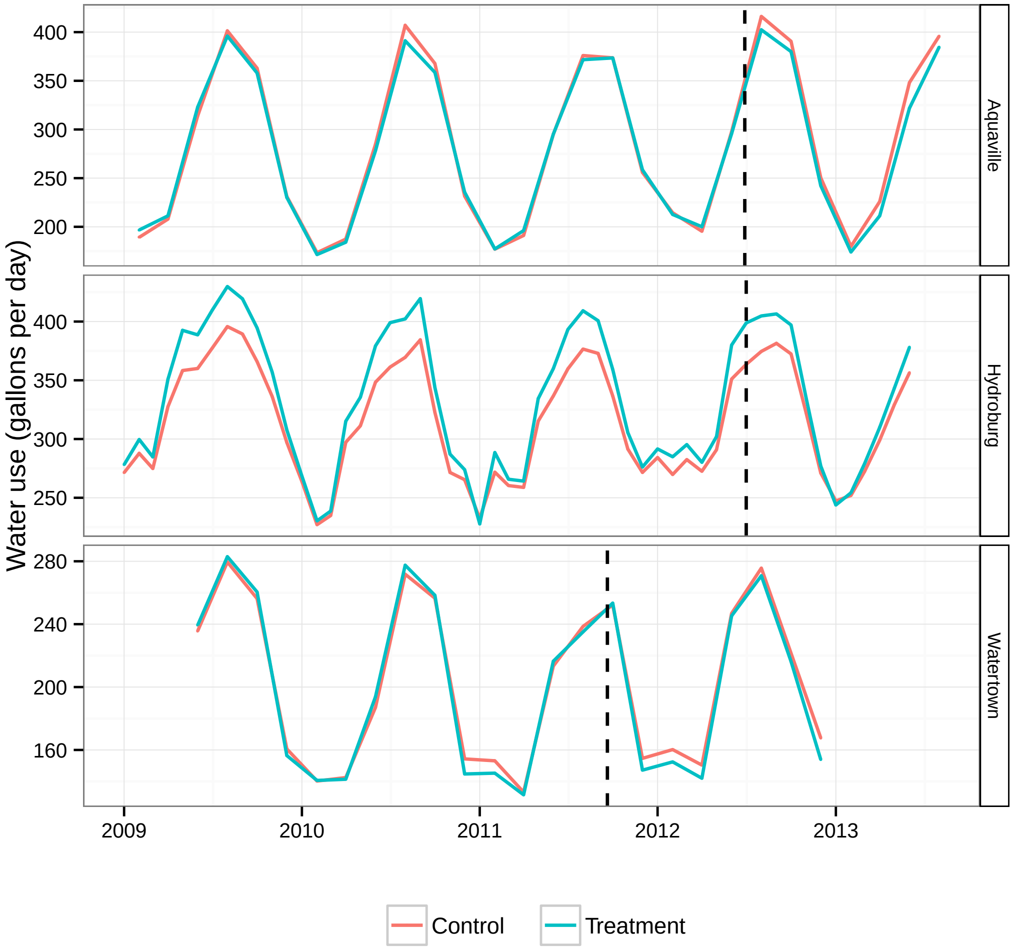

Notes: The graph contains the mean water used per day in gallons for both the treatment and control groups over time. The vertical dashed line indicates the start of the program for each pilot.

Table 3.4: Sample Sizes

Pilot	Start Date	End Date	N: Obs	N: Post-Treat	HHs	Treated HHs
Watertown	2011-09-20	2013-01-01	44,822	16,266	2,575	992
Aquaville	2012-06-28	on going	93,664	18,583	3,096	1,547
Hydroburg	2012-07-01	on going	666,981	115,649	11,710	1,184

of the population. The aforementioned literature highlights some of advantages of unconditional QTEs that provide inference on the marginal distribution of the outcome variable, as opposed to the conditional distribution. The conditional quantile does not have the same interpretation as the conditional mean, and in general the marginal distribution is more relevant in a policy context. The interpretation of the QTE is subject to the rank preservation assumption that presumes treatment assignment does not affect the relative ranking of households consumption. If the assumption holds the process of treatment assignment does not systematically determine the consumption quantile for a household. Assuming rank preservation holds a control household at a given quantile is an appropriate counterfactual for the corresponding treated household. Therefore the QTE represents quantiles of the treatment effect, and provides inference on the distribution of treatment effects across households at different quantiles. Without this assumption, we can only interpret the results from quantile regressions as the QTE on the outcome distribution [55], which still may have economic and policy relevance.

It is not possible to directly test this assumption, but the literature [20, 21, 44] suggests an ad hoc assessment that examines the difference in means within quartiles across treatment status for key variables.[15] Many significant differences in means make the rank preservation assumption difficult to justify since the distribution of observables that affect water demand vary across treatment status. However, even good balance within quartiles does not prove that the rank preservation assumption holds. Table 3.5 reports the p-values for difference in means for several variables within quartiles for each pilot. 25% of the differences have p-values less than 0.05, suggesting that the quantile treatment effects may provide a decent approximation for the quantiles of the treatment effect. In the next section we present our estimation strategy and empirical results.

---

[15] All the variables in this test are pre-determined relative to treatment assignment, and this primarily motivates that households at a given quantile have similar characteristics.

Table 3.5: Motivating the Rank Preservation Assumption

Pilot	Quantile	Baseline Water	Green Index	Lot Size	Sq.Ft.	Year Built
Aquaville	1.0	(0.352)	(0.102)	(0.013)	(0.111)	(0.146)
Aquaville	2.0	(0.003)	(0.006)	(0)	(0.747)	(0.138)
Aquaville	3.0	(0)	(0.624)	(0.001)	(0.939)	(0.131)
Aquaville	4.0	(0.742)	(0.861)	(0.109)	(0.552)	(0.959)
Hydroburg	1.0	(0)	(0.303)	(0.195)	(0.001)	(0.368)
Hydroburg	2.0	(0)	(0.187)	(0.895)	(0)	(0.04)
Hydroburg	3.0	(0.002)	(0.001)	(0.318)	(0.086)	(0.812)
Hydroburg	4.0	(0.262)	(0.942)	(0.581)	(0.489)	(0)
Watertown	1.0	(0)	(0)	(0.005)	(0.127)	(0.243)
Watertown	2.0	(0)	(0)	(0)	(0.029)	(0.273)
Watertown	3.0	(0)	(0)	(0)	(0)	(0.82)
Watertown	4.0	(0.782)	(0.049)	(0.415)	(0.572)	(0.149)

Notes: Table contains the p-values for difference in means between treatment and control groups within each quartile. The rows designate a pilot-quartile pair and the the columns are variables important to water consumption.

3.5 Estimation Strategy

3.5.1 Average Treatment Effects

The primary goal of this research is to estimate the average treatment effect and explore treatment effect heterogeneity. Because households were randomly selected into treatment, a simple t-test for difference in water use means after treatment between the treatment group and the control group provides a valid statistical tool for the analysis. However, water demand has a large variance and in order to identify relatively small effect sizes with the conventional significance levels we need to control for covariates.[16] A Hausman test that is robust to heteroskedastic and clustered standard errors as suggested by [28] provides evidence for the fixed effects model written as

$$\ln(w_{it}) = \alpha_i + \gamma T_i \times P_t + \theta P_t + \beta X_{it}' + \xi_{it} \tag{3.1}$$

---

[16] The conventional for a small effect size is 0.2 (standardized to the variance of the outcome variable) and in our setting we are attempting to identify effect sizes on the order of 0.05, which is only possible while controlling for covariates.

where w_{it} is water consumption for household i at time t, T_i is a dummy that identifies the treatment group, P_t is a dummy for treatment period, X_{it} is a set of time-varying covariates, and ξ_{it} is an idiosyncratic error term. In our regressions the covariates include the average number of rainy days and cooling degree-days; as well as year-period fixed effects. Huber-White standard errors [139] are clustered at the household level to account for serial correlation within a household over time. The regression results for each utility are presented in Table 3.6.

Table 3.6: Average Treatment Effects

	(1) Watertown	(2) Aquaville	(3) Hydroburg
Treatment Effect	-0.0534***	-0.0313***	-0.0170**
	(0.0146)	(0.0119)	(0.0073)
Rainy Days	24.1253	-0.3020***	-8.5027***
	(19.8758)	(0.0367)	(0.3942)
Cooling Degree Days	0.2633	0.0067***	0.0013
	(0.6321)	(0.0009)	(0.0047)
Constant	4.6138***	5.1315***	5.5719***
	(0.0972)	(0.0130)	(0.0046)
Household FEs	Yes	Yes	Yes
Year-Period FEs	Yes	Yes	Yes
Adjusted R^2	0.157	0.251	0.190
Households	2,575	3,093	11,696
Observations	44,763	83,342	664,075

Notes: The dependent variable is the natural logarithm of water consumption in gallons per day. Rainy days and cooling degree days are the average during the reading period, generated from daily weather data. Robust standard errors clustered at the household level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

We find that sending home water reports reduces average monthly water use in the billing periods after treatment by 5.34% in Watertown, 3.13% in Aquaville, and 1.70% in Hydroburg.[17] The results from Table 3.6 display the differences in treatment response across pilots, and we

---

[17] Relative effect of a dummy variable with estimated coefficient α in a log-linear regression is equal to $e^{\alpha} - 1$ (Halvorsen and Palmquist, 1980).

run a pooled regression with pilot-level interactions for all variables to test if these differences are statistically significant. An F-test rejects the null for a constant ATE across pilots at the 1% level. In grouped F-tests between each pair of pilots we reject the null for equal ATEs for Watertown and Hydroburg at the 5% level and for Hydroburg and Aquaville at the 10% level. In determining whether to run a pooled model with pilot-level interaction terms we must evaluate whether it is appropriate to pool the error term across all the pilots. One way to evaluate if unobserved heterogeneity is equal across pilots is to examine the observed heterogeneity. All the interaction terms for weather variables and many of the interactions for year-period fixed effects are statistically significant. The presence of significant observed heterogeneity signals a likelihood that unobserved heterogeneity such as utility policies and regional macro-level shocks exists that will cause the error variances to differ across pilots. The results suggest that it will be difficult to identify pilot-level effects in a pooled model, and therefore we restrict our econometric analysis to within-pilot heterogeneity.

3.5.2 Durability

Our first foray into the heterogeneity looks at the treatment effect over time. Ex ante there are several valid hypotheses explaining the dynamic pattern of consumer response to the treatment. One potential pattern of results is that initially the program increases the salience of water consumption and households make temporary behavioral adjustment to their use. In this setting, after an initial strong response, the treatment effect will wane over the course of the program. Another conceivable outcome is that the program induces capital investments or permanent changes in behaviors that cause consistent reductions throughout the program. The series of period-level treatment effects over time also provides insight into the mechanisms through which the HWRs cause consumers to reduce water use, including the role of seasonality and outdoor water demand. Both [54] and [8] examine the effect of social norm messaging after the treatment ceases. [8] distinguish ”durability”, the effect over time during the treatment period from ”persistence”, how the effect remains after treatment ends. They find that, in energy, the effect of social norms is both durable and persistent, though the persistence wanes with time from the end of treatment. In [54] the social norm is only sent in one summer so analysis of durability is not possible. They see a similar tapering in the persistence of the treatment effect over time as [8]. Our experimental design allows us to examine durability but not persistence, because the control group also begins receiving the treatment (the HWR) at the end of the experiment. In order to examine the durability of the treatment effect *within* the active administration of the program we run separate regressions for each treatment period. Thus the regression equation is now

$$\ln(w_{it}) = \alpha_i + \gamma T_i \times P_{t,h} + \theta P_{t,h} + \beta X'_{it} + \xi_{it} \tag{3.2}$$

where $h = 1, ..., H$ is a specific treatment period as opposed to the whole course of the program. In this setting we drop all observations for other treatment periods, P_{-h}, so that we compare a single treatment period to all controls as opposed to other treatment periods. The results from equation 3.2 for all treatment periods and for each pilot are shown graphically in Figure 3.4. The left vertical axis represents the coefficient and 95% confidence intervals and the horizontal axis is the series of treatment periods over time. To highlight the interaction with seasonality we include average pilot-level evapotranspiration, which measures the consumptive use of turf grass in millimeters, on the right vertical axis.

The graphs reveal variation between pilots in the pattern of savings over the course of the program. Note that Watertown and Aquaville have meters read every two months so each period corresponds to two calendar months, whereas in Hydroburg the time series is observed every month. The first home water report for each pilot is sent out in the middle of the reading period, after the previous period's data is processed in order to generate the metric on relative consumption. Therefore, a portion of consumption in the first period is actually untreated, which attenuates the treatment effect. This explains the smaller savings in the first period relative to most other periods observed in all pilots. As a robustness check we re-run the regressions from equation 3.1 dropping all observations in the first treatment period for each pilot. These ATEs, as seen in Table 3.7, are very similar to the base regressions in Table 3.6 but are slightly larger in absolute value(-5.46%, -3.31%, and -2.02% for Watertown, Aquaville, and Hydroburg, respectively).

The treatment effect for both Watertown and Aquaville is relatively stable after the first period (Figure 3.4), suggesting that consumers either made early investments in water efficiency or

Figure 3.4: Durability of Treatment Effects - Percentage

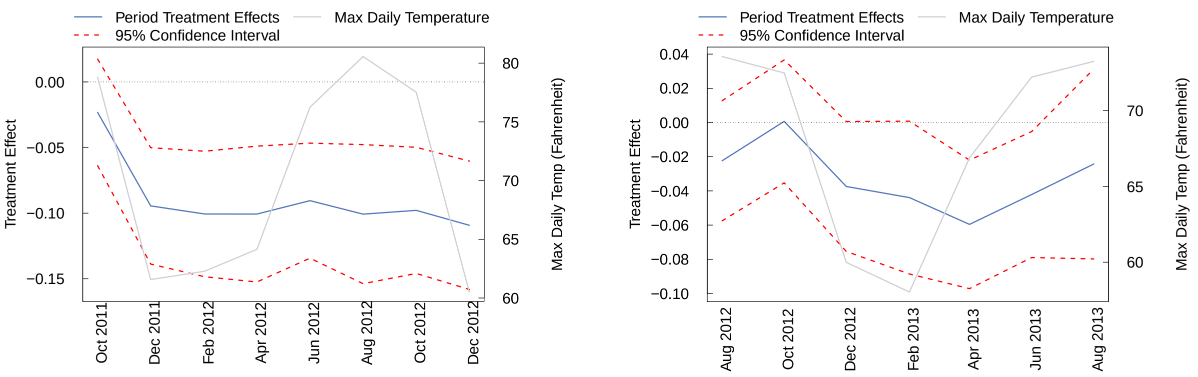

Watertown

Aquaville

Hydroburg

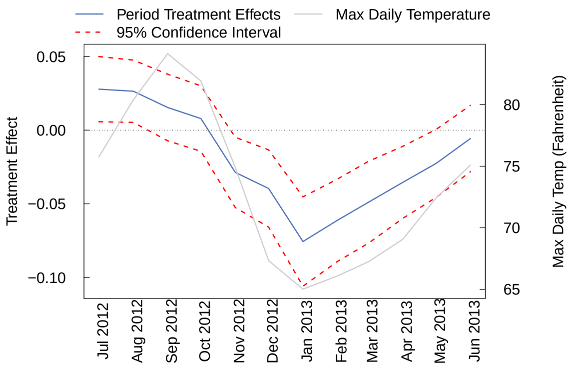

Notes: The solid line is set of point estimates for a regressions on an individual treatment period. All data from other treatment periods are omitted and all regressions contain household and year-period fixed effects as well as weather controls. The dashed line represents the 95% confidence interval constructed from robust standard errors clustered at the household level. The grey line, on the right vertical axis, is the average maximum daily temperature in degrees Celsius for the pilot to highlight the correlation of individual-period treatment effects and weather.

Table 3.7: Average Treatment Effects: Exclude first period

	(1) Watertown	(2) Aquaville	(3) Hydroburg
Treatment Effect	-0.0561***	-0.0337***	-0.0204***
	(0.0148)	(0.0118)	(0.0076)
Rainy Days	23.8145	-0.3048***	-7.6478***
	(19.8971)	(0.0370)	(0.3997)
Cooling Degree Days	0.2959	0.0069***	0.0121**
	(0.6425)	(0.0009)	(0.0047)
Constant	4.8397***	5.1302***	5.5307***
	(0.0804)	(0.0129)	(0.0047)
Household FEs	Yes	Yes	Yes
Year-Period FEs	Yes	Yes	Yes
Adjusted R^2	0.164	0.253	0.183
Households	2,504	3,092	10,490
Observations	42,394	80,845	589,030

Notes: The dependent variable is the natural logarithm of water consumption in gallons per day. All observations from the first treatment period for each pilot are omitted. Rainy days and cooling degree days are the average during the reading period, generated from daily weather data. Robust standard errors clustered at the household level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

established permanent behavioral changes. The treatment effect in Hydroburg took several months to appear (perhaps because there was a lag between initially receiving the HWR and undertaking conservation efforts), and then wanes over the treatment period. The treatment effect is essentially zero in Hydroburg 11 or 12 months after the HWRs were first sent. This lack of durability in the treatment in this pilot should not be confused with persistence; the HWRs continued to be sent to households in Hydroburg through June 2013.

Learning whether the program is effective at reducing outdoor water use is critical. This is the peak demand period and often corresponds with supply shocks. If the main saving mechanism was in reducing outdoor water use we would expect a negative correlation between temperature and the treatment effect because larger reductions in water use would correspond with higher temperatures.

Examining the pattern of individual treatment effects in conjunction with the weather data indicates that the social norm does not generate savings primarily through outdoor use (Figure 3.4). This result is not surprising given that Watertown has the most moderate climate and relatively less outdoor use than the other pilots. However, even in Hydroburg, the hottest and driest pilot, the correlation is positive with peak savings in January- exactly the opposite result that would support savings in outdoor water use. Similarly, the U-shaped curve for Aquaville suggests that more conservation occurs in the winter months with peak savings occurring in March and April.

Though we control for seasonality with year-period fixed effects and weather variables, there may be concern that the logarithmic transformation masks some of the aggregate savings in gallons. Since water use is higher in the summer the same amount of gallons saved during the winter is larger relative to average in percentage terms. In order to test the sensitivity to our specification of the dependent variable we re-run the regressions in equation 3.2 with water use in gallons per day. Investments made in indoor water efficiency during the winter should still be visible in later summer months. The results are graphed in Figure 3.5. In Aquaville we observe a steadily increasing treatment effect in absolute value (except for a slight decrease in July 2013), supporting the notion that households gradually invest in water efficiency after receiving the HWR. However, regressions on water as opposed to log water do not change the pattern in Hydroburg, suggesting a different experience. The return to a zero treatment effect after several months suggests transient behavioral savings as opposed to investments in efficient consumer durables. Hydroburg's existing rate structure already reflects relative water use and is very punitive for high relative water users. Consumers in this pilot may have already made investments in water efficiency that were not yet undertaken in the pilots with less aggressive rate structures.

Figure 3.5: Durability of Treatment Effects - Gallons per day

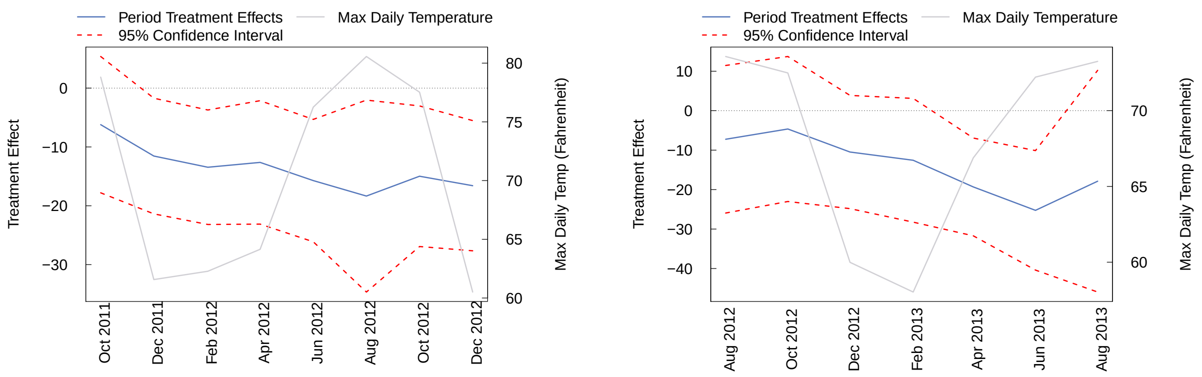

Watertown **Aquaville**

Hydroburg

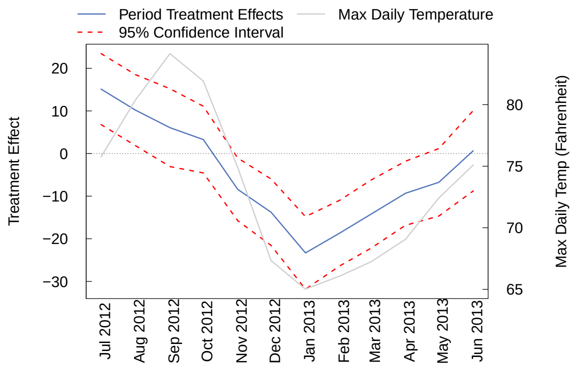

Notes: The solid line is set of point estimates for a regressions on an individual treatment period. All data from other treatment periods are omitted and all regressions contain household and year-period fixed effects as well as weather controls. The dashed line represents the 95% confidence interval constructed from robust standard errors clustered at the household level. The grey line, on the right vertical axis, is the average maximum daily temperature in degrees Celsius for the pilot to highlight the correlation of individual-period treatment effects and weather.

