Title
Essays on Environmental and Resource Economics

Author
Uz, Dilek

Publication Date
2016

Peer reviewed|Thesis/dissertation

Essays on Environmental and Resource Economics

by

Dilek Uz

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

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Sunding, Chair
Professor Sofia Berto Villas-Boas
Professor Gordon Rausser
Professor Catherine Wolfram

Spring 2016

Essays on Environmental and Resource Economics

Abstract

Essays on Environmental and Resource Economics

by

Dilek Uz

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor David Sunding, Chair

In this dissertation, I present three essays that empirically study water and energy economics issues in California.

The objective of the first chapter is to investigate whether and to what extent farmers' crop choice decision is affected by the irrigation water salinity. Using a highly granular land use data and random coefficients logit method, the effect of irrigation water salinity on crop choice is studied in the context of Sacramento-San Joaquin River Delta— California's major water source and home to prime agricultural farmlands. The results show that though the effect of salinity was statistically significant during the past decade, highest and most significant coefficients were those of crop class indicators and weather. This finding suggests that it is essential to reach out to the farmer community to ensure that they are fully capable of coping with expected salinity increases in medium to long run. Additionally, there is evidence for heterogeneity in farmers' response to salinity even though the area studied is relatively small. Ignoring the heterogeneity can result in misleading coefficient estimates especially for those researchers who wish to study farmer behavior in larger regions. Finally, revenue losses are simulated under baseline salinity and potential future salinity scenarios due to building a water conveying facility around the Delta, which suggests an expected revenue loss of about 19%.

In the second chapter, together with Steven Buck, I question the wisdom of selecting a forecast model based on a within-sample goodness-of-fit criterion in the context of commercial and industrial (C&I) water demand in the Southern California. Initially, a set of about 350 thousand regression models are estimated using retailer level panel data featuring water consumption, price, employment, weather variables, and GDP. Out-of-sample forecasting performances of those models that rank within the top 1 % based on various in and out-of-sample goodness-of-fit criteria were compared. We found that the models that provide the best in-sample fit are not necessarily the most favorable ones when it comes to forecasting water demand. The results indicate that on average, these models have a significantly higher absolute forecast error and a larger gap between the highest and lowest

forecasts that they generate compared to the models that rank high based on out-of-sample fit criteria we defined.

Finally, the third chpater investigates the effect of the 2000 California energy crisis on the take up of an engineering audit program funded by the Department of Energy, aiming operational improvements in various domains, including energy efficiency, at small and medium sized firms. Using a detailed data set containing information on both firm characteristics and the specifics of the recommendations made, a linear probability model is estimated using difference-in-difference strategy. In order to keep the treatment and the control groups as comparable as possible to ensure credible identification, the firms that applied to be audited and made the take up decision before the crisis are compared to those that applied right before the crisis and had to decide after the crisis started. The results show that the 2000 California energy crisis was associated with a 16% increase in the take up of the IAC energy efficiency recommendations. The coefficient estimate is statistically significant and robust to different model specifications.

To my mother Sevim Uz and my father Kadir Uz.

Acknowledgements

I would like to thank my dissertation committee: David Sunding, Sofia Berto Villas-Boas, Gordon Rausser, and Catherine Wolfram for the feedback and advice they gave me.

I also thank Steven Buck, Meredith Fowlie, Peter Berck, Patrick Kline, as well as the participants of SABE 2014 conference, workshop organized by Women in Economics and Business group of UC Berkeley, and department seminar at the Agricultural and Resource Economics Department of UC Berkeley for their helpful comments on Chapter 1 of this dissertation.

I received a generous summer grant from the Berkeley Graduate Division which I gratefully acknowledge.

I am forever in Prof. Maureen Kilkenny's debt. Ms. Threatt shared her beautiful house with me. Funkydoor yoga studio helped me stay healthy.

I am thankful to my wonderful friends Maria Klein (my-fair-lady), Olivia Hyun Young, Mia, Hatice Yıldız, May Kosba, Simona Schneider, Sikandra Christian, Lydia Ashton, Gina Waterfield, Lilia Chaidez, Gavin McCormick, and Megan Stevenson. Without them I would be lost.

Last but definitely not least, I thank my family for supporting me in every way imaginable.

Chapter 1: Water Quality and Farming Revenues: The Case of California Bay Delta

1.1 Introduction

Excessive salinity induced by irrigation is one of the biggest and oldest environmental challenges faced by the farmers from all around the world. Naturally occurring salts are dissolved in and carried through surface and ground flows. When used in irrigation, the water itself evaporates but the salts accumulate in the plants' root, gradually impeding its ability to grow. Salinity causes billion dollar losses in farming revenues, renders millions of acres of agricultural land unproductive, and poses a significant threat to food security (Pitman and Läuchli (2002)). How the farmers will adapt in the face of this environmental problem and how their revenues and the crops grown will be affected as a result is crucial especially in a world where drought is expected to become more prevalent due to changing climate.

The objective of this chapter is to investigate whether and to what extent farmer's crop choice decision is affected by the irrigation water salinity. In other words, do the farmers switch to more salt tolerant crops when they face increases in salinity and how does the importance of the role played by salinity compare to other factors?

Using a highly granular panel data on Sacramento-San Joaquin River Delta (the Delta, hereafter) agricultural activity, together with ownership information, weather, and crop revenues per acre, a discrete choice model for agricultural crops is estimated in order to investigate the effect of irrigation water salinity changes on agricultural land use. Subsequently, expected land acreages and corresponding farming revenues are simulated using the parameter estimates and salinity scenarios obtained from DSM II –"a [computerized] river, estuary and land modeling system" used by the California Department of Water Resources.

The Delta is a main water source of California (see Figure 1), providing drinking water for millions of Californians and irrigation water for millions of acres of land. During the past 150 years, it has been altered via draining the marsh areas into agricultural islands that are protected with an extensive levee system. Two major projects, State Water Project (SWP) and Central Valley Project (CVP) that export water to arid regions of California further contributed to the alteration of natural water flowing patterns in the Delta. Though the diverted water brought wealth and prosperity to its destination, aggressive diversions have caused gradual degradation in the ecosystem leading to court decisions that limit the amount of export water to protect the Delta ecosystem.

To meet the "coequal goals" of ensuring the reliability of water supply as well as restoring the damaged ecosystem, the Bay Delta Conservation Plan (BDCP) was formulated. In addition to many habitat conservation measures, the plan also includes construction of an isolated water conveying facility to carry water from the northern part of the Sacramento River directly to the pumping facilities in the southern Delta. The plan raises concerns regarding the quality of the water flowing through the Delta which is used, among other things, for irrigation by the Delta farmers.

There is an extensive literature that studies farmer decision making. An important portion of these studies involve areas such as land use and technology adoption as a function of factors like input prices and environmental factors. This chapter contributes to the current body of knowledge by studying farmers' decision making process as a function of salinity, a factor that is not studied from a behavioral point of view. Additionally, this study expands the existing literature in the following ways: First, instead of cross sectional data, a highly disaggregated panel dataset of crop choice featuring corresponding field and weather characteristics is utilized in the estimation. Having observations from multiple years both makes it possible to control for aggregate year to year changes in crop preferences and to account for path dependence in the crop choice decision. Second, farm level ownership data is incorporated into the analysis which allows utilizing mixed (random coefficients) logit. Mixed logit offers advantages over standard multinomial logit in various ways. In addition to freedom from the "independence from irrelevant alternatives" (IIA) property that imposes restrictive substitution patterns across the alternatives, mixed logit also allows testing whether there is heterogeneity in preferences inflicted by unobserved factors. A variety of forces might bear upon farmers' responses to water salinity changes such as capability to mitigate salinity effectively through other means, whether they need to deliver a certain type of crop under binding contracts, or the way in which they interpret past observations while making current decisions. Hence, a rigorous analysis of farmer behavior calls for accounting for these unobserved sources of heterogeneity and random coefficients logit method provides the necessary econometric machinery for that. To my knowledge, this is the first study that utilizes micro panel data and features mixed logit analysis to estimate farmers' land use decisions.

Estimation results show that although its effect is statistically significant, salinity has not been the primary factor that drove Delta farmers' crop decision in the previous decade. The highest and most significant coefficients are those of crop class indicators and the maximum temperature. From the policy standpoint, it is important to understand the nature of the relative indifference to water salinity among farmers before implementing a major construction that will likely have large and permanent effects on the irrigation water quality. If the reason for farmer's reluctance to react to salinity is due to being fully capable of mitigating the effects, then the damage from impending salinity surges will be more manageable. However, if it is due to suboptimal decision making, they are likely to incur losses and a successful policy implementation should involve effective outreach to the farmer community in order to facilitate optimal decision making.

Second, there is evidence for heterogeneity in farmers' response to salinity even though the area studied is relatively small. Therefore, researchers who wish to study farmer behavior in larger regions, should be aware of potential errors ignoring heterogeneity in behavior may cause.

The results are robust to a different measure of salinity (average of previous 3 consecutive years instead of only previous year's salinity). Additionally, three different measures of goodness-of-fit are provided.

Finally, revenue losses are simulated under baseline and policy salinity scenarios considering possible yield declines. The simulation exercise predicted about a 19% decline in the farming revenues compared to the base case salinity scenario. The losses will vary depending on the farmers' actual crop preferences and their willingness to switch to more salt tolerant species due to high water salinity within an estimated interval of 10 to 30%.

This chapter proceeds as follows: Section 1.2 gives an overview of the salinity problem and an institutional background; section 1.3 provides a review of the related literature; section 1.4 explains the empirical model; section 1.5 provides the details of the data set used; section 1.6 presents and discusses the results of the estimation; and section 1.7 concludes. All the tables and figures are provided in section 1.8.

1.2 Background

Irrigation has been the major driving force of agricultural development, transforming the economies all around the world by providing food security and supporting many industries. The benefits of irrigation are accompanied by major costs such as salinity, water logging, soil erosion, and spread of diseases. Salinity is the most severe environmental challenge faced by the farmers all around the world causing billion-dollar losses in farming revenues and agricultural land (Pitman and Läuchli (2002)). Irrigation water taken from streams carry tons of naturally occurring diluted salts and minerals. Salts remain in the root zone while the majority of the water evaporates either through heat or biological process (evapotranspiration). The effects of this physical process on civilizations can be traced back to ancient Mesopotamia (Jacobsen and Adams (1958)).Unless there is proper drainage to leach these salts away from the root zone, they may accumulate to a point where the salinity starts to impede the growth of the plant. Sometimes, the saline water does leach below the root zone down to the ground water causing the ground water to get gradually saltier and the water table to rise. Rising salty water table eventually hits the root zone and drowns the crops in salty water, a problem called water logging (Oosterbaan (1988)). Even with the existence of drainage facilities, the issue persists as the saline water drained out of the farmland still will need to be properly disposed.

Farmers from all over the world are challenged by salinity induced by irrigation. 30% of the agricultural land in the western side of Andes along the Pacific Ocean in Peru is threatened by water logging and salinity (De La Torre (1987)). In Pakistan, despite millions of rupees were spent to reclaim over 4.5 million acres of land under the threat of serious salinity, the goals could only be partially met (Bhatti (1987)). Massive irrigation following the construction of Assuan Dam forced the Egyptians to install drainage facilities over millions of hectares of agricultural land (Abdel-Dayem (1987)). In Australia, most of the wetland, damp land, forests, and at least 450 species are under serious threat and without a massive amount of intervention, they are destined to extinction (Pannell (2001)).

20 to 25% of irrigated land in US faces yield reduction due to salinity (El-

Ashry, van Schilfgaarde, and Schiffman (1985)). Imperial Valley located in the south east of California has been irrigated since 1901. Accumulated salts and rising water table reached threatening levels within 20 years. In 1922, a $2.5 million bill was passed to build the drainage system that will channel the saline agricultural drainage waters to the Salton Lake.

In 1961, water quality became a major issue in the Colorado River due to increased usage of river water within US and highly saline drainage water from the Wellton-Mohawk Irrigation and Drainage District. Salinity of the water crossing from the Mexican border went up from 800ppm to 1500ppm causing an international crisis between the US and Mexico (Brownell and Eaton (1975)).

The geographical context of this study is the California Bay Delta. The Delta is the most important water resource in California.[1] The construction of the State Water Project and the (federal) Central Valley Project made the Delta the water hub of California. It supplies drinking water for 25 million Californians and irrigation water to 4.5 million acres of agricultural land. Reduced water quality and quantity due to aggressive fresh water exports together with altered flowing patterns significantly degraded the Delta ecosystem. Several fish species including the Smelt, a key species in the Delta ecosystem, came under the risk of going extinct. In 2007, a Federal Court order mandated certain water levels in order to protect the fish.[2] The order limited the amount of water that can be exported from the Delta, especially in dry years.

With water reliability and ecosystem concerns in mind, Governor Jerry Brown recently revived the peripheral canal idea the California voters defeated in 1982. His proposal of constructing a $16B water conveyance facility, also known as the Bay Delta Conservation Plan (BDCP), to divert Sacramento River water underneath the Delta is recently incorporated into the Delta Plan. The facility will divert the fresh water before entering the Delta and will ensure the reliability the water exports to the Bay Area and Southern California. Since it will be "isolated", it will also protect the conveyed water from risks such as levee collapse and sea water intrusion. Additionally, because the natural water flow pattern will no longer be disturbed by the strong pumping facilities, it will allegedly allow the fish population to restore.[3]

---

[1] For an excellent introduction to California water issues visit Water Education Foundation at `http://www.watereducation.org/topic-delta-issues`. "Managing California's Water: From Conflict to Reconciliation" by Hanak et.al. is an invaluable resource providing a detailed institutional history and reviewing California's water (Hanak (2011)).

[2] "Court Finalizes Order to Protect Bay-Delta, Smelt and Water Supply for Millions of Californians" `http://www.nrdc.org/media/2007/071214.asp`

[3] The plan has aspects other than solely changing the the way water is exported from the Delta. Overall, it is a Habitat Conservation Plan (HCT) and a Natural Community Conservation Plan (NCCP) aiming to restore the habitats. It reportedly aims a holistic approach to mitigate the degradation in Delta ecosystem which is threatened by multiple stressors (Snow (2010)). It is developed by the California Department of Water Resources and with assistance from several other agencies including California Natural Resources Agency, California Department of Fish and Game, US Fish and Wildlife Service, national Marine Fisheries Service, US Army Corps of Engineers, State Water Resources Control Board, and the US Environmental Protection Agency (US EPA). $216 million has been spent to develop the draft plan, conduct the environmental review, host numerous public outreach activities, and complete preliminary engi-

However, there are still major objections to the plan. First of all, the financing of the project is critical. The actual costs will likely to be larger than the budget amount when the costs to the stakeholders are factored in. Even if there are no environmental lawsuits, some of the funding depends on the passage of water bonds by the voters. Also, there is a concern that too much water will be taken out of the rivers and will cause drying up of downstream habitats especially in low precipitation years though the authorities assert that the total quantity of the water exported will be within 10% of the average annual amount.

While the construction of isolated conveyance facilities will insure the quality and the reliability of the exported water, taking away the fresh water before entering the Delta is expected to increase the salinity within the Delta. Compromised water quality will likely have ramifications from agricultural, ecological, and municipal aspects. This study focuses on the agricultural implications from the point of view of the Delta farmers.

Delta is composed of many fertile agricultural islands that were originally drained from marshland and are protected from flood via an extensive 1000 mile long levee system. According to California Farmland Mapping and Monitoring Program (FMMP) 80% of the Delta's 500 thousand acre farmland are top tier "Prime Farmland". High salinity has already been issue for the Delta farmers due to combination of San Joaquin River's high salinity combined with rising sea level. Carrying Sacramento River's fresh water before it enters the Delta, raises serious concerns regarding further degradation in water quality in the area.

1.3 Literature Review

Salinity is a well documented subject in the agricultural economics literature. Earlier studies estimated production functions using the data gathered from agricultural experiments. Dinar, Rhoades, Nash, and Waggoner (1991) uses data from experiments to estimate production functions relating yield to water quantity, quality, and salinity for wheat, sorghum and tall wheatgrass. Then, the estimated parameters are used to simulate crop yields in San Joaquin and Imperial Valleys under different salinity conditions. Datta, Sharma, and Sharma (1998) estimates the production function for wheat under different water quality and quantities holding other inputs constant using experimental data from India.

Another set of studies aim to shed light to the prescription of optimal decision making. Dinar, Letey, and Knapp (1985) use the production function parameters for corn and cotton (representative sensitive and tolerant crops, respectively), to compute optimal applied water and associated profits. Using a short run, single-crop optimization model Yaron and Bresler (1970) determines the least cost combinations of water quality-quantity given climate, soil and land use conditions. Yaron, Bresler, Bielorai, and Harpinist (1980) uses a dynamic model for optimal pre-plant leaching and irrigation schedule under hypothetical water supplies with varying salinity levels. Their findings suggest that frequent applications of small

neering and design of the proposed conveyance facilities (`http://baydeltaconservationplan.com/AboutBDCP/YourQuestionsAnswered.aspx#CostFinancing`).

quantities of water is a better method than applying larger quantities at a time with lower frequency. Francois (1982) conducted a field experiment to investigate the possibility of increasing the density of cotton growing on saline soils. Since the plants tend to be smaller under high salinity his idea was that more can be planted per acre of land and he derived optimal water quality, quantity and crop densities given the biological yield relationship. Knapp, Stevens, Letey, and Oster (1990) uses a dynamic optimization model to prescribe the optimal irrigation systems and derives the conditions under which it is optimal to pay for the drainage system.

Agricultural land use in discrete choice setting is studied under climate adaptation context. The idea is to estimate the changes in crop choices as a response to changes in climate variables to be able to quantify the effects of global climate change in the various regions of the world. Using a multinomial logit model Seo and Mendelsohn (2008) studies a cross section of 7000 farmers from 7 countries in South America to estimate which crops are more likely to be chosen by the farmers. Sanghi and Mendelsohn (2008) focuses on India and Brazil while Kurukulasuriya, Mendelsohn, et al. (2007) looks at Africa in a similar fashion. Surprisingly, however, very little is done so far to study irrigation water salinity as one of the environmental decision variables.

In this chapter, farmers' land use decision making as a function of salinity is studied. There is a positive correlation between the crop salinity sensitivity and revenue per acre (see Fig. 2). In other words, the crops that are more sensitive to salinity tend to yield more revenue per acre harvested. This raises the question of whether the increases in the irrigation water salinity will result in a shift in the crop mix towards more salt tolerant hence (usually) less valuable crops.

Farmers who face excessive water salinity can take actions in order to mitigate the negative effect on the yield. The methods include applying excess water in order to induce leaching, switching to micro irrigation, drainage, adjusting fertilizer, crop rotation, and fallowing. The extent to which the farmers will benefit from each of these options will depend on a number of conditions such as the costs of the alternatives, soil quality, and expected prices of different crops. Therefore, the link between yield (or revenue) and irrigation water salinity may not always be given by the relationships that are estimated under experimental conditions and will ultimately depend on what the farmers choose to plant. For this reason, it is more appropriate to study this subject from a behavioral perspective.

To my knowledge, the conference manuscript by MacEwan and Howitt (2012) is the only academic study that looks at farmers' land use decision making as a function of salinity using a behavioral approach. Using a cross sectional data set on California's Kern County, they first estimate a standard multinomial logit model where the salinity coefficients are statistically meaningful. Then under the profit maximization assumption, they find that there is a significant difference between the "behavioral" patterns and the "experimental" patterns of yield changes as a response to salinity. This finding implies that farmer's switching behavior needs to be taken into account when estimating yield changes rather than relying solely on the physiological relationship, in order to get more accurate results.

The methodology in this study differs from MacEwan and Howitt (2012). It features a panel dataset of crop choice and salinity as well as control variables such as maximum temperature, owner size indicators, and information of what was planted the year before on the same land. Additionally, this chapter uses a random coefficient logit (mixed logit) method and the results are contextualized with simulation.

1.4 Empirical Analysis

A major empirical challenge in estimating the effect of salinity on agricultural crop choice is that hard to observe farm characteristics that play into crop choice may correlate with salinity levels. If we can assume that farm characteristics tend to stay relatively stable over time, and irrigation water salinity varies from year to year, in order to be able to credibly identify salinity response parameters, we need to allow salinity level to vary within a farm and observe the decision making under different conditions. Observing the farms for multiple years will also allow accounting for year to year aggregate changes in farmers preferences via year indicators and the path dependence in farmers' decision making.

Another potential challenge is that farmers might differ in their response to changing environmental conditions. In this study, ownership information is incorporated into the analysis[4] in order to control for owner specific unobservable factors that could be potentially correlated with the variable of interest (water salinity) and the outcome (crop choice).[5]

As a preliminary analysis, two sets of regressions are run prior to the mixed logit estimation. The first one involves regressing the salinity sensitivity of the chosen crop (proxied by the slope of the salinity-crop yield curve) on the previous year's water salinity. The yield of a crop that is more sensitive to salinity will decline faster as the water salinity goes up and hence will have a steeper slope (larger in absolute value). For most crops, the salinity-yield relationship is established by agricultural scientists. The studies from which this information is obtained are Hanson, Grattan, and Fulton (1999), Hoffman (2010), and Maas and Hoffman (1977) (See the appendix for the crops and the salinity tolerances).The regression results can be found in Table 2. We see that the salinity sensitivity of the chosen crop tends to decline as the previous year's irrigation water salinity goes up. In other words the farmers tend to choose more salt tolerant crops as they observe increases in the irrigation water salinity. All three models include owner indicators. As we sequentially add the year and the conservation zone indicators, we see that both the magnitude and the significance of the salinity coefficient go up while the magnitude and the significance of other variables remain somewhat stable.

A similar analysis is presented in Table 3, where the revenue per acre is the dependent variable. Here we do not observe a clear pattern in the sign, magnitude, and the significance of the salinity coefficient. These findings suggest that the

[4]It is assumed that farmlands owned by the same entity/person are managed by the same decision maker.

[5]When estimating random coefficient logit, this is done by setting the coefficient for the variable of interest same for same decision makers in the optimization algorithm.

farmers respond to irrigation water salinity by switching to more salt tolerant crops but it is hard to claim with any reasonable of confidence that the crops they switch to are necessarily low value. In other words, data suggests that farmers watch out for water salinity in crop choice while at the same time protecting their revenues.

The regressions are followed by a discrete choice model estimation. Crop choice is conceptualized with a simple decision framework. The farmer first observes salinity, weather, and crop prices before she makes a decision. Then, she chooses the crop class that maximizes the benefits given the physical characteristics of the land and her observations.

Let U_{ijknt} denote i^{th} farmer's net benefits from choosing crop class j, for farmland n, in conservation zone k at time t.

The researcher does not observe the farmers' individual preferences but farm level physical characteristics, weather, revenue per acre, and salinity are observed. The total value of the crop to the farmer is partitioned into observed (V_{ijknt}) and unobserved (ε_{ijknt}) components such that $U_{ijknt} = V_{ijknt} + \varepsilon_{ijknt}$.