3.5.3 Treatment Effect Heterogeneity

There are many ways to explore the treatment effect heterogeneity for interventions using social norms to promote sustainable behavior. [5] use quantile regression and interactions with baseline electricity consumption quantile dummies, [53] explore the baseline consumption and differences due to home prices, and [36] utilize voting data to examine differences in treatment effect due to ideology. We build on the prior research by combining these approaches in a setting where we observe considerable variation in both ideology and treatment response.

Quantile Treatment Effects

We estimate unconditional quantile regressions in the presence of covariates through the methodology of [56]. Estimating conditional quantile regressions in the vein of [85] is less appealing in a policy framework when the unconditional distribution is of primary interest. Interpreting a coefficient from a quantile regression when conditioning on many variables is difficult and the results can vary substantially depending on the data generating process [22]. Since our data measure water use over time, the identification of the treatment effect requires controls for time varying covariates through weather and time fixed effects.[18] The widely employed estimator from [55] only introduces covariates to remove selection bias in estimating quantile treatment effects, a problem that does not exist in our case of random assignment into treatment. Therefore we use the estimator of [56] in order to identify treatment effects by controlling for time varying covariates while still providing unconditional quantile treatment effects, the quantiles of the treatment effect over the distribution of all households under current assumptions.

The basic logic behind the estimator draws from the influence function (IF) describing the influence of an individual observation on a given distributional statistic. A convenient feature is that the recentered influence function (RIF) is simply the IF plus the distributional statistic; the expectation of the RIF is the distributional statistic itself. In the case of quantiles the RIF for the τth quantile is $q_\tau + (\tau - \mathbb{1}\{Y \leq \tau\}/f_Y(q_\tau)$. The RIF regression function with covariates is $E[RIF, Y, q_\tau|X] = m_\tau(X)$, which [56] prove represents an unconditional quantile regression. Given

---

[18] If our samples sizes were larger we should be able to identify the treatment effects with a simple difference in means, or difference in quantiles, but we require covariates to identify statistically significant effects.

the assumption on the validity of the control group as a counterfactual the coefficient representing the treatment effect can be interpreted as the QTE. With the rank preservation assumption, the coefficient can be interpreted as the quantile of the treatment effect.

Even without the rank preservation assumption, learning about the difference in marginal distributions of water consumption between treated and untreated groups may be relevant to policymakers. One important clarification is that the estimator does not account for the panel structure of the data and contains no household fixed effects. Exploring quantile treatment effect models that incorporate the panel data structure and household fixed effects is left for future research and will follow [112].

We estimate the QTE using the RIF approach at each quantile of water use from 0.05 to 0.95 in 0.05 intervals with standard errors that are robust to heteroskedasticity. Figure 3.6 shows the results for the quantile treatment effects. The thick blue line represents the coefficient estimate for the treatment effect at each quantile of water use. The thick red dashed lines are the 95% confidence interval. The corresponding OLS estimates and 95% confidence interval are also shown on the graph by the horizontal green dashed and dotted lines respectively. The QTE is relatively constant across quantiles for Hydroburg. Aquaville and Watertown show a similar pattern to [5] and [53] in that the QTE increases at the highest quantiles. Watertown is unique in that there is a significantly non-zero QTE at low quantiles as well. This is a different pattern than seen in the literature on the effect of social norms and argues for the effectiveness of social norms at the low end of the consumption distribution. Another key finding is that at no point in the distribution is the QTE above zero, which is important for policymakers cautious of provoking the perverse effect of increased consumption in response to messages of relatively low water use.

In Figure 3.6 Watertown does experience savings at the lower end of the consumption distribution with a significant QTE at $tau = 0.1$, however the results suggest that targeting large water users will generate larger savings per household. Based on the empirical results WaterSmart is concentrating an expansion in one pilot exclusively for high water users.[19]

---

[19] The targeted expansion is occurring in the population that does not have email addresses registered with the utility company, requiring a more costly printed version of the HWR to be sent by conventional mail as opposed to email.

Figure 3.6: Quantile Treatment Effects

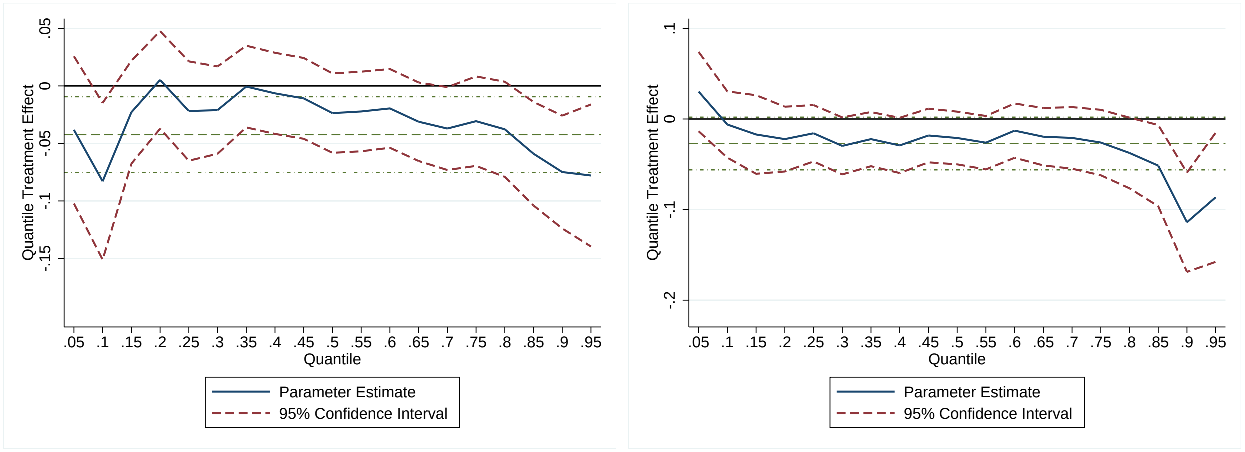
Watertown

Aquaville

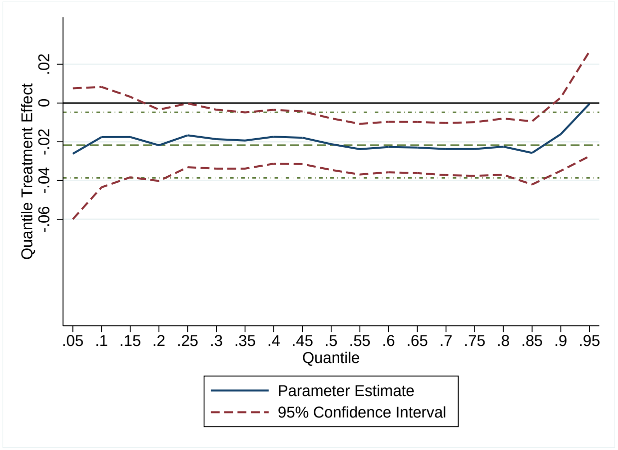
Hydroburg

Notes: The thick blue line is the estimate for unconditional quantile treatment effect by quantile in 0.05 intervals, with red thick dashed line representing 95% confidence intervals based on heteroskedastic-robust standard errors. The OLS regression and 95% confidence intervals are the flat horizontal dashed and dotted lines respectively. Pooled OLS is used for comparison, as opposed to the FE model, because the quantile regressions do not account for the panel structure of the data.

Conditional Average Treatment Effects

For a more intuitive regression model in addition to the unconditional quantile regressions above, we interact the treatment effect with quintiles of baseline water consumption for each utility. These regressions also attempt to parse out the effect of treatment across the distribution of water use, but focus on conditional average treatment effects (CATE) for each quintile as opposed to QTE. In this case the treatment effect of higher quantiles are only relevant for household types that populate that end of the distribution. Baseline deciles are calculated by taking the average pre-treatment household usage across the whole utility and ordering it into equal groups of five. We perform a similar technique to create quintiles of the GII and housing values for each utility. Note that higher quintiles of baseline consumption relate to higher baseline usage and higher GII quintiles correspond to stronger environmental preferences. The regressions include all the variables as presented in equation 3.1 in addition to interactions between the treatment effect and the quintiles of water and GII. The regressions have the form

$$\ln(w_{it}) = \alpha_i + \gamma T_i \times P_t + \sum_{d=1}^{5} \delta_d (T_i \times P_t \times D_d) + \theta P_t + \beta X'_{it} + \xi_{it} \tag{3.3}$$

where D_d is the d^{th} quintile of either baseline consumption, the GII, or housing values. The parameters of interest are the linear combination of γ and δ_d.

Table 3.8 shows the effect of treatment for households in different quintiles of baseline water use. The estimates in Table 3.8 are for $\gamma + \delta_d$, the linear combination of the treatment effect and the interaction of the treatment effect with each quintile of baseline consumption.[20] The significance levels are determined by p-values from test for joint significance using F-statistics. The full regression output is provided in Table C.18 in the appendix. In Watertown the highest three quintiles are all negative and statistically significant, but the lower two quintiles are not significant. This is an indication that the majority of savings comes from households with above average water consumption. Similarly, Hydroburg has negative and significant coefficients on the top three quintiles, but also has a positive coefficient on the second quintile. This suggests that

---

[20] The third quintile, containing the median, is the omitted category and in this case the raw coefficient and stand error are shown.

some households may actually increase consumption due to the treatment. For Aquaville the only statistically significant quintile is in the middle of the distribution. The lowest quintile in Aquaville is positive suggesting a similar relationship seen in the other pilots that the program is less effective in households that start out using less water. This raises the question of whether the social norms are effective for households that already are consuming below the average prior to the intervention. This results are slightly different from those conclusions drawn from the unconditional quantile regressions, with starker differences between the two methodological results for Aquaville and Hydroburg.

Table 3.8: Heterogeneity: Baseline Water Use

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Watertown	-0.00502	-0.0104	-0.0537*	-0.0853***	-0.110***
	(0.0289)	(0.0233)	(0.0237)	(0.0211)	(0.0195)
Aquaville	0.0213	-0.0345	-0.0732***	-0.0361	-0.0366
	(0.0225)	(0.0219)	(0.0195)	(0.0191)	(0.0210)
Hydroburg	0.0208	0.0439**	-0.0509***	-0.0449***	-0.0415**
	(0.0192)	(0.0158)	(0.0138)	(0.0132)	(0.0149)

Notes: The coefficients reported here are the linear combinations of the coefficients from the regression results in Table C.18. Robust standard errors clustered at the household level for these linear combinations are reported in parentheses. Stars denote significance from F-tests evaluating the difference of the linear combination from zero. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Unlike the impact of baseline consumption, the effect of ideology and housing values on treatment effect heterogeneity is not consistent across the three pilots. Tables 3.9 and 3.10 show the linear combinations of of treatment effect and interaction terms for regressions based on quintiles of the GII and housing values. Watertown displays relatively consistent CATE across all quintiles of the GII and housing values, though there are insignificant effects at the second quintile for both variables. Aquaville and Hydroburg show opposite effects for both variables, though the interactions are generally insignificant and it is difficult to discern a trend. For ideology there is a positive, but insignificant effect at the second quintile in Aquaville and the fifth quintile in Hydroburg. Interpreting the results with caution due to the noise in the estimate, it appears that more environmentally-conscious households respond more to social norms in Aquaville while in Hydroburg the least eco-friendly households are more responsive. One caveat is that the ideology data are at the census block level as opposed to the household level as in [36].

In terms of housing values, which serve as a proxy for income, Aquaville has negative and significant effects at the lowest two quintiles whereas Hydroburg has a significant effect at the fourth quintile. The results from CATEs by housing value Aquaville are the starkest, and also contradict the finding of [53] that households in expensive homes are more responsive. Overall it is difficult to make any conclusions about how environmental ideology and housing values drive heterogeneity in the response to treatment. Most of the interaction terms as seen in Tables C.19 and C.20 in the Appendix are insignificant, and the trends run in opposite directions for Aquaville and Hydroburg. This suggests that the biggest driver of heterogeneity is baseline water consumption. Another issue is that we are only exploiting intra-pilot variation in ideology and housing values. There is more variation across pilots for these two variables, and that may be part what is determining the difference in magnitudes of the pilot-level ATEs. In a purely descriptive analysis however, the ATE of social norm messaging in Hydroburg the most conservative and affluent area in our study, is lower than that of the other two pilots. Watertown, the least affluent and most environmentally friendly utility had the strongest ATE. We must note here that Hydroburg's long standing rate structures aggressively tied to relative water use confound even these descriptive conclusions.

Table 3.9: Heterogeneity: Ideology

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Watertown	-0.0682**	-0.0271	-0.0533	-0.0634*	-0.0930**
	(0.0252)	(0.0262)	(0.0281)	(0.0317)	(0.0311)
Aquaville	-0.0152	0.0196	-0.0458*	-0.0258	-0.0516
	(0.0211)	(0.0215)	(0.0231)	(0.0225)	(0.0291)
Hydroburg	-0.0234	-0.0679***	-0.0148	-0.0184	0.00879
	(0.0167)	(0.0163)	(0.0156)	(0.0171)	(0.0210)

Notes: The coefficients reported here are the linear combinations of the coefficients from the regression results in Table C.18. Robust standard errors clustered at the household level for these linear combinations are reported in parentheses. Stars denote significance from F-tests evaluating the difference of the linear combination from zero. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3.10: Heterogeneity: Housing Values

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Watertown	-0.0662*	-0.0171	-0.0593**	-0.0621**	-0.0603**
	(0.0272)	(0.0268)	(0.0216)	(0.0232)	(0.0219)
Aquaville	-0.0704**	-0.0701**	-0.0372	-0.0108	0.0196
	(0.0232)	(0.0252)	(0.0216)	(0.0185)	(0.0170)
Hydroburg	-0.0176	-0.0276	-0.0251	-0.0295*	-0.00622
	(0.0173)	(0.0180)	(0.0155)	(0.0138)	(0.0157)

Notes: The coefficients reported here are the linear combinations of the coefficients from the regression results in Table C.18. Robust standard errors clustered at the household level for these linear combinations are reported in parentheses. Stars denote significance from F-tests evaluating the difference of the linear combination from zero. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

3.5.4 Interactions with Existing Conservation Programs

Since the HWRs provide information about utilities' water conservation programs and generally increase the awareness of water consumption, we may expect those in the treatment group to participate in those programs at a higher rate. On the other hand, the deluge of information about water conservation may inundate consumers and social norms may crowd out existing conservation efforts. Using data on participation in conservation programs that spanned both the pre and post-intervention periods, we estimate the impact of the social norm message on the probability of engaging in an additional utility program.

In the three pilots only Watertown, the longest running pilot, had enough post-treatment participation data to estimate a statistical model. We hope to get sufficient data from the other pilots in the future.[21] We estimate a logit model where the dependent variable is a dummy equal to one if a household participates in a program in a given period and zero otherwise. We run several specifications of the dependent variable. The first specification is to pool all the different programs together to create a single indicator for participation, defined as all programs. We also create separate indicators for the completion of a home visit by an expert to provide personalized advice on

---

[21] Aquaville and Hydroburg had many observations that were missing a program participation date. We will likely be able to run these models as the utilities enter recent participants into their database.

water conservation, the receipt of a high-efficiency toilet rebate, or the receipt of a high-efficiency clothes washer rebate.[22] Since our goal is to assess the impact of the social norm campaign on participation and we only observe 384 instances of program participation during our sample, we simplify the temporal dimension to before or after treatment. Our regression then takes the form of a difference-in-difference logit model where we control for any time and group level effects as well as static household characteristics that may influence participation.

$$\Pr(c^l_{i,t} = 1 | X_{i,t}) = \frac{\exp(\beta X_{i,t})}{1 + \exp(\beta X_{i,t})} \tag{3.4}$$

where $X_{it} = T_i + P_t + T_i \times P_t + Z_i$. The treatment group and treatment period indicator are defined as T_i and P_t respectively, i designates the household, time, $t = \{1, 2\}$, is pooled into pre and post intervention periods, l specifies the type of conservation programs and Z_i is a vector of household characteristics.

The results from equation 3.4 are presented in Table 3.11 as marginal effects, with standard errors calculated by the delta method. The treatment effect defined as $T_i \times P_t$ increases the probability of participation in conservation programs in the pooled specification as well as the completion of a home water audit, but is statistically insignificant for each of the individual rebates. These results provide evidence that social norms are a complement to existing programs and do not crowd out other conservation efforts.

In order to connect increased participation rates to water use we estimate the impact of the conservation programs on water demand. The regressions take the following form

$$\ln(w_{it}) = \alpha_i + \gamma T_i \times P_t + \pi_1 C^l_{i,t} + \pi_2 \tilde{C}^l_{i,t} + \theta P_t + \beta X'_{it} + \xi_{it} \tag{3.5}$$

where $C^l_{i,t}$ is the cumulative sum of programs that household i has participated in at time t, for l defined by the categories: all programs, home visit, toilet rebate, and clothes washer rebate. $\tilde{C}^l_{i,t}$ is the sum of programs for those in the treatment group that were initiated after the start of the social norm campaign. Since households can participate in multiple programs of the same type $C^l_{i,t}$ and

---

[22] There are two programs, greywater systems and landscape conversion rebates, that are included in "all programs" category but not estimated separately due to the small number of observations.

Table 3.11: Logit Regressions for Conservation Programs - Watertown

	Pooled		Rebates	
	(1) All Programs	(2) Home Survey	(3) Toilet	(4) Clothes Washer
Treatment Effect	0.1605***	0.1178***	0.0102	0.0001
	(0.0359)	(0.0363)	(0.0140)	(0.0085)
Treatment Group	0.0093	-0.0005	0.0019	0.0004
	(0.0080)	(0.0058)	(0.0016)	(0.0029)
Treatment Period	-0.1381***	-0.0308***	-0.0484***	-0.0339***
	(0.0161)	(0.0090)	(0.0146)	(0.0079)
Single Family Home	0.0102	0.0021	0.0032	0.0002
	(0.0094)	(0.0053)	(0.0020)	(0.0040)
Bathrooms	0.0245***	0.0065**	0.0042***	0.0045**
	(0.0048)	(0.0028)	(0.0016)	(0.0019)
Home Size (Sq. Ft.)	-0.0017	0.0023	-0.0010	0.0013
	(0.0098)	(0.0055)	(0.0020)	(0.0040)
Year Built	-0.0001	-0.0000	-0.0001**	-0.0000
	(0.0002)	(0.0001)	(0.0000)	(0.0001)
Lot Size (Sq. Ft.)	-0.0091**	-0.0112***	-0.0025**	-0.0007
	(0.0037)	(0.0029)	(0.0011)	(0.0008)
Home Value	0.0075	0.0024	0.0005	0.0039**
	(0.0048)	(0.0029)	(0.0011)	(0.0018)
Pseudo R2	0.0849	0.0859	0.156	0.106
Observations	4,015	4,015	4,015	4,015

Notes: The dependent variable is a dummy for participation in a given water utility conservation program. The data are pooled temporally by pre/post WaterSmart intervention. The marginal effects are reported along with standard errors obtained from the delta method.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

$\tilde{C}^l_{i,t}$ not only capture the presence but also the intensity of program participation. The term $\tilde{C}^l_{i,t}$, essentially an interaction term, represents the marginal effect of the normative message on water savings beyond the baseline impact of conservation programs and identifies whether information campaigns improve the efficiency of existing campaigns. [17] argues that conservation programs like rebates do not pass the test of additionality because those that take advantage of rebates may have undertaken the conservation activity regardless of the presence of the rebate. If a social norm campaign prompts households to engage in a toilet rebate who otherwise would not have replaced their toilet, then the conservation programs should be more effective in the population treated by the social norm. Due to random assignment into treatment we do not expect differences in natural appliance depreciation rates across groups. Therefore, if the social norm spurs additional program participation it is likely that these are households that otherwise would not have replaced inefficient appliances. In other words, a positive coefficient on $\tilde{C}^l_{i,t}$ means that households that receive a social norm utilize conservation programs more effectively than the untreated population. The results from the regressions in equation 3.5 for each of the four specifications of water conservation programs are presented in Table 3.12.[23]

In the pooled specification, participating in a conservation program reduces water use by almost 7%. Home surveys and toilet rebates each lead to a drop in consumption of over 10%. The clothes washer program is not statistically significant on its own, but it is jointly significant with the treatment interaction term at the 5% level. None of the interaction terms are statistically significant, but there is still valuable information in the sign and size of the coefficients.[24] Home surveys are essentially an additional information treatment by providing households with tailored advice to reduce consumption. The coefficient on the home survey interaction is positive but insignificant. It is difficult to posit why an interaction of two information treatments would produce savings greater than the sum of the parts. The sign of coefficients on toilet and clothes washer rebates is congruent with the hypothesis that social norms attract households to appliance rebates that otherwise would not have made those investments. However, the noise in the estimates require caution when making the argument that social norms alleviate concerns over additionality in rebate programs. Examining the results in Tables 3.11 and 3.12 together also demonstrates one of the methods through which

---

[23] While there is not a sufficient sample size post-treatment to report the regression results in this paper, the conservation programs have a similar effect on water demand in Aquaville and Hydroburg. Tables presenting those results are available upon request.

[24] The lack of significance on the individual treatment programs is likely due to the sparseness of post-treatment participation data for each of the individual programs, whereas the coefficient on all programs is combining positive and negative effects.

Table 3.12: Interactions with Conservation Programs - Watertown

	Pooled		Rebates	
	(1) All Programs	(2) Home Survey	(3) Toilet	(4) Clothes washer
Treatment Effect	-0.0501***	-0.0500***	-0.0530***	-0.0534***
	(0.0147)	(0.0147)	(0.0146)	(0.0147)
All Programs	-0.0674***			
	(0.0192)			
Treat*All	0.0128			
	(0.0294)			
Home Survey		-0.1274***		
		(0.0453)		
Treat*Survey		0.0656		
		(0.0505)		
Toilet Rebate			-0.1146***	
			(0.0382)	
Treat*Toilet			-0.0841	
			(0.0768)	
Clothes Washer Rebate				-0.0534
				(0.0351)
Treat*Clothes Washer				-0.0524
				(0.0639)
Household FEs	Yes	Yes	Yes	Yes
Year-Period FEs	Yes	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes	Yes
Adjusted R^2	0.158	0.158	0.158	0.157
Households	2,575	2,575	2,575	2,575
Observations	44,763	44,763	44,763	44,763

Notes: The dependent variable is the natural logarithm of water consumption in gallons per day. Robust standard errors clustered at the household level are reported in parentheses. Each regression equation includes a dummy variable if a household has participated in the conservation program by time t, as well as an interaction of the conservation variable with the treatment effect. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

social norm programs cause reductions in water consumption - by increasing participation rates in conservation programs. A back of the envelope calculation based on the coefficients in column

(1) of Tables 3.11 and 3.12 suggests that approximately 20% of the ATE is due to participation in additional utility conservation programs. This is a strain of research that is important to pursue as data arrive from the other pilots. The interactions with existing programs may contribute to the overall ATE in a utility and could help describe the variation in ATE across pilots.

3.6 Future Work

This research is a portion of the broader project on the role of information in municipal water campaigns. WaterSmart is working with several other utilities that provide new opportunities for research. Several of the utilities do not want to implement randomized field experiments in the hopes of maximizing conservation gains immediately by targeting specific problem populations or by providing universal access to the normative message. Analyzing the data for utilities targeting subsets of their customer base re-introduces selection concerns and the identification of suitable control observations. This will be left for future research. An ongoing partnership with WaterSmart also provides opportunities to set up randomized experiments testing alternative messages across subgroups of those new pilots providing universal access. This has the same attractive features of WaterSmart's first three programs analyzed above by avoiding concerns of endogenous sample selection. The results of this paper suggest that there may be conservation gains available by targeting consumers to treatment based on their observables. Conducting a natural field experiment will test if additional conservation gains can be achieved by targeting the message based on observables. There may be important interactions with water rates and normative messaging. These are difficult to compare across utilities because there are many unobserved shocks that vary across utility. Randomizing subgroups within a utility to exploit interactions between the social norms and water rates is another promising area of research. We hope to explore several of these avenues in future research.

3.7 Conclusion

Behavioral economic is a particularly attractive policy instrument to address environmental externalities where using Pigouvian taxes is politically intractable. Municipal water demand is a domain that must address issues of scarcity associated with rising populations, the exhaustion of cheap supply sources, and stronger demand for water devoted for ecosystem services. Climate change is expected to increase the variability of supply and increase the probability of extreme droughts. Zero profit constraints, concerns over equity, and the notion of access to water as a basic human right, limit the ability for public utilities to manage demand through the price mechanism. Utilities therefore seek a plethora of non-pecuniary demand side management practices. Among the suite of options, information campaigns using social norms are gaining momentum as an effective conservation instrument.

We analyze the data from three randomized field experiments that harness social norms as a water conservation tool. There is significant heterogeneity in the treatment effects both within and across pilots. While there are many dimensions that may impact response to treatment, we focus on the distribution of water use, environmental ideology, and housing values as observable drivers of heterogeneity. At an aggregate level Watertown, a small utility with a moderate climate and low baseline water use experiences the largest ATE of 5.34%. This is statistically different than the ATE for Hydroburg, with the smallest ATE of 1.7%. Interacting the treatment with quintiles of the baseline consumption corroborates the results of other social norm campaigns that treatment effects are largest for households with high baseline consumption.

Several results augment and address gaps in the prior literature. Contrary to [53] who find a higher CATE for houses with assessed values above the median; the only significant heterogeneity in terms of housing values that we observe shows that higher housing values is correlated with lower savings. Across utilities there is monotonically decreasing relationship between average income and the ATE. Though it is difficult to draw inference from the three aggregate results in our data, this does provide evidence that social norms are effective with low-income households. We also observe two effects of ideology that work in opposite directions. In Aquaville stronger environmental preferences increase the CATE while in Hydroburg they correspond with smaller CATEs. In a series of unconditional quantile regressions we find that Watertown exhibits a QTE larger than the median at both the low and high end of the consumption distribution. Prior research [5, 53] finds a monotonic relationship between consumption quantiles and the QTE. This opens the possibility that social norms can be effective at reducing consumption for households that are already using water efficiently. These households may be more conscious of the financial and/or environmental costs associated with water consumption.

In addressing the mechanisms of how WaterSmart's HWR translate into water savings we look at the time pattern of individual-period treatment effects. We find little correlation of the treatment effect with weather variables that increase outdoor water use, suggesting that the HWRs do not have a large impact on outdoor use. Rather, the durability of savings over time is consistent with indoor water efficiency investments or systematic behavioral changes. Additionally, the time pattern of water savings varies across utilities, which indicates that demographics, climate, and utility policies affect the actions prompted by the normative message. One mechanism that we do identify is increased participation rates in additional utility conservation programs. This contributes up to 20% of the ATE in Watertown and demonstrates that social norms act as a complement rather than a substitute to existing conservation programs. To address water scarcity in the presence of regulatory and political pricing constraints utilities are seeking new policy instruments to reduce demand. Due to the heterogeneity in treatment response, both within and across pilots, utilities must address underlying interactions in order to maximize the effectiveness of social norms as a water conservation instrument.

BIBLIOGRAPHY

[1] Rimjhim M. Aggarwal, Subhrajit Guhathakurta, Susanne Grossman-Clarke, and Vasudha Lathey. How do variations in Urban Heat Islands in space and time influence household water use? The case of Phoenix, Arizona. *Water Resources Research*, 48(6):1–13, June 2012.

[2] Donald E. Agthe and R. Bruce Billings. Dynamic models of residential water demand. *Water Resources Research*, 16(3):4764880, 1980.

[3] Donald E. Agthe, R. Bruce Billings, John L. Dobra, and Kambiz Raffiee. A Simultaneous Equation Demand Model for Block Rates. *Water Resources Research*, 22(1):1, 1986.

[4] Hunt Allcott. Consumers' Perceptions and Misperceptions of Energy Costs. *American Economic Review*, 101(3):98–104, 2011.

[5] Hunt Allcott. Social norms and energy conservation. *Journal of Public Economics*, 95(9-10):1082–1095, 2011.

[6] Hunt Allcott and Michael Greenstone. American Economic Association Is There an Energy Efficiency Gap? *American Economic Review*, 26(1):3–28, 2012.

[7] Hunt Allcott and Sendhil Mullainathan. Behavioral Science and Energy Policy. *Science*, 327(6970):1204–1205, 2010.

[8] Hunt Allcott and Todd Rogers. How Long Do Treatment Effects Last? Persistence and Durability of a Descriptive Norms Intervention's Effect on Energy Conservation. 2012.

[9] Hunt Allcott and Todd Rogers. The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation. 2012.

[10] RG Allen, ME Jensen, JL Wright, and RD Burman. Operational estimates of reference evapotranspiration. *Agronomy Journal*, 1989.

[11] John M. Anderies. Robustness , institutions , and large - scale change in social - ecological systems : the Hohokam of the Phoenix Basin. *Journal of Institutional Economics*, 2(02):133, July 2006.

[12] R. C. Balling, P. Gober, and N. Jones. Sensitivity of residential water consumption to variations in climate: An intraurban analysis of Phoenix, Arizona. *Water Resources Research*, 44(10):1–11, October 2008.

[13] Javier A. Barrios García and José E. Rodríguez Hernández. Housing demand in Spain according to dwelling type: Microeconometric evidence. *Regional Science and Urban Economics*, 38(4):363–377, July 2008.

[14] Bryson Bates, Zbigniew W. Kundzewicz, Shaohong Wu, and Jean Palutikof. Climate Change and Water. Technical report, Technical Paper of the Intergovernmantal Panel on Climate Change, IPCC Secretariat, Geneva, 2008.

[15] David R Bell and Ronald C Griffin. Urban Water Demand with Periodic Error Correction. *Land Economics*, 87(3):528–544, 2011.

[16] Shlomo Benartzi and Richard H Thaler. Heuristics and Biases in Retirement Savings Behavior. *Journal of Economic Perpectives*, 21(3):81–104, 2007.

[17] Lori S Bennear, Jonathan M Lee, and Laura O Taylor. Municipal Rebate Programs for Environmental Retrofits : An Evaluation of Additionality and Cost-Effectiveness. *Journal of Policy Analysis and Management*, 32(2):350–372, 2013.

[18] R Bruce Billings and Donald E Agthe. Price Elasticities for Water: A Case of Increasing Block Rates. *Land Economics*, 56(1):73–84, 1980.

[19] Erik Biørn, Rolf Golombek, and Arvid Raknerud. Environmental Regulations and Plant Exit. *Environmental and Resource Economics*, 11:35–59, 1998.

[20] Marianne P Bitler, Jonah B Gelbach, and Hilary W Hoynes. What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments. *American Economic Review*, 96(4):988–1012, September 2006.

[21] Marianne P. Bitler, Jonah B. Gelbach, and Hilary W. Hoynes. Distributional impacts of the Self-Sufficiency Project. *Journal of Public Economics*, 92(3-4):748–765, April 2008.

[22] Bijan J Borah and Anirban Basu. Highlighting Differences Between Conditional and Unconditional Quantile Regression Approaches Through an Applicatino to Assess Medication Adherence. *Health economics*, 22:1052–1070, 2013.

[23] Stuart H. Burness and James P. Quirk. Appropriative Water Rights and the Efficient Allocation of Resources. *American Economic Review*, 69(1):25–37, 1979.

[24] Oscar R. Burt. Optimal Resource Use Over Time with an Application to Ground Water. *Management Science*, 11(1):80–93, September 1964.

[25] Van Butsic, David J Lewis, and Lindsay Ludwig. An Econometric Analysis of Land Development with Endogenous Zoning. *Land Economics*, 87(3):412–432, 2011.

[26] Van Butsic and Noelwah R. Netusil. Valuing Water Rights in Douglas County, Oregon, Using the Hedonic Price Method. *Journal of the American Water Resources Association*, 43(3):622–629, June 2007.

[27] Ricardo J. Caballero and Eduardo M. R. A. Engel. Beyond the Partial-Adjustment. *American Economic Review*, 82(2):360–364, 1992.

[28] A. Colin Cameron and Pravin K. Trivedi. *Microeconometrics: Methods and Applications*. Cambridge University Press, 2005.

[29] Philip H. Carver and John J. Boland. Short- and long-run effects of price on municipal water use. *Water Resources Research*, 16(4):609, 1980.

[30] By Anita Castledine, Klaus Moeltner, and Michael K Price. Free to choose : Promoting Conservation by Relaxing Outdoor Watering Restrictions. In *Association for Public Policy Analysis & Management*, page 44, Washington, DC, 2011.

[31] Gary Chamberlain. Analysis of Covariance Data with Qualitative. *The Review of Economic Studies*, 47(1):225–238, 1980.

[32] Gary Chamberlain. Multivariate Regression Models For Panel Data. *Journal of Econometrics*, 18:5–42, 1982.

[33] Pat S Chavez. Image-Based Atmospheric Corrections - Revisited and Improved. 62(9):1025–1036, 1996.

[34] S. V. Ciriacy-Wantrup. Concepts Used as Economic Criteria for a System of Water Rights. *Land Economics*, 32(4):295–312, 1956.

[35] Katharine Coman. Some unsettled problems of irrigation. *The American Economic Review*, 1(1):1–19, 1911.

[36] Dora L. Costa and Matthew E. Kahn. Energy Conservation “Nudges” and Environmentalist Ideology: Evidence From a Randomized Residential Electricity Field Experiment. *Journal of the European Economic Association*, 11(3):680–702, June 2013.

[37] Jan P Crouter. Hedonic Estimation Applied to a Water Rights Market. *Land Economics*, 63(3):259–271, 1987.

[38] Gordon B Dahl. Mobility and the Return to Education: Testing a Roy Model with Multiple Markets. *Econometrica*, 70(6):2367–2420, 2002.

[39] Jasper M. Dalhuisen, Raymond J. G. M. Florax, Henri L. F. de Groot, and Peter Nijkamp. Price and Income Elasticities of Residential Water Demand: A Meta-Analysis. *Land Economics*, 79(2):292, May 2003.

[40] Graeme Dandy, Tin Nguyen, and Carolyn Davies. Estimating Residential Water Demand in the Presence of Free Allowances. *Land Economics*, 73(1):125–139, 1997.

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

[42] Xinshen Diao, Ariel Dinar, Terry Roe, and Yacov Tsur. A general equilibrium analysis of conjunctive ground and surface water use with an application to Morocco. *Agricultural Economics*, 38(2):117–135, March 2008.

[43] WE Diewert. Duality approaches to Microeconomic Theory. In K.J. Arrow and M.D. Intrilligator, editors, *Handbook of Mathematical Economics*. Amsterdam, 1978.

[44] Habiba Djebbari and Jeffrey Smith. Heterogeneous impacts in PROGRESA. *Journal of Econometrics*, 145(1-2):64–80, July 2008.

[45] Paul Dolan and Robert Metcalfe. Better neighbors and basic knowledge : a field experiment on the role of non-pecuniary incentives on energy consumption . 2011.

[46] Julie S Downs, George Loewenstein, and Jessica Wisdom. Strategies for Promoting Healthier Food Choices. *American Economic Review: Papers & Proceedings*, 99(2):159–164, 2009.

[47] Jeffery A. Dubin and Daniel L. McFadden. An Econometric Analysis of Residential Electric Appliance Holdings and Consumption. *Econometrica*, 52(2):345–362, 1984.

[48] Jon Elster. Social Norms and Economic Theory. *Journal of Economic Perspectives*, 3(4):99–117, 1989.

[49] M. Espey, J. Espey, and W. D. Shaw. Price elsasticity of residential demand for water: A meta-analysis. *Water Resources Research*, 33(6):1369–1374, 1997.

[50] Y. H. Farzi, K. J. M. Huisman, and P. M. Kort. Optimal timing of technology adoption. *Journal of Economic Dynamics and Control*, 22:779–799, 1998.

[51] John Faux and Gregory M. Perry. Estimating Irrigation Water Value Using Hedonic Price Analysis: A Case Study in Malheur County, Oregon. *Land Economics*, 75(3):440–452, August 1999.

[52] By Paul J Ferraro, Juan Jose Miranda, and Michael K Price. The Persistence of Treatment Effects with Norm-Based Policy Instruments : Evidence from a Randomized Environmental Policy Experiment. *American Economic Review*, 101(3):318–322, 2011.

[53] Paul J. Ferraro and Juan José Miranda. Heterogeneous Treatment Effects and Mechanisms in Information-Based Environmental Policies: Evidence from a Large-Scale Field Experiment. *Resource and Energy Economics*, Forthcomin, April 2013.

[54] PJ Ferraro and MK Price. Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. *Review of Economics and Statistics*, 95(1):64–73, 2013.

[55] Sergio Fipro. Efficient Semiparametric Estimation of Quantile Treatment Effects. *Econometrica*, 75(1):259–276, 2007.

[56] Sergio Firpo, Nicole M. Fortin, and Thomas Lemieux. Unconditional Quantile Regressions. *Econometrica*, 77(3):953–973, 2009.

[57] Anthony C. Fisher, W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker. The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment. *American Economic Review*, 102(7):3749–3760, 2012.

[58] M Frölich and B Melly. Quantile treatment effects in the regression discontinuity design. *Journal of Econometrics*, 168:382–395, 2012.

[59] Markus Frölich and Blaise Melly. Unconditional Quantile Treatment Effects Under Endogeneity. *Journal of Business & Economic Statistics*, 31(3):346–357, July 2013.

[60] S. Gaudin. Effect of price information on residential water demand. *Applied Economics*, 38(4):383–393, March 2006.

[61] Andrew Gelman and Jennifer Hill. *Data Analysis using Regression and Multilevel/Hierarchical Models*. Cambride University Press, New York, NY, 2007.

[62] J Geweke. Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments. In J.M . Bernardo, J.O. Berger, A.P. Dowd, and A.F.M. Smith, editors, *Bayesian Statistics*. Oxford Universit, 4 edition, 1992.

[63] Kenneth Gillingham. Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices. 2012.

[64] Anatoly A Gitelson. Remote estimation of crop fractional vegetation cover: the use of noise equivalent as an indicator of performance of vegetation indices. *International Journal of Remote Sensing*, 34(17):6054–6066, 2013.

[65] Uri Gneezy and John a List. Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments. *Econometrica*, 74(5):1365–1384, September 2006.

[66] Patricia Gober, Anthony Brazel, Ray Quay, Soe Myint, Susanne Grossman-Clarke, Adam Miller, and Steve Rossi. Using Watered Landscapes to Manipulate Urban Heat Island Effects: How Much Water Will It Take to Cool Phoenix? *Journal of the American Planning Association*, 76(1):109–121, December 2009.

[67] Markus Goldstein and Christopher Udry. The Profits of Power : Land Rights and Agricultural Investment in Ghana. *Journal of Political Economy*, 116(6):981–1022, 2008.

[68] Lawrence H. Goulder, Antonio M. Bento, Mark R. Jacobsen, and Roger H. von Haefen. Distributional and Efficiency Impacts of Increased US Gasoline Taxes. *American Economic Review*, 99(3):667–699, 2009.

[69] R. Quentin Grafton and Michael B. Ward. Prices versus Rationing: Marshallian Surplus and Mandatory Water Restrictions*. *Economic Record*, 84(September 2008):S57–S65, September 2008.

[70] R. Halvorsen and R. Palmquist. The interpretation of dummy variables in semilogarithmic equations. *American Economic Review*, 70(3):474–75, 1980.

[71] W. Michael Hanemann. Discrete/Continuous Models of Consumer Demand. *Econometrica*, 52(3):541–561, 1984.

[72] Glenn W Harrison and John A List. Field Experiments. *Journal of Economic Literature*, 42(4):1009–1055, 2004.

[73] Kevin A. Hassett and Gilbert E. Metcalf. Energy tax credits and residential conservation investment: Evidence from panel data. *Journal of Public Economics*, 57(2):201–217, June 1995.

[74] James Heckman. Shadow Prices , Market Wages , and Labor Supply. *Econometrica*, 42(4):679–694, 1974.

[75] James Heckman. *The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process and some Monte Carlo evidence.* University of Chicago Center for Mathematical studies in Business and Economics, 1987.

[76] Julie A. Hewitt and Michael W. Hanemann. A Discrete/Continuous Choice Approach to Residential Water Demand under Block Rate Pricing. *Land Economics*, 71(2):173–192, 1995.

[77] TA Howell, JA Tolk, AD Schneider, and SR Evett. Evapotranspiration, yield, and water use efficiency of corn hybrids differing in maturity. *Agronomy Journal*, 1998.

[78] A R Huete, H Q Liu, K Batchily, and W Van Leeuwen. A Comparison of Vegetation Indices over a Global Set of TM Images for EOS-MODIS. *Remote Sensing of Environment*, 59:440–451, 1997.

[79] Koichiro Ito. Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing. 2012.

[80] Koichiro Ito. How Do Consumers Respond to Nonlinear Pricing? Evidence from Household Water Demand. 2013.

[81] Eric J Johnson and Daniel Goldstein. Do Defaults Save Lives ? *Science*, 302:1338–1339, 2003.

[82] David R. Kamerschen and David V. Porter. The demand for residential, industrial and total electricity, 1973–1998. *Energy Economics*, 26(1):87–100, January 2004.