The unobservable portion is assumed to be independently and identically distributed Type I extreme value (Gumbel) the density of which is given by
$f(\varepsilon_{ijknt}) = e^{\varepsilon_{ijknt}} e^{e^{-\varepsilon_{ijknt}}}$.

The difference of two extreme value variables is a logistic variable and the probability of alternative j being chosen by decision maker i is given by[6]:

$$P_{ijknt} = \frac{e^{V_{ijknt}}}{\sum_{l \in J} e^{V_{ilknt}}},$$

where J represents the choice set (Train (2009)). This is also known as the "logit probability".

The observed component of the total benefit of farmer i is modeled as

$$V_{ijknt} = \beta_{\mathbf{1j}}\mathbf{F_n} + \beta_{\mathbf{2j}}\mathbf{O_i} + \beta_{\mathbf{3j}}\mathbf{W_{n(t-1)}} + \beta_{\mathbf{4ij}}\mathbf{S_{n(t-1)}} + \beta_{\mathbf{5}}\mathbf{C_{jn(t-1)}} + \mathbf{T_{jt}} + \mathbf{K_{jk}}$$

for each crop class j in the choice set where,
$\mathbf{F_n}$ is the vector of farm characteristics such as soil quality, size, elevation, and slope;
$\mathbf{O_i}$ is the vector of owner size categories;
$\mathbf{W_{n(t-1)}}$ is the previous year's weather measurement;
$\mathbf{S_{n(t-1)}}$ is the previous year's salinity measurement;
$\mathbf{C_{jn(t-1)}}$ is the vector of time varying crop choice characteristics such as revenue per acre and path dependence. Path dependence here is an indicator variable that is equal to 1 if option j was picked by the farmer in the previous period and zero

[6]The reason why we take the difference is straightforward. When making a voice, only the differences in the utilities matter. Therefore, in order to simplify the analysis, one of the alternatives' (which constitutes the baseline) utility is set to zero by subtracting it's utility form each of the alternative including itself. As a result, instead of N (number of alternatives) errors we are left with N-1 error differences.

otherwise.
T_{jt} is the year specific constant;
K_{jk} is the conservation zone specific constant. Conservation zones are regions in the Delta with similar ecological characteristics.[7]

This model is estimated both using a random coefficients logit (mixed logit) and the standard (fixed coefficients) multinomial logit model. The latter is a widely used method to estimate discrete choice decision parameters in the economics literature. Studies that utilized this method to analyze agricultural crop choice include Sanghi and Mendelsohn (2008), Seo and Mendelsohn (2008), Kurukulasuriya, Mendelsohn, et al. (2007), and MacEwan and Howitt (2012). The properties are well understood and the implementation is straightforward. However, mixed logit offers major advantages relative to standard multinomial logit making it a worthwhile modeling option. First of all it provides freedom from rather restricted substitution patterns also known as the independence from irrelevant alternatives (IIA) property caused by the ratio of probabilities of two alternatives being only dependent on their own attributes (Train (2009)).

Second, logit model forces the parameters in the value function to be same for every decision maker. Therefore, a change in a certain attribute will necessarily affect every decision maker's valuation of the alternative in the same direction and by the same amount. With mixed logit on the other hand, the parameter is treated as a random variable with a certain distribution and the mean and the standard deviation of this distribution are estimated in the maximum simulated likelihood procedure.

In this study, the salinity parameters are modeled to be normally distributed across the farmers. In other words, $\beta_{4ij} = \bar{\beta}_{4j} + \xi_{ij}$ and $\xi_{ij} \sim N(0, \sigma_j)$. It should be interpreted as follows: the contribution of salinity to the value that farmer i gets from alternative j is a normally distributed random variable with mean $\bar{\beta}_{4j}$ and standard deviation σ_j.

The estimation procedure can be summarized as follows: Asume, for simplicity, that we only have one parameter, β. The logit probability formula is valid for a given value of β. Now suppose, β itself is a random variable, then the logit probability will be a function of this random variable. Hence, the "unconditional" probability of a certain alternative will be given by $\int L(\beta) f_\theta(\beta) d\beta$ where $L(\beta)$ is the logit probability and $f_\theta(\beta)$ is the density of β given the parameter vector θ. We can think of the mixed logit probability as the weighted average of the logit probabilities, where weights are determined by the probability density function of β evaluated at each point. In other words, we are *mixing* different logit

[7] Notice that in order for the coefficients of the variables that do not vary across alternatives to be identified, they need to be specified as alternative specific. For example, a farm specific attribute will contribute to the benefit of each option differently.

distributions via a normal distribution.[8]

By the same token, conditional on β, an individual's decision sequence will have the probability $S(\beta) = \prod_t L_t(\beta)$. The integral of this probability over the density of β, $\int S(\beta) f_\theta(\beta) d\beta$, will give the unconditional probability. Since the integral of normal distribution does not have a closed form, this quantity needs to be simulated. The simulated probability is $SP_\theta = (1/R) \sum_r S(\beta_r;\theta)$, where R is the number of draws taken from the underlying distribution. This clearly is an unbiased estimate of the actual unconditional probability. It is smooth which paves the way for maximum (simulated) likelihood estimation. Under regularity conditions, maximum simulated likelihood is consistent and asymptotically normal (Hajivassiliou and Ruud (1994)) when the number of draws increases faster than the square root of the sample size. Also it is asymptotically equivalent to maximum likelihood estimator.

In the estimation procedure, the simulated log-likelihood function is constructed by adding up the natural logarithms of the individual simulated sequence probabilities and population parameters are estimated by maximizing this log-likelihood over θ.[9] Since this is a non-linear transformation, the expectation the log of simulated probabilities does not equal the log of actual probabilities. The bias vanishes as the number of repetitions increases (Revelt and Train (1998)). The standard errors are calculated using the BHHH methodology (named after the authors of Berndt, Hall, Hall, and Hausman (1974)) by taking the square root of the diagonals of the "Hessian" matrix calculated by the outer product of the gradients of the log-likelihood function with respect to the parameters. This method offers a significant advantage as it does not require additional computing besides the (already computed) gradients (Greene (2008)).

Allowing the coefficients to come from a random distribution incorporates the heterogeneity across the farmers. Note, however, that the researcher stays agnostic to the source of this heterogeneity. It could be caused by a number of factors such as differences in farming approaches, management practices, binding contracts, and the way the expectations are formed based on past observations. Therefore, it is worthwhile to use an econometric method which is capable of handling heterogeneity inflicted by these hard to observe characteristics.[10]

The salinity measures taken during the previous year's irrigation season is the variable of interest here. The identifying assumption is that after controlling for year, conservation zone, farmer size, weather and the farm attributes like elevation, soil quality, and size, the previous year's average irrigation season salinity level

[8]Note that the random coefficients do not necessarily have to be normally distributed. For example, if the researcher does not desire negative (positive) values for the parameter, she can truncate the support or perhaps pick a distribution with a strictly positive (negative) support. Ben-Akiva and Bolduc (1996) prefers the term "Multinomial Probit with Logit Kernel" to express specifically that the preferences are normally distributed but the idiosyncratic errors are extreme value.

[9]Programming codes written in MATLAB by Prof. Kenneth Train are modified to this project.

[10]The virtue of including ownership information also becomes evident in this context as the estimation machinery uses the same coefficient for the same owners in the sample when calculating the likelihood function, the maximand of the maximum likelihood.

is as good as random. Table 4 summarizes how the salinity varies across and within conservation zones. The sizes of the within conservation zone standard deviations of the salinity relative to the means imply a significant amount of variation in salinity levels within each conservation zone. Table 5 shows the results of regressing salinity on the controls. We can see that after regressing the irrigation water salinity on all the controls previously mentioned, about 20% of the variation in the salinity remains unexplained. In figures 5 and 6 we see the histogram of the size of the absolute residuals from the regressions presented in Table 5 relative to the mean salinity levels at farm and conservation zone level, respectively. These figures suggest that the size of the identifying variation in salinity across farms is greater than 20% of the mean values observed both at farm and conservation zone levels. In other words, we have a sizable variation in salinity with which the salinity parameters are identified.

Note that panel data allows observing the same farm over different years and varying salinity levels. However, farm fixed effects are not being utilized here as the number of parameters to be estimated increases rather fast with the number of right hand side variables. The issue is resolved by including time invariant farm level characteristics such as slope, soil quality, elevation, and size which will allow accounting for fixed farm characteristics that are expected to effect crop choice.

1.5 Data

This study features a unique and rich data set that is compiled from a variety of different sources. The farmland level crop choice data came from the county agricultural commission offices. California Department of Pesticide Regulation requires farmers to report any pesticide use to the county offices. The counties digitally map the farm lands to form a mosaic of agricultural fields which gives the coordinates of the centroids and the size of the farmland. This data provides information on agricultural activity at the highest possible granularity for the vast majority of the Delta. For the small percentage of the fields where pesticide were not used, the crop choice was estimated using satellite remote sensing data from National Agricultural Statistics Service (NASS).

The study is confined to those cases where there is actually change in the crop species. In the cases where the crop is repeated it is assumed that no active choice had been made. For example, if the farmer has an almond orchard, she will keep them for a long time and clearly no crop decision takes place for many years. Also if a farmer has a binding contract that forces her to grow a certain kind of crops, once again it is safe to assume that the no active crop choice is taking place at the time of planting. Those cases where the farmer actively decides to keep planting the same crop are hard to isolate so all the repeated crop cases removed from the data sample. With this elimination, the total number of choice situations went from 24,639 (285 thousand unique acres) down to 7,650 (166 thousand unique acres).

The choice set is defined based on the agronomic classes which are deciduous, field, grain, pasture, truck, and vineyard. The classification is taken from the Economic Sustainability Plan (Delta Protection Commission, 2012). Though it is not strict, the crops in the same agronomic class could be considered to require similar management approaches and technologies.[11] The list of the crops, their agronomic classes, salt tolerance classes (obtained using Hanson, Grattan, and Fulton (1999), Hoffman (2010), and Maas and Hoffman (1977)) can be found in the appendix.

While the correspondence between salinity tolerance classes and the agronomic crop classes is not one-to-one, the breakdown of the agronomic classes into salinity tolerance classes by their acreages show that the deciduous crops in this study are mostly sensitive; field, truck, and vineyard crops are medium sensitive; grains are mostly medium tolerant; and pasture crops are mostly tolerant (see Figure 4).

Parcel level land ownership data is available through the UC Berkeley Earth Sciences and Map Library in GIS format. The data is collected from the counties and commercialized by Boundary Solutions Inc. In this dataset, for the majority of the records the assessment year was indicated as 2004. Ownership is assumed to be fixed throughout the entire time period that this chapter studies (2003-2010). In order to verify the validity of this assumption the land transaction data compiled by Dataquick is used for a cross-check. According to the data, total amount of farmland that changed hands between 2001-2008 in the San Joaquin county, which constitutes the biggest chunk of land in this study, was about 10 thousand acres. The total size of farmland from the San Joaquin county in this study is about 210 thousand acres. So less than 5% of the agricultural land in San Joaquin county was sold within about a decade.

Ownership size is a categorical variable constructed by adding up the acres by owners and dividing the owners into three bins based on their total size. Using the 33rd and the 66th percentile, owners with less than a total of 48 acres, between 48 and 156 acres, and greater than 156 acres were classified as 'small', 'medium', and 'large', respectively. 90% of the farmers included in this study own less than 520 acres.

Salinity (measured in electro-connectivity, micro Siemens/cm) data comes from 40 water-monitoring stations that are maintained by different organizations like Interagency Ecological Program (IEP), California Department of Water Resources, the US Bureau of Reclamation, and US Geological Survey. The average value of the observed salinity between the months of May and August is used, as this is the period when the crops are most sensitive to the changes in salinity. Then, the salinity values are mapped to individual crop fields by averaging the salinity values measured from the stations that are within 3-mile radius of the center of the field. If a field does not have a monitoring station within 3 miles then the salinity from the nearest station was used.

The soil Storie index is created by the National Resources Conservation Service (NRCS). It varies from 1 to 100, where a score of 100 represents the highest quality. The measure includes factors such as permeability—the measure of the soil's ability to hold water— and its acidity (pH).

Weather data comes from the PRISM Climate Group at the Oregon State

[11] Personal communication with Michelle Leinfelder-Miles, PhD, farm advisor at University of California Division of Agricultural and Natural Resources.

University. The numbers are provided in a continuous grid estimated from point measurements. Annual data at 800 meter resolution was used. Slope and elevation data comes from US Geological Survey extracted based on the coordinates of the centroids of the farm lands.

Revenue per acre values are the weighted averages calculated using county crop reports where the weights are determined by the acreages or each crop in a given category. All the values are converted to 2010 dollars using the PPI for farm products from the Bureau of Labor Statistics. Acres and salinity are scaled by 1/100, and revenue per acre is scaled by 1/1000.

Farms that were less than 10 acres and those that were in remote conservation zones were excluded from this study. (Conservation zones are areas with similar ecological characteristics.) These areas correspond to about 10 percent of the total choice situations but in terms of acreage, they account for less than 1% of the farmlands.

As of 2010, a total of about 278 thousand acres were planted. There are 49 different different crops in the mix (not counting the subtypes of the crops). Field crops had the largest acreage with 156 thousand acres followed by truck crops with 43 thousand acres. Top 5 crops were corn (88 thousand acres), alfalfa (59 thousand acres), vineyards (24 thousand acres), tomatoes (21 thousand acres), and wheat (20 thousand acres).

The total size of the farm land included in the study is about 166 thousand acres (down from a total of 284 thousand acres after the elimination of the cases where the crop choice is repeated). Overall, the data comes from 3 counties (San Joaquin, Yolo, Sacramento)[12] that are made of 7 conservation zones, from a total of 3,245 farmlands, owned by 772 distinct farmers, spanned from 2003 to 2010. A total of 99 parameters are estimated using 7650 choice situations.

1.6 Results

The model is estimated using a maximum simulated likelihood method in MATLAB.[13] Afterwards, the expected acreages of each crop class under different salinity scenarios (baseline vs. policy) provided by the Delta Simulation Model II (DSM II) are simulated.

The mixed logit results are presented in Table 6. Most variables studied here do not vary over alternatives, therefore the estimation had to be done with the interaction of the variable and the alternative indicator. Hence, for each variable there are 5 different parameter estimates i.e. the coefficients are alternative specific.[14]

These numbers are the estimates of the parameters of the value function of each crop class, pasture being the outside option. In other words, they show how

[12] San Joaquin County: 2002-2010, Sacramento: 2010, and Yolo: 2007-2010.

[13] For a rigorous treatment of maximum simulated likelihood see chapters 6 and 10 of Train (2009).

[14] Although path dependence does vary over the alternatives, the coefficients are also estimated to be alternative specific because doing so significantly improved the maximum likelihood estimation.

each factor's effect on farmers' benefit from each option *relative* to pasture class. The salinity parameter is assumed to be randomly distributed across the Delta farmers. This means that for a given alternative, the contribution of salinity to the value function follows a normal distribution over the farmers. The mean and the standard deviation of this distribution are the parameters of interest which are estimated in the maximum simulated likelihood algorithm. The statistically significant standard deviation estimates show that the data supports this view of heterogeneity.

Previous year's salinity was a statistically significant factor in Delta farmers' crop decision during the previous decade. The relative magnitude of the parameters indicate, however, that salinity has not been the most influential factor. Factors such as farm size (for deciduous crops), weather (maximum temperature), and most importantly, unobserved crop characteristics which are captured by alternative constants were more important determinants. Estimation procedure yielded very large negative coefficients for the alternative constants. These numbers likely represent the effect of the cost of planting each crop category. Larger coefficients for deciduous and vineyard are consistent with the sizable upfront costs of planting and maintaining these valuable crops.

Revenue per acre varies over the alternatives which obviates the need for the estimation to be done in an alternative specific way. It is positive and highly significant, and for most crops, it is a more important factor relative to salinity. The relatively small magnitude can likely be attributed to the presence of alternative specific constants.

Path dependence is an important factor determining the farmers' crop choice. We see that for farmers who already planted a deciduous crop are slightly more likely to plant another deciduous crop when they switch crops. For field, grain, and truck crops, that they are less likely to pick a crop in the same category. This could be due to deciduous crops requiring different and very valuable expertise while for the latter group a crop rotation regimen might be preferred. The estimation of this parameter for the vineyard class was not possible due to lack of observations pertaining to staying in the same category.

Table 7 presents the estimates of the standard multinomial logit (MNL) version of the same model. In this version, all coefficients, including salinity, are modeled to be fixed across the farmers. The coefficients estimates for salinity here are very different than the ones we see in the mixed logit table. For example, for field crops and grains, the standard MNL estimates the salinity parameter to be statistically indistinguishable from zero for all farmers while in the mixed logit results suggest that this coefficient is positive for some farmers while it is negative for others. Furthermore, the magnitude of the of these coefficient estimates turned out to be very different.

The results point to an important amount of heterogeneity regarding the effect of water salinity on farmers' choice in a relatively small area with relatively homogenous physical characteristics. This suggests that studies focusing on larger areas with diverse physical and climatic properties should take this into consideration to ensure the reliability of their results. This issue is particularly relevant as we expect to see more research being done studying farmer behavior in the verge of impending dramatic environmental changes.

Next, the expected acreages are simulated under the two scenarios generated by DSM II. For this simulation exercise, only choice situations from 2010 are used. Additionally, those farms that have a high likelihood of being lost to urbanization according to UC Berkeley Resilient and Sustainable Infrastructure (RESIN) data are excluded from the simulation. Figure 3 depicts the average salinity figures that are observed in each cross sectional unit and the salinity scenarios that are estimated by the hydrological model. The salinity projections are within the realm of the already observed figures for the farms that are included in this study.[15]

Columns 2 and 3 of Table 9 show the expected acreages under two scenarios. We see that they are virtually identical due to estimated salinity coefficients being very small in size. In this simulation exercise, in order to also construct simulated confidence intervals[16], taking draws are done in two steps. First, 1000 draws are taken from the variance-covariance matrix generated in the maximum simulated likelihood estimation process. I call these the "outer" draws. Each time, we get a draw from a multivariate normal distribution with mean and variance given by the vector of the parameter estimates and their variance covariance matrix, respectively.[17] These include draws for the means and the standard deviations of the coefficients that are modeled to be random. Afterwards, for the random coefficients 500 "inner" draws are taken from the distribution implied by each outer draw. Each one of these draws will give an expected acreage and revenue for each choice situation through the logit probabilities. For revenue per acre, 2010 county crop reports are used[18]. The revenues are averaged over these inner draws so that we are back to one set of numbers (acreages, revenues etc.) for each outer draw. The acreages shown at columns 2 and 3 and the revenue at column 5 of Table 9 are the averages. The 95% confidence interval is constructed by taking the 2.5^{th} and 97.5^{th} percentiles of these 1000 figures corresponding to each outer draw. The bottom row of Table 9 gives these simulated confidence intervals.

In an effort to take into consideration the yield declines associated with the salinity increases I used the average percentage yield decline for crop classes after a threshold salinity is passed. In general, sensitive crops have low thresholds—a point of salinity level beyond which there will be yield declines— and high slopes, percentage decline in the yield associated with each unit increase in salinity. The

[15]There are locations in the Delta with really high expected increases in the water salinity due to the water canal. They stayed within the group of data excluded due to no change in the crop choice.

[16]Constructing the confidence intervals analytically, using the Delta method is not preferred as it is an immediate application of first order Taylor series expansion. The method would be plausible for a linear model but at a highly nonlinear setting such as this, computer simulation emerged as a more viable option.

[17]Such draws can be obtained from $\beta + \mathbf{\Gamma e}$, where β is the vector of parameter estimates, $\mathbf{\Gamma}$ is the Choleski factor of the variance covariance matrix, and $\mathbf{e}$ is a vector of draws from the standard normal distribution.

[18]The possible price impacts of change in the crop mix due to salinity changes is ignored as the area that is covered in this study is relatively small portion of California's agriculture market for each of the crop classes.

slopes are interacted with the *change* in the salinity level from baseline to policy scenario for each choice situation in the following fashion.

$$\% \text{ Decline}_j = \begin{cases} (s_p - s_b)Sl_j & \text{if } s_p,\, s_b > T_j \\ (s_p - T_j)Sl_j & \text{if } s_p > T_j \ \&\ s_b < T_j \\ -(s_b - T_j)Sl_j & \text{if } s_p < T_j \ \&\ s_b > T_j \\ 0 & \text{if } s_p < T_j \ \&\ s_b < T_j \end{cases}$$

for each choice situation, where, s_p is the salinity for the "policy" case, s_b is the salinity for the base case, Sl_j is the slope of the yield decline curve for crop class j, T_j is the threshold level of salinity beyond which the yield starts to decline. The threshold and slopes can be found in the appendix.

These decline figures can be considered as the upper bound of the revenue declines as it is assumed that the farmers are not going to take any action besides switching crops.[19]

We see that the confidence intervals are fairly large. The higher revenue and loss figures correspond to the states of the world where most farmers have high preferences for sensitive (i.e. higher value crops) and low willingness to switch in the face of salinity. By the same token, low revenue-yield loss cases are likely to happen if most farmers choose highly tolerant species and likely to switch more with salinity. Note that revenue per acre, the salinity levels and weather conditions are treated as if they are certain across the scenarios. The only sources of uncertainty accounted for here are the farmer's preferences for the sake of isolating the effect of salinity associated with the policy as well as the fact the the estimations are done based on a finite sample of farms and periods. The simulation exercise shows that if the farmers choose to plant sensitive species and do not mitigate the salinity the expected revenue loss is about 19% of the expected revenue under the baseline salinity scenario.

In order to check the robustness of the results, instead of using the previous period's salinity, the average of previous 3 years salinity was used as an alternative measure for water quality. The idea is that it may not be very realistic to conclude that the farmers are insensitive to changes in salinity when the measure is taken on an annual basis. The results with this alternative measure turned out to be qualitatively and quantitatively very similar. Table 8 has the results.

There are several points one needs to keep in mind when interpreting these numbers. First of all, the analysis here only includes possible revenue changes from switching crops due to salinity changes. It does not take into account possible adaptation strategies by farmers in terms of adoption of new irrigation technologies or switching to some other form of land use. Second, any loss of farmland to the construction of the new facility is ignored. Third, these numbers only reflect the revenues from crop sales. Overall values of different crop classes are different as some crops generate extra revenue and jobs if they are further processed in the

[19] The yield loss due to increased salinity is limited to 50% while the yield gain due to reduced salinity is limited by 20%.

local facilities. Therefore, a full evaluation of a policy such as this one requires a broader consideration of the impacts on various stakeholders.

Finally, note that the annual revenue losses simulated here should be considered as short run changes in the revenues as the parameters are estimated on based on annual variation in salinity rather than long term changes. Additionally, by time the project is built, the farmer profile may change and the preferences and the management approaches of the future farmers may differ from their contemporary counterparts.