[83] H Allen Klaiber and V Kerry Smith. Recovering Household Valuation of Urban Heat Island in the Presence of Omitted Variables across Spatial Scales. 2011.

[84] H. Allen Klaiber, V. Kerry Smith, Michael Kaminsky, and Aaron Strong. Measuring Price Elasticities for Residential Water Demand with Limited Information. 2012.

[85] Roger Koenker and Gilbert Basset. Regression Quantiles. *Econometrica*, 46(1):33–50, 1978.

[86] Gary Koop. *Bayesian Econometrics*. Wiley, West Sussex, 2003.

[87] Craig E. Landry, Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp. Toward an understanding of the economics of charity: Evidence from a field experiment. *Quarterly Journal of Economics Journal of Economics*, (May), 2006.

[88] Lung-Fei Lee. Generalized Econometric Models with Selectivity. *Econometrica*, 51(2):507–512, 1983.

[89] Steven D. Levitt and John a. List. Field experiments in economics: The past, the present, and the future. *European Economic Review*, 53(1):1–18, January 2009.

[90] Gary D. Libecap. Institutional Path Dependence in Climate Adaptation: Coman's "Some Unsettled Problems of Irrigation". *American Economic Review*, 101(February):64–80, 2011.

[91] JA List and D Lucking-Reiley. Demand reduction in multiunit auctions: Evidence from a sportscard field experiment. *The American Economic Review*, 90(4):961–972, 2000.

[92] Charles F. Manski. Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3):531, July 1993.

[93] Charles F Manski. Economic Analysis of Social Interactions. *Journal of Economic Perspectives*, 14(3):115–136, August 2000.

[94] Erin T. Mansur and Sheila M. Olmstead. The value of scarce water: Measuring the inefficiency of municipal regulations. *Journal of Urban Economics*, 71(3):332–346, May 2012.

[95] Roberto Martínez-Espiñeira and Céline Nauges. Is all domestic water consumption sensitive to price control? *Applied Economics*, 36(15):1697–1703, September 2004.

[96] Robert Mendelsohn and Ariel Dinar. Climate, Water, and Agriculture. *Land Economics*, 79(3):328, August 2003.

[97] Robert Mendelsohn, William D Nordhaus, and Daigee Shaw. The Impact of Global Warming on Agriculture: A Ricardian Analysis. *American Economic Review*, 84(4):753–771, 1994.

[98] C. Mood. Logistic Regression: Why We Cannot Do What We Think We Can Do, and What We Can Do About It. *European Sociological Review*, 26(1):67–82, March 2009.

[99] Yair Mundlak. On the Pooling of Time Series and Cross Sectino Data. *Econometrica*, 46(1):69–85, 1978.

[100] Antonio Musolesi and Mario Nosvelli. Long-run water demand estimation: habits, adustment dynamics and structural breaks. *Applied Economics*, 43(17):2111–2127, 2011.

[101] Soe W Myint, May Yuan, Randall S Cerveny, and Chandra P Giri. Comparison of Remote Sensing Image Processing Techniques to Identify Tornado Damage Areas from Landsat TM Data. *Sensors*, 8:1128–1156, 2008.

[102] Shanthi Nataraj and W. Michael Hanemann. Does marginal price matter? A regression discontinuity approach to estimating water demand. *Journal of Environmental Economics and Management*, 61(2):198–212, March 2011.

[103] Céline Nauges and Alban Thomas. Long-run Study of Residential Water Consumption. *Environmental and Resource Economics*, 26:25–43, 2003.

[104] Noelwah R. Netusil and Matthew T. Summers. Valuing instream flows using the hedonic price method. *Water Resources Research*, 45(W11429):1–7, November 2009.

[105] Michael L Nieswiadomy and David J. Molina. Comparing Residential Water Demand Estimates under Decreasing and Increasing Block Rates Using Household Data. *Land Economics*, 65(3):280–289, 1989.

[106] Sheila M. Olmstead, W. Michael Hanemann, and Robert N. Stavins. Water demand under alternative price structures. *Journal of Environmental Economics and Management*, 54(2):181–198, September 2007.

[107] Raymond B Palmquist. Land as a Differentiated Factor of Production: A Hedonic Model and Its Implications for Welfare Measurement. *Measurement*, 65(1):23–28, 1989.

[108] Raymond B Palmquist and Leon E Danielson. A Hedonic Study of the Effects of Erosion Control and Drainage on Farmland Values. *American Journal of Agricultural Economics*, 71(1):55–62, 1989.

[109] Ragan A Petrie and Laura O Taylor. Estimating the Value of Water Use Permits: A Hedonic Approach Applied to Farmland in the Southeastern United States. *Land Economics*, 83(3):302–318, 2007.

[110] Ellen M Pint. Household Responses to Increased Water Rates during the California Drought. *Land Economics*, 75(2):246–266, 1999.

[111] R.A. Pollak. Conditional demand functions and consumption theory. *The Quarterly Journal of Economics*, 83(1):60–78, 1969.

[112] David Powell. A New Framework for Estimation of Quantile Treatment Effects: Nonseperatble Disturbance in the Presence of Covariates. 2013.

[113] Adrian E Raftery and Steven Lewis. How Many Iterations in the Gibbs Sampler? In J.M . Bernardo, J.O. Berger, A.P. Dowd, and A.F.M. Smith, editors, *Bayesian Statistics*. Oxford University Press, 4 edition, 1992.

[114] Adrian E Raftery, David Madigan, and Jennifer A Hoeting. Bayesian Model Averaging for Linear Regression Models. *Journal of the American Statisitical Association*, 92(437):179–191, 1997.

[115] Mary E Renwick and Sandra O Archibald. Demand Side Management Policies for Residential Water Use: Who Bears the Conservation Burden? *Land Economics*, 74(3):343–359, 1998.

[116] Mary E. Renwick and Richard D. Green. Do Residential Water Demand Side Management Policies Measure Up? An Analysis of Eight California Water Agencies. *Journal of Environmental Economics and Management*, 40(1):37–55, July 2000.

[117] David Revelt and Kenneth Train. Mixed Logit with Repeated Choices: Households' Choices of Appliance Efficiency Level. *Review of Economics and Statistics*, 80(4):647–657, November 1998.

[118] Sherwin Rosen. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82(1):34, January 1974.

[119] Todd Schatzki. Options, uncertainty and sunk costs: an empirical analysis of land use change. *Journal of Environmental Economics and Management*, 46(1):86–105, July 2003.

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

[121] Wolfram Schlenker, W. Michael Hanemann, and Anthony C. Fisher. Water Availability, Degree Days, and the Potential Impact of Climate Change on Irrigated Agriculture in California. *Climatic Change*, 81(1):19–38, January 2007.

[122] Carl P. Schmertmann. Selectivity bias correction methods in polychotomous sample selection models. *Journal of Econometrics*, 60:101–132, January 1994.

[123] Jan C. Semenza, Carol H. Rubin, Kenneth H. Falter, Joel D. Selanikio, Dana W. Flanders, Holly L. Howe, and John L. Wilhelm. HEAT-RELATED DEATHS DURING THE JULY 1995 HEAT WAVE IN CHICAGO. *The New England Journal of Medicine*, 335:84–90, 1996.

[124] Steven E Sexton. Energy Price Salience and Residential Energy. 2010.

[125] Jeong-Shik Shin. Perception of Price When Price Information is Costly: Evidence from Residential Electricity Demand. *The Review of Economics and Statistics*, 67(4):591–598, 1985.

[126] Kent A Sovocool, Mitchell Morgan, and Doug Bennett. An in-depth investiagation of Xeriscape as a water conservation measure. *Journal of American Water Works Association*, 98(2):82–93, 2006.

[127] William L. Stefanov and Maik Netzband. Assessment of ASTER land cover and MODIS NDVI data at multiple scales for ecological characterization of an arid urban center. *Remote sensing of Environment*, 99(1-2):31–43, 2005.

[128] William L. Stefanov, Michael S. Ramsey, and Philip R. Christensen. Monitoring urban land cover change: An expert system approach to land cover classification of semiarid to arid urban centers. *Remote Sensing of Environment*, 77(2):173–185, 2001.

[129] Richard H Thaler and Cass R Sunstein. Libertarian Paternalism. *American Economic Review: Papers & Proceedings*, 93(2), 2003.

[130] Yacov Tsur and Theodore Graham-Tomasi. The buffer value of groundwater with stochastic surface water supplies. *Journal of Environmental Economics and Management*, 21(3):201–224, November 1991.

[131] USBR. Interim Comprehensive Basin Operating Plan. Technical report, Yakima Field Office, Yakima, WA, 2002.

[132] USBR. Yakima River Basin Study. Technical report, U.S. Bureau of Reclemation, 2011.

[133] USBR. Yakima River Basin Study - Proposed Integrated Water Resource Management Plan Volume 1. Technical report, United States Bureau of Reclamation; Washington State Department of Ecology, Yakima, WA, 2011.

[134] USBR. Yakima River Basin Study - Water Needs for Out-of-Stream Uses. Technical report, United States Bureau of Reclamation, Yakima, WA, 2011.

[135] Julie A. Vano, Michael J. Scott, Nathalie Voisin, Claudio O. Stöckle, Alan F. Hamlet, Kristian E. B. Mickelson, Marketa McGuire Elsner, and Dennis P. Lettenmaier. Climate change impacts on water management and irrigated agriculture in the Yakima River Basin, Washington, USA. *Climatic Change*, 102(1-2):287–317, May 2010.

[136] Andrés Viña, Anatoly a. Gitelson, Anthony L. Nguy-Robertson, and Yi Peng. Comparison of different vegetation indices for the remote assessment of green leaf area index of crops. *Remote Sensing of Environment*, 115(12):3468–3478, December 2011.

[137] Elizabeth A. Wentz and Patricia Gober. Determinants of Small-Area Water Consumption for the City of Phoenix, Arizona. *Water Resources Management*, 21(11):1849–1863, February 2007.

[138] Western Region Climate Center. Monthly Climate Summaries, 2010.

[139] Jefferey M. Woolridge. *Econometric Analysis of Cross Section and Panel Data.* MIT Press, Cambridge, MA, 2002.

[140] Andrew C. Worthington and Mark Hoffman. an Empirical Survey of Residential Water Demand Modelling. *Journal of Economic Surveys*, 22(5):842–871, July 2008.

[141] Feng Xu, Ron C. Mittlehammer, and L. Allen Torrell. Modeling nonnegativity via truncated logistic and normal distributions: an application to ranch land price analysis. *Journal of Agricultural and Resource Economics*, 19(1):102–114, 1994.

[142] Yakima County Superior Court. Draft Schedule of Rights: Acquavella Surface Water Rights Adjudication. Technical Report 17, Yakima County Superior Court, Yakima, WA, 2012.

Appendix A

APPENDIX FOR CHAPTER 1

A.1 Bayesian Model Specification

The basic structure of the linear regression with a general covariance matrix is taken from Koop (2003). The regression function is $y = X\beta + \epsilon$, where y is the real log sale price per acre, X is a matrix of covariates, β is a coefficient vector and ϵ is a heteroskedastic error term distributed $\epsilon \sim N(0, \sigma^2\Omega)$. I first outline the model with general covariance matrix $\sigma^2\Omega$, and then with the restrictions I impose to aid estimation. The likelihood function is

$$p(y|\beta, \sigma^2, \Omega) = (2\pi\sigma^2)^{-\frac{n}{2}}|\Omega|^{-\frac{1}{2}}\left[\exp\left(-\frac{1}{\sigma^2}(y - X\beta)'\Omega^{-/frac12}(y - X\beta)\right)\right] \tag{A.1}$$

and the priors are

$$p(\beta) \sim N\left(\underline{\beta}, \underline{\mathrm{V}}\right) \tag{A.2a}$$

$$p(\sigma^2) \sim I\Gamma\left(\frac{\nu}{2}, \frac{\nu s^2}{2}\right) \tag{A.2b}$$

$$p(\Omega) \sim p(\Omega) \tag{A.2c}$$

Prior values are $\underline{\beta} = 0$, $\underline{\mathrm{V}} = 1000^2 I_k$, $\underline{\nu} = 1$, and $\underline{s^2} = 1^{-1000000}$. This leads to a joint posterior of the following form

$$\begin{aligned} p(\beta, \sigma^2, \Omega|y) \propto p(\Omega)\left[\exp\left(-\frac{1}{2}\left[\frac{1}{\sigma^2}(y - X\beta)'\Omega^{-1}(y - X\beta) + (\beta - \underline{\beta})'\underline{\mathrm{V}}^{-1}(\beta - \underline{\beta})\right]\right)\right] \\ \left(\frac{1}{\sigma^{n+\nu-2}}\right)\exp\left(\frac{\underline{\nu}}{2\sigma^2\underline{\mathrm{s}}^{-2}}\right) \end{aligned} \tag{A.3}$$

Since this posterior is not of standard form we draw from the conditional posterior distributions for β, σ^2, given in equation (8). I will not assume that I know the structure of the heteroskedasticity, but will make some parametric assumptions to aid in the computation. The structure of Ω is given by the following priors $\Omega = \begin{bmatrix} \omega_1 & 0 & \cdots & 0 \\ 0 & \omega_2 & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \cdots & 0 & \omega_n \end{bmatrix}$ and $\lambda \equiv (\lambda_1, \lambda_2, \ldots, \lambda - n)' \equiv (\omega_1^{-1}, \omega_2^{-1}, \ldots, \omega_n^{-1})'$

$$p(\lambda) = \prod_{i=1}^{n} \Gamma(\lambda_i | 1, \nu_\lambda) \tag{A.4a}$$

$$p(\nu_\lambda) = \Gamma(\underline{\nu_\lambda}, 2) \tag{A.4b}$$

We use a value of $\nu_\lambda = 15$, and the conditional posterior for λ_i and ν_λ are given in equation (8). Figure A.1 shows a histogram for the posterior estimates of ν_λ in the top panel and the M-H acceptance rate in the bottom panel.

Figure A.1: Posterior Estimates for Heteroskedasticity Dispersion Parameter

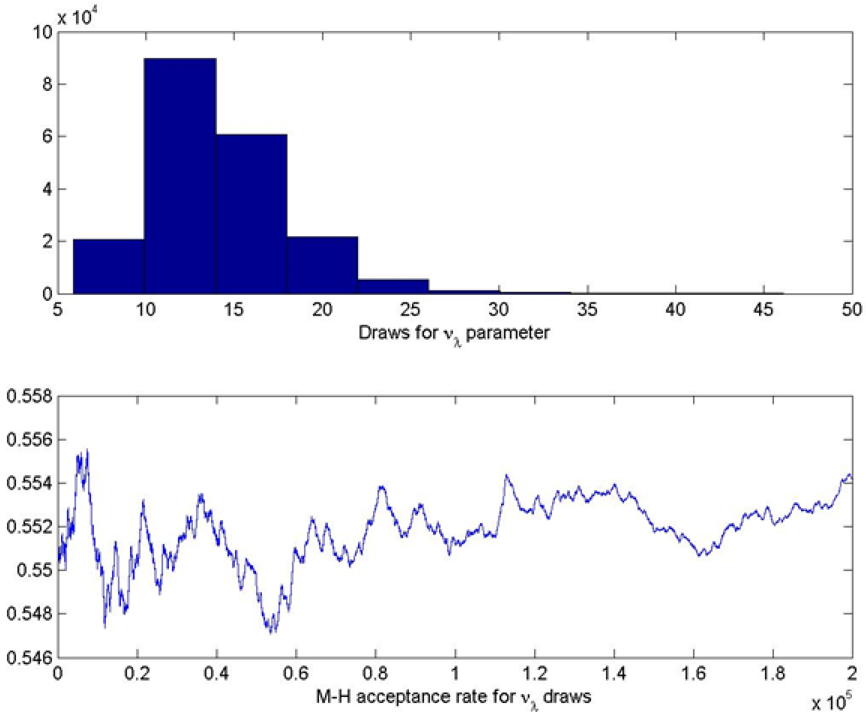

*Notes:*The top panel is a histogram for the draws of the degree-of-freedom parameter, ν_λ, that determines the form of heteroskedasticity in the base regression. The bottom graph shows the acceptance rate for the Metropolis-Hastings algorithm.

A.2 MCMC Convergence Diagnostics

The Gibbs sampler is an MCMC procedure where arbitrary initial values may bias the results. There are several diagnostic tools used to assess the convergence of the Gibbs sampler to the true joint posterior distribution, ensuring that the effect of the starting values has worn off. We employ three tools that all indicate that the Gibbs sampler reached convergence. The I-statistic is the ratio of the number of draws required for given accuracy level to the number of draws necessary if the chain was i.i.d., developed by [113]. For an accuracy level of 1% the I-statistic is 1.047, safely below the recommended threshold of 5. The autocorrelation of draws in the parameter chain is another metric to determine if the Gibbs sampler is drawing from the true distribution. The low level of serial correlation in the Gibbs draws as shown in Table A.1 provides evidence that the draws represent and independent sample. Lastly we show the results for Geweke χ^2 test for equality in means for two regions of the Gibbs sampler we use the first 20% and the last 50% of the Gibbs draws. If the Gibbs sampler reached convergence then any subset should represent the true joint posterior and there should be no difference in parameter means for different regions. Table A.2 shows the p-values for the χ^2 test for the null of equal means. In all cases the test fails to accept the null at the 90% level. These diagnostics tool suggest that the Gibbs sampler has reached convergence; not a surprise given that running 220,000 draws with 20,000 burn-in draws is extremely circumspect.

Table A.1: MCMC Convergence Diagnostics - Autocorrelations for parameter chains

Variable	Lag 1	Lag 5	Lag 10	Lag 50
Senior	0.129	-0.002	0	0
Time	0.146	-0.003	0.006	0.002
Time$\hat{2}$	0.126	-0.008	0.002	0
Acres	0.189	0.004	0.001	-0.003
Improvements	0.176	0.008	0.008	0.009
Rolling Avg	0.144	0.007	0.003	0.001
Class 1	0.028	0.002	0.002	0.003
Class 2	0.028	0.002	0.003	0.003
Class 3	0.028	0.002	0.002	0.003
Class 4	0.012	0.001	-0.001	0
Class 5	0.029	0.002	0.002	0.003
Distance City	0.126	0.001	0.001	0.001
Distance Stream	0.127	0	-0.003	0.002
Inv Dist UGA	0.142	0.003	0.001	0.002
Inv Dist Yak Riv	0.161	-0.001	0.004	0.003
Right	0.117	0.003	0	0
Residential	0.115	0.001	-0.004	0.001
Kittitas	0.137	0.001	0	0.002
Benton	0.128	0.002	0	0.001
Senior	0.129	-0.002	0	0

Notes: Autocorrelation measures of the posterior estimates based on the draws of the Gibbs sampler. 220,000 initial draws were taken with 20,000 omitted resulting in 200,000 draws.

Table A.2: Geweke Chi-square Test for Equality of Means

Variable	NSE	NSE 4%	NSE 8%	NSE 15%
Senior	0.8868	0.9231	0.9258	0.9221
Time	0.5542	0.6586	0.6529	0.6123
Time$\hat{2}$	0.7283	0.7314	0.716	0.6604
Acres	0.8886	0.9201	0.9238	0.9238
Improvements	0.1471	0.4905	0.4748	0.4434
Rolling Avg.	0.8088	0.8693	0.8708	0.8571
Class 1	0.6951	0.7084	0.6804	0.6245
Class 2	0.778	0.7867	0.7621	0.7191
Class 3	0.7679	0.7744	0.7517	0.7068
Class 4	0.5468	0.5438	0.5388	0.5182
Class 5	0.7487	0.7602	0.7345	0.691
Distance City	0.4681	0.5105	0.4703	0.3905
Distance Stream	0.8367	0.8584	0.8442	0.8369
Inv Dist UGA	0.6766	0.729	0.7411	0.7439
Inv Dist Yak River	0.6679	0.6425	0.5978	0.5735
Right	0.6345	0.6199	0.5521	0.4226
Residential	0.8324	0.8646	0.8659	0.8569
Kittitas	0.2728	0.3412	0.3022	0.1148
Benton	0.7977	0.8142	0.8173	0.786
Intercept	0.7737	0.7627	0.7389	0.6847

Notes: Results are p-values for the Geweke chi-square test for difference in means for two intervals of Gibbs draws. I use the first 20% and the last 50% of draws as the two intervals. 220,000 initial draws were taken with 20,000 omitted resulting in 200,000 draws.

Appendix B

APPENDIX FOR CHAPTER 2

B.1 Water Rates in Phoenix

Figure B.1 shows the water rate structure in the City of Phoenix during the sample period.

Figure B.1: Phoenix Water Rate Structure

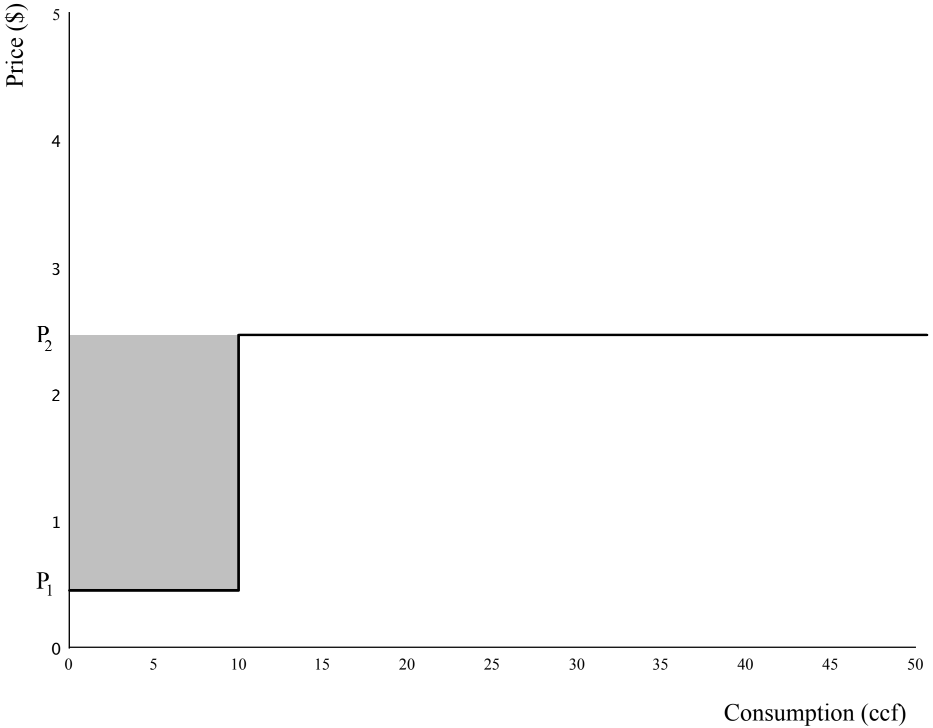

Notes: The marginal price is p_1 if consumption is less than $\bar{w}$ and p_2 when consumption is greater than $\bar{w}$. The total bill is the area under the line and never includes the shaded area.

B.2 Processing Landsat Data

The Landsat data is publicly available in a raw form and requires processing in order to apply the data for analysis. Processing is particularly important when comparing two or more images over time or across space. I follow the steps presented in [101] to process the Landsat data in Erdas Imagine. The first step is to find valid scenes via the USGS Glovis system. I only select images with less than 10% cloud cover and an overall quality score of at least 9 out of 10. I try to select at least two scenes for each year in between the months of June and August, though some scenes are also take from May and September. Combining data from two scenes helps to alleviate idiosyncratic shocks due to weather. I end up with 25 scenes between the years of 1998 and 2009. I repeat the following steps for each of the 25 scenes.

Since the Landsat data is stored in separate bands I import the data and stack all the layers on top of each other. Next, I subset the image to limit the geographic area to Phoenix metro. Each Landsat scene is 185km x 117km so limiting the image to the study area greatly increases computational speed and the digital space required for storage. Each image is registered to a base image in order to reduce locational errors using 14 ground control points and ensuring root mean squared errors of less than 0.1. This process ensures that all the images line up properly and that a parcel has the same geo-reference in each image over time. In order to account for differences in atmospheric reflectance and solar radiation I apply the Cos(t) method of radiometric correction of [33]. Once theses steps are complete the images are suitable to be compared over time and I calculate the Normalized Difference Vegetation Index (NDVI).

NDVI is the calculated from the visible and near-infrared bands in the Landsat data[1]. Healthy green vegetation absorbs visible light and reflects infrared light so the difference performs well in identifying healthy vegetation. The formula used to calculate the index is $NDVI = (NIR - VIS)/(NIR + VIS)$, where NIR is the near infrared band and VIS is the visible red band. This formula results in an index ranging from -1 to 1 with higher values representing more robust vegetation. A sample image of NDVI for the Phoenix metro area, along with the border of the water utility's service area is shown in Figure B.2. Figure B.3 shows an area just northwest of Arizona State University to give an example of how different values of NDVI correspond to land use features.

The NDVI data are then merged with parcel boundaries in Geographic Information System software. Each pixel of the Landsat data is 30m x 30m and is often larger, or matches imperfectly, with the parcel boundaries. An example of the problems that can arise from merging NDVI at the parcel level are displayed in Figure B.4. It is clear that smaller parcels create challenges for spatially merging NDVI data. To reduce the noise in the spatial merge I downscale each NDVI pixel to nice 10m x 10m pixels and take the spatially weighted average of the pixels within the parcel. The final results is a panel dataset at the parcel level that contains the time series variation

---

[1] More information on NDVI is available at
http://earthobservatory.nasa.gov/Features/MeasuringVegetation/measuring_vegetation_2.php.

of NDVI for each of the 25 Landsat scenes.

Figure B.2: Phoenix Metro & NDVI Coverage

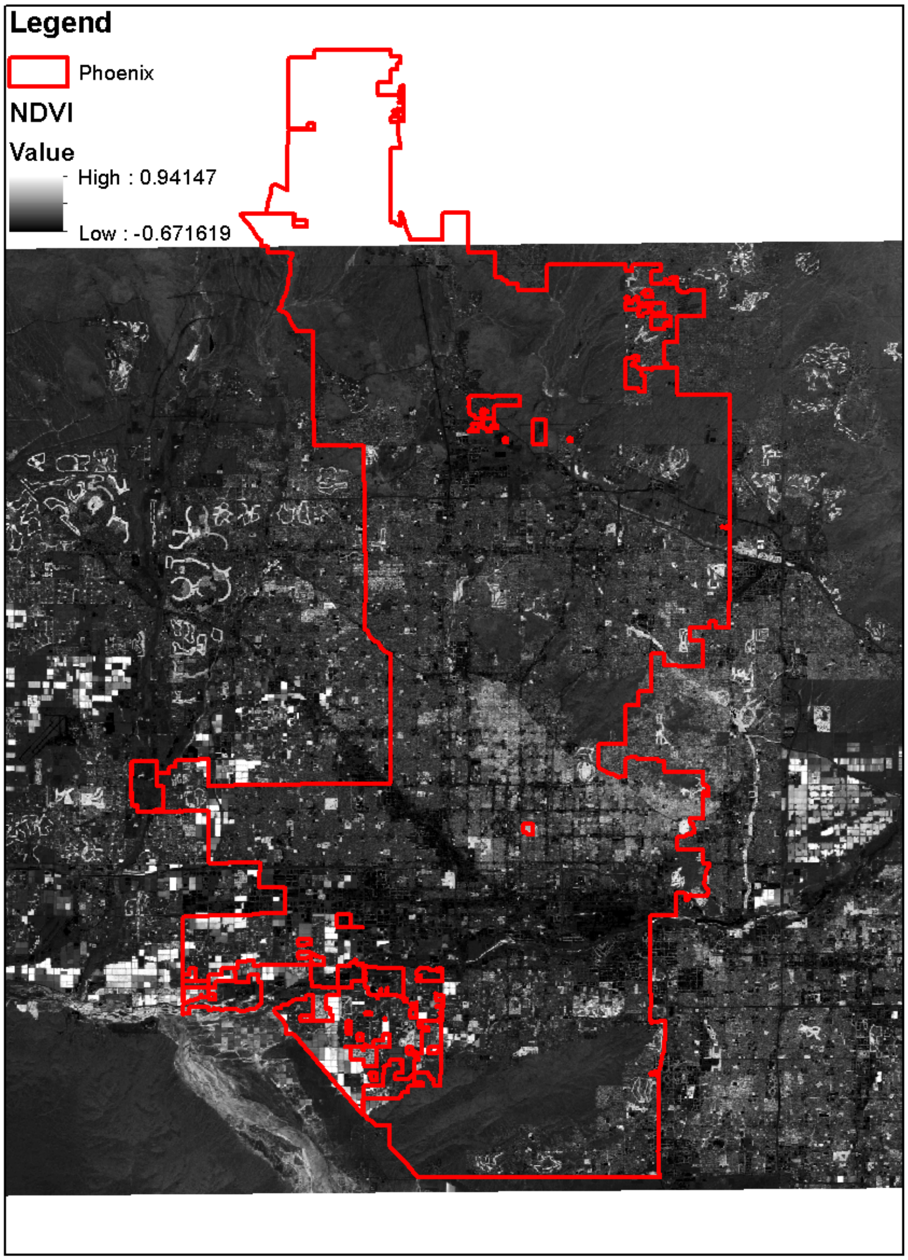

Note: This is the geographical sample space for the NDVI data, along withe the border of the utility's service area. There is a small section of northern Phoenix that is not captured by NDVI but there is relatively little development there.

Figure B.3: NDVI of Known Land Features

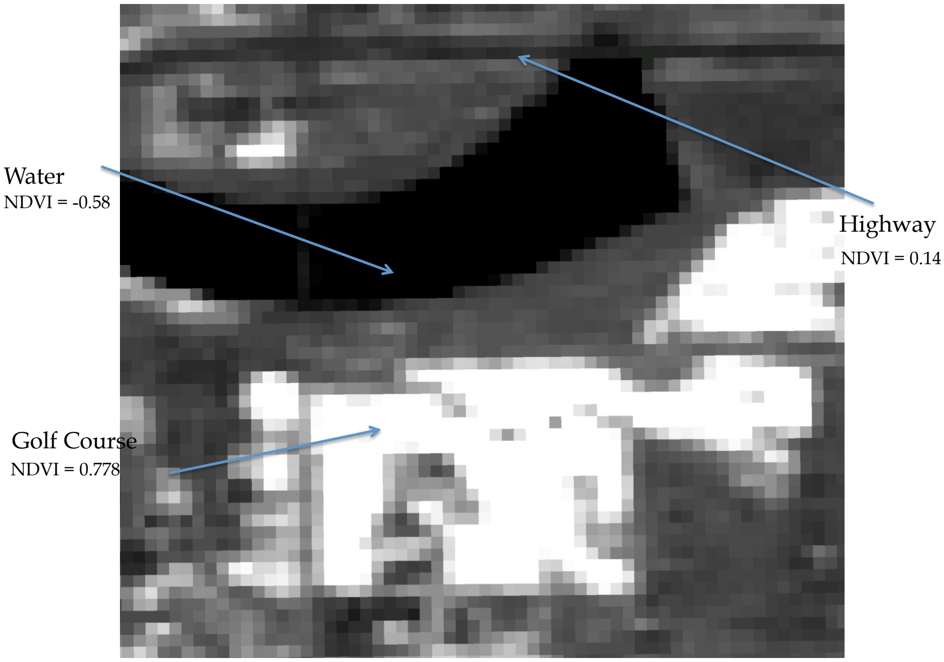

Note: This is a sample area just northwest of Arizona State University, whose land use features are known.

Figure B.4: Merging Parcels and NDVI

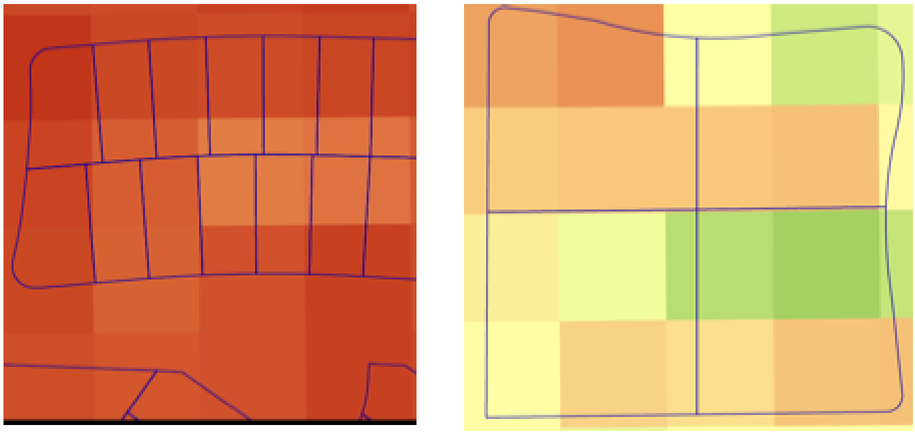

Note: The color gradient for the images is the same, but is purely for illustrative purposes. Each pixel is $900m^2$ and is actually composed of nine $100m^2$ homogeneous pixels that improve the spatially weighted average of parcel-level NDVI.

B.3 Weather Normalization

While the processing steps described above alleviates concerns due fluctuations over time in how the satellite captures images, I also need to account for the impact of natural variations in weather on NDVI. Irrespective of human watering practices NDVI will vary based on the weather conditions in the area. In order to minimize these variations and focus on the water-added component of landscape I normalize the NDVI for weather. Since the images are taken at different times of the month I match daily weather data based on the image date. From this I construct variables representing weather conditions for each of the four weeks prior to the image. In addition since Phoenix is very dry and often will not have rained within four weeks in the summer I generate a variable for the number of days since the last precipitation event. Next I regress NDVI on these weather variables and keep the residuals as weather-normalized values of NDVI.

Table B.1 presents the results of the weather normalizations with lags of up to four weeks. I select residuals from the regression in column (4) as my preferred measure, though robustness checks using other normalizations produce very similar results. The results in Table B.1 are mostly intuitive with the cumulative effect of higher soil temperatures and evapotranspiration leading to lower values of NDVI. Great precipitation generally increases NDVI and the longer dry periods decrease NDVI. Overall weather explains between 10-14% of NDVI suggesting that most of the variation is spatial, due to different landscaping practices across the city. The regressions appear to maintain the spatial variation because analyzing the residuals shows that the standard deviation in normalized NDVI within a given year is very similar to the the standard deviation of the raw data. Figure B.5 shows the difference in the raw and normalized NDVI between two years. The differences the overall NDVI is less discernible in the normalized NDVI, facilitating the comparison of NDVI values across years.

Table B.1: Weather Normalization Regression

VARIABLES	(1)	(2)	(3)	(4)
Max Soil Temp (1 week ago)	-0.00415***	-0.00863***	-0.0101***	-0.0127***
	(1.01e-05)	(3.69e-05)	(3.79e-05)	(4.67e-05)
Max Soil Temp (2 weeks ago)		0.00402***	0.00103***	0.00694***
		(3.53e-05)	(4.15e-05)	(6.15e-05)
Max Soil Temp (3 weeks ago)			0.00451***	0.00329***
			(3.70e-05)	(5.47e-05)
Max Soil Temp (4 weeks ago)				-0.00399***
				(3.94e-05)
Total Rain (1 week ago)	0.00585***	0.00381***	0.00921***	0.000613**
	(7.14e-05)	(9.21e-05)	(0.000102)	(0.000276)
Total Rain (2 weeks ago)		-0.00301***	-0.000735***	-0.00312***
		(1.92e-05)	(3.29e-05)	(7.81e-05)
Total Rain (3 weeks ago)			-0.000890***	-6.05e-06
			(1.12e-05)	(1.47e-05)
Total Rain (4 weeks ago)				-0.00128***
				(1.51e-05)
Evapotranspiration (1 week ago)	-0.000680***	0.00197***	0.00337***	0.00147***
	(1.24e-05)	(1.62e-05)	(1.81e-05)	(2.17e-05)
Evapotranspiration (1 week ago)		-0.00285***	-0.00239***	-0.00179***
		(1.16e-05)	(1.19e-05)	(1.57e-05)
Evapotranspiration (1 week ago)			-0.000942***	-0.00194***
			(7.97e-06)	(1.10e-05)
Evapotranspiration (1 week ago)				-3.77e-05*
				(2.02e-05)
Days Since Rain	-0.000330***	-1.76e-05***	5.80e-05***	-1.26e-05***
	(1.64e-06)	(2.36e-06)	(2.61e-06)	(3.24e-06)
Constant	0.317***	0.332***	0.270***	0.470***
	(0.000948)	(0.000949)	(0.00124)	-0.00166
Monthlty Dummies	Yes	Yes	Yes	Yes
Observations	4,671,486	4,671,486	4,671,486	4,671,486
R-squared	0.119	0.132	0.138	0.147

Notes: Dependent variable is parcel-level NDVI for a given scene. Robust standard errors are in parenthesis. ** p<0.01, * p<0.05

Figure B.5: Differences in Raw and Normalized NDVI

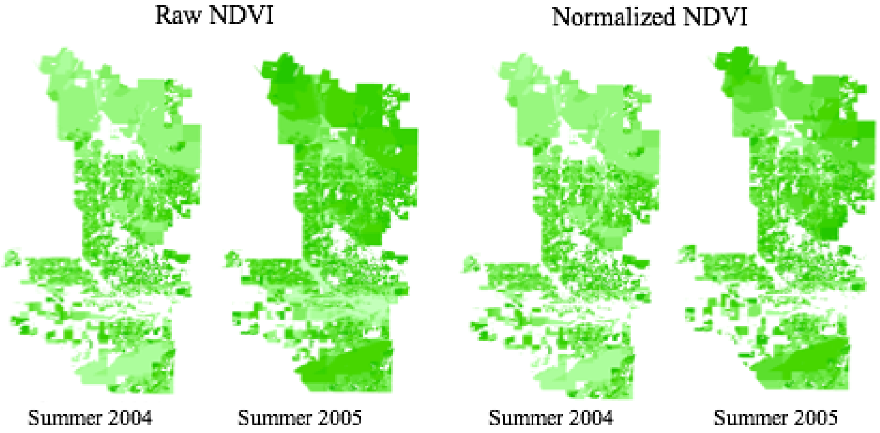

Note: Color gradient is based on deciles of NDVI over the years 1998-2009 with darker colors representing higher deciles.

B.4 Landscape Classification Diagnostics

In order to test the feasibility of using NDVI to classify different landscape varieties I compare the quantiles of NDVI to data from a widely cited remote sensing paper. [128] classifies 11 different types of land use for Phoenix using 1998 data, including mesic and xeric residential. I merge the parcels with 1998 NDVI with the classification from [128] keeping all parcels identified as either mesic or xeric. Table B.2 presents the percentage of parcels that were correctly identified using various quantiles of NDVI. The columns show the thresholds for NDVI quantiles to make a classification. Parcels with NDVI above the higher threshold are classified as wet in a given year and parcels with NDVI less than the low threshold are defined as dry. Therefore decreasing the high threshold and increasing the low threshold relaxes the conditions to observe a conversion. NDVI does a relatively better job classifying dry landscapes and for that reason I use an asymmetric threshold for defining landscape groups For example in the conditional demand models I designate a the Wet group by households that have NDVI above the 80th quantile every year, and the Dry group by households that have below the 30th quantile every year. The results in Table B.2 contribute to establishing a relatively less stringent threshold for the Dry group.

Table B.2: Landscape Diagnostics

NDVI Quantiles	90/10	80/20	70/30	60/40
<u>% Correct</u>				
Wet	82%	77%	73%	70%
Dry	91%	88%	84%	80%

Notes: The columns designate the quantile of NDVI to compare with wet and dry landscapes. The higher quantile is used to determine wet parcels and the lower quantile designates dry parcels. The percentage correct takes the data from [128] as the true value. We only compare single family residential households that were classified as xeric or mesic.

B.5 Landscape Conversion Robustness

I do not have validation data to verify that my definition of a landscape conversion is indeed correct. Adding landscape conversions into the water demand model provides some justification that I am indeed identifying conversions, but I also perform robustness checks by using alternative definitions of a conversion. The first check is to use the raw NDVI data to ensure that the weather normalization process did not introduce and bias. In this setting I compare the household level NDVI observation to the quantiles of the distribution of NDVI for the sample in that year. Next I relax the threshold for classifying a yearly NDVI observation as either wet or dry. This translates to a yearly NDVI observation that is above the 60th quantile being classified as wet as opposed to the base case where the threshold is above the 70th quantile. Likewise the dry classification is relaxed from being below the 40th quantile to below the 50th quantile. Table B.3 displays the results from the fixed effects logit under all combinations for generation of the landscape conversion variable. The columns represent regressions using different lags of the water rate, similar to Tables 4 and 5. All other control variables that are used in Table 5 are also present in these regressions, and the coefficients are omitted to condense space. The different panels (a)-(d) represent different designations of the landscape conversion variable as described above. Though the magnitudes do vary somewhat the pattern sign of the coefficients is the same, and the relative magnitude across equations is also consistent. This provides support that the landscape classification methodology is not driving the results.