1.6.1 Goodness of Fit

For the logit model, a prominent way to study the goodness of fit is via the likelihood ratio method. Likelihood ratio index is given by the following formula (Train (2009)) :

$\rho = 1 - \frac{LL(\hat{\beta})}{LL(0)}$

where, $LL(\hat{\beta})$ represents the log likelihood function evaluated at the estimated parameters, while $LL(0)$ the log likelihood evaluated with all parameters are set to zero. The log likelihood at the estimated parameters is -8217 and the function evaluates to -13,706 when all parameters are zero. This gives a likelihood ratio index of 0.40.[20]

The t-test results obtained from the parameter estimation procedure suggest that the salinity parameters are significant (with the exception for the field crops). To further investigate the importance of these parameters, likelihood ratio test is conducted on them. In order to test the hypothesis that "the salinity parameters are zero", the model is estimated once again without the salinity parameters. The test statistic is given by $-2(LL(\hat{\beta}_0^H) - LL(\hat{\beta}))$ which is distributed χ^2 with degrees of freedom equal to the number of parameters that the null hypothesis suggests to equal zero (which is 10 in this case, i.e. the means plus the standard deviations of the salinity parameters). The log likelihood from the estimation of the model after excluding the parameters was -8254. The test statistic is 74.5. The cut off point for 10 degrees of freedom at 1 percent significance is 23.21 which suggests an overwhelming statistical evidence for the importance for these parameters.

The final metric is the percent correctly guessed. In this exercise we calculate the logit probability of each option in the choice set for every single choice situation and see if the option with the highest probability is in fact the selected option. The model estimated here correctly guessed the option 49% of the time. Since we have 6 alternatives, randomly guessing the correct crop class would five an expected correct guess of 1/6 (about 16%).

[20] Likelihood ratio is an index varies between 0 and 1 and represent the extent to which estimated parameters improve our estimate of each option's probability. With no parameter estimation our best guess would be assigning each option an equal chance which is what we would precisely get by setting all the parameters to zero at the logit probabilities formula. If the model is no better, we would get a likelihood ratio index of zero and if the model perfectly estimates each option's probability by assigning 1 to the chosen option and 0 to everything else we would get a 1.

1.7 Conclusion

California's history in the past couple of centuries is the history of subduing nature and Sacramento-San Joaquin River Delta is no exception. The beneficiaries of the exports, enjoyed abundant fresh water thanks to which the population and the economy in the arid regions of California flourished over the past half century against all odds. After 150 years of alterations with an extensive levee system and massive water diversion facilities, the Delta, California's most important water resource, host to a unique ecosystem of wide variety of plants and animals, and a big hub for agricultural activity providing food, water, and jobs for millions of people, has reached a point of collapse. Drought, failing levees, court decisions limiting water export to protect endangered and threatened species limits the amount and the reliability of the water conveyed from the Delta. About 30 years after being defeated by the Californians, the peripheral canal idea is brought back by Gov. Brown in an effort to meet the "coequal' goals of water supply and habitat restoration as status quo is no longer sustainable.

The government authorities formulated various measures including the construction of an isolated water conveyance facility. This, however, raises the concerns about increased salinity in the irrigation water and hence the future of agricultural activities in the Delta.

Using a unique and rich panel dataset on Delta agriculture, a discrete choice model for agricultural crops is estimated using random coefficients logit method. The parameter estimates, then, are used to simulate expected acreages and associated farming revenues under two salinity scenarios generated by a hydrology model used by the California Department of Water Resources. Utilization of panel data structure together with the land ownership information is what differentiates this study from other studies, allowing a more credible identification of the salinity parameters, accounting for time varying aggregate changes in preferences, path dependence, and heterogeneity in response to salinity across the farmers.

The results indicate that salinity was a statistically significant factor but weather and alternative specific factors which are captured by alternative dummies have been more important than irrigation water salinity during the last decade for the Delta farmers. This has important policy implications for farmers' welfare. If the farmers are capable of mitigating salinity and protecting their yields without having to switch to salt tolerant species then they are likely to be able to keep their high value crops in the future when the canal is built. However if the relative reluctance to respond to changing salinity levels is due to suboptimal decision making by the farmers, then it is important to understand the source and address any problems accordingly. It is especially important in a case where substantial amount of public funding is about to be channeled to a project that is likely to affect their farming practices.

The simulation results point to a 19% revenue losses *relative* to the baseline scenario where the simulated confidence interval of losses ranges from 30% to 10%. This is the expected value of the loss in the short run if the farmers do not take any mitigating action besides switching crops and let the salts burn their crops, which is highly unlikely given that the Delta farmers have already been facing salinity. The high losses are associated with the states of the world where the farmers have high preferences for high value sensitive crops and low willingness to switch to salt tolerant species in the face of salinity.

Another important take-away form this study is that it is important to take farmer heterogeneity into consideration especially for those researchers who want to study land use and climate adaptation in larger regions. Even in a region as small as studied here, the heterogeneity turned out to be statistically meaningful, and the salinity coefficient estimates of the standard MNL where the heterogeneity component was shut off were quite different. Ignoring this part of the equation can lead to misleading coefficient estimates.

Future work should include engaging with the Delta farmers and see what salinity mitigation options they have and costs of utilizing them. A more comprehensive consideration of Delta water quality from ecological and municipal points of view is also in order to get the full picture of the likely impacts.

1.8 Figures and Tables

Figure 1: Sacramento San Joaquin River Delta Map.[21]

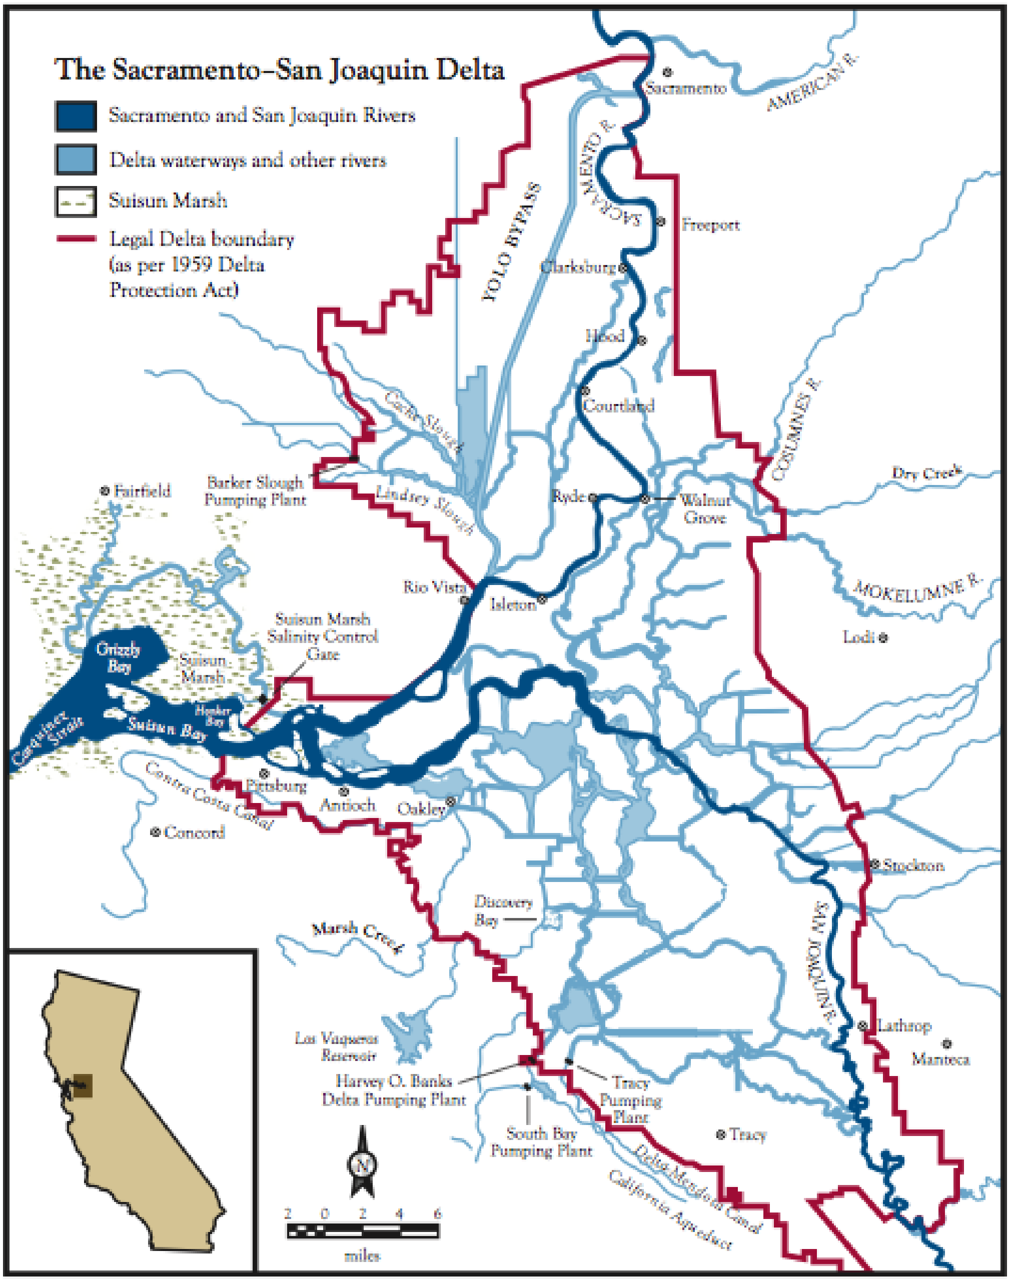

[21]Source: `http://users.humboldt.edu/ogayle/hist383/SacramentoDeltaMap.png`

Figure 2: Salinity vs. Revenue/acre[22]

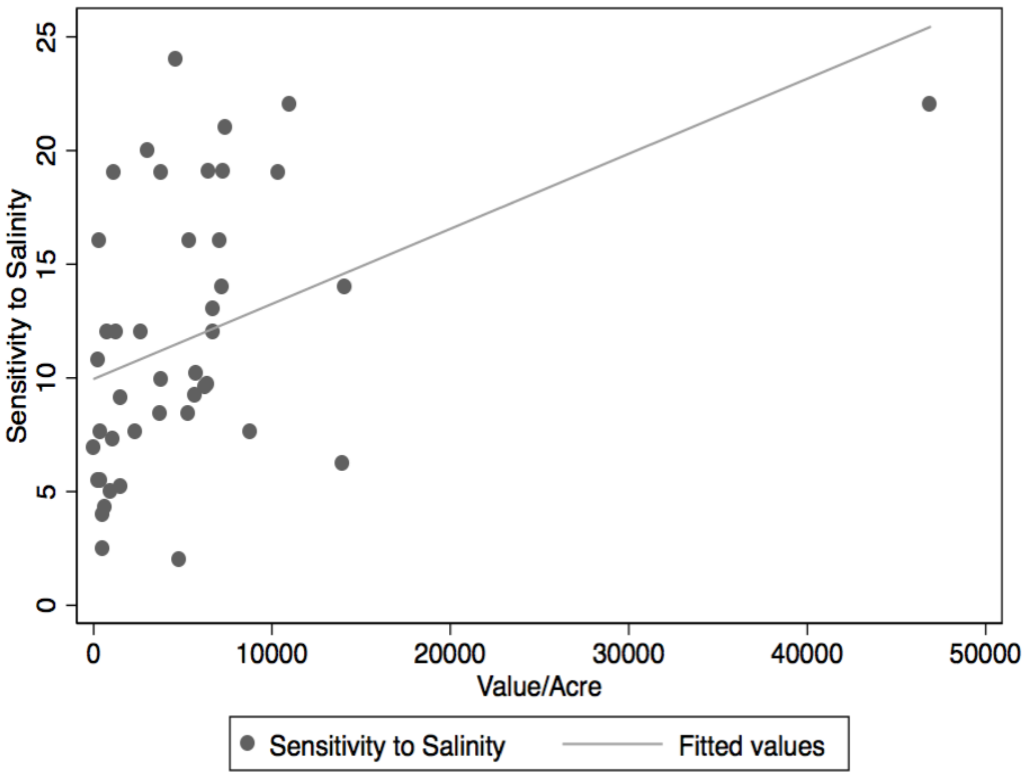

[22]In this figure, we see the positive correlation between the crop revenue per acre and crop sensitivity. Namely, the crops that are more sensitive to salinity are also the ones that yield higher revenue per acre, in general. Here, crop sensitivity is measured in terms of how fast the crop yield declines in percentage terms as a result of a unit increase in salinity.

Figure 3: Comparing the Actual Average and Scenario Salinity Levels[23]

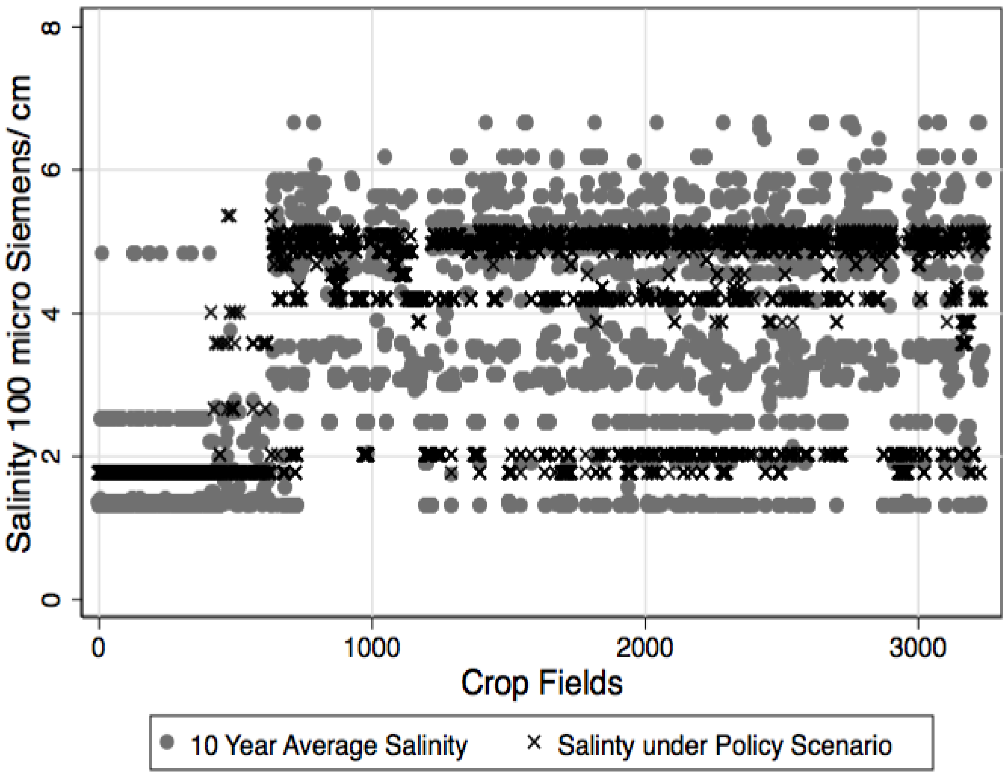

[23]Figure 3 shows the actual salinity levels and the salinity levels used in the simulations which are generated by DSM II —a river, estuary, and land modeling program used by California Department of Water Resources. We see that the scenario salinity levels are all within the already observed ranges across the farms.

Figure 4: Breakdown of the Salinity Tolerance Among the Crop Classes by Acreage[24]

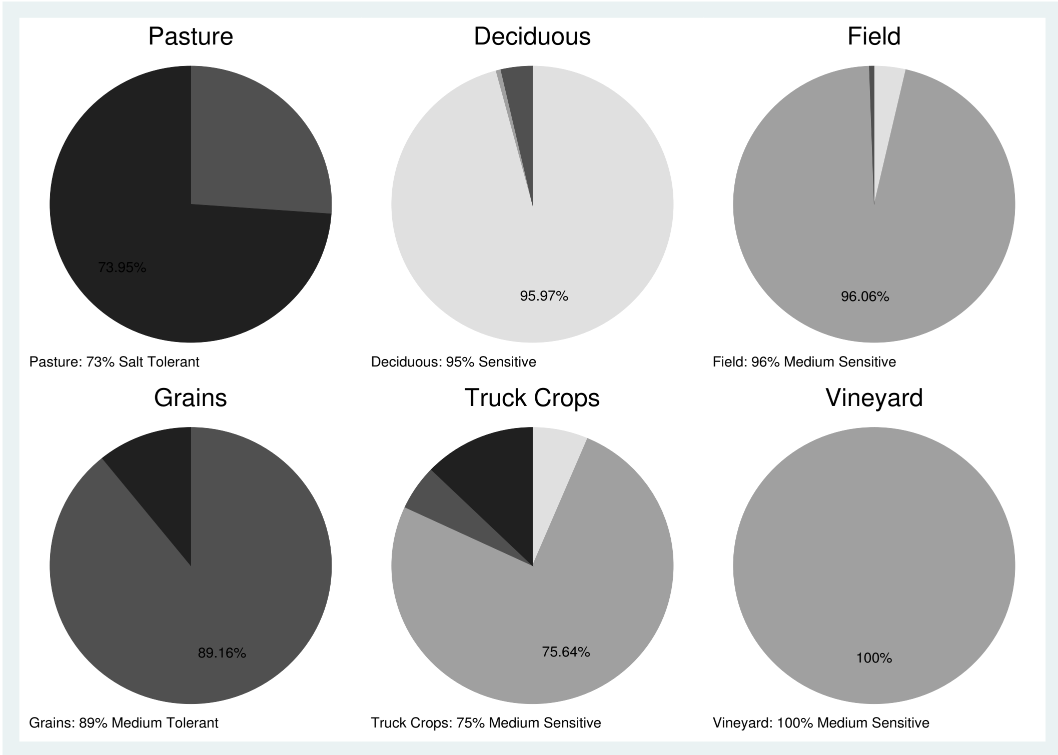

[24]Figure 4 presents that even though there is not a one to one correspondence between the salinity classes and the agronomic crop classes, in terms of acreages, we see that most deciduous crops planted in the Delta in 2010 were sensitive, field, truck, and vineyard crops were medium sensitive, grains were medium tolerant, and the pasture crops were tolerant. The salinity tolerance information were gathered from Hanson, Grattan, and Fulton (1999), Hoffman (2010), and Maas and Hoffman (1977).

Figure 5: Relative Size of the Residuals from Regressing Salinity on Controls (Within Farms)[25]

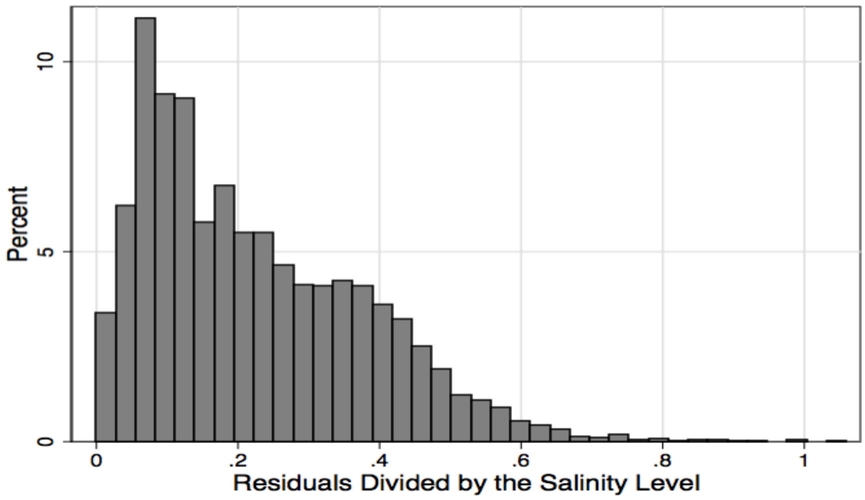

Figure 6: Relative Size of the Residuals from Regressing Salinity on Controls (Within Conservation Zones)

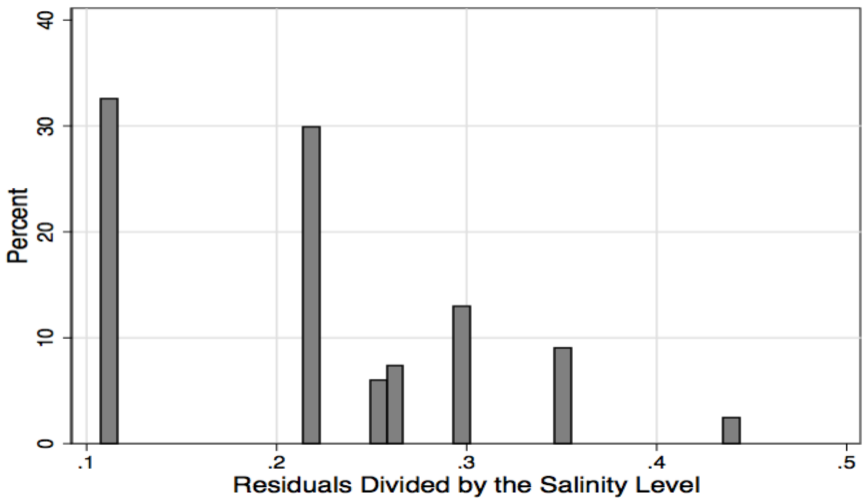

[25]Figures 5 and 6 present the histograms of the size of the absolute value of the residuals from the regressions presented in Table 5 relative to the mean salinity levels at farm and conservation zone level, respectively. These pictures suggest that the size of the identifying variation in salinity across farms is greater than 20% of the mean values of the observed salinity levels both at farm and conservation zone level. In other words, we have quite sizable variation in salinity with which the salinity parameters are identified.

Table 1: Summary Statistics of the Variables

Variable	Mean	Std. Dev.	N
Farm Size (100 Acres)	0.05	0.05	7650
Soil Quality (0-100 Index)	47.85	14.98	7650
Elevation (Feet)	2.62	6.72	7650
Max. Temp (Celcius)	24.14	0.31	7650
Field Slope (Decimal Degrees)	0.13	0.55	7650
Medium Farmer	0.12	0.33	7650
Large Farmer	0.83	0.38	7650
Salinity (100 mS/cm)	3.63	1.84	7650
Deciduous Rev/Acre (1000 USD)	4.63	0.84	7650
Field Rev/Acre (1000 USD)	1.13	0.16	7650
Grain Rev/Acre (1000 USD)	0.51	0.12	7650
Truck Rev/Acre (1000 USD)	2.91	0.64	7650
Vineyard Rev/Acre (1000 USD)	3.25	0.71	7650

Table 2: Regressing Sensitivity to Salinity on Water Salinity and Controls[26]

	Yield Decline	Yield Decline	Yield Decline
Salinity (100 mS/cm)	-0.0663	-0.156*	-0.188**
	(0.047)	(0.061)	(0.067)
Farm Size (100 Acres)	-1.352	-1.439	-2.102
	(1.148)	(1.144)	(1.155)
Elevation (Feet)	0.0165	0.0126	-0.00154
	(0.025)	(0.025)	(0.025)
Soil Quality (0-100 Index)	-0.00438	-0.00415	-0.00505
	(0.004)	(0.004)	(0.004)
Precipitation	-0.0138*	-0.0708*	-0.0757*
	(0.007)	(0.029)	(0.031)
Constant	10.19***	12.90***	11.85***
	(0.331)	(1.121)	(1.180)
R^2	0.266	0.275	0.278
Owner FE	Yes	Yes	Yes
Year FE	No	Yes	Yes
Con. Zone FE	No	No	Yes

Standard errors in parentheses

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

[26]Table 2 provides the results of regressing salinity sensitivity of the selected crop on previous year's salinity level and other covariates used in the mixed logit. The estimations suggest that the salinity sensitivity of the chosen crop tends to decline as the previous year's irrigation water salinity goes up. We infer that the farmers tend to choose more salt tolerant crops as they observe increases in the irrigation water salinity. *Sensitivity* here refers to how fast the crop yield declines per unit increase in salinity. All three models include owner dummies. As we sequentially add the year and the conservation zone indicators, we see that both the magnitude and the significance of the salinity coefficient goes up while the magnitude and the significance of other variables remain somewhat stable.