Table B.3: Robustness of Landscape Conversion

	(1)	(2)	(3)	(4)	(5)
(a) Normalization - 70/40					
Price	12.34***	12.50***	13.28***		
Price (t-1)		-0.244	-5.759***	7.439***	3.453***
Price (t-2)			10.84***		9.428***
Additional Controls	Yes	Yes	Yes	Yes	Yes
(b) No Normalization - 70/40					
Price	14.16***	16.86***	15.96***		
Price (t-1)		-3.528***	-7.268***	7.180***	3.321***
Price (t-2)			10.06***		10.23***
Additional Controls	Yes	Yes	Yes	Yes	Yes
(c) Normalization - 60/50					
Price	8.697***	9.366***	8.455***		
Price (t-1)		-1.043	-3.625***	5.980***	2.844***
Price (t-2)			9.989***		10.87***
Additional Controls	Yes	Yes	Yes	Yes	Yes
(d) No Normalization - 60/50					
Price	7.735***	7.820***	6.731***		
Price (t-1)		-0.117	-1.090	5.755***	3.849***
Price (t-2)			6.223***		7.497***
Additional Controls	Yes	Yes	Yes	Yes	Yes

Notes: These are the same regressions for the fixed effect logit model of landscape conversion presented in Table 5. All other controls that appear in those regressions are included in all regression in Table A.3. The columns represent regressions using different lags of price and panels (a)-(d) designate different classification techniques for a landscape conversion. Normalization and No Normalization designate whether the NDVI data was weather normalized. 70/40 and 60/50 correspond to the low and high thresholds of NDVI quantiles that classify a yearly observation as wet or dry. Bootstrapped standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

B.6 Raw % Changes in Consumption and Conversions

In Figure 2.3 I overlay landscape conversions with quantiles of changes in consumption over time. Figure B.6 presents the same data in a different format. The color gradations now are the percentage changes in consumption over the course of sample and the hatched regions represent census blocks where at least 5% of the households in the sample converted their landscape from wet to dry. The percentage change is calculated at the census block level and is the change from the first two years of the sample to the last two years.

Figure B.6: Landscape Conversions & Consumption over Time

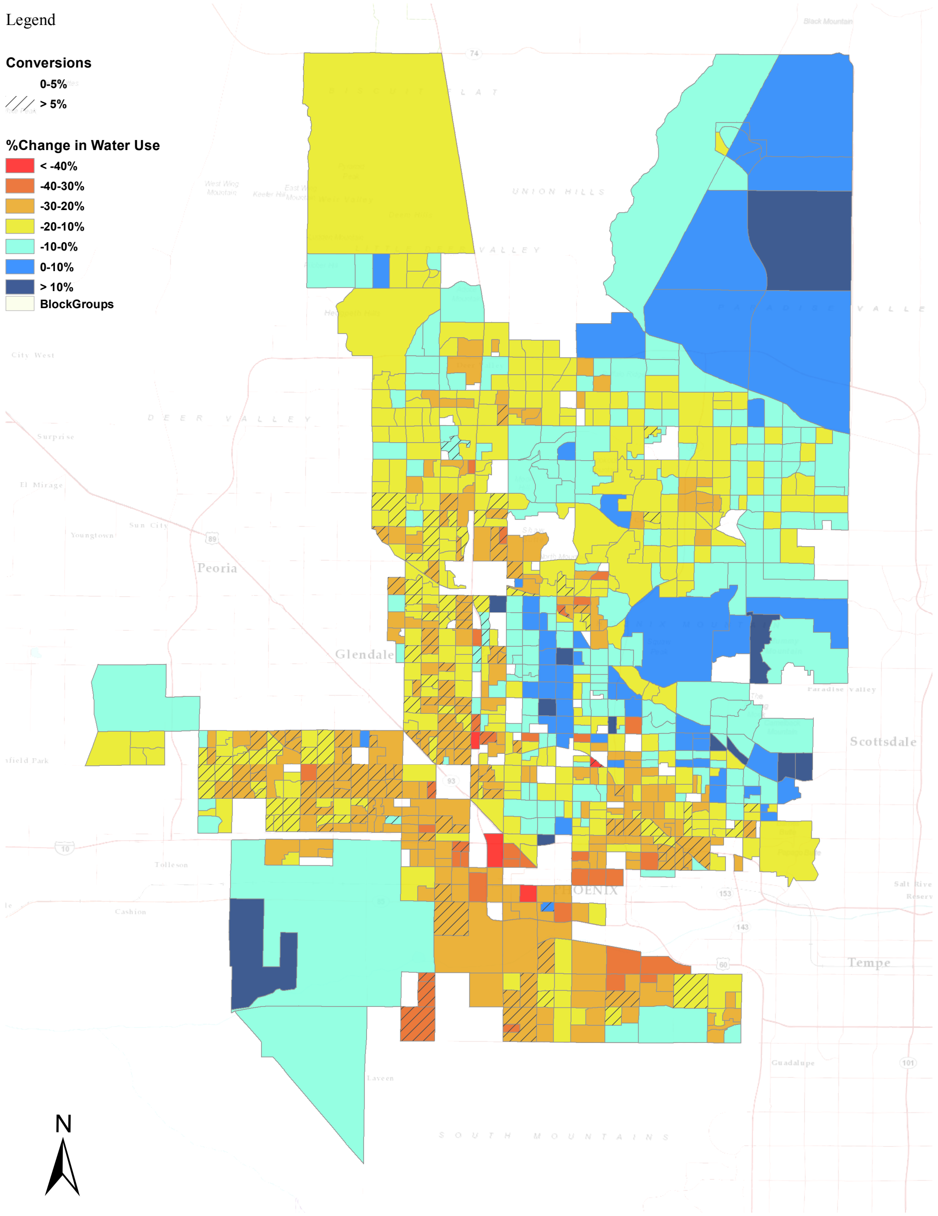

Notes: The shaded coloring are based on the average change in water consumption during the course of the sample. Specifically it measures the average of consumption in 2008-2009 divided by average consumption in 1998-1999. Census blocks with less than 20 houses in the sample are removed for confidentiality concerns. The hatched blocks are where more that 5% of the households in the sample experienced landscape conversions.

Appendix C

APPENDIX FOR CHAPTER 3

C.1 Pilots

WaterSmart partners with water utilities to send Home Water Reports, which provide information to the household on their water use relative to a similar cohort of households as well as personalized conservation recommendations. All of WaterSmart's current pilot programs are with California utilities; the names of these utilities have been disguised to protect the privacy of the water utilities' customers. Three of these programs have a randomized experimental design and have been active long enough to analyze for a combined initial sample of 805,467 total water use observations.[1] WaterSmart's first pilot, Watertown, has completed its experimental phase; all customers in this utility have since begun to receive home water reports. The experiments in the other two pilots are on going; water observations used in this paper for these pilots are as recent as August 2013. Sample sizes, the date of the first home water report sent, experiment end dates, and the number of households treated in each pilot can be found in Table C.1.

Table C.1: Initial Sample Sizes

Pilot	Start Date	End Date	N: Obs	N: Post-Treat	HHs	Treated HHs
Watertown	2011-09-20	2013-01-01	44,822	16,266	2,575	992
Aquaville	2012-06-28	on going	93,664	18,583	3,096	1,547
Hydroburg	2012-07-01	on going	666,981	115,649	11,710	1,184

C.2 Irregular Period Lengths and Staggered Reads

WaterSmart provides water meter reads for each household in their pilot program, though the data are originally generated from the water utilities themselves. The individual utilities define regularly spaced official read periods. For example, in Aquaville and Watertown, a new read period begins every two months from an initial date of January 1st. Hydroburg, however, begins a new official read period every month. With some exceptions, a water read is taken for each household

---

[1] This initial dataset is after very preliminary data cleaning. Observations before January 2008 have been dropped due to the limited number of households with water reads spanning back before that time. We've also dropped water read observations with extreme outliers: those beyond the median ±28×interquartile range by pilot or those equal to zero.

during every read period. For individual households however, the time period between one read and then next, at which point the accumulation of water used since the last read is recorded, is irregular.

Individual read period lengths are not uniform across households, even within a given pilot. Table C.2 provides selected quantiles of period length by pilot. We determine what defines a short or long period relative to the distribution of period lengths within the pilot. Currently, the high end cutoff for each pilot is set by the median period length plus the inter-quantile range between the 85th and 15th percentile. The low end cutoff is the median less this inter-quantile range. The specific quantiles for these ranges were chosen by inspection of the distributions in Figure C.1. These cutoffs are given in Table C.3 below. We consider short read periods to be a more serious deviation from the norm and run versions of our regressions excluding them. We assume that a utility must have had a reason, such as a leaking pipe, to return after a short period to take another water read and restart the meter for the next period.

Table C.2: Period Length: high and low quantiles by pilot

	Watertown	Aquaville	Hydroburg
0th	12.0	1.0	1.0
1st	54.0	35.0	27.0
5th	56.0	58.0	28.0
25th	60.0	60.0	29.0
50th	60.0	61.0	30.0
75th	63.0	62.0	32.0
95th	66.0	64.0	33.0
99th	66.0	67.0	35.0
100th	155.0	143.0	50.0

Figure C.2 shows the distribution of water reads across the days of the month on which the meter was read by pilot. Notice that in Aquaville the density is much higher towards the end of the month, whereas in Hydroburg and Watertown meter reads are centered near the middle of the month. The staggered and irregular nature of bimonthly meter read periods leads to an even month preference in Watertown, see Table C.4. Only a handful of treatment and control observations are read during odd months. Due to the small number of these odd month meter

Figure C.1: Period Length Distributions: Identifying short and long periods

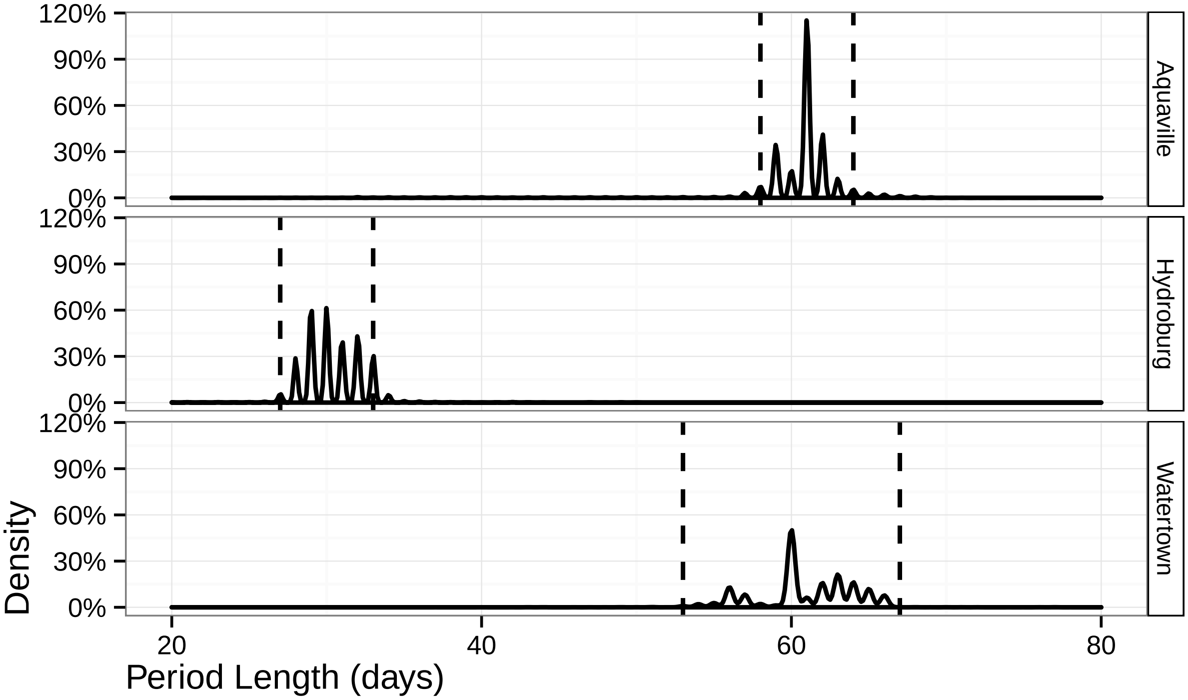

Notes: The graph provides the distribution of water read period lengths. The vertical dashed lines indicate the cutoffs we use to define short or long read periods.

Table C.3: Short and long read periods relative to pilot

Pilot	Cutoff: long period	Cutoff: short period
Aquaville	64	58
Hydroburg	33	27
Watertown	67	53

reads, these observations are easily dropped. This avoids potential problems that may arise if the water reads pulled during odd months are systematically different from reads taken during even months.

Table C.4: Observation counts by month and pilot.

	Control Observations			Treated Observations		
Month	Aquaville	Hydroburg	Watertown	Aquaville	Hydroburg	Watertown
1	4,050	6,626	9	4,046	6,540	9
2	4,215	6,646	2,612	4,184	6,551	2,884
3	3,961	6,676	4	3,954	6,563	2
4	4,423	6,690	2,606	4,379	6,591	2,874
5	4,031	6,716	0	4,007	6,605	1
6	4,503	6,708	3,479	4,450	6,607	3,823
7	4,025	5,628	7	4,046	5,542	7
8	3,827	5,641	3,475	3,758	5,560	3,821
9	3,310	5,642	1	3,302	5,572	1
10	3,784	5,664	3,440	3,746	5,580	3,776
11	3,219	5,673	1	3,201	5,594	3
12	3,639	5,682	3,402	3,604	5,606	3,738

The time based fixed effects are important to control the effect of seasonality on water use. However, for pilots with bimonthly meter reads, if meters read on odd months are systematically different from water reads drawn during even months in a way that is correlated to the treatment group or other cross-sectional covariates of interest, then the year-month dummies will capture some of this cross-sectional variation and corrupt the regression coefficients.

Fortunately, we can avoid these difficulties by instead using utility-defined reading periods. Table C.5 shows that the number of readings are more balanced across reading periods.

Figure C.2: Frequency of water reads by pilot and day of the month

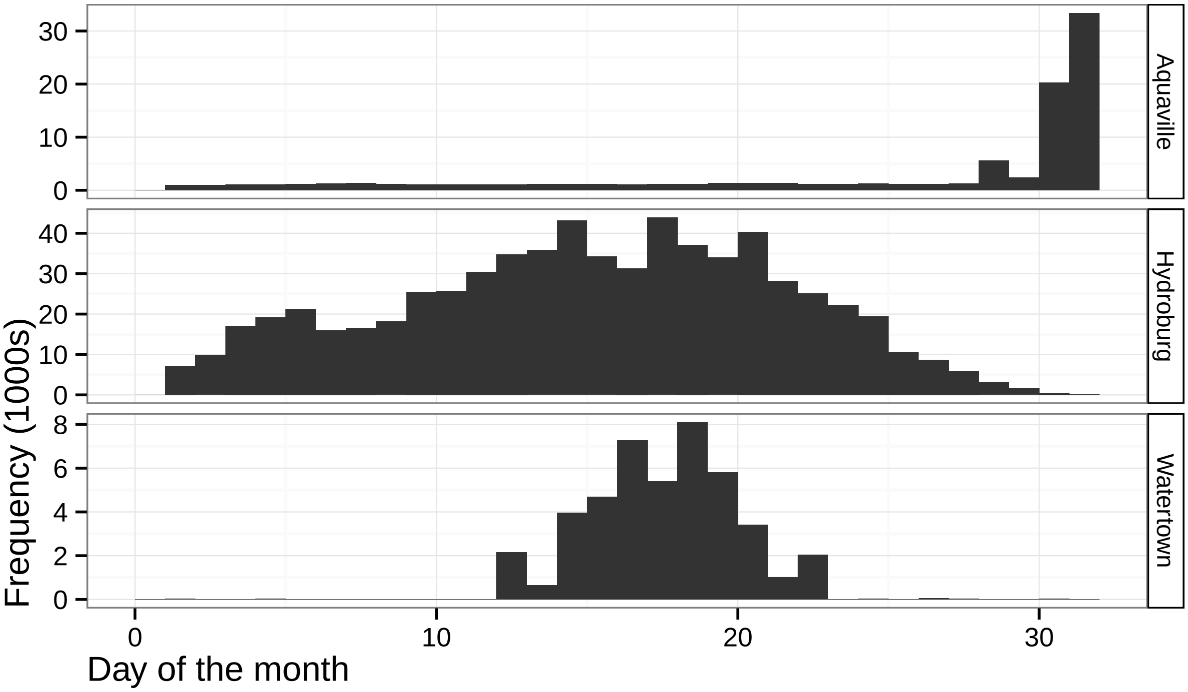

Table C.5: Observation counts by reading period and pilot.

	Control Observations			Treated Observations		
Period	Aquaville	Hydroburg	Watertown	Aquaville	Hydroburg	Watertown
1	8,265	6,626	2,621	8,230	6,540	2,893
2	8,384	6,646	2,610	8,333	6,551	2,876
3	8,534	6,676	3,479	8,457	6,563	3,824
4	7,852	6,690	3,482	7,804	6,591	3,828
5	7,094	6,716	3,441	7,048	6,605	3,777
6	6,858	6,708	3,403	6,805	6,607	3,741
7	0	5,628	0	0	5,542	0
8	0	5,641	0	0	5,560	0
9	0	5,642	0	0	5,572	0
10	0	5,664	0	0	5,580	0
11	0	5,673	0	0	5,594	0
12	0	5,682	0	0	5,606	0

To further ensure balanced samples for panel regressions using year-period fixed effects, we must ensure that each household has only one water read per year-period to avoid overweighting the contribution of individual households during certain periods. Removing short periods will remove many duplicates, but not all. We randomly drop one water read in each household-year-period pair where there are duplicate observations. Table C.6 provides the number of additional observations dropped to balance the panel samples.

Table C.6: Ensuring a balanced panel: identify multiple obs within year-period.

Pilot	Total Obs	Duplicates	Duplicates (no short periods)
Aquaville	93,664	911	324
Hydroburg	666,981	1,995	927
Watertown	44,822	59	6

C.3 Matching Weather Data to Irregular Meter Read Periods

Given the use of year-period fixed effects, the inclusion of weather covariates is needed to control for within read period variation in water use due to seasonality. This is especially necessary for bimonthly pilots. However, given that meter reads are taken at different times throughout the month and are of irregular lengths, matching water reads to monthly or even weekly weather would require many assumptions. We instead pull daily maximum temperature and daily precipitation from National Oceanic and Atmospheric Administration's National Climatic Data Center, specifically from the Global Historical Climatology Network-Daily (GHCN-D) database. We match households to the nearest weather station with available data. Then by looping through each irregular read period, we find the average maximum daily temperature and the average daily precipitation. We also estimate the average cooling degree days (CLDD; the accumulated degrees above 65 degrees Fahrenheit) per day using the daily maximum temperature. We use an average measure of CLDD since CLDD has a better predictive power than average temperature, and we also need to control for the variation in period length, making CLDD per meter read inappropriate. With the same argument we created the average number of rainy days (days with more than 0.01 inches of rain) per day. Exploring other options and data sources, we collected daily evapotranspiration (ET) at the zipcode level from the California Irrigation Management Information System (CIMIS) through their Spatial CIMIS system. ET is the combined loss of water from evaporation and transpiration (like evaporation but also the consumptive water use of plant tissues). For healthy plants including turf grass and gardens, this water needs to be replaced each day. Figure C.3 provides the monthly accumulation of ET from daily averages across the zipcodes covered by the respective pilots. The variation in ET is driven more by solar radiation, relative humidity, and cloud cover, and less by average daily temperature (Hidalgo et al., 2005). For this reason, we suspect it is a better measure to use than average temperature. However, daily ET is missing for several full weeks and months in Watertown, the utility with the fewest households and observations. See Table C.7 for the percent of water reads with missing weather data. For this reason we use average CLDD and average days with precipitation in meter reading periods in our regressions to control for within read period variation in water use. See Figures C.4 and C.5 for pilot level time series of monthly averages of temperature and precipitation. Table C.8 reports pilot level monthly averages of all weather data collected.

Figure C.3: Evapotranspiration

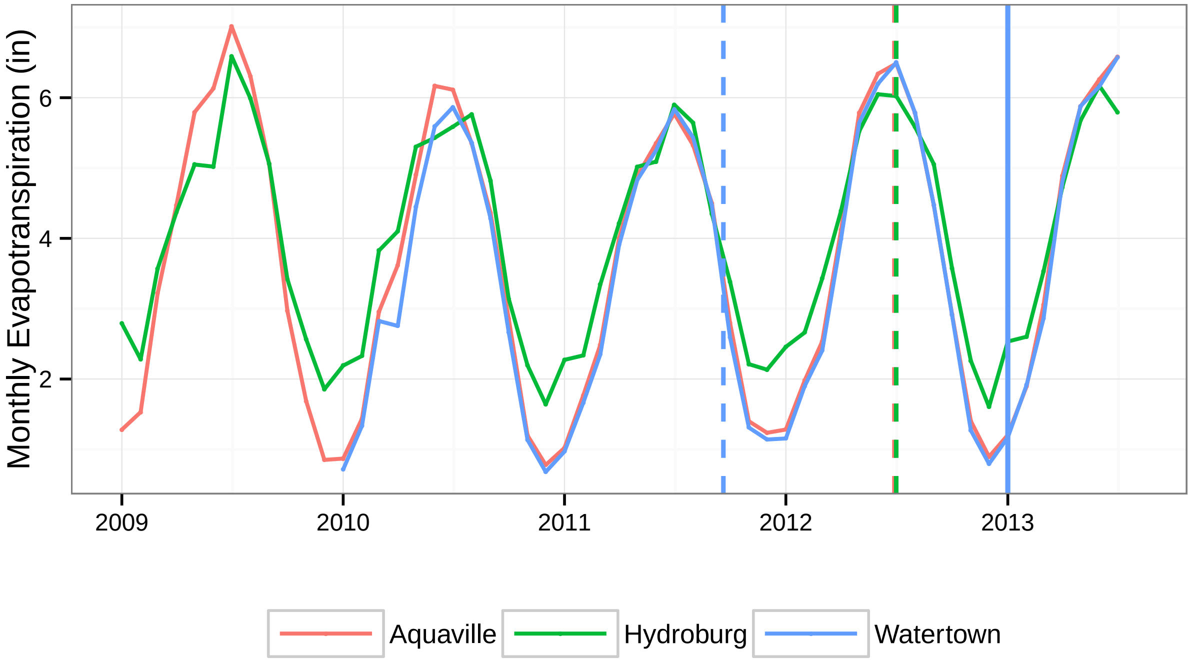

Notes: Evapotranspiration (ET) is the combined loss of water from evaporation and transpiration (like evaporation but from plant tissues). For healthy plants including turf grass and gardens, this water needs to be replaced each day. The graph provides the monthly accumulation of ET from daily averages across the zipcodes covered by the respective pilots. The variation in ET is driven more by solar radiation, relative humidity, and cloud cover, and less by average daily temperature **(SITE: Hidalgo et al. 2005)**. The vertical dashed lines indicate the start of the program for each pilot. The solid line indicates the end of the experiment for Watertown.