Table 3: Crop Revenue Per Acre on Water Salinity and Controls[27]

	Revenue/acre	Revenue/acre	Revenue/acre
Salinity (100 mS/cm)	127.0***	5.460	-26.22
	(29.592)	(37.738)	(41.136)
Farm Size (100 Acres)	476.6	-24.08	-562.9
	(703.031)	(695.395)	(702.770)
Elevation (Feet)	-6.892	-11.77	-22.31
	(16.564)	(16.517)	(16.796)
Soil Quality (0-100 Index)	-2.089	-1.569	-2.282
	(2.547)	(2.522)	(2.566)
Precipitation	-4.948	-6.774	-18.14
	(4.225)	(18.269)	(19.567)
Constant	1784.7***	1955.9**	1465.0*
	(206.610)	(713.438)	(744.135)
R^2	0.294	0.312	0.317
Owner FE	Yes	Yes	Yes
Year FE	No	Yes	Yes
Con. Zone FE	No	No	Yes

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

[27] Table 3 shows the results of regressions similar to the ones in Table 2 but this time, revenue per acre appears as the dependent variable instead. Unlike the previous case, we do not observe a clear pattern in the sign, magnitude, and the significance of the salinity coefficient. These findings suggest that the farmers seem to respond to irrigation water salinity by switching to more salt tolerant crops but it is hard to make a claim with any degree of confidence that the crops they switch to are necessarily low value. Hence, farmers seem to pay attention to for water salinity while protecting their revenues.

Table 4: Information by Conservation Zones [28]

Zone	Years	County	Number of obs.	Number of farms	Mean Salinity (100mS/cm)	St. Dev.
2	2007-2010	Yolo	184	85	1.84	0.97
3	2007-2010	Sacr., Yolo	688	379	1.70	0.45
4	2003-2004, 2006-2010	Sacr., San Joaquin	559	261	1.63	0.53
5	2003-2004, 2006-2010	Sacr., San Joaquin	988	407	1.74	0.58
6	2003-2004, 2006-2010	San Joaquin	2,287	819	3.67	1.24
7	2003-2004, 2006-2010	San Joaquin	2,491	1,115	5.16	1.11
8	2003-2004, 2006-2010	San Joaquin	453	179	3.63	0.99

[28] Table 4 presents a summary of how the salinity varies within and across the conservation zones. Conservation zones refer to the regions in the Delta with similar ecological characteristics. This is useful for seeing the nature of salinity variation, because in the analysis uses conservation zone dummies to account for the time invariant factors that effect crop choice.

Table 5: Regressing Salinity on Controls[29]

	Salinity	Salinity	Salinity	Salinity	Salinity
Soil Quality	0.0199***	0.0197***	0.000233	0.0000142	0.000237
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Elevation	0.103***	0.109***	0.00567**	0.00358	0.00388*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Farm Size	-1.341***	-1.043***	-0.998***	-0.911***	-0.985***
	(0.339)	(0.312)	(0.209)	(0.210)	(0.210)
Max. Temp.	3.378***	4.106***	1.062***	1.073***	1.074***
	(0.051)	(0.061)	(0.059)	(0.059)	(0.059)
Constant	-79.09***	-95.79***	-22.94***	-23.17***	-23.05***
	(1.243)	(1.459)	(1.414)	(1.410)	(1.408)
Observations	7650	7650	7650	7650	7650
R^2	0.450	0.536	0.795	0.797	0.798
Year FE	No	Yes	Yes	Yes	Yes
Conzone FE	No	No	Yes	Yes	Yes
Owner-size FE	No	No	No	Yes	Yes
Salinity Rating FE	No	No	No	No	Yes

Standard errors in parentheses

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

[29] In order to further deliberate on the identifying variation in salinity, Table 5 presents the results of regressing the salinity on the control variables used in the study. We can see that after regressing the irrigation water salinity on all the controls, still a sizable portion (about 20%) of the variation in the salinity remains unexplained.

Table 6: Random Coefficients Logit Estimation Results Using Previous Year's Salinity [30]

Fixed Coefficients					
Crop Choice	**Deciduous**	**Field**	**Grains**	**Truck**	**Vineyard**
Acres	−16.83**	0.91	−0.13	−0.34	3.18
	(7.77)	(2.41)	(2.45)	(2.47)	(5.42)
Soil	0.01	0.02***	0.02**	0.02***	0.05***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Elevation	0.14***	0.03	0.01	0.02	−0.35***
	(0.04)	(0.02)	(0.02)	(0.02)	(0.13)
Max. Temp.	6.16***	4.40***	3.20***	3.55***	8.88***
	(1.36)	(0.63)	(0.64)	(0.65)	(1.96)
Slope	−0.16	0.09	0.10	−0.03	0.38
	(0.46)	(0.18)	(0.18)	(0.18)	(0.31)
Medium Farmer	−2.21**	0.09	−0.57	0.28	−1.03
	(1.03)	(0.52)	(0.52)	(0.56)	(1.11)
Large Farmer	−0.89	−0.09	−0.93*	0.36	0.01
	(0.79)	(0.48)	(0.48)	(0.51)	(0.97)
Same as last year	0.67	−1.30***	−0.85***	−0.92***	–
	(1.03)	(0.07)	(0.09)	(0.08)	–
Alternative Constant	−147.27***	−103.88***	−73.54***	−84.26***	−213.69***
	(32.14)	(14.78)	(15.01)	(15.31)	(46.69)
Revenue/Acre	0.50***	0.50***	0.50***	0.50***	0.50***
	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)
Random Coefficients					
Salinity	−1.84***	−0.01	−0.14	−0.29**	−3.38***
	(0.34)	(0.13)	(0.13)	(0.13)	(0.59)
Salinity Standard Dev.	−0.63***	0.24***	0.27***	0.44***	−1.55***
	(0.13)	(0.01)	(0.02)	(0.02)	(0.25)

[30] Table 6 presents the random coefficients logit estimation results. Each column represents the alternative specific coefficient estimate and its standard error *relative* to the outside option which is pasture crops. Significance stars next to the standard deviation estimates at the bottom suggest the existence of heterogeneity in the corresponding salinity parameters across the decision makers.

Table 7: Standard Multinomial Logit Results

Crop Choice	Deciduous	Field	Grains	Truck	Vineyard
Acres	-17.842^{*}	1.568	0.371	-2.529	4.098
	(7.042)	(2.328)	(2.363)	(2.352)	(3.767)
Soil	0.017	0.013^{**}	0.015^{**}	0.017^{***}	0.035^{***}
	(0.011)	(0.006)	(0.006)	(0.006)	(0.010)
Elevation	0.102^{***}	0.034^{*}	0.018	0.013	-0.284^{***}
	(0.027)	(0.019)	(0.020)	(0.020)	(0.084)
Max. Temp.	5.541^{***}	3.476^{***}	3.449^{***}	2.870^{***}	4.452^{***}
	(1.029)	(0.580)	(0.585)	(0.582)	(1.191)
Slope	-0.138	0.067	0.110	-0.015	0.230
	(0.352)	(0.175)	(0.176)	(0.177)	(0.269)
Medium Farmer	-1.440^{*}	-0.102	-0.409	0.061	-0.921
	(0.733)	(0.494)	(0.494)	(0.502)	(0.775)
Large Farmer	-0.764	-0.278	-0.685	0.417	-0.750
	(0.625)	(0.457)	(0.457)	(0.464)	(0.689)
Same as last year	2.002^{***}	-1.059^{***}	-0.465^{***}	0.088	–
	(0.586)	(0.059)	(0.076)	(0.057)	–
Alternative Constant	-133.637^{***}	-81.821^{***}	-80.169^{***}	-68.427^{***}	-109.676^{***}
	(24.421)	(13.682)	(13.821)	(13.722)	(28.135)
Revenue/Acre	0.484^{***}	0.484^{***}	0.484^{***}	0.484^{***}	0.484^{***}
	(0.094)	(0.094)	(0.094)	(0.094)	(0.094)
Salinity	-1.052^{***}	-0.105	-0.056	-0.245^{*}	-0.524^{*}
	(0.196)	(0.114)	(0.115)	(0.115)	(0.214)

[30] Table 7 presents the results of estimating same parameters on Table 6 with standard multinomial logit instead of random coefficients logit. The difference here is that unlike the mixed logit, the salinity parameters are set to be fixed across all decision makers by setting the standard deviation of the salinity parameters to zero. We see that salinity parameters differ to a great extent. This means that ignoring the presence of heterogeneity will result in misleading parameter estimates.

Table 8: Random Coefficients Logit Estimation Results Using Previous 3 Year's Salinity [31]

Fixed Coefficients					
Crop Choice	**Deciduous**	**Field**	**Grains**	**Truck**	**Vineyard**
Acres	-17.50^{**}	0.84	-0.13	-0.83	3.76
	(7.93)	(2.38)	(2.45)	(2.44)	(4.93)
Soil	0.02	0.02^{***}	0.02^{**}	0.02^{***}	0.05^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Elevation	0.10^{***}	0.02	0.01	0.01	-0.37^{***}
	(0.03)	(0.02)	(0.02)	(0.02)	(0.13)
Max. Temp.	4.05^{***}	3.85^{***}	3.20^{***}	2.71^{***}	6.56^{***}
	(1.09)	(0.63)	(0.64)	(0.65)	(1.77)
Slope	-0.15	0.09	0.10	-0.01	0.32
	(0.43)	(0.18)	(0.18)	(0.18)	(0.31)
Medium Farmer	-1.90^{*}	0.11	-0.57	0.41	-0.93
	(0.97)	(0.52)	(0.52)	(0.55)	(1.05)
Large Farmer	-0.78	0.06	-0.93^{*}	0.40	0.05
	(0.76)	(0.49)	(0.48)	(0.51)	(0.97)
Same as last year	1.01	-1.31^{***}	-0.85^{***}	-0.87^{***}	–
	(0.93)	(0.07)	(0.09)	(0.08)	–
Alternative Constant	-97.58^{***}	-91.12^{***}	-73.54^{***}	-64.45^{***}	-159.38^{***}
	(25.76)	(14.77)	(15.01)	(15.25)	(42.05)
Revenue/Acre	0.54^{***}	0.54^{***}	0.54^{***}	0.54^{***}	0.54^{***}
	(0.11)	(0.11)	(0.11)	(0.11)	(0.11)
Random Coefficients					
Salinity	-1.57^{***}	-0.01	-0.15	-0.32^{***}	-2.76^{***}
	(0.33)	(0.13)	(0.13)	(0.13)	(0.57)
Salinity Standard Dev.	-0.56^{***}	0.26^{***}	0.29^{***}	0.41^{***}	-1.43^{***}
	(0.12)	(0.02)	(0.02)	(0.02)	(0.26)

[31] Table 8 shows the results of the random coefficients logit estimates with the average of the previous three years' salinity as the salinity measure instead of previous year's salinity. We see that the results are robust to such a change in salinity measure.

Table 9: Results of the Simulation Exercise[32]

Crop Classes	Baseline Acreage	Gross Scenario Acreage	Net Scenario Acreage	Rev./Acre	Baseline Revenues	Net Scenario Revenues	%Change
Pasture	414	414	414	\$116	47,966	48,003	0.1%
Deciduous	6,027	6,027	4,367	\$4,612	27,795,681	20,139,665	-27.5%
Field	67,137	67,138	50,009	\$780	52,366,928	39,007,336	-25.5%
Grain	61,092	61,091	54,837	\$426	26,025,065	23,360,381	-10.2%
Truck	66,492	66,492	55,122	\$3,903	259,516,975	215,142,113	-17.1%
Vineyard	5,591	5,591	4,182	\$3,566	19,938,905	14,913,477	-25.2%
Total Rev.					\$385,691,521	\$312,610,975	-19%
95% C.I.					(\$92,560,000 - \$929,770,000)	(\$82,380,000 - \$680,150,000)	(-30% to -10%)

[32] Table 9 shows the results of the simulation exercise. The first and the second columns of numbers are the expected values of the total acreages of each alternative while the third column presents the acreages after taking into consideration possible yield declines due to salinity changes. The declines are calculated based on the average threshold and yield decline slope of the crops in each class taken from the agricultural experiments studies (see the appendix). Column 4 is the revenue per acre measures. Columns 5 and 6 show the expected revenues under baseline and the policy scenario, respectively. Latter is after the yield declines are accounted for. The last column shows the percentage change in the expected revenues due to the policy. The overall expected decline in revenues is about 19%. At the bottom of the table are the confidence intervals of the revenues under the two scenarios and the associated percentage of revenue losses.

Chapter 2: Forecasting Commercial, Institutional and Industrial Water Demand in the Southern California (with Steven Buck)

2.9 Introduction

The main purpose of this study is to compare the out-of-sample forecasting performances of different regression models that perform the best with respect to various in and out-of-sample goodness-of-fit criteria in the context of commercial and industrial[33] (C&I) water demand in the Southern California. The majority of the water demand forecast studies focus on the residential sector as it accounts for the largest portion of the publicly supplied water demand.[34] Water use data published by the US Geological Survey (Kenny, Barber, Hutson, Linsey, Lovelace, and Maupin (2014)) indicates that public supply made up about 17% of all water withdrawals in California in 2010. According to the annual survey conducted by the Bay Area Water Supply Conservation Agency, residential demand accounts for about 60% of the publicly supplied water while commercial and industrial makes up about 20%.[35]

The C&I water demand forecast still requires rigorous scrutiny. This sector makes up a considerable portion of water consumption in other parts of the world. For example, in Europe industrial sector accounts for 23% of total water demand and even within the EU countries we see a significant variability in water consumption patterns (de Bono, Del Pietro, Giuliani, Harayama, Le Sourd, and Diana (2004)). Furthermore, changing trends in the economic landscape of both developed and developing countries (a shift towards commercial activities in developed countries while industrial activities gaining pace in the developing nations) is likely to increase the relative importance of the C&I water demand. The majority of the water used in industrial activities in California for example, is self supplied while commercial water is supplied by the public utilities. So if we see a shift from industrial towards commercial activities in the economic mix, it is reasonable to expect that the C&I sector will take up a greater share within that publicly supplied water. Most importantly, it has been gradually more evident that the commercial sector can be an important avenue to save water through various rebate programs. Potential quests to evaluate such programs will call for accurate water demand forecasts in the C&I sector which will also improve the

[33]By the US Geological Survey definition, commercial water use includes consumption by "motels, hotels, restaurants, office buildings, other commercial facilities, military and non-military institutions, and (for 1990-1995) off stream fish hatcheries". Industrial water demand captures the water used "for fabrication , processing, washing, and cooling. Includes industries such as chemical and allied products, food mining, paper and allied products, petroleum refining, and steel."(`http://water.usgs.gov/watuse/wuglossary.html#CO`)

[34]Public supply refers to water withdrawn by public and private water suppliers that provide water to at least 25 people or have a minimum of 15 connections (Kenny, Barber, Hutson, Linsey, Lovelace, and Maupin (2014))

[35]`http://www.bawsca.org/docs/BAWSCAFY2012-13AnnualSurvey_Water%20Consumption%20by%20Class.pdf`

effectiveness of water plans and budgets. Water related issues are likely to be more pressing as the water is becoming progressively scarce in the face of global warming, rising population as well as the surge of per capita demand due to globally increasing life standards.

When choosing a model for forecasting, it is often perceived as customary to select a model that provides the best in-sample fit (usually high coefficient of determination i.e. R-squared). Forecasting by definition requires stretching the estimation out-of-sample. Therefore, an investigation into how the out-of-sample forecast performances of models selected solely based on a high R-squared is warranted. How do these models compare to the models that would get selected under different criteria? Using models that are chosen based on in-sample fit can be suboptimal when forecasting out-of-sample (McCracken and West (2002)). In this study, we undertake a discussion of this problem in the C&I water demand context.

Using a model space of 352,112 models, we look at how the models that yield the best results based on R-squared and/or other popular in-sample-fit criteria (Adj. R-Squared, AIC, BIC) perform when forecasting out-of-sample. We then compare them to those models that would be selected under three different out-of-sample criteria we define. While, by construction, the models that are selected based on out-of-sample forecasting performance should do better than those based on any other criterion in forecasting out-of-sample, the main goal is to demonstrate by how much.

The out-of-sample fit criteria are defined following Auffhammer and Steinhauser (2012) which are also used commonly used in evaluating models in the field of machine learning. Models are generated mostly through inclusion and exclusion of different key covariates and the actual versus logged values of the dependent variable. The dependent variable is the water retailer (utility) level total annual C&I demand. Covariates include median tier price; manufacturing and service sector employment in the service area of the retailer; and weather variables (maximum temperature, cooling degree days, and precipitation). In order to account for the fact that C&I water demand is a derived demand together with other factors of production, we also included real US GDP as a proxy for the overall purchasing power in the economy.

The results indicate that selecting models solely based on in-sample fit will yield a poor performance when it comes to forecasting out-of-sample. This is a well known fact in the machine learning discipline and in a way, this study brings it to the attention of water planners. We also demonstrate that the predictions that are generated by the highest R-squared models are highly dispersed around the actual value relative to those that the lowest absolute aggregate error models generate.

These findings suggest that the policy makers and planners should make sure that the models to forecast demand are selected taking out of sample prediction performance into account.

This chapter proceeds as follows: a review of the related literature in section 2.10 is followed by the summary of the data in section 2.11. Section 2.12 describes the model selection procedure and presents the results and section 2.13 concludes. All figures are tables are provided in section 2.14

2.10 Literature Review

Though not as numerous as residential water demand studies, C&I studies have been conducted both in academia as well as by government organizations.

Academic studies are usually centered around estimating the price and/or output elasticity of the water demand in the industrial sector using a sample of disaggregated (plant or building level building) panel data which include observations on price of water, quantity consumed, and the output. Studying the water use from 30 large plants, De Rooy (1974) find that a 1% increase in water price is associated with a 0.89 % decline in the quantity. Ziegler and Bell (1984) and Renzetti (1992) use instrumental methods for price to avoid the canonical simultaneity problem in their estimations. Using a system of simultaneous equations method, Babin, Willis, and Allen (1982) and Renzetti (1992) examine the relationship of water intake and the utilization of other outputs (the degree of substitutability). Using data from 51 industrial plants in France and seemingly unrelated regression and feasible generalized least squared methodology, Reynaud (2003) finds that the elasticity of water demand varies across the water sources. As an alternative to econometric method, Calloway, Schwartz, and Thompson (1974) develops a linear programing model on order to analyze the effects water quality policy, on the use of water in ammonia production and on the cost of ammonia.

Commercial water demand in the forecasting context has not been as widely studied as residential water demand in academia. This is most likely due to lack of data, and the complexity of C&I demand because of the diversity in the nature of water use in this sector. The result of the studies done suggest that although the demand in commercial sector is somewhat sensitive to water price, it is not as sensitive as the residential and industrial consumption (Schneider and Whitlatch (1991)).

The studies that are conducted by the government organizations are more geared towards forecasting for planning purposes. Water Resources Municipal and Industrial Needs model (IWR-MAIN) has been used historically to forecast the water demand.The size of each C&I sector is estimated using total employment and C&I water use is estimated based on the Standard Industrial Classification (SIC) sectors. The method used regression analysis to determine water intensity of each sector where the explanatory variables were number of employees, the price of the water and sewer services, and whether or not there was a water conservation program (Boland (1997)). Later on a nationwide survey of over 3 thousand establishments and surveys from manufacturers from the US Census Bureau and the California Department of Water Resources were utilized to improve the model (Dziegielewski and Boland (1989)). The main intuition of this approach is to estimate a “water use coefficient” for each sector and multiply that with the forecasted size of the sector and sum up over all the sectors. A nice treatment of the historical progression of the IWR-MAIN model can be found in Morales, Martin, and Heaney (2009). Even though IWR-MAIN itself is no longer in use, the forecast methods that are currently utilized for planning use similar methodology. The approach inspired some other government studies as well. For example, using the establishment level water billing and employment data Cook, Urban, Maupin, Pratt, and Church (2001) calculate the standard industrial classification (SIC) level water demand employment coefficients, which is basically the weighted average of the per employee water consumption for the SICs. Then, under various growth scenarios and employment forecasts, they project the water consumption into the future periods. As a major improvement to these models, Morales, Martin, and Heaney (2009) presents a C&I water demand estimation methodology explaining the availability of a rich database of parcel level consumer attributes and water use billing from Florida.

The path followed in this study is similar to the one used in Auffhammer and Steinhauser (2012) for forecasting CO_2 emissions. They use 41 years of state level data to test about 27,000 models and compare the out-of-sample forecasting performances of benchmark models from the related literature and the ones that they find to be best under the aggregate error criterion. They find that benchmark models which are calibrated against in-sample performance criterion are likely to overestimate the CO_2 emissions which might be consequential in climate policy and international agreements.

We do not really have any "benchmark" models in the field of aggregate C&I water demand. What we do in this study is to compare the out of sample performances of the models that would be selected under various in and out-of-sample criteria given the available dataset. Our findings are in tune with Auffhammer and Steinhauser (2012) in the sense that the model selection criterion matters for the forecasting performance.

2.11 Data

The geographical scope of this study is the boundaries of Metropolitan Water District of Southern California (MWDSC) (see Figure 7). The dataset used here is a subset of a larger dataset collected for an ongoing study about forecasting single family residential (SFR) sector water demand. Data collection effort, therefore, was focused on the retailers that reported more than 3,000 single-family residential accounts as it is estimated that these retailers account for about the 99% of this sector. 153 retailers were contacted within the realm of the study. C&I data was obtained from 75 retailers and has 709 observations from 25 of the 26 member agencies that are under MWDSC. The only unrepresented member agency is San Marino which likely has one of the lowest C&I sectors of all member agencies. Table B2 in the appendix lists the agencies and the associated retailers. The water retailers in the study are located in Los Angeles, Orange, Riverside, San Bernardino, San Diego, and Ventura counties.

The rate schedules were received directly from retailers while the consumption figures are mostly based on monthly data reported to the Public Water System Statistics (PWSS) augmented with data received from retailers and is aggregated to the calendar year.[36] For the price measure, we use the median tier of the rate schedule. Rates are reported in year 2000 real dollars.[37]

Data on average precipitation were obtained through the use of the geographical information and mapping software system, ArcGIS. Spatially referenced boundaries of state and private water districts were obtained from the Cal-Atlas geospatial clearinghouse.[38] These boundaries allowed visualization of each water district polygon using ArcGIS. The points at the centroid of each water system polygon were then geo-referenced. Based on the resulting set of points the local precipitation data were extracted from rasters provided by the PRISM Climate Group.