Figure C.4: Average Maximum Daily Temperature

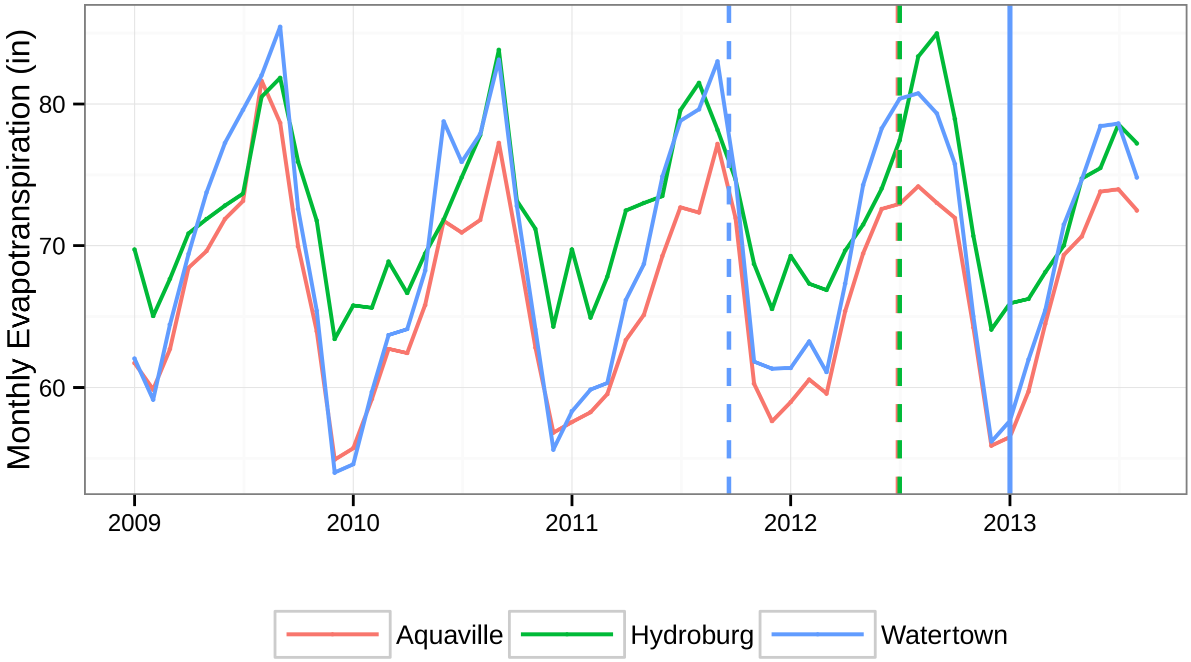

Notes: The graph provides the average daily maximum temperature in the given months across multiple weather stations within the separate pilot locations. The vertical dashed lines indicate the start of the program for each pilot. The solid line indicates the end of the experiment for Watertown.

Figure C.5: Monthly Precipitation

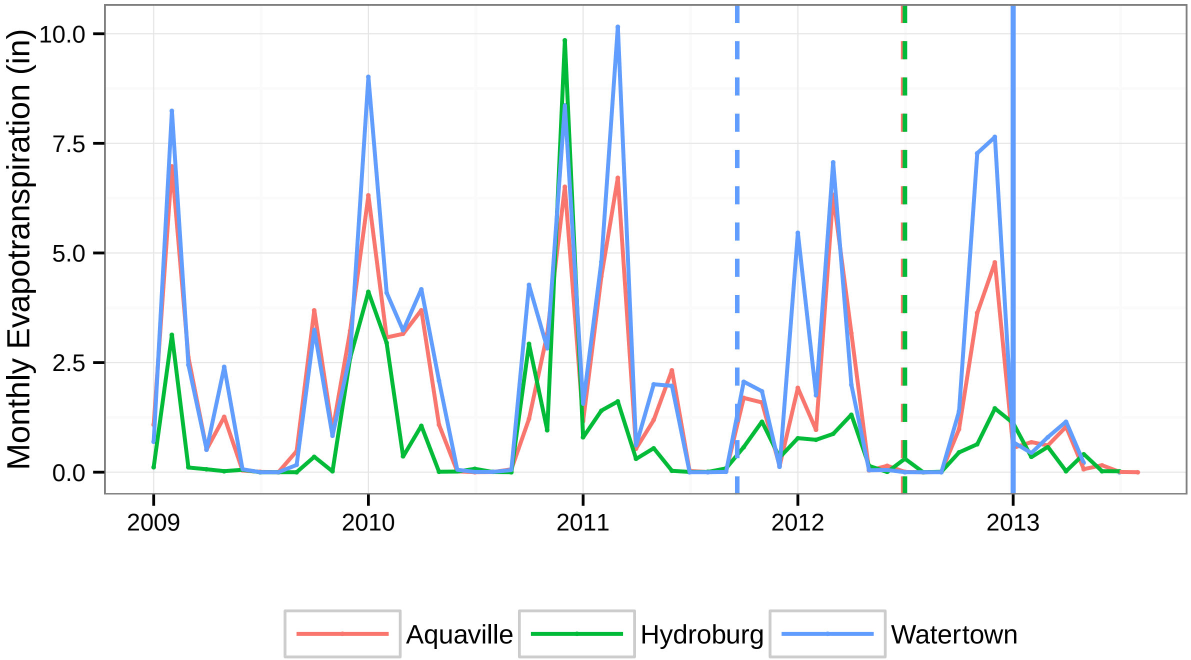

Notes: The graph provides the monthly accumulation of precipitation from the average daily precipitation observed from multiple weather stations within the separate pilot locations. The vertical dashed lines indicate the start of the program for each pilot. The solid line indicates the end of the experiment for Watertown. In Dec 2010, Southern California experienced record storms causing flood and land slides.

Table C.7: Percent of Water Read Observations Missing Corresponding Weather Data

Pilot	Temperature	Evapotranspiration	Rain
Watertown	0.0	13.7	0.0
Aquaville	10.1	0.0	10.1
Hydroburg	0.1	0.1	0.1

Table C.8: Summary Statistics By Pilot: Baseline water use and monthly weather

Pilot	Water Use (GPD)	Temp (F)	CLDD	Rainy Days	ET (in)	Rain (in)
Watertown	202	70	138	5.5	3.5	2.2
Aquaville	282	67	49	7.0	3.7	1.7
Hydroburg	348	72	216	3.9	4.0	0.8

Notes: The water use reported is the mean gallons per day before for treatment across all households; Temperature is the average monthly maximum temperature in degrees Fahrenheit; CLDD stands for cooling degree days and is the average sum of degrees exceeding 65 degrees per month; Rainy days is the average number of days with more than 0.01 inches of rain per month; ET stands for evapotranspiration and is the combined loss of water from evaporation and transpiration (like evaporation but from plant tissues); Rain is the monthly average precipitation in inches.

C.4 Cleaning other outside data sources

We use block level median housing values from Zillow, an on-line real estate company with a database of millions of homes and sales transactions across the US.[2] This is meant to proxy for household income which is only available at the block-group level from the five-year American Community Survey. Also from Zillow's extensive database, we pull median year built and lot size at the census block level to replace missing values of these variables provided at the individual level from WaterSmart. With the exception of lot size in Hydroburg, a vast majority of the values used are at the individual level from WaterSmart. Table C.9 provides the percent of non-missing observations of these covariates that were provided by WaterSmart. We clean all home feature covariates by replacing values with missing if the value exceeds the median ±10×the standard deviation within pilot.

The last external data source are voting records from the Statewide Database hosted by the

---

[2] See http://www.zillowblog.com/research/2012/01/21/zillow-home-value-index-methodology/ for information over the Zillow Home Value Index.

Table C.9: Percent of Non-missing Home Feature Data Provided by WaterSmart

Pilot	Year Built	Lot Size
Aquaville	99.8	99.8
Hydroburg	100.0	0.1
Watertown	90.6	92.5

Notes: The remainder of the non-missing data are referred from block level medians using the Zillow datsbase.

University of California, Berkeley to create ideological indices capturing the environmental sympathies of individual block groups.[3] The Green Ideology Index (GII) is formed using six census block level measures from the 2008 and 2010 California elections:

- the percent who voted Democrat for U.S. Senator in 2008 and Governor in 2010
- the percent who voted yes on Proposition 7 (2008), which would have required California utilities to produce half their electricity from renewable resources by 2025
- the percent who voted yes on Proposition 10 (2008), which would have allocated \$5 billion as cash incentives for high fuel economy and alternative fuel vehicles and R&D for and education on renewable energy and alternative fuel technologies
- the percent who voted yes on Proposition 21 (2010), which would have increased vehicle license fees in the state by \$18 in order to raise \$500 million a year dedicated to California State Parks
- the percent who voted no on Proposition 23 (2010), which would have suspended the state's 2006 Global Warming Solutions Act until the state's unemployment rate decreased to 5.5% for four consecutive quarters, a condition not met for decades.

For each measure, we assign a score near 100 to the top 0.1% of blocks with the greatest percentage of individuals voting along green-friendly lines and a score near zero to the least green-friendly census block. To ensure data quality, we only assign and use scores to blocks with more than thirty votes. Averaging these ordinal scores together yields the GII and maintains the 0 to

---

[3] The data are publicly available at http://statewidedatabase.org/.

100 scale across all of California (the greenest areas are consistently green across all votes). To ensure data quality, we only maintain score averages (index values) where the census block has at least five of the six scores from the separate measures. We match individual households to the GII via their census block as a proxy for their own ideology under the assumption that like-minded individuals tend to cluster around green-related amenities that cater to their preferences.

Tables C.10 and C.11 provide pilot level summaries and data coverage for these descriptive covariates.

Table C.10: Summary Statistics By Pilot: Baseline water use and HH characteristics

Pilot	Water Use	Green Index	Income	Home Value	Year Built	Lot Size
Watertown	202	71	57,333	358,394	1987	5,227
Aquaville	282	62	96,696	670,242	1952	5,600
Hydroburg	348	34	120,943	815,231	1994	5,100

Notes: The water use reported is the mean gallons per day before for treatment across all households; Green Index is the median Green Index across pilot households, ranging from 0 to 100. The household index value is inferred from the Green Index of the census block. See the text for the creation of the block level index; Income is the median income across households in the pilot program. Individual household income is inferred from the average reported income in the census block; Home value is the median home value across households in the pilot program. Individual household home value is inferred from the median home value in the census block; Year build is the median year homes were built across pilot households. When not provided by WaterSmart, year built is inferred from the census block median; Lot size, in square feet, mirrors the logic of year built.

Table C.11: Percent of Households Missing Data

Pilot	Green Index	Income	Home Value	Year Built	Lot Size
Watertown	28.3	0.0	5.4	0.1	8.5
Aquaville	36.2	0.0	2.7	0.0	0.1
Hydroburg	22.3	0.1	12.1	0.1	13.9

Notes: Non-missing data for year built and lot size include those inferred from regional aggregates.

C.5 Water Use

C.5.1 Distribution and Outliers

There are many gallons per day observations that seem beyond reasonable. Table C.12 provides water use by pilot and by selected quantiles. We consider observations to be outliers, and run versions of our regression results excluding them, if the gallons per day exceed the $median + 8 \times interquartile\ range$ (within pilot and season). The assumption is that this procedure only drops water reads due to leaks or broken pipes and not the observations of heavy-handed water users. The outlier counts and water distributions by season are given in Figure C.6. Table C.13 provides seasonal average water use by pilot.

Table C.12: Water Use: High and low quantiles by pilot

	Watertown	Aquaville	Hydroburg
0th	0	5	19
1st	16	25	48
5th	33	51	97
25th	89	123	199
50th	150	196	299
75th	250	338	434
95th	531	811	748
99th	967	1521	1169
100th	4484	5800	7012

Notes: Water use is reported in gallons per day.

Table C.13: Median Water Use by Season

Pilot	Summer	Winter	Ratio
Aquaville	282.0	156.8	1.8
Hydroburg	374.0	249.3	1.5
Watertown	193.5	125.0	1.5

Notes: Water use is reported in gallons per day.

We also consider identifying outliers within households. In this case, there are too few within HH and season observations as of now to determine without household outliers by season. Also,

Figure C.6: Distributions of gallons per day by season

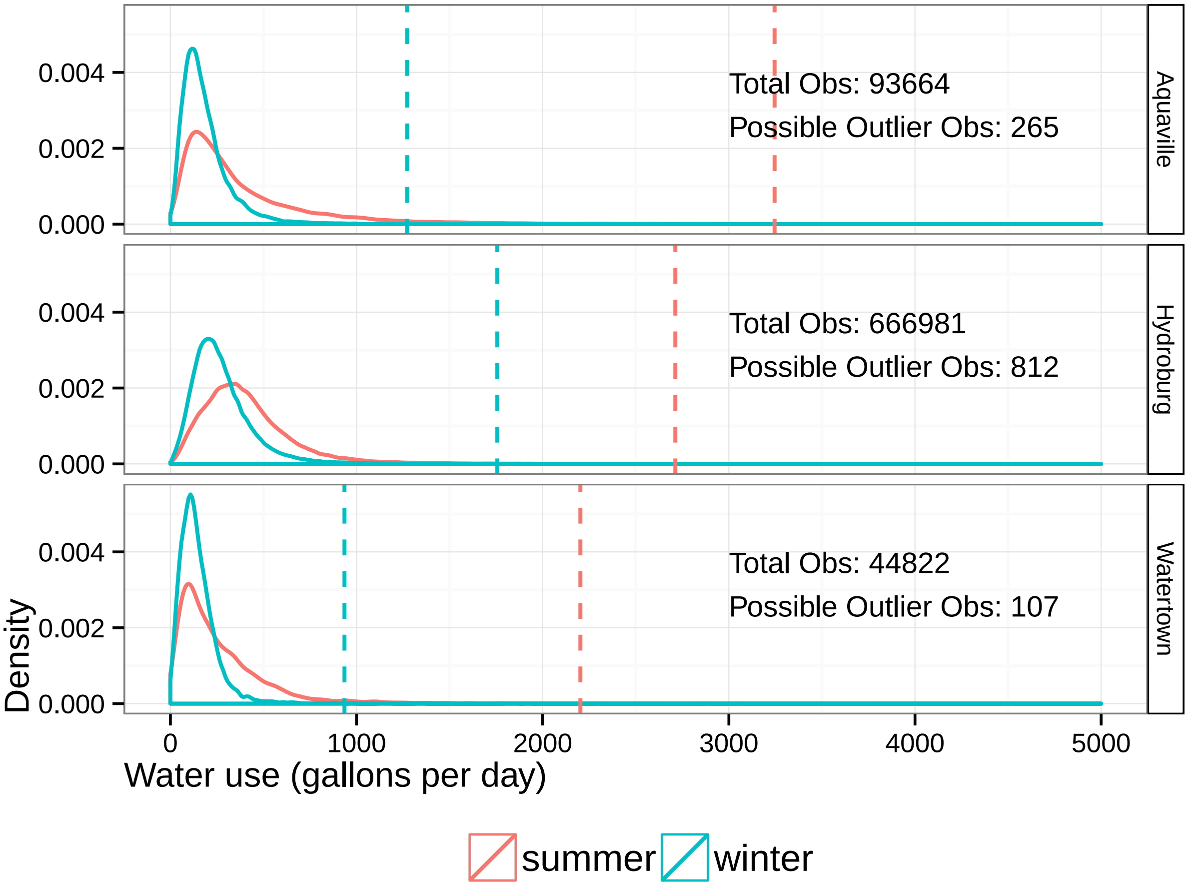

within HH distributions are wider (less n) so we also dial back the requirement. By this we mean that we consider daily water use to be too high or too low for the household if the water use exceeds the household median gallons per day ±6× the interquartile range of the household. Table C.14 and C.15 provide the number of observations defined as outliers for each of the rules described above as well as the number of observations identified under multiple rules.

Table C.14: Outliers in High Water Use (observation counts)

Pilot	Tot Obs	High 4 HH	High 4 Pilot	High 4 Pilot & High 4 HH
Aquaville	93,664	191	265	15
Hydroburg	666,981	821	812	73
Watertown	44,822	157	107	16

Table C.15: Outliers in Low Water Use (observation counts)

Pilot	Tot Obs	Below 20 GPD	Low 4 HH	Below 20 Gallons & Low 4 HH
Aquaville	93,664	217	18	6
Hydroburg	666,981	1	4	0
Watertown	44,822	1,007	13	9

C.5.2 Balance in Baseline Water Use Between Treatment and Control

In all three pilots, water accounts were randomly selected into the treatment group. A randomized field experiment facilitates estimation of the treatment effect by allowing us make the following assumptions. First, no selection bias exists: the receipt of a Home Water Report must not depend on the change in water use before and after treatment. Households in neighborhoods selected to receive HWRs are no more likely than control households to invest in water saving infrastructure or engage in conservation behaviors in the absence of treatment. This could mean that relative to control households treatment households are equally susceptible to peer pressure, equally sympathetic to environmentalist causes, or equally lacking costly information on their costs of water use reduction. Randomization ensures the independence between treatment and any observed or unobserved covariates that may confound the estimation of treatment effects.

However, we may need to perform additional cleaning on the randomized pilots when analyzing heterogeneity in the response to treatment according to home values, environmental idealism, and pre-treatment water use. This is because randomization was not conducted within such subgroups, only across the entire population leaving room for imbalance within subgroups and potential spurious correlations (Ferraro, 2013). Balance refers to the degree in which the covariates' multidimensional distribution of the treatment group resembles the multidimensional distributions of the control group. With more than three covariates, it is very difficult to inspect these multi-dimension distributions. Also, because balance is a property of the observed sample and not of of a hypothetical population there is no hypothesis test for balance. While we can never be sure we've achieved perfect balance, we can find evidence of imbalance by examining the first moments of individual covariates. Table C.16 shows that baseline water use is significantly larger in the control group than in the treatment by almost 14 gallons per day. Home values between treatment and control groups are also significantly different in Watertown, but due to the smaller sample size, Watertown is not a good candidate for matching procedures.

On Hydroburg, we employ a matching algorithm developed by Gary King[4] to produce a control sample as similar to treated households in pretreatment water use and select covariate space as possible to reduce or eliminate the bias from imbalance between treatment and control groups. This algorithm samples from the full set of control observations to minimize the distance in covariate space between the sampled control set and the original treatment group. The specific covariates used for matching are the same observed variables that would facilitate the controlled prediction of the response variable, i.e. they directly effect the response variable as well as effecting treatment. In the case of this analysis which employs Difference-In-Difference (DID) regressions to estimate impacts on water, the response variable is ultimately water use after treatment. Important predictors then include pre-treatment water use (captured by household average baseline water use as well as seasonal averages in baseline water use), irrigable area, and year built (a proxy for the age

---

[4] Documentation can be found at http://gking.harvard.edu/files/gking/files/cem.pdf.

Table C.16: Differences in Means Between Control and Treatment Groups

(a) Variables used to explore heterogeneity

Pilot	Baseline Water	Green Index	Home Value	Income
Aquaville	-1.211 (0.861)	0.3 (0.48)	-12,107 (0.946)	-84 (0.77)
Hydroburg	13.89 (0)	2.7 (0.606)	109,701 (0.979)	3,337 (0.011)
Watertown	-6.541 (0.993)	0.5 (0.195)	-19,930 (0)	-530 (0.21)

(b) Structural Features of the House

Pilot	Sqft	Lot Size	Single Family Home	Year Built
Aquaville	-11 (0.753)	-178.4 (0.116)	0 (0.112)	1.7 (0.538)
Hydroburg	208 (0.292)	217 (0.738)	-0.015 (0)	8.5 (0.495)
Watertown	-32 (0.123)	-124.6 (0.854)	-0.029 (0.126)	1.2 (0.079)

Notes: Values reported are the control mean minus the treatment mean. P-values for the difference from zero are in parenthesis. While we don't use income explicitly to explore heterogeneity, the use of home value is used as a proxy because of its finer geographical granularity.

and water-intensity of appliances). While other covariates may also be important, this set is more readily available for individuals in the Hydroburg pilot. Before running the algorithm, we drop control households without sufficient observations to overlap with the observations drawn from the treatment group.