In those cases where the retailer level district boundaries were not available, zip codes were used as a geographical proxy. Retailers were assigned to representative zip codes on a case by case basis. The centroid of each zip code polygon were geo-referenced, and based on the resulting set of points local precipitation data were extracted. The precipitation variable in our dataset is in millimeters of rainfall per year.

Data on temperature were obtained in the same manner as the precipitation data, described above. Rasters for the temperature data (in degrees Celsius) were obtained from PRISM Climate Group.[39] The year round maximum and minimum temperatures are used to calculate retailer specific cooling degree days.

Total employees within a retailer are computed based on two data sources. Historical annual employment is provided by the metropolitan water district (MWD) at the member agency level from 1990 to 2010. To calculate employment at the retailer level we use the Census Zip Code Business Statistics (ZCBS), which reports historical employment estimates at the zip code level from 2004 to 2010. The ZCBS only provides employment numbers based on the majority of sectors (largely excludes non-service oriented government positions) so total employment is not complete. Therefore, we only use the ZCBS to calculate the share of employment within a member agency due to a particular retailer. We calculate the relevant share using a crosswalk between zip codes and retailer level boundaries, and zip codes and member agency level boundaries again, we are able calculate the share of each member agency's employment due to a particular retailer. Finally, to compute a historically based retailer level total employment measure, we multiply the share of employment within a member agency estimated from the ZCBS by

[36]PWSS is a database kept by the Department of Water Resources (DWR), containing the annual voluntary surveys of a subset of California utilities. The data contains numbers of connections, water deliveries for consumer classes such as single family residential, multifamily residential, commercial/institutional, industrial, landscape irrigation and other at a monthly level.

[37]We do not use instrumental variables to address the simultaneity bias in the price coefficient. However, note that the purpose of this study is to compare the forecasting performances of models that rank high under different goodness-of-fit criteria given the data available rather than credibly identifying a price parameter.

[38]Cal-Atlas Geospatial Clearing House, accessible: `http://atlas.ca.gov/download.html#/casil/boundaries`

[39]PRISM Climate Group, "Near-Real-Time High-Resolution Monthly Average Maximum/Minimum Temperature for the Conterminous United States", raster digital data, accessible: `http://www.prism.oregonstate.edu/`

the total employment in the member agency obtained from MWD (based on Employment Development Department data). For years prior to 2004 when ZCBS is unavailable, we assume the retailer level average employment shares from 2004 to 2006. That is, for each retailer we assume their share of total employment within a member agency is constant between 1994 and 2003.

GDP is also included in the regressions. Unlike the residential sector, water demand in the C&I sector is derived, together with other inputs, as a part of the production process. In other words, water demand in C&I is *indirectly* caused by the demand for the goods and services that these sectors offer to consumers. Therefore, the water demand should not only depend on its own price but also to the total demand in the economy which ultimately depends on income. National, rather than regional GDP figure is utilized as these the C&I water customers are likely to supply to a larger region than the state they are physically located. The real GDP data here is obtained from the publicly available international macroeconomical data series provided at the USDA website.[40] [41] Figures are converted to 2000 dollars. Tables 10 and 11 give the summary statistics of the variables both in the training and the forecasting subsamples, respectively. Training sample is composed of data from years 2000-2005, while the forecast sample is the data from years 2006-2010.

2.12 Model Selection and Comparison

In the regression models, we follow a similar approach to studies that forecast the residential sector water demand with the following general regression template:

$q_{tar} = \beta \cdot price_{tar} + \mu \cdot man.emp_{tar} + \sigma \cdot serv.emp_{tar} + \tau \cdot tmax_{tar} + \pi \cdot precip_{tar} + \gamma \cdot cool_{tar} + \alpha_a + \eta_t + \epsilon_{tar}$

where, q_{tar} is the annual water demand in the C&I sector in year t served by the retailer r that is under agency a;
$price_{tar}$ is the median tier price charged;
$man.emp_{tar}$ is the number of manufacturing employees;
$serv.emp_{tar}$ is the the total number of service employees;
$tmax_{tar}$ is the average maximum temperature;
$precip_{tar}$ is the average annual precipitation;
$cool_{tar}$ is the cooling degree days;
α_a is the agency fixed effects;
and η_t represent one of the time indicators.[42]

The model space was created using different permutations of dependent and independent variables (and their actual and logged values). There are three main avenues through which new models are added to the model space. First is the inclusion versus exclusion of the main variables: price, number of employees in

[40] `http://www.ers.usda.gov/data-products/international-macroeconomic-data-set.aspx`

[41] Note that to avoid perfect collinearity only one of year fixed effects and annual GDP covariates could be used at a time in a regression equation.

[42] Year dummies or GDP. Note that the year dummies and GDP can not be used at the same time due to perfect collinearity.

the manufacturing and service sectors in the retailers' boundaries, maximum temperature, cooling degree days, precipitation, GDP as well as lagged dependent variables (up to two lags). Second variation source pertains to the inclusion of the variables needed in order to account for the heterogeneity with respect to the time and the institutions that govern the water of different locations. These include agency indicators, the permutations of time trend (up to cubic time trend), and year indicators. Finally, further variations are generated using logged vs. level dependent variables as well as total quantity versus per employee quantity as the dependent variable. Table 12 summarizes the details. All these permutations initially resulted in a total of 497,724 models. In the forecasting procedure it was observed that some of these models yielded negative forecast values. This happened when the coefficient of trend, trend squared, or trend cubed variable turned out to be negative. When extrapolated out of sample the negative trend lead to negative forecasts. After these models are omitted from the model space, we ended up with 352,112 models.

For the regressions, the data set is divided into two subsets: training and forecast samples. Data from the years 2000-2005 are used to train (estimate) the models and the years 2006-2010 are used as the forecast sample to measure the out-of sample performances.[43]

For every single model, both popular in-sample (R-squared, adjusted R-squared, Akaike information criterion (AIC), Bayesian information criterion (BIC)) and three out of sample performance measures (explained below) are calculated.

Afterwards, the models are sorted based on their performances with respect to each these criteria, and the out-sample performances of the top 1% in each category are compared.

Many different performance criteria can be chosen based on the forecasting and planning goals such as aggregating absolute or squared errors across different geographical or institutional boundaries. Here we studied three different out-of-sample performance measures: mean squared forecasting error at retail and agency levels, and the overall absolute aggregate error. Table 13 provides the formulations for the in and out-of-sample performance criteria.

First out-of-sample criterion is the retail level mean squared forecast error (third one from the bottom of Table 13). Here, q_{tar} the annual C&I water demand in year t of retailer r that belongs to agency a in the forecasting sample. $\hat{q}_{tar}$ is the forecasted quantity for the same data point. R_{ta} is the number of retailers for which the data was available in agency a in year t, A_t is the number of agencies in the sample in year t, and N is the total number of the data points in the forecast sample (N= 310).

The second criterion is the counterpart of the first one at the agency level. We first aggregate the differences between the actual numbers and forecasted numbers at agency level for each forecast year. Afterwards, we take the mean of the squared forecasted error over this *collapsed* sample (M=101).

[43] The data was available starting from year 1994 however these years were omitted from the dataset for a healthier analysis as prior to 2000 the number of available observations per year are less then 20.

Final out-of-sample performance criterion is the absolute aggregate forecasting error. All the quantities (both forecasted and actual) are aggregated over the forecast sample for each year, the aggregate of the forecasts are subtracted from the aggregate of the actual numbers and the the average of the absolute value of the aggregate error is taken over the years.

One important detail to note is the comparability of the performance criteria across the models with different dependent variables (i.e. level vs. logged). It is important to establish this comparability of the goodness-of-fit measures across the models to be able to make meaningful statements about their relative performances. In order to do that, the performance measures for the models with logged dependent variable had to be transformed in the following manner.

After the models with logged dependent variable are estimated, the fitted values are exponentiated. Then, the actual quantities are regressed (without a constant term) on these exponentiated values. Once again the fitted values are calculated from this second regression. These fitted values (Wooldridge (2002))[44] are used to calculate the prediction errors. The square of the correlation coefficient between the actual and the ffitted within the training sample is our comparable R squared. The adjusted R squared is calculated from this R squared by using the usual formula : $1 - \frac{n-1}{n-k}(1 - R^2)$.

Another layer of adjustment is done due to the existence of models both with total quantity and quantity per employee as dependent variable. AIC and BIC scores are calculated using the sum or squared errors (see Table 13 for the formula). Therefore, unlike R-squared and adjusted R-squared the magnitude will depend on the scale of the variables. For this reason the scale of the error needs to be adjusted for a fair comparison of different models. All AIC and BIC scores are calculated using the deviance from the actual total quantity and the implied total quantity by the model. In other words if the model is logged, the AIC and BIC scores are calculated from the squared errors obtained from the fitted values described above. On the other hand if the dependent variable is per employee water quantity, predicted total quantity is obtained by multiplying the fitted value by the total employee number.

Next, the models are ranked based on each one of the criteria in our list. Tables 14 and15summarize and compare the performances the top 1% of the models in each criterion. Every column (except for the first column) refers to a certain subset of all the models in our model universe. Each row gives the mean and the standard deviation of the top 1% based on the criteria listed in that row within the subset given by the column. For example, the numbers in the first row of the second column give the mean and the standard deviation of the "Retail Level MSFE" of the models that rank in the top 1% in terms of "Retail Level MSFE" category among the models that only use levels of the dependent variable (rather than the logs). This categorization gives an association between the inclusion of certain

[44] Another popular method to recover the forecasted value for y when $logy$ is the dependent variable is using the $\hat{y} = exp(\frac{\hat{\sigma}^2}{2})exp(\hat{logy})$ formula. But this transformation relies on the assumption that the errors are normally distributed. Since we are running a large number of regressions in this study, we chose a method that is robust to the error distributions.

variables in the models and the performance as a result. We see in Table 14) that within the models where the dependent variable is total quantity (as opposed to per employee quantity) log models (3rd versus 1st and 2nd columns) showed better out of sample performance on average while in sample performances were almost the same for the models in all categories.

One notable result is models without any lagged variables did much worse than overall (comparing columns 1 versus 5) in almost all criteria. This perhaps is not surprising given the serially correlated nature of water consumption. Additionally we see that though it may reduce the noise, adding agency fixed effects did not improve the forecasting performance.

Since the data is annual, we had to pick between year dummies and the (lagged) per capita GDP as including more than one of these covariates at the same time would result in perfect collinearity. In the models that use year fixed effects the projection needs special consideration as we do not have a clear way to forecast the year fixed effects for the future years. For simplicity, we treated all years in the forecast sample as the end year of the training sample. Comparing the final two columns we see that in fact, instead of using some kind of a proxy for the year fixed effects for the future years, we could use the GDP forecast as the performance of the models with year fixed effects and per capita GDP are fairly comparable for both in and out-of-sample criteria. So in addition to providing a proxy for the size of the economy in forecasting this indirect demand for water, this is another motivation for including the GDP in the forecast models.[45] The results are very similar for quantity per employee models (Table 15).

Tables 16 and 17 compare the absolute aggregate error of the models ranked within the top 1% of our criteria. For example, the number on the second row and the first column of Table 16 is the mean of the absolute aggregate error (in thousand acre feet) of the models that are in the top 1% based on the "Retailer Level MSFE" criteria. We see that the the models that score high based on in-sample-criteria did poorly compared to the models that get picked based on the out-of -sample criteria in aggregate forecast. This result is expected given the way we choose the models and constructed the criteria. The key point here is the difference between the mean of the aggregate error under different categories. The models that score high in the out-of-sample performance criteria yielded much lower absolute errors (10.09 for the absolute aggregate error (in 1000-Acre feet)), and a narrower width (standard deviation of 1.37) whereas the models that were among the highest R-squared, for example, on average relatively did very poorly (mean absolute aggregate error: 405.51 (in 1000-Acre feet) and the dispersion of the performance was much larger (standard deviation of 539.27). The results are similar for the comparison of the models where the dependent variable was quantity per employee (Table 17). Here we see that models with highest R-squared and adjusted R-squared was more comparable, although still worse, while lowest AIC and BIC models performed really poorly.

The panels in Figures 8 and 9 help visualize to the point made in Tables 16 and

[45] It is important, however, to note that the GDP figures here are the actual numbers. When forecasting in reality, the forecasted GDP not the actual figures will have to be utilized instead.

17. In these graphs we see the actual aggregate, "best","lowest", and the"highest" of the models that ranked among the top 1% of the criteria studied here. In this context the "best" model means the model with the lowest absolute aggregate forecast error among the top 5% of the given criterion. The "highest' and the "lowest" models are the ones with the smallest and the largest mean aggregate forecast error, respectively. A large error means that the model underestimated the actual quantity whereas a small error (which would be a large negative number) means there is an over estimation. Notice the wide gap between the lowest and the highest model in the graph for the top R-squared and AIC models (panels (a) and (b) of Figures 8) and 9. We see clearly in these figures that the forecasts generated by the models that are selected based on in-sample criteria are much widely dispersed compared to those that are selected based on out-of-sample forecast criteria. The models selected using the out-of-sample criteria, on the other hand, (panels (c) and (b) of Figure 8) generate a much narrower spectrum of forecasts.

To provide a further visual insight, each panel in Figures 10 and 11 shows the actual aggregate (represented with the black spikes) and the histogram of the aggregate of the forecasts for the models that are within the top 1% of the R-squared and the absolute aggregate error criteria by the year. We see that in every single year the forecasts that are generated by the models that are in the top 1% of the latter criterion are more closely gathered around the actual value. This implies again that selecting a model based on the out-of-sample prediction performance will make it more likely to have a close forecast.

Finally, the panels in Figures 12 and the future projections of aggregate C&I demand using the top models under the select criteria and the projections of the covariates. Note that the models with lagged dependent variables as well as squared and cubed time trend had to be eliminated. The projected employment numbers were available at the agency level and disaggregated to the retailer level using some assumptions. For the employment shares of each retailer in the future years, the employment figure in the most recent year in which the data was available were used. For the future GDP numbers, a 1% annual growth was assumed. For the weather variables at the retailer level, the agency level averages were used. The red lines are the average quantities of the top 50 models in each category. We see that for the models where total quantity were used. For both category of models (total quantity as well as per employee quantity) the forecasts are still somewhat more disperse for the top R-squared models. More importantly models selected by different criteria provide different views of the future. We see that according to top R-squared models both total quantity and per employee quantity demanded in the C&I sector will continue growing. On the other hand, lowest aggregate forecast error models suggest that the boom in the total quantity will be more with a more stable per employee quantity.

2.13 Conclusion

Even though the commercial and industrial water demand is a relatively small portion of the publicly supplied water demand in the US, the share is larger in some other countries and it is likely to get larger in the US as the composition of the economic activity shifts. The emergence of the commercial sector as a potential source of water conservation further increased its relative importance. Especially in the existence of dramatic droughts driven by the changing climate, the accuracy of the forecasts becomes increasingly important for the ability to plan ahead more effectively by the institutions that govern water.

Using over 352,000 models, out-of-sample forecasting performances of different models are compared in the context of C&I water demand in the Metropolitan Water District of Southern California. Additionally the water demand trajectories implied by models picked based on different criteria are depicted side by side.

We found that selecting models based on high in-sample goodness-of-fit (usually R-squared) value may not result in the best forecasts out of sample. Furthermore, we saw a much higher variance in the forecasts among these models relative to the top (smallest) aggregate forecast error models.

The water policy makers and planners who rely on water demand forecasts, therefore, should pay attention to the out-of-sample performance of the models that are being utilized in their analyses. Commonly, water demand forecasting done for planning purposes are performed by "black box" models and it is hard to see the underlying modeling methodology and standard errors of the parameter estimates. If a water governing body, instead, chooses to use econometric methods using the data available from the local region, they should avoid selecting models based on in-sample-fit as this may result in suboptimal results in the quality of the forecasts.

Future work should use instrumental variables method for water price for a credible parameter identification as questionable parameters may compromise the overall forecasting ability of the model. Precipitation in the area where the water of Southern California comes from (Sierra Nevada or California Bay Delta for example) could be used as an instrument for the price as it plausibly affects the supply but should have little or no effect on the water demand in the Southern California.

Furthermore, sophisticated machine learning algorithms such as artificial neural networks as well as data at different granularity could also be used for prediction as well. Current computing and monitoring technologies make it easier than it has ever been to collect and process large amounts of data. A recent study done on the municipal water demand in Cyprus showed that, neural networks yielded superior results to multiple regression models (Adamowski and Karapataki (2010)). Water planners who seek to improve their forecasting precision will certainly benefit from these techniques.

2.14 Tables and Figures

Table 10: Summary Statistics of the Training Sample[46]

	Mean	Std.	Min.	Max.
Water Quantity (1000 Acre-feet)	7.73	23.62	0.07	182.24
Real GDP(in Year 2000 $ Bil.)	16992.49	764.27	16102.66	18249.53
Price (in Year 2000 $)	1.33	0.42	0.44	2.48
Manufacturing Employment	9472.27	27108.52	31.56	225540.53
Service Employment	29425.89	93287.30	206.14	682037.75
Max Temp. (C)	24.19	1.73	19.90	28.85
Cooling Degree Days	1.17	0.41	1.00	3.27
Precip. (mm Per Year)	373.92	157.34	102.81	900.13
Observations	326			

Table 11: Summary Statistics of the Forecast Sample[47]

	Mean	Std.	Min.	Max.
Water Quantity (1000 Acre-feet)	7.21	22.30	0.05	181.84
Real GDP(in Year 2000 $ Bil.)	18857.88	214.49	18486.07	19069.42
Price (in Year 2000 $)	1.47	0.51	0.47	3.73
Manufacturing Employment	8078.80	21815.74	28.53	171248.09
Service Employment	31226.67	94688.02	216.71	719254.69
Max Temp. (C)	24.44	1.68	20.46	28.24
Cooling Degree Days	1.44	0.62	1.00	3.44
Precip. (mm Per Year)	339.76	182.32	102.15	909.27
Observations	310			

[46] Region: Metropolitan Water District of Southern California, Years: 2000-2005
[47] Region: Metropolitan Water District of Southern California, Years: 2006-2010

Table 12: Variables Used to Define the Model Universe

Main variables
Price
Number of employees in the manufacturing sector (both linear and squared form)
Number of employees in the service sector (both linear and squared form)
Average maximum temperature
Cooling degree days
Precipitation
GDP
Lagged dependent variable (up to two lags)
Temporal and insitutional heterogeneity
Agency fixed effects
Time trend (up to cubic)
Year fixed effects
Additional
Levels vs. logs of both dependent and independent variables
Total quantity vs. quantity per employee as the dependent variable

Table 13: Summary of the In-Sample and Out-of-sample Criteria Used for Model Selection

In sample performance criteria	
R- Squared	$1-\frac{\sum_{t=2000}^{2005}\sum_{a=1}^{A_t}\sum_{r=1}^{R_{ta}}\epsilon_{tar}^2}{\sum_{t=2000}^{2005}\sum_{a=1}^{A_t}\sum_{r=1}^{R_{ta}}(q_{tar}-\bar{q}_{tar})^2}$
Adjusted R-Squared	$1-\frac{N-1}{N-k}(1-R^2)$
AIC	$ln(\sum_{t=2000}^{2005}\sum_{a=1}^{A_t}\sum_{r=1}^{R_{ta}}\frac{\epsilon_{tar}^2}{N})+\frac{2k}{N}$
BIC	$ln(\sum_{t=2000}^{2005}\sum_{a=1}^{A_t}\sum_{r=1}^{R_{ta}}\frac{\epsilon_{tar}^2}{N})+\frac{k}{N}ln(N)$
Out-of-sample performance criteria	
Retail level MSFE	$\frac{1}{N}\sum_{t=2006}^{2010}\sum_{a=1}^{A_t}\sum_{r=1}^{R_{ta}}(q_{tar}-\hat{q}_{tar})^2$
Agency level MSFE	$\frac{1}{M}\sum_{t=2006}^{2010}\sum_{a=1}^{A_t}(\sum_{r=1}^{R_{ta}}(q_{tar}-\hat{q}_{tar}))^2$
Total forecast error	$\lvert\sum_{t=2006}^{2010}\sum_{a=1}^{A_t}\sum_{r=1}^{R_{ta}}(q_{tar}-\hat{q}_{tar})\rvert$

Table 14: Summary Statistics of the Results from Top 1% of the Models - Dependent Variable: Total Quantity[48]

	All Models	Levels Only	Logs Only	No Agency FE	No Lags	Year FE	Lagged GDP
Retail Level MSFE	1.512	2.880	1.473	1.517	9.165	1.578	1.593
	(0.121)	(0.314)	(0.115)	(0.106)	(1.181)	(0.0750)	(0.187)
Agency Level MFSE	5.117	15.91	4.970	5.046	31.84	5.342	5.842
	(0.432)	(4.543)	(0.401)	(0.437)	(4.072)	(0.290)	(0.616)
Absolute Aggregate Error	10.09	26.25	9.673	10.93	15.30	10.87	10.63
	(1.367)	(3.324)	(1.336)	(1.073)	(3.659)	(0.939)	(2.192)
R Squared	0.988	0.988	0.985	0.985	0.988	0.988	0.988
	(0.000311)	(0.000473)	(0.00105)	(0.000484)	(0.0000367)	(0.000407)	(0.0000751)
Adj. R Squared	0.987	0.987	0.983	0.984	0.987	0.987	0.987
	(0.000327)	(0.000524)	(0.00115)	(0.000468)	(0.0000344)	(0.000454)	(0.0000507)
AIC	15.90	15.88	16.34	16.02	15.91	15.90	15.91
	(0.0241)	(0.0380)	(0.0390)	(0.0274)	(0.00454)	(0.0347)	(0.00368)
BIC	16.17	16.12	16.46	16.14	16.24	16.19	16.18
	(0.0332)	(0.0198)	(0.0223)	(0.0276)	(0.0110)	(0.0310)	(0.0358)

[48] Each row in Tables 15 and 14 corresponds to the mean and the standard deviation of the subset of models described in the column based on the criteria on the row. For example, the numbers in the first row and the forth column from left give the mean and the standard deviation of the retail level MSFE of the top 1% of the models that used logged dependent variable.

Table 15: Summary Statistics of the Results from Top 1% of the Models - Dependent Variable: Quantity Per Employee

	All Models	Levels Only	Logs Only	No Agency FE	No Lags	Year FE	Lagged GDP
Retail Level MSFE	1.523	1.540	1.518	1.510	10.36	1.542	1.751
	(0.106)	(0.127)	(0.0989)	(0.112)	(0.891)	(0.0855)	(0.186)
Agency Level MFSE	5.242	5.843	5.084	5.242	38.67	5.469	7.518
	(0.501)	(0.596)	(0.468)	(0.534)	(5.280)	(0.330)	(0.619)
Absolute Aggregate Error	10.02	10.37	9.863	10.09	13.94	10.20	15.37
	(1.406)	(1.296)	(1.455)	(1.256)	(3.747)	(1.248)	(2.656)
R Squared	0.884	0.880	0.885	0.877	0.470	0.881	0.885
	(0.000983)	(0.000251)	(0.000777)	(0.000421)	(0.00246)	(0.000660)	(0.000846)
Adj. R Squared	0.874	0.871	0.875	0.874	0.415	0.871	0.874
	(0.000690)	(0.000165)	(0.000600)	(0.000603)	(0.00168)	(0.000232)	(0.000617)
AIC	15.99	15.91	16.08	16.12	16.38	15.89	16.14
	(0.0831)	(0.0806)	(0.0834)	(0.0301)	(0.0722)	(0.0658)	(0.0587)
BIC	16.26	16.22	16.41	16.25	16.63	16.22	16.34
	(0.0381)	(0.0377)	(0.0467)	(0.0204)	(0.0530)	(0.0382)	(0.0346)

Table 16: Aggregate Forecast Error (in 1000-Acre-Feet) of the Top 1% of the Models[49]

	Mean	Std.	Min.	Max.
Absolute Aggregate Error	10.09	1.37	4.32	11.54
Retailer Level Average MSFE	13.17	2.21	7.38	21.81
Agency Level MFSE	12.60	1.84	7.38	19.08
R Squared	405.51	539.27	71.84	2613.93
Adjusted R Squared	240.23	351.31	71.84	2613.93
AIC	163.27	140.02	66.73	1253.92
BIC	114.41	91.02	34.82	773.57

Table 17: Aggregate Forecast Error (in 1000-Acre-Feet) of the Top 1% of the Models. Dependent Variable: Quantity Per Employee

	Mean	Std.	Min.	Max.
Absolute Aggregate Error	10.02	1.406	3.44	11.80
Retailer Level Average MSFE	13.23	4.320	6.74	30.40
Agency Level MFSE	11.75	2.126	6.24	19.55
R Squared	47.57	50.886	6.74	372.72
Adjusted R Squared	28.25	13.176	6.74	97.89
AIC	555.05	845.184	22.57	3425.30
BIC	760.53	1049.162	11.43	3539.80

[49] Each row in Tables 16 and 17 represents the summary statistics for the top 1% of the models in the criteria that is labeled by the row. For example the first row gives the mean, standard deviation, minimum, and the maximum of the aggregate forecasting error of the models that ranked in the top 1% in terms of their absolute aggregate error. We see that the models that rank high in terms of in-sample criteria yielded a relatively high aggregate forecasting error compared to the ones that rank in out-of-sample criteria.

Figure 7: Map of Metropolitan Water District of Southern California

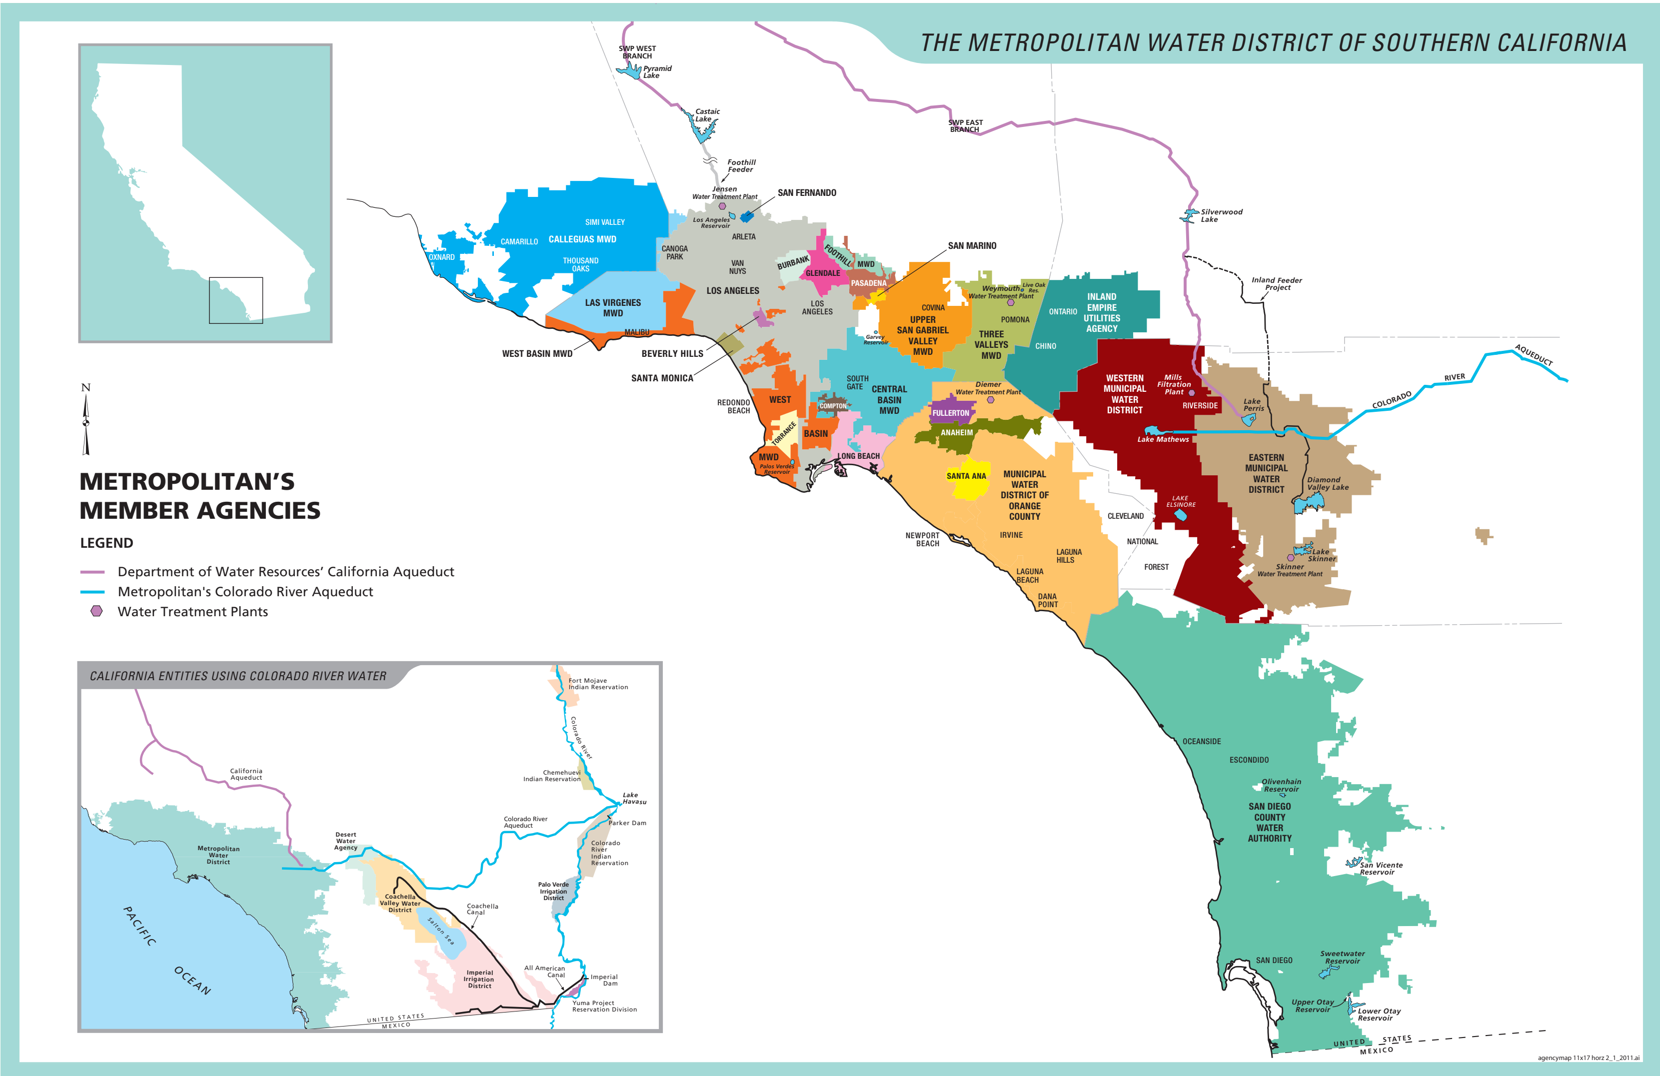

Source: http://www.mwdh2o.com/mwdh2o/pages/news/MA_map.pdf

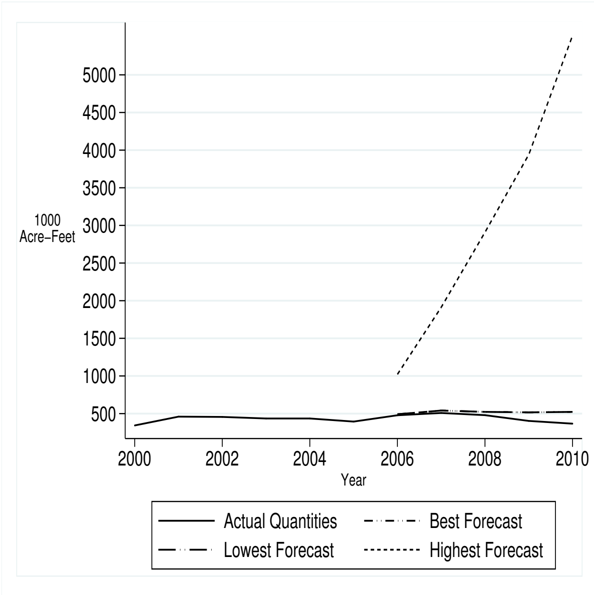

(a) Top R-Squared Models

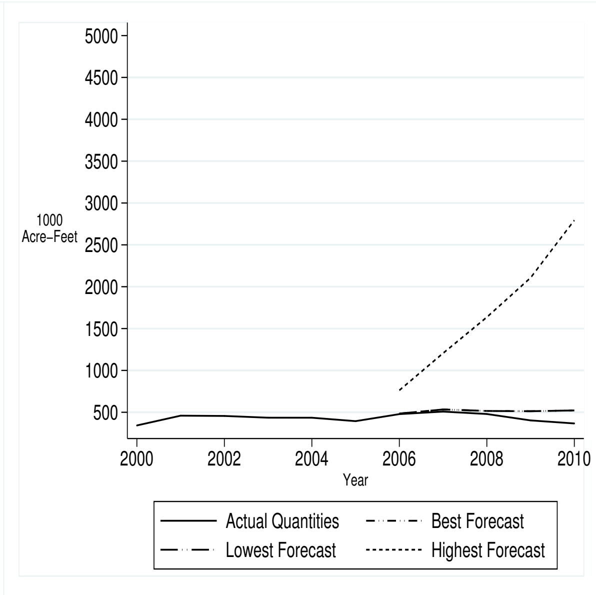

(b) Top AIC Models

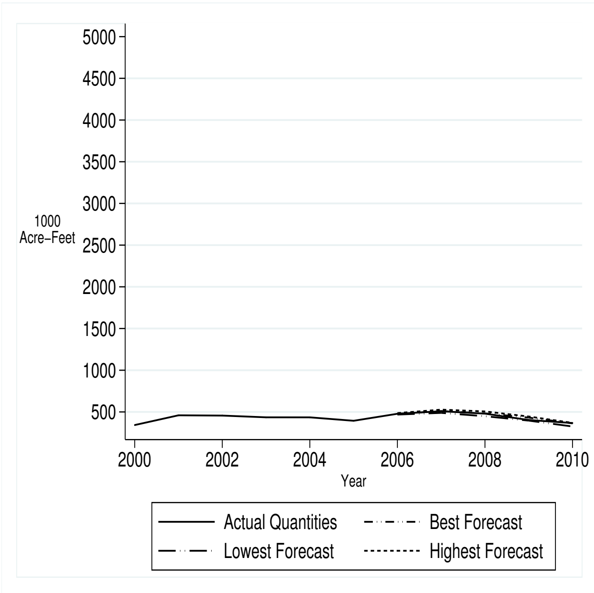

(c) Top MSFE Models

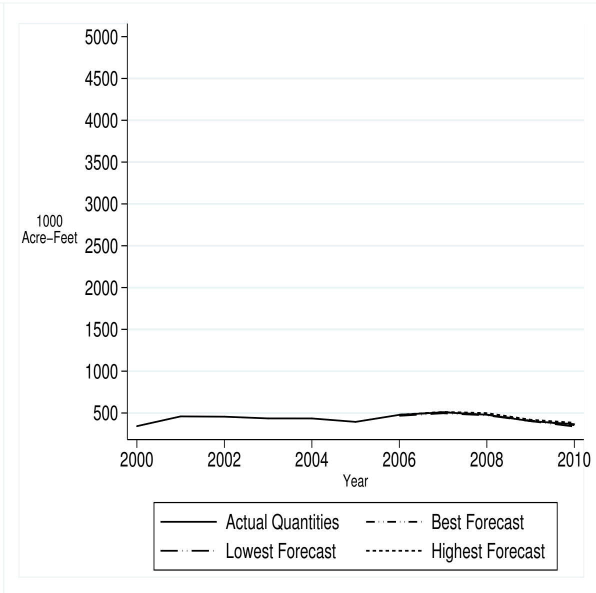

(d) Top Absolute Aggregate Error Models

Figure 8: Highest, Lowest, and the Best Projections. Dependent Variable: Total Quantity[50]

[50] The panels of Figures 8 and 9, show the actual quantity consumed as well as the best, highest, and lowest forecasts generated for the models that rank within the top 1% of R-Squared, AIC, mean squared forecast error (MSFE) and absolute aggregate error criteria, respectively. We see in panel (a) that the models that ranked high based on the R-squared criterion, the highest forecast turned out to be extremely high compared to the actual numbers. And the pattern is very similar with panel (b) which represents the top AIC score models. Contrasting panels (c) and (d), we see that using out of-sample-fit as a selection criterion yields a much narrower range of forecasts. The results look similar in Figure 9.

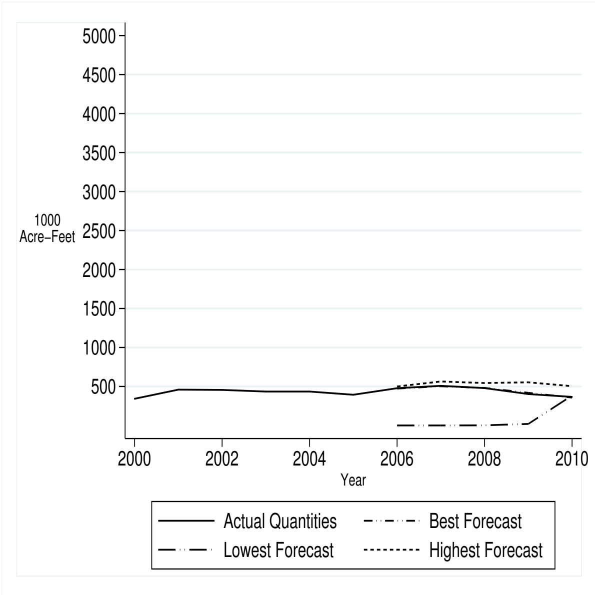

(a) Top R-Squared Models

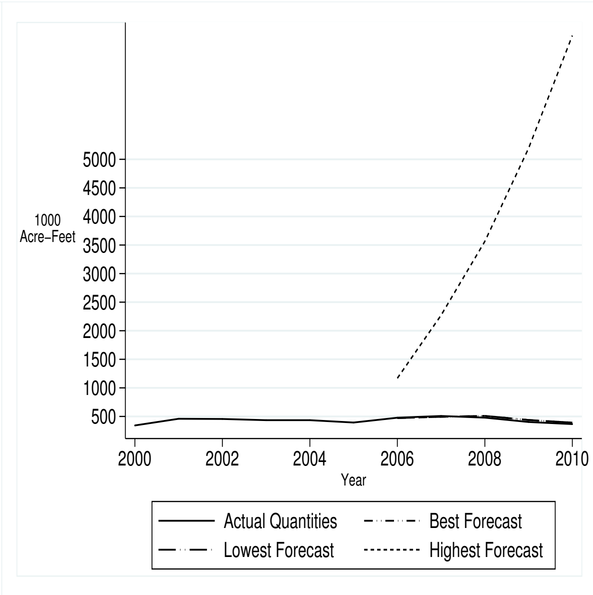

(b) Top AIC Models

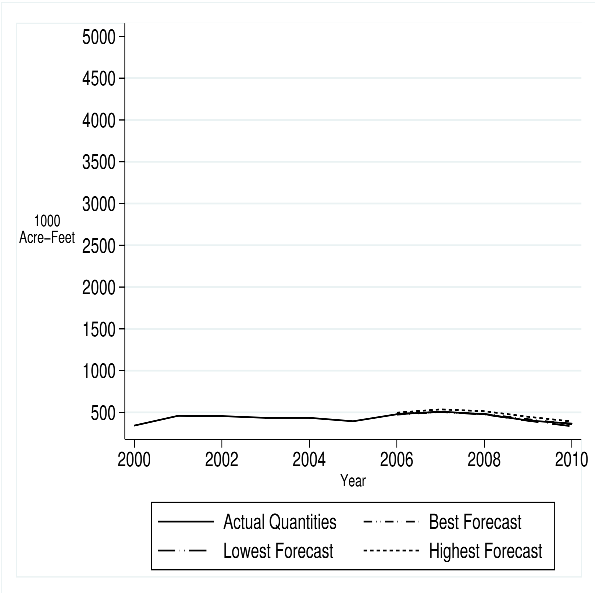

(c) Top MSFE Models

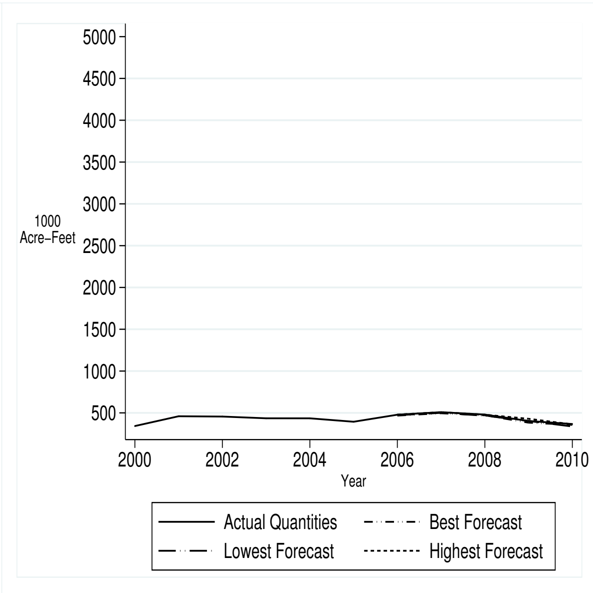

(d) Top Absolute Aggregate Error Models

Figure 9: Highest, Lowest, and the Best Projections. Dependent Variable: Quantity Per Employee

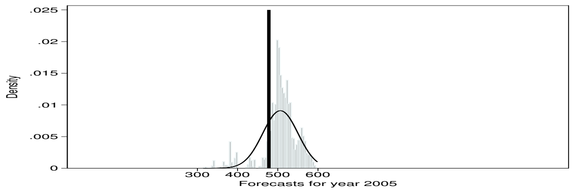

(a) 2006 Forecasts of Top R-Squared Models

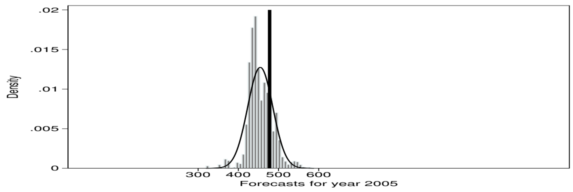

(b) 2006 Forecasts of Top Abs. Agg. Error Models

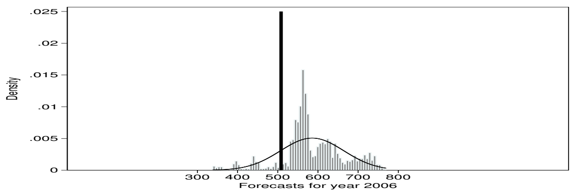

(c) 2007 Forecasts of Top R-Squared Models

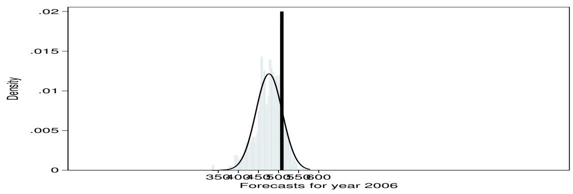

(d) 2007 Forecasts of Top Abs. Agg. Error Models

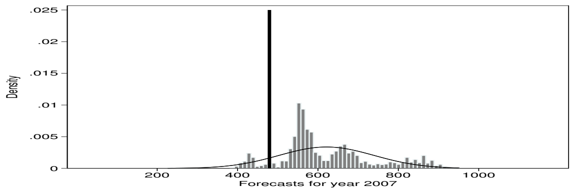

(e) 2008 Forecasts of Top R-Squared Models

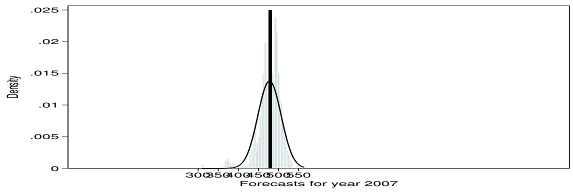

(f) 2008 Forecasts of Top Abs. Agg. Error Models

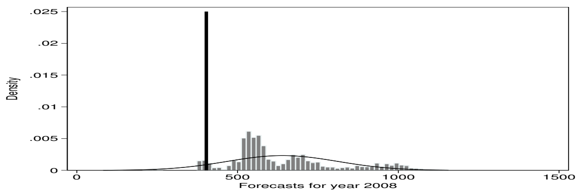

(g) 2009 Forecasts of Top R-Squared Models

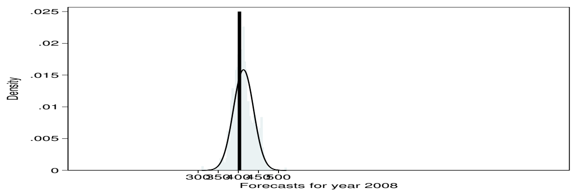

(h) 2009 Forecasts of Top Abs. Agg. Error Models

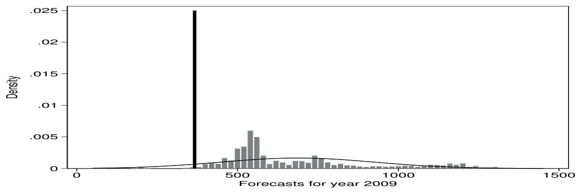

(i) 2010 Forecasts of Top R-Squared Models

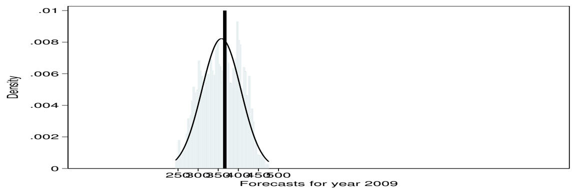

(j) 2010 Forecasts of Top Abs. Agg. Error Models

Figure 10: Distribution of the Forecasts Around the Actual Value Criteria: R-Squared vs. Aggregate Absolute Error Criteria. Dependent Variable: Total Quantity[51]