QQplots help to illustrate the balance achieved by the matching algorithm and can be found in Figure C.7.

After applying this matching algorithm the difference between treatment and control group baseline water use remains, but the magnitude of this difference at the very least is reduced to only 0.151 gallons, see Table C.17.

We also need to assume there are no treatment spill overs between treated and non-treated neighbors. The best we can do to avoid this in these random experiments is to drop the control households who have somehow found their way to the WaterSmart web portal feature without receiving their own HWR. This occurs for only a few households in Watertown. We know this because WaterSmart provides information on their site traffic.

Figure C.7: Quantile-quantile plots to explore effect of matching algorithm on balance

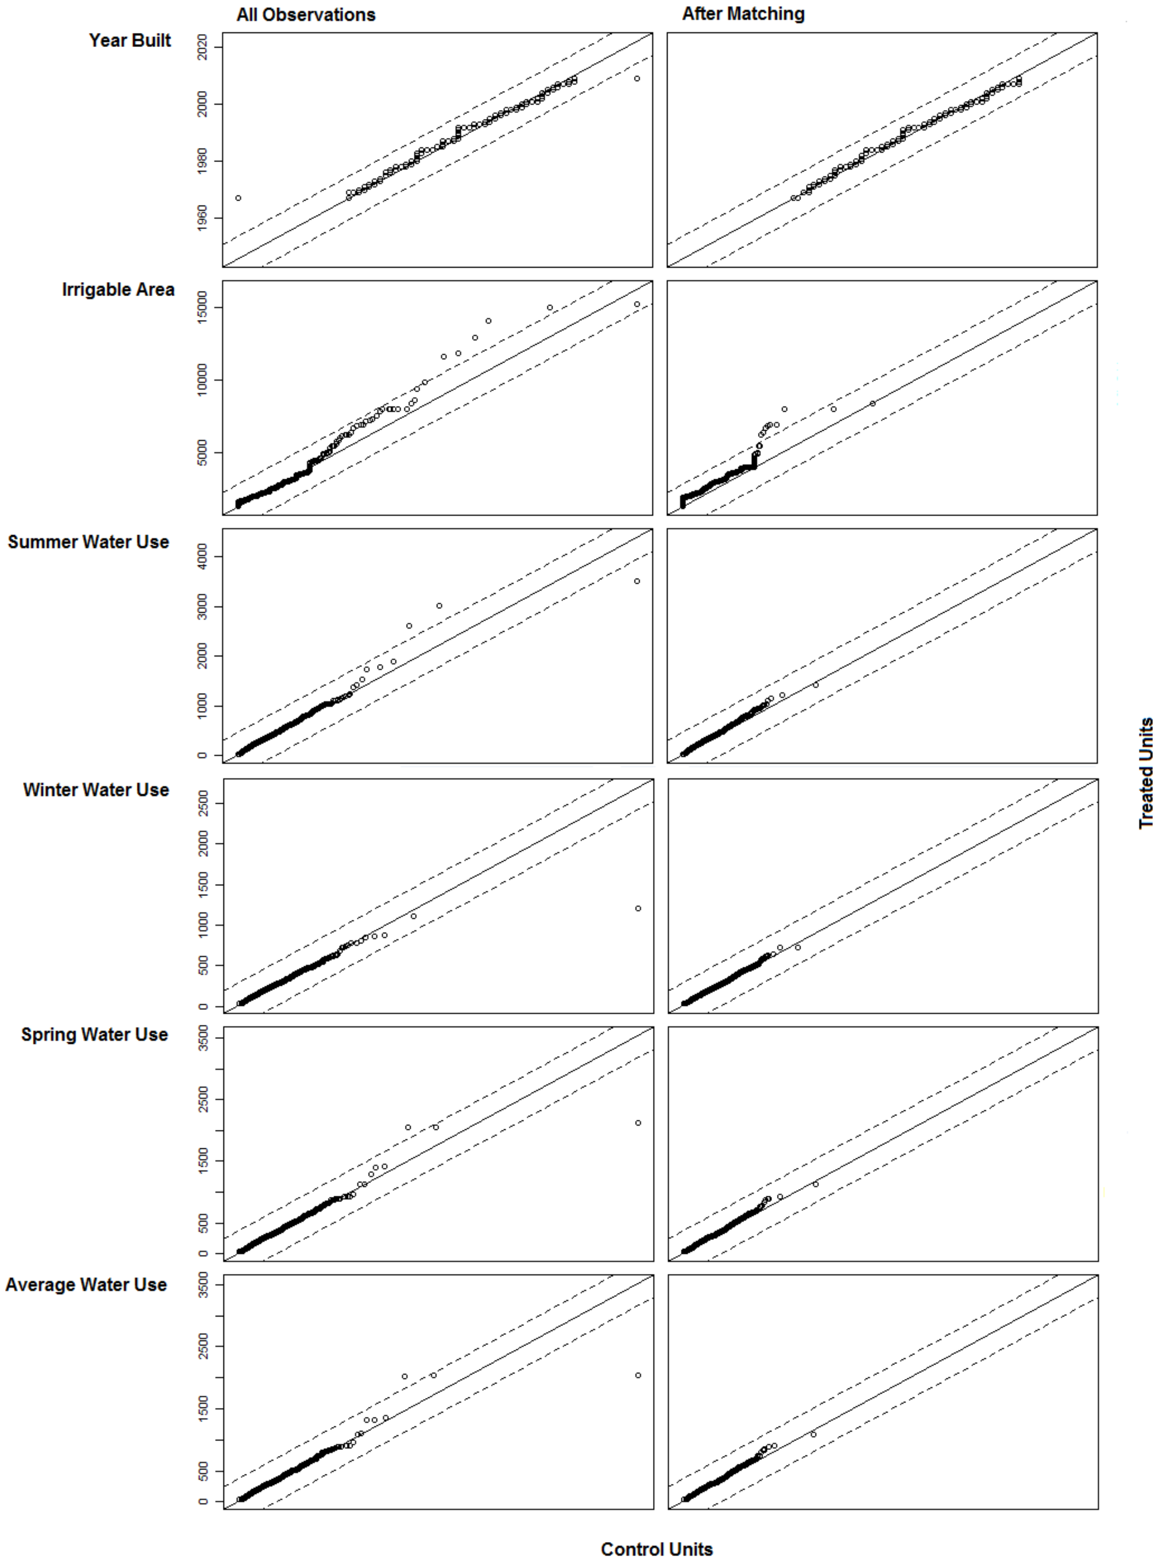

Table C.17: Differences in Means Between Control and Treatment Groups for Hydroburg after Matching

(a) Variables used to explore heterogeneity

Pilot	Baseline Water	Green Index	Home Value	Income
Hydroburg	0.151 (0)	2.1 (0.319)	98,329 (0.651)	2,796 (0)

(b) Structural Features of the House

Pilot	Sqft	Lot Size	Single Family Home	Year Built
Hydroburg	177 (0.204)	75 (0.88)	-0.014 (0)	7.6 (0.843)

Notes: Values reported are the control mean minus the treatment mean. P-values for the difference from zero are in parenthesis. While we don't use income explicitly to explore heterogeneity, the use of home value is used as a proxy because of its finer geographical granularity.

C.5.3 Water Use Time Series

Figures C.8 provides time series of water use in gallons per day and Figure C.9 provides time series of year-over-year percentage changes in water use for each pilot.

Figure C.8: Water Use by Pilot

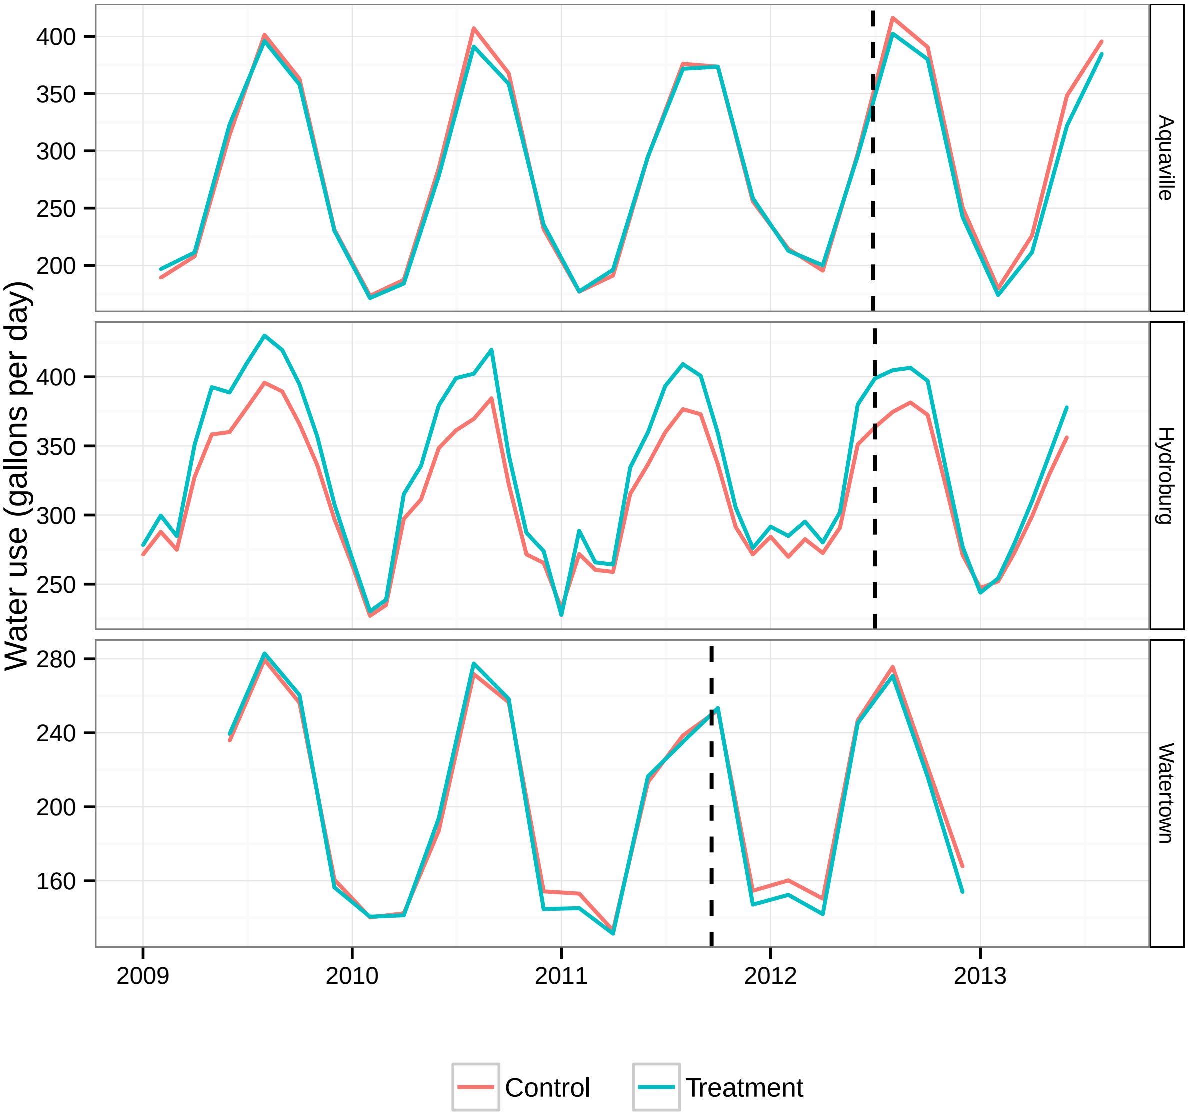

Notes: The graph contains the mean water used per day in gallons for both the treatment and control groups over time. The vertical dashed line indicates the start of the program for each pilot.

Figure C.9: Year-Over-Year Change in Water Use by Pilot

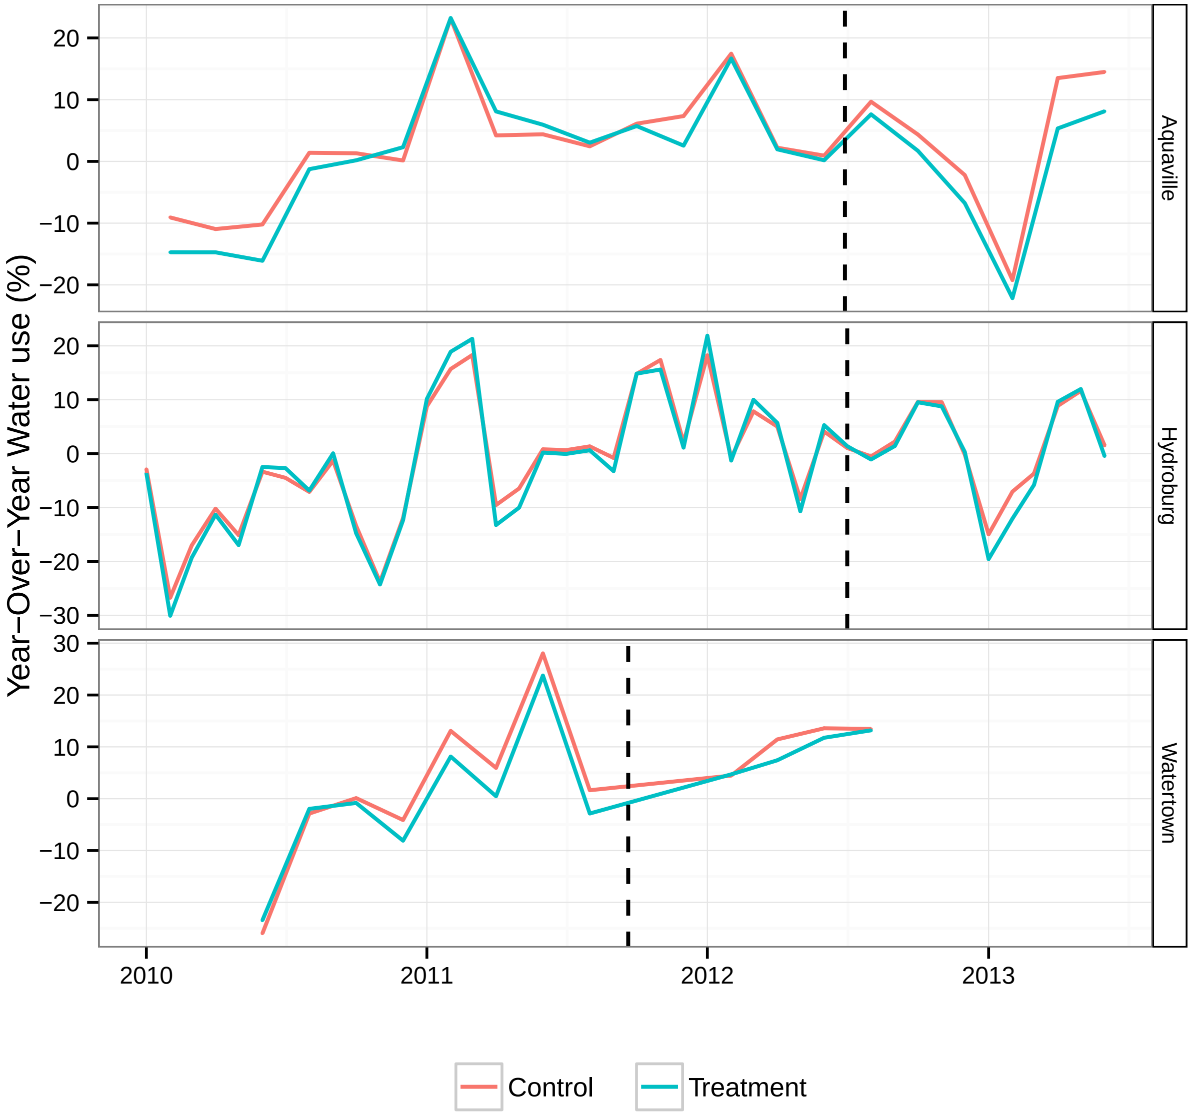

Notes: The graph contains the year-over-year change in water use in percentage terms for both the treatment and control groups over time. The vertical dashed line indicates the start of the program for each pilot.

C.6 Conditional Average Treatment Effects

This section presents the full output for regressions that interact the treatment effect with covariates to identify heterogeneity in conditional average treatment effects. Tables C.18, C.19, and C.20 contain the regression that produce Tables 3.8, 3.9, and 3.10 respectively that display linear combinations for the treatment effect and the treatment effect interaction.

Table C.18: Heterogeneity: Baseline Water Use

	(1) Watertown	(2) Aquaville	(3) Hydroburg
Treatment Effect	-0.0537**	-0.0732***	-0.0509***
	(0.0237)	(0.0195)	(0.0138)
TE*Quintile 1	0.0487	0.0945***	0.0716***
	(0.0338)	(0.0274)	(0.0234)
TE*Quintile 2	0.0433	0.0387	0.0947***
	(0.0291)	(0.0269)	(0.0206)
TE*Quintile 4	-0.0316	0.0371	0.0060
	(0.0274)	(0.0247)	(0.0187)
TE*Quintile 5	-0.0561**	0.0366	0.0094
	(0.0262)	(0.0262)	(0.0200)
Rainy Days	22.1677	-0.3019***	-8.5055***
	(19.9632)	(0.0367)	(0.3943)
Cooling Degree Days	0.2068	0.0067***	0.0013
	(0.6671)	(0.0009)	(0.0047)
Household FEs	Yes	Yes	Yes
Year-Period FEs	Yes	Yes	Yes
Adjusted R^2	0.162	0.251	0.190
Households	2,232	3,091	11,696
Observations	43,541	83,340	664,075

Notes: The dependent variable is the natural logarithm of water consumption in gallons per day. Robust standard errors clustered at the household level are reported in parentheses. Interaction variables are the treatment effect multiplied by quintiles of water consumption averaged over the entire pre-treatment period. The 3rd quintile is omitted, so the base effect can be interpreted as the effect at the median of baseline consumption. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table C.19: Heterogeneity: Ideology

	(1) Watertown	(2) Aquaville	(3) Hydroburg
Treatment Effect	-0.0533*	-0.0458**	-0.0148
	(0.0281)	(0.0231)	(0.0156)
TE*Quintile 1	-0.0149	0.0306	-0.0086
	(0.0330)	(0.0280)	(0.0225)
TE*Quintile 2	0.0262	0.0654**	-0.0531**
	(0.0338)	(0.0284)	(0.0222)
TE*Quintile 4	-0.0101	0.0200	-0.0037
	(0.0383)	(0.0291)	(0.0228)
TE*Quintile 5	-0.0397	-0.0058	0.0236
	(0.0377)	(0.0344)	(0.0258)
Rainy Days	25.7393	-0.3302***	-8.2944***
	(24.1518)	(0.0446)	(0.4408)
Cooling Degree Days	-0.1411	0.0075***	-0.0030
	(0.7071)	(0.0011)	(0.0049)
Household FEs	Yes	Yes	Yes
Year-Period FEs	Yes	Yes	Yes
Adjusted R^2	0.162	0.290	0.203
Households	1,841	1,958	9,024
Observations	32,110	54,869	516,534

Notes: The dependent variable is the natural logarithm of water consumption in gallons per day. Robust standard errors clustered at the household level are reported in parentheses. Interaction variables are the treatment effect multiplied by quintiles of the green ideology index measured at the census block level. The 3rd quintile is omitted, so the base effect can be interpreted as the effect at the median ideology. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table C.20: Heterogeneity: Housing Values

	(1) Watertown	(2) Aquaville	(3) Hydroburg
Treatment Effect	-0.0593***	-0.0372*	-0.0251
	(0.0216)	(0.0216)	(0.0155)
TE*Quintile 1	-0.0069	-0.0333	0.0075
	(0.0308)	(0.0294)	(0.0229)
TE*Quintile 2	0.0422	-0.0330	-0.0026
	(0.0304)	(0.0310)	(0.0234)
TE*Quintile 4	-0.0028	0.0264	-0.0044
	(0.0273)	(0.0259)	(0.0204)
TE*Quintile 5	-0.0010	0.0568**	0.0189
	(0.0262)	(0.0248)	(0.0217)
Rainy Days	25.2592	-0.3032***	-8.9716***
	(20.2693)	(0.0374)	(0.4406)
Cooling Degree Days	0.4395	0.0068***	0.0168***
	(0.6802)	(0.0009)	(0.0050)
Household FEs	Yes	Yes	Yes
Year-Period FEs	Yes	Yes	Yes
Adjusted R^2	0.162	0.255	0.210
Households	2,417	3,011	10,170
Observations	42,368	81,034	584,217

Notes: The dependent variable is the natural logarithm of water consumption in gallons per day. Robust standard errors clustered at the household level are reported in parentheses. Interaction variables are the treatment effect multiplied by quintiles of housing values averaged at the census block level. The 3rd quintile is omitted, so the base effect can be interpreted as the effect at the median average housing value. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$