[51]Figures 10 and 11 show the distribution of the aggregate forecasts around the actual values for years 2006-2010 for the models that ranked in the top 1% of the R-squared (displayed on the left) and the absolute aggregate error (displayed on the right), respectively (for total quantity and quantity per employee models). Each graph shows the actual aggregate value (represented by the black spikes) and a histogram of the forecasts generated by each model picked under different criteria for the given year. We see that for all years, the forecasts generated by the

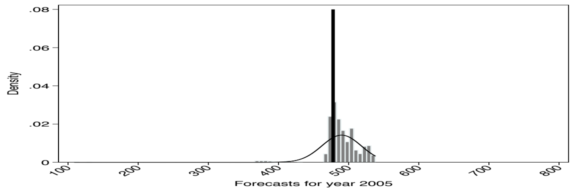

(a) 2006 Forecasts of Top R-Squared Models

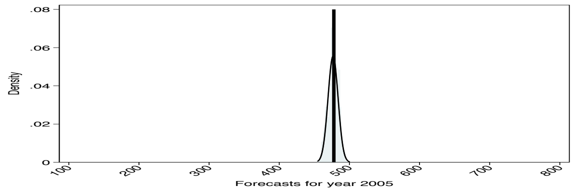

(b) 2006 Forecasts of Top Abs. Agg. Error Models

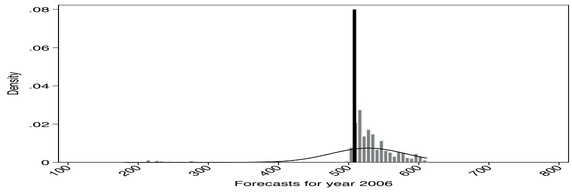

(c) 2007 Forecasts of Top R-Squared Models

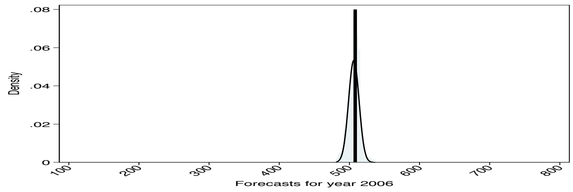

(d) 2007 Forecasts of Top Abs. Agg. Error Models

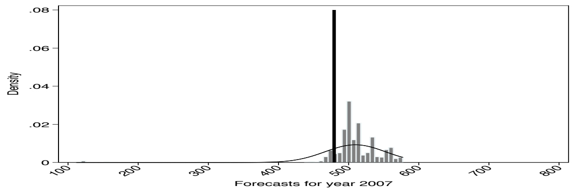

(e) 2008 Forecasts of Top R-Squared Models

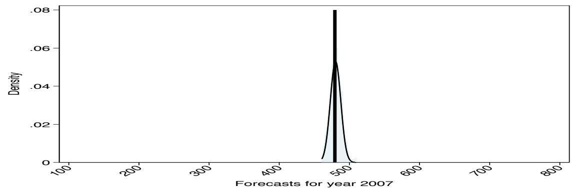

(f) 2008 Forecasts of Top Abs. Agg. Error Models

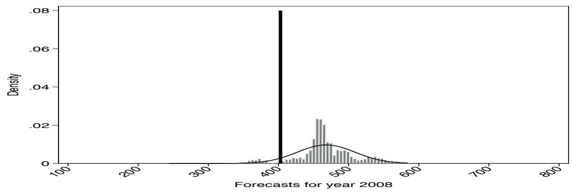

(g) 2009 Forecasts of Top R-Squared Models

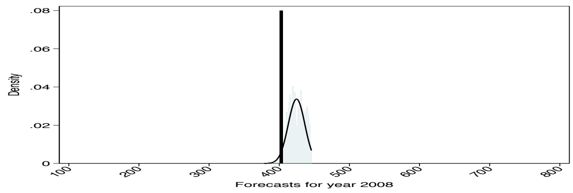

(h) 2009 Forecasts of Top Abs. Agg. Error Models

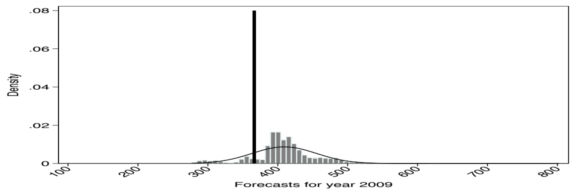

(i) 2010 Forecasts of Top R-Squared Models

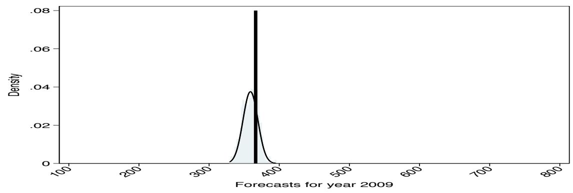

(j) 2010 Forecasts of Top Abs. Agg. Error Models

Figure 11: Distribution of the Forecasts Around the Actual Value Criteria: R-Squared vs. Aggregate Absolute Error Criteria. Dependent Variable: Quantity Per Employee

models that are picked based on the absolute aggregate error are more densely distributed close to the actual value compared to the top R-squared models for all years.

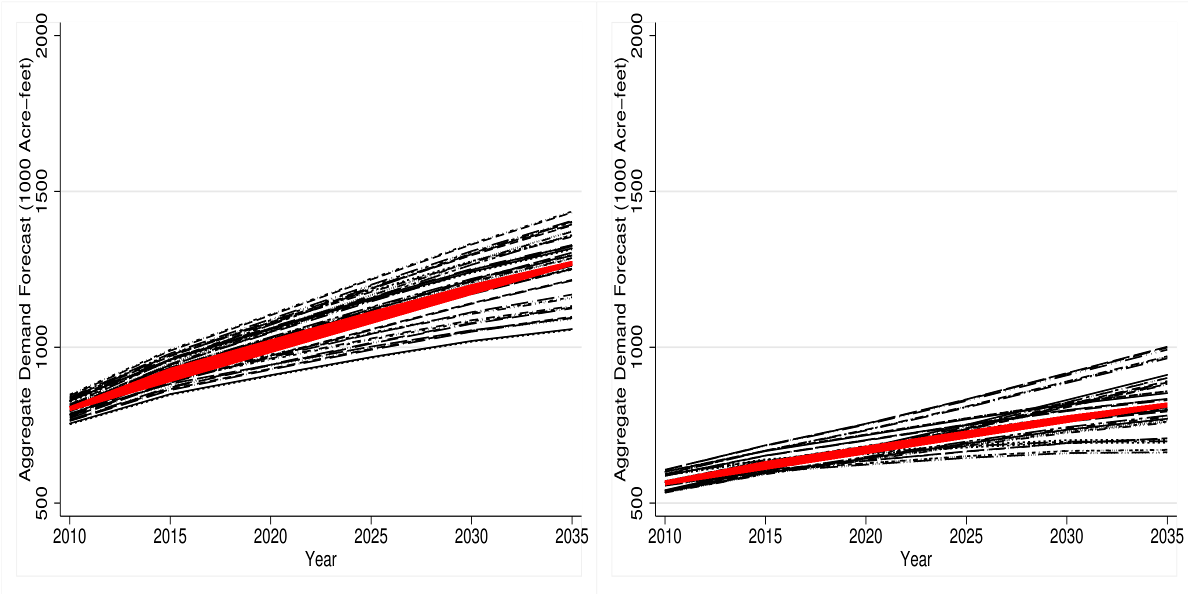

(a) R-Squared Models **(b)** Absolute Aggregate Error Models

Figure 12: Future Projections Generated by the Top Models of Different Criteria. Dependent Variable: Total Quantity[52]

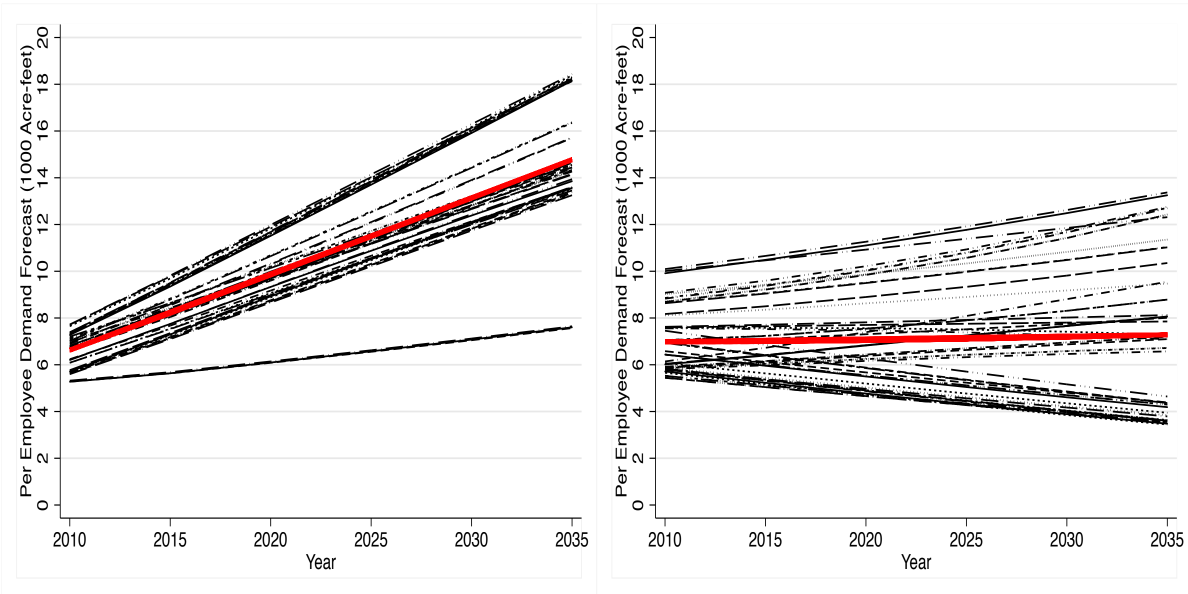

(a) R-Squared Models **(b)** Absolute Aggregate Error Models

Figure 13: Future Projections Generated by the Top Models of Different Criteria. Dependent Variable: Quantity Per Employee

[52]Figures 12 and 13 present the future projections by the models selected under different criteria. We see the similar pattern of dispersion as before.

Chapter 3: Did The California Energy Crisis Increase The Take Up Of The Industrial Assessment Centers (IAC) Program?

3.15 Introduction

Industrial Assessment Centers (IAC) is an energy audit program funded by the Department of Energy targeting small and medium sized firms. It involves engineering teams from partnering institutions across the US conducting on-site visits to participating firms and providing them with tailored recommendations regarding a variety of operational improvements including energy efficiency.

Using a detailed data set containing information on both firm characteristics and the specifics of the recommendations made, this study investigates whether the 2000 California energy crisis had an effect on the take up of the recommendations pertaining to energy efficiency.

After decades of chronic high retail prices, California's electricity market went through a major restructuring starting 1998. With this deregulation attempt, the retail electricity prices were fixed while the wholesale prices were allowed to be determined freely by the market forces expecting that the wholesale price would remain below the retail price. Within less than two years into the restructuring, however, the wholesale prices skyrocketed due to reasons including limitations in hydro generation caused by unfavorable weather conditions, disappearing extra generation capacity associated with a fast growing economy, increased pollution permit prices, and most importantly, supply disruptions that are schemed by corrupt wholesalers who sought to manipulate the prices and exercise market power. The crisis resulted in rolling blackouts, rate hikes, bankruptcies, and law suits costing the Californians billions of dollars.

During the crisis, even though it was initially set to be fixed, the electricity price and the volatility thereof were higher than usual. Price and its volatility are opposing forces that act on an investment project's value. Higher energy prices imply higher savings associated with energy efficiency investments and, therefore, will provide an incentive to invest in efficiency. On the other hand, theory suggests that higher volatility in the prices will increase the uncertainty on the returns of the energy efficiency investments. Higher uncertainty in any investment project will reduce its appeal. While it is impossible to tease apart the magnitude of the individual factors without observing the managers' expectations on the price and its volatility, this study focuses on the net effect of such movements in market signals occurred in the context of an energy crisis. The main goal is to contribute to the knowledge of firm behavior regarding energy efficiency adoptions.

A graphical illustration of how the project take up rates moved over time both in and outside of California can be found in Figures 14 and 15. We see an increase in the three-month moving average of the implementation rates of those recommendations that are evaluated by the firms during crisis in California while it stays somewhat stable everywhere else in the US. Although this may suggest that the take up increase was due to the crisis, the exceedingly variable nature of the monthly take up rate calls for more rigorous analyses.

According to linear probability model estimation results estimated with difference-in-difference method, the crisis was associated with about a 16% increase in the probability that a project is implemented. In order to keep the treatment and the control groups as comparable as possible to ensure credible identification, the firms that applied to be audited and made the decision before the crisis are compared to those that applied right before the crisis and had to decide after the crisis started in and out of California (see Table 21). The coefficient estimate is robust to a variety of different specifications including the incorporation of numerous fixed effects. The results can be found in Table 25. Further robustness checks are provided subsequently.

The organization of the chapter is as follows: next section provides a background on the California energy crisis. A review of the related literature is provided in section 3.17, followed by a description of the data that has been used and the construction of the treatment and control groups in section 3.18. Section 3.19 lays out the model, the estimation framework, and the results. A robustness exercise is presented in section 3.20, and section 3.21 concludes. The table and figures can be found in section 3.22.

3.16 Background: California Energy Crisis

California's electricity sector suffered chronic above-national-average electricity prices for decades. The building of nuclear power plants that eventually turned out to be far more expensive than expected and long term binding contracts that forced the utilities to purchase power at higher-than-market prices were the root cause. Deeming the market as "fragmented, outdated, arcane and unjustifiably complex", the Public Utilities Commission (PUC) voted to restructure the market by opening it up to competition in December of 1995 (Vogel (2000)).

During the restructuring, the investor owned utilities were offered incentives to divest out of electricity generation. The electricity retail price was fixed at 6 cents per kilowatt hour (kwh) while the wholesale price was allowed to be determined in the market. The expectation was that the latter would remain below the fixed retail level, while the difference would help finance the recovery of the stranded costs from the previous decades.[53] In 1998 the deregulation legislation started to take effect.

Electricity industry has unique properties that played a major role in the unfolding of the events. First, it is prohibitively costly to store electricity at large quantities. This means that the electricity consumed at any given moment needs to be generated at the same time. Additionally, the demand for electricity is highly inelastic. Coupled with the inherently rigid consumption patterns, the inelastic demand is further reinforced by the lack of price signals that would reflect the differences in costs of generation throughout different hours of the day, despite the technological feasibility (i.e. the smart meters). Finally, the electricity supply is

[53]After the recovery or March of 2002, whichever comes first, the utilities would just pass through the wholesale electricity prices.

also highly inelastic both in the short run and the long run. In the short run, it is virtually impossible to generate beyond the already established capacity without inflicting irreversible damage on the transmission system. In the long run, supply is tight due to high capital costs.

In addition to these, California itself had special conditions contributed to the severity of the crisis. First of all, excess generation capacity that would buffer the sharp increases in the demand disappeared quickly due to a fast booming economy. Additionally, because of hotter and dryer than average weather conditions, the hydro generation plummeted. Also, the cost of nitrogen oxide emission permits increased by more than an order of magnitude. Last but not least, the supply became increasingly short due to more-often-than-usual power outages which turned out to be a scheme implemented by some of the whole-sellers who wanted to manipulate prices by cutting supply and exercise market power.

With all these factors playing their parts, a major market dysfunction was inevitable. On May 22nd of year 2000, California Independent System Operator (CAISO) declared the first Stage 2 power alert when heavy usage contributed to power reserves dropping to 5 percent. Rolling blackouts started to take effect in June in the San Francisco Bay Area affecting about 97,000 customers of Pacific Gas & Electricity (PG&E), one of the three major utilities in California. The average wholesale prices skyrocketed to more than twice the level the year before. Since San Diego Gas & Electric (SDG&E) had already completed the recovery of the stranded costs, it was free to pass through the increased wholesale prices to the retail customers whose bills tripled during that summer. However the rate hikes were not welcome by the consumer groups and soon enough, the legislature passed AB 265 which froze the rates for SDG&E (Weare (2003)).

The crisis lasted for about a year resulting in rolling blackouts, increase in electricity prices, bankruptcies, and law suits leaving the Californians with billions of dollars in losses. The turmoil subsided a year later thanks to mild weather conditions, federally mandated price caps on the wholesale prices, and price declines in the natural gas market (Borenstein (2002)).

3.17 Literature Review

The energy efficiency gap which refers to the underinvestment in energy efficiency capital, is a widely studied topic in the energy economics literature. It is a multifaceted economical problem not only because it involves inefficient allocation of resources, but also the environmental externality implications of energy production. As climate change receives unprecedented amount of media coverage and public attention, energy efficiency takes place in the middle of many public policy debates.

Jaffe and Stavins (1994) provide an extensive list of possible explanations for slow diffusion of energy efficiency investments despite high rates of return. Listed among the potential non-market-failure explanations are: possible value of delaying due to uncertainty of the returns combined with the irreversibility of the investment, different qualitative attributes of the technologies, costs associated with learning in addition to the monetary cost, and inertia. Underprovision of information due to its public good nature, incentives to wait in order to learn from new adopters' experience, and possible principle-agent problems are articulated as the issues which the markets potentially might be failing to address. On uncertainty and irreversibility, Hassett and Metcalf (1992) show that if the prices and the cost of the investment follow a random walk, then the optimal timing of the investment (provided that the investor can choose a time) is when the annual savings is greater than a certain multiple of the cost where the multiplier is determined by the discount rate, volatility, and the drift parameters of the random walks. However, testing this theory empirically, Sanstad, Blumstein, and Stoft (1995) find that the magnitude of this option value effect is relatively small.

There are studies investigating the possible systematic patterns of energy efficiency investments at the firm level. For example, DeCanio and Watkins (1998) show a positive significant relationship between certain firm characteristics including financial performance of both the firms and the industry group they belong to and their likelihood of joining the EPA's Green Lights program. Analyzing the Department of Energy funded Industrial Assessment Center (IAC) Program, Tonn and Martin (2000) indicate how this program triggers the firms' gradual evolution from being completely new to the idea to a point where they develop an active energy efficiency agenda by staying in touch with the alumni that took part in this study (mostly through employment), and/or receiving further information from IAC websites. Another study that analyzed the IAC program was done by Anderson and Newell (2004). They find that while firms are responsive to economic attributes of the projects, certain technologies are more likely to be adopted than others. Furthermore, the firms are more responsive to costs than savings and this difference is exacerbated with increasing costs.

Another aspect worth examining is how the decision makers adjust their energy consumption as a response to exogenous shocks. Reiss and White (2008) showed that after the unprecedented price increase during the California energy crisis, San Diego households dropped their electricity consumption by 13% within 2 months. Even after the price cap imposed in September 2000, they kept conserving energy which is attributed to media coverage implying possible opportunities for effective non tax interventions to reduce energy consumption.

This study contributes to the knowledge of energy efficiency investments by studying the behavior of small and medium sized manufacturing firms in the context of an energy crisis. Even though the prices did end up increasing, the majority of the burden of the increased prices were carried by the utilities which eventually had to file for bankruptcy. Instead of attempting to quantify the magnitude of the individual forces such as price and volatility increase as well as loss of reliability, this study provides an insight into the net effect of a large energy crisis on energy efficiency investment behavior.

3.18 Data

The IAC provides a publicly available data set on energy audits of small and medium sized manufacturing firms.[54] The audits are performed by 26 IAC centers from 31 universities with engineering programs. The teams normally pay a one-day on-site visit to an industrial plant. During these visits, the plant is toured and various operational parameters are measured. The visits are followed by a written report by the team regarding energy utilization, waste handling and other manufacturing procedures along with a list of recommendations. After giving the firms between six to nine months, the centers call the firms to follow up on the recommendations to find out if the recommendations were implemented. All this data gathered from assessment and the recommendations are formatted into a spreadsheet boilerplate. The database includes firm specific information on sales, sector (SIC code), plant size, and number of employees, as well as information pertaining to each recommendation such as implementation cost, energy savings, and whether or not the project was implemented. The dollar values represent the best estimates of the engineering teams of current costs or projected costs for the coming year.[55] The result is a database with more than 12,000 assessments and 87,000 recommendations.[56]

This study utilizes the data from all audits that took place between January 1998 and June 2000 pertaining to electricity consumption. A total of 2094 such recommendations were made to 610 unique firms from 24 different states by 13 participating institutions.[57] The number of energy recommendations for a given firm range from 1 to 15. 391 of these recommendations were made by 2 California centers (SDSU and SFSU) to 110 unique California firms. Table 18 has the summary statistics.

The variable of interest is the "implementation status"— a binary variable that equals 1 if the project was implemented within 6 to 9 months of the visit, zero otherwise. The database provides information about both firms and recommendations. During the criteria check prior to conducting site visits, IAC teams access firm level data including annual sales, number of employees, plant size, and annual energy expenditure. Variables specific to the projects include estimated implementation costs[58], annual savings and payback. Payback is calculated by

[54] The criteria to qualify for this service are very specific. The firms must have gross annual sales of $100 million or less, consume energy at a cost greater than $100,000 and less than $ 2.5 million per year, have no more than 500 employees, and do not employ a technical staff whose primary duty is energy analysis.

[55] Dr. Michael Muller IAC Field Manager, personal communication

[56] The data documentation could be reached at `http://iac.rutgers.edu/manual_database.php`.

[57] This number does not include the recommendations where there was a rebate issued. The cost data on these recommendations were reported net of the rebate, however the dataset does not contain information on how much the rebate was or when and how it was issued. Because the existence of rebate can be interpreted differently across the decision makers and the data does not grant a good grasp of the nature of these rebates, the recommendations with rebate were omitted. Additionally, observations with missing values on key variables like implementation cost and savings were not included as well as the ones with negative payback ratios.

[58] Firms are provided with an estimate of capital and non-capital costs. Firms do their own

dividing the cost by the annual savings to calculate roughly how many years it takes the project to pay itself back.

Table 19 summarizes the estimated annual savings and the implementation costs of the projects. The mean is larger than even the 75th percentile for both variables implying highly skewed distributions. About 90% of the projects costed under $20K while 9 projects had implementation costs of over a million dollars. As with the savings, more than 90% of the projects had savings less than $15K while 4 projects had estimated savings of over a million dollars. A scatter plot of the savings against the cost can be seen in Figure 16 which suggests a high positive correlation. This might be an artifact of the characteristics of the firms interested in energy audits. In other words, these firms might have already implemented possible projects with low costs and high savings.

One important figure, which is exemplary of the energy efficiency gap discussed in the literature, is the low implementation rates despite the low payback figures. Only about 47% of those recommendations were implemented. The implementation costs and estimated annual savings[59] of the projects have wide ranges and so do the firm characteristics such as sales, number of employees, floor area, and annual energy expenditures.

Table 20 gives a further break down of the details with respect to the types of projects. The majority of the observations are "Motor systems" (e.g. operation, maintenance and repair of motor systems) and "Building and grounds" (e.g. lighting, heating, ventilation and air conditioning (HVAC), and infiltration) followed by "Thermal systems" (e.g. operation and maintenance of heat treating, recovery, and containment systems; and cooling), "Electrical power" (e.g. scheduling and generation of power), "Operations" (e.g. process and material specific efficiency improvements), and "Combustion systems" (e.g. operation and maintenance of furnaces, ovens, and boilers or possible fuel switching). The positive correlation between the costs and the savings of the projects is evident from the table.

California firms that decided whether to implement during the crisis were defined as the treatment group. Selection bias is a potential empirical challenge, as the firms voluntarily participate in these programs. The firms that are interested in such a program under different circumstances (i.e. crisis vs. not) could differ in their project valuation which can potentially be a confounding factor. In order to keep the treatment and control groups as comparable as possible, it is necessary to look at the firms that chose to be audited under similar conditions, but faced an exogenous change in circumstances while deciding whether or not to implement the recommendations. With this idea in mind, the treatment group is determined in the following manner. First, starting point of the crisis was marked as June 2000 as this is when the blackouts and the extensive media coverage started. The firms are contacted 6-9 months after the audit for follow up. This means that the California firms that are visited between January 2000 and June 2000 (inclusive) signed up for the audits without any expectation of the crisis and were contacted

estimation of the cost before implementing. This study only uses firms' own estimates of the cost because of many missing values in the cost numbers produced by the IAC teams

[59] All the dollar values reported here are inflation adjusted for 2000 dollars using the producer price index (finished goods, series WPUSOP3000) from the US Bureau of Labor Statistics (2010).

for follow up sometime during the crisis.[60] Therefore the final decision recorded in the database for these recommendations must be made sometime during the crisis. On the other hand, firms that signed up and are audited before August 1999 were clearly not affected (control) by the shock during their decision making period *and* were not contacted for follow up sometime during the crisis. The firms that are audited between September and December 1999 may or may not have been contacted during the crisis. These observations are dropped because the follow up date is not certain.

Hence, to summarize, time period between January 1998 and August 1999 is labeled as "pre-crisis", the period between September to December of 1999 is dropped, and the the period of January to June 2000 is labeled as "crisis". The observations before January 1998 are not included to prevent any other unobservable factors from convoluting the sample.

Table 22 compares the covariates of the California firms that were in the crisis group to those of the firms that were either pre-crisis group (both in an outside of California) and the firms in the crisis group but outside of California. P-values of the t-test results of their differences are given at the last column. Despite the fact that these factors are controlled for in the ensuing regressions, the overlap in the observables across the groups provides an advantage for the reliability of these regressions.

Tables 23 and 24 compares the breakdown of the project types and standard industrial classification (SIC) across the groups, respectively. We see that the share of the projects recommended and the industries represented by the groups are similar. This provides further reassurance for the comparability of the firms in the "treatment" group (California firms visited during the "crisis" period) with the other firms in the sample.

3.19 Modeling and Estimation Results

In the field of energy economics, there is a considerable amount of concern over the inadequacy of the energy efficiency investments despite the seemingly high potential savings. By what is referred to as the "energy paradox", it is suggested that the decision makers apply surprisingly high implicit discount rates to the future savings of the energy conserving projects. Previous studies have estimated discount rates as high as 20-50% (Hausman (1979) and Train (1985)).

Hassett and Metcalf (1992) suggest that the irreversible nature of energy efficiency investments together with the uncertainty associated with the energy price and the cost of the investment will deter such decisions and could offer a plausible explanation for what puzzled the energy economists.

To briefly summarize their model and its implications, suppose a firm constantly utilizes 1 unit of energy and has an opportunity to invest in equipment that will save a fraction, δ, of energy consumption. Once implemented, the capital invested in this project is completely sunk. Also let the energy price and the

[60] The month of June is included because based on personal communication with Prof. Asfaw Beyene of San Diego IAC engineering team, it takes engineers more or less a month to get to the firms' plants once it is established that the firm is interested.

cost of the energy conserving equipment follow geometric Brownian motion, i.e. continuos time random walk. The motion processes are specified by the following equations:

$$dP_{it} = \mu_p P_{it} dt + \sigma_p P_{it} dz_p$$
$$dK_{it} = \mu_k K_{it} dt + \sigma_k K_{it} dz_k$$

where i is the index for the decision maker and t for time; z is the standard Brownian motion process for the corresponding variable which has a change of dz with zero mean and unit variance. The change in P_{it} over time t has mean $\mu_p t$ and variance $\sigma_p^2 t$. They show that the optimal time to invest in the energy efficiency project is when:

$$\delta_i P_{it} > \frac{b}{b-1}(\gamma_i - \mu_p) K_{it}$$

where:

$$b = \frac{0.5\sigma_0^2 - \alpha + \sqrt{(0.5\sigma_0^2 - \alpha)^2 + 2(\gamma_i - \mu_c)\sigma_0^2}}{\sigma_0^2};$$

γ_i is the discount rate;

σ_0^2 and α are, respectively, the variance and the drift of the geometric Brownian motion process given by $\frac{P_{it}}{K_{it}}$.

Note that the term $\frac{b}{b-1}$ will approach 1 as σ_0^2 goes to zero.[61] In other words, without uncertainty, the investment happens exactly when the net present value (NPV) of the savings equals the cost of investment. We see from the equation that the higher energy price will increase while higher uncertainty will decrease the incentive for the manager to implement the project. For positive values of σ_0^2, optimal investment timing occurs not only when the project is "in the money", but when it is "deep in the money". When the returns on the investment are uncertain, a rational investor will require a larger present value for compensation to give up the "option to invest", and terminate the ability to wait and see the realization of future costs and prices to make a more informed decision in the future. Put differently when there is uncertainty, the project's net present value needs to be large enough to cover both the investment cost and the opportunity cost of forgoing the option to wait.

In the case of the IAC recommendations, arguably every project is an irreversible investment with a risky return. A manager will implement the project if the following condition holds:

$$Savings_{ij} > \frac{b}{b-1}(\gamma_i - \mu_p) Cost_{ij}$$

In order to be able to make meaningful statements regarding the discount rates

[61] This is true as long as $\alpha + \mu_c < \gamma$ In other words if the growth of $\frac{P_{it}}{K_{it}}$ is slower than the discount rate. Otherwise it is never optimal to undertake the investment.

with which the firms evaluate the energy savings, we need to observe the individual project lifetimes, the exact time that the project was implemented, and the managers' estimate for the mean and the variance the price and investment cost movements. Due to lack of data on these variables, no such estimations can be made at this time. However, we do observe visit dates, estimated costs and savings as well as whether or not the project was implemented within 6-9 months of the visit. Therefore we are able to do analysis on the effect of the energy crisis using a reduced form approach.

The effect of the California energy crisis on the implementation rates of the IAC recommendations is estimated using the following linear probability model:

$$Imp_{ij} = \beta_0 + \beta_1 S_{ij} + \beta_2 C_{ij} + \beta_\mathbf{3} \mathbf{X_i} + \beta_4 Ca_i + \beta_5 Crisis_{ij} + \beta_7 Ca_i \times Crisis_j + \varepsilon_{ij}$$

where:
Imp_{ij} is a binary variable equals 1 if the project is j is implemented by firm i;
S_{ij} is the logarithm of expected annual savings;
C_{ij} is the logarithm of expected total implementation cost;
$\mathbf{X_i}$ is the vector of firm characteristics such as firm's size and energy intensity;
Ca_i is a binary variable equals 1 if firm i is located in California;
$Crisis_{ij}$ is an indicator of whether or not the California energy crisis was going on during firm i's implementation decision of project j;
$Ca \times Crisis_{ij}$ is the interaction of Ca_i and $Crisis_{ij}$;
and ε_{ij} represents the unobserved factors that affected the implementation.

Plant area, logarithm of the annual sales, and number of employees are used to proxy for firm size. Energy intensity is calculated by dividing the annual energy expenditure by total sales which is then standardized by dividing by the industry average by state. This could be thought of as the share of energy input in the total output of the firm. Although the theory implicitly assumes that these characteristics should not matter, in practice firm characteristics may play a role in their project evaluation. ε_{ij} which includes firm's overall value of the energy efficiency projects, opportunity cost of not waiting until some of the uncertainty settles etc.

In an empirical setting where the dependent variable is binary, logit could have been chosen to model the decision making process. An important advantage of logit is, unlike linear probability model, it insures that the probability estimate is between zero and one. Yet in this study, linear probability model was chosen to be the method of estimation for an important reason. Logit model hinges on the assumption that the utility function can be broken into observed and unobserved components. The unobserved component is modeled as a random variable that is independently and identically distributed with type I extreme value distribution (Train (2009)). Note that in the representation of the utility function, the scale of the utility does not matter. In other words, the utility of all the options could be divided by the same number and this would not change the option with the highest utility. Traditionally, with the standard logit model, the utility is scaled in a way that the error variance is $\pi^2/6$. So the parameters are identified up to the scale parameter. The coefficient that we obtain is the ratio of the actual utility coefficient to the scale parameter, neither of which can individually be identified. While this does not pose a problem when the errors are homoscedastic, a different model needs to be estimated for all the groups if the unobserved portion of the utility is heteroskedastic.[62]

In this study, the data comes from a wide variety of geographic locations and industries from across the US. Since assuming homogeneity in the unobserved factors would be rather unrealistic, dividing the data into more homogenous subgroups and estimating a separate model each one of them would be required. This, however, would obviate the tractability of esimation process and the interpretability of the results. Linear probability model emerges as a simpler yet a more feasible approach. It allows including all the data from the period to be included in the same model which paves the way for the utilization of the difference-in-difference strategy.

The models are estimated using ordinary least squares regression method with robust standard errors. Table 25 has the results. The parameter of interest is on the third row. The analysis started with a simple model with only group indicators. Since the crisis is exogenous, this specification by itself should be enough to identify the parameter of interest. However in order to reduce the noise, various factors are gradually controlled for the progression of which can be found in the table. We see that the parameter estimate and its significance stays stable as the state indicators, project characteristics (payback time, cost, savings), firm characteristics (energy intensity, plant area, and number of employees), time trend, quarter indicators, and industry indicators are gradually added to the model. The models unanimously estimate that the crisis is associated with about a 16% increase in the probability of adoption. Results also indicate that California on average had a lower rate of adoption during the time of the study. The project cost, as expected, has a negative significant coefficient. For every percent increase in the project cost, the model predicts about 5% decrease in implementation rate. Once the cost is accounted for the payback period loses it's significance which implies that the size of the cost itself is the biggest determinant without much regard to the savings opportunities, a finding that is consistent with that of Anderson and Newell (2004). This result may suggest a potential liquidity or credit constraint faced by the small and medium sized firms.

3.20 Robustness

In addition to estimating the model with different variables, the robustness is checked in the time dimension as well with what I call "placebos". The records of California recommendations in the IAC dataset used for this study date from April 1991 and to September 2009. Starting from the beginning of this time period, the regression estimation procedure was repeated with 1 month increments in the following fashion: The original regressions are run using data that spreads over a 30-month period (see Table 21 for details) which is divided into 3 portions: pre-

[62]For further details with an example see: Train (2009) pp26.

crisis, drop-out, and crisis, respectively. Starting April 1991, 30-month horizon is set aside and divided into 3 segments that are of same length as in the original regressions using which the final regression equation on Table 25 is estimated. The same procedure is repeated in one-month increments, in other words, the beginning and the ending of the regression period was moved up by one month at a time until the end of the data availability was reached. After each regression estimation, the t-score of the coefficient of the *hypothetical* treatment coefficient is recorded. This procedure yielded a total of 192 regression estimations, and only 19 (about 10%) of which had the t-score greater then or equal to that of the original equation. In other words the probability of getting a t-score that is as big as the one that we got in the original regression by chance is 1 out of 10. The histogram of the t-scores can be found in Figure 17.

3.21 Conclusion

IAC is a Department of Energy funded industry audit program that targets small and medium sized firms to provide them with engineering recommendations in an effort to improve their operations in a variety of domains including energy efficiency. Using s detailed data set containing information on both firm characteristics and the specifics of the recommendations made, the effect of the 2000 California energy crisis on the take up of the recommendations pertaining to energy efficiency is measured. California energy crisis refers to the a sequence of developments in California's energy sector originating from the attempt to restructure the industry which resulted in highly publicized rolling blackouts, skyrocketing electricity prices, utility bankruptcies, and lawsuits. The time series graphs of implementation rates in and outside of California can be seen in Figures 14 and 15. The effect of the crisis is estimated using linear probability model with a difference-in-difference strategy. The results suggest that the crisis was associated with about a 16% increase in the implementation rate and are robust to a variety of different specifications including the incorporation of different time, state, and industry fixed effects. The robustness of the estimation is checked in the time dimension as well with the reestimaton of the regression for other time periods and studying the t-scores of the hypothetical treatment coefficients.

Future research projects that seek to improve the take up of the this audit program could potentially deploy experimental methods to measure the effectiveness of loan availability. Also, further studies are needed to understand the sources and the nature of differences across different states and industries.

3.22 Figures and Tables

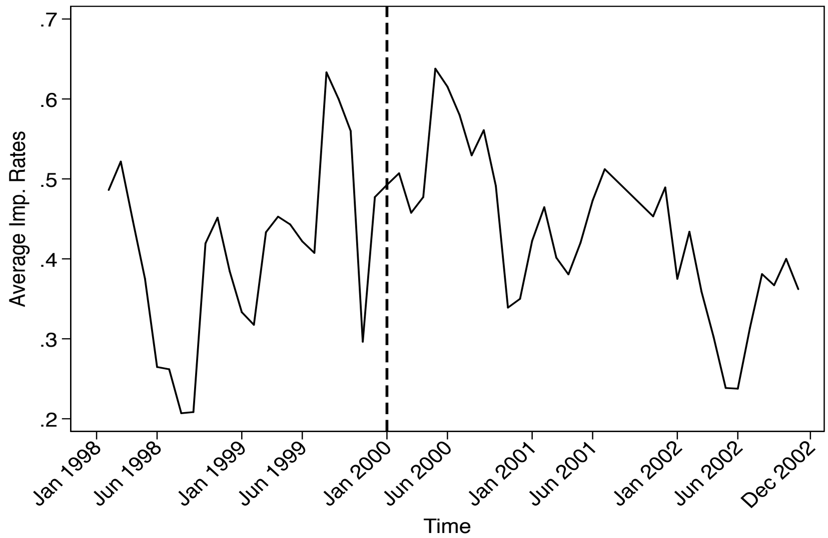

Figure 14: Implementation Rates in California Over Time

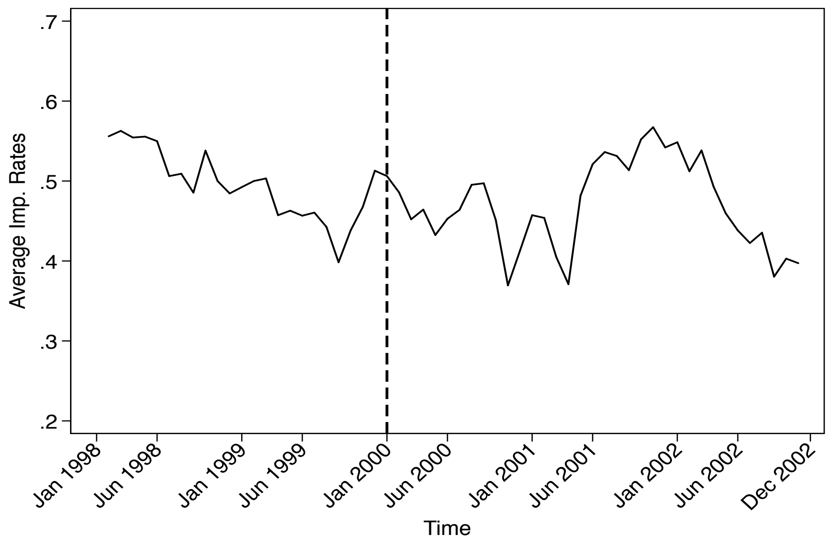

Figure 15: Implementation Rates out of California Over Time
[63]

[63]Figures 14 and 15 show how the 3-month moving average implementation rate of the IAC recommendations behaved over time in and out of California between January 1998 and December 2002. The dashed vertical black line represents the date of the recommendations for which the follow up calls are made when the crisis started. We see an increase in the implementation rates right after the crisis while it was somewhat stable out of California. However, the significant variability in the implementations rates over time calls for more rigorous quantitative analyses which can be found in the subsequent sections.

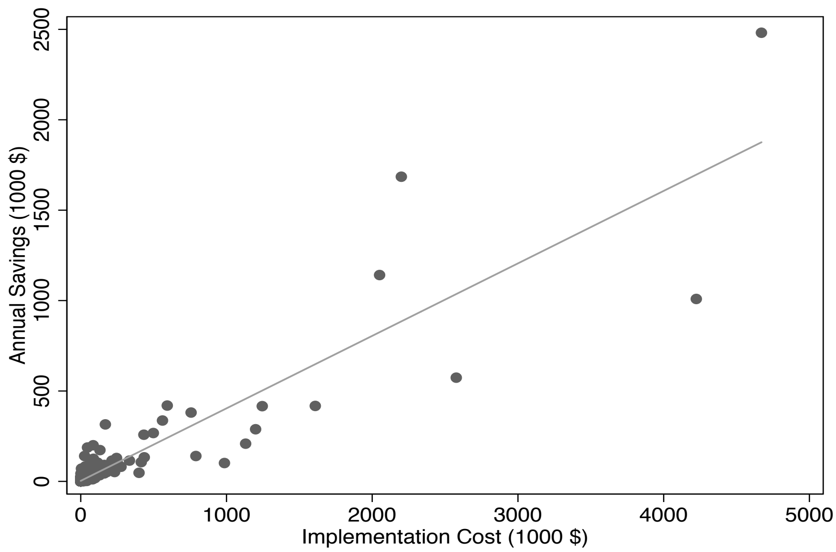

Figure 16: Scatterplot of Estimated Annual Savings and Implementation Costs
64

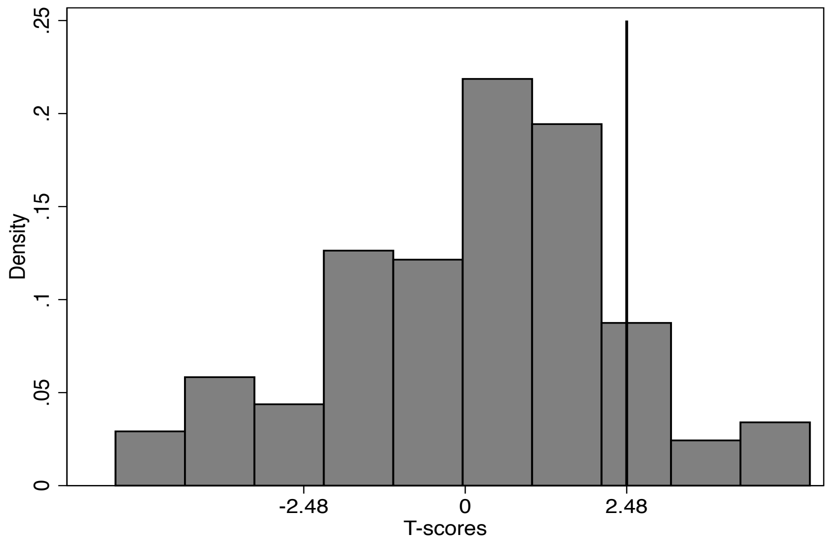

Figure 17: T-scores of the Placebo Regressions

65

[64] Figure 16 shows the scatterplot of estimated annual savings and the implementation costs of the projects.

[65] In order to test the robustness of the regression results in the time dimension, the regressions were reestimated using different time segments within the data where the time segment is incremented by one month at a time. Figure 17 shows the histogram of the t-scores of the hypothetical treatment coefficients of these regressions. Only 1 out of 10 regressions yielded a t-score that is as big as the ones that is obtained in the original regression.

Table 18: Summary Statistics[66]

	Mean	StD	Min	Max
Imp. Status	0.47	0.50	0.00	1.00
Payback Time (Years)	1.37	1.78	0.00	42.63
Log Imp. Cost	7.33	1.95	2.36	15.36
Log Ann. Savings	7.70	1.49	2.69	14.72
Log Ann. Sales	16.77	0.94	13.85	20.03
Energy Intensity	0.90	0.90	0.02	7.84
Plant Size (Mil. Sqf.)	0.91	9.95	0.00	224.25
Log Ann. Energy Exp.	12.58	1.03	9.21	16.27
No of Employees	164.66	148.34	1.00	1500.00
Crisis	0.24	0.43	0.00	1.00
California	0.19	0.39	0.00	1.00
Observations	2094			

Table 19: Summary of Annual Savings and Implementation Costs of the Projects[67]

	25th Perc.	Median	75th Perc.	Mean	StD
Annual Savings (1000 $)	0.77	2.11	5.58	11.08	79.05
Implementation Cost (1000 $)	0.42	1.29	5.64	20.32	178.02
Observations	2094				

[66] Table 18 has the summary statistics. In total, data from 2094 recommendations made to 611 unique firms are used. 388 of these recommendations were made to 110 unique California firms. The dates vary from January 1998 to June 2000. The "crisis" is a binary variable equal to one if a firm got the audit right before the crisis started but had to decide during the crisis (for more information on the dates and treatment assignment, see Table 21). "California" represents whether the firm is located in California.

[67] Table 19 summarizes the estimated annual savings and the implementation costs of the projects. We see that these figures are highly skewed. About 90% of the projects costed under $20K while 9 projects had an implementation cost of over a million dollars. As with the savings more than 90% of the projects had savings less than $15K while 4 projects had estimated savings of over a million dollars.

Table 20: Implementation Rates, Payback, Cost, and Savings by Project Type[68]

	No of Proj.	Imp. Rate	Payback	Cost (1000 USD)	Savings (1000 USD)
Ancillary costs	2	1.00	3.81	21175.54	5608.24
Building and grounds	759	0.40	1.58	5706.11	3956.73
Combustion systems	90	0.17	1.35	34046.59	27019.20
Electric power	94	0.34	1.65	259997.57	110878.79
Industrial design	3	0.67	2.10	159815.04	77236.19
Motor systems	977	0.57	1.17	6682.25	5341.68
Operations	53	0.47	0.69	1667.74	3766.64
Thermal systems	116	0.34	1.78	30806.22	14538.44
Total	2094	0.47	1.37	20318.61	11081.82

[68]Table 20 breaks down the important characteristics (implementations rate, cost, savings, and payback period) of the recommendations with respect to the types of operations being targeted. The majority of the recommendations are pertaining to “Motor systems” (e.g. operation, maintenance, and repair of motor systems) and “Building and grounds” (e.g. lighting, HVAC, and infiltration) followed by “Thermal systems” (e.g. operation and maintenance of heat treating, recovery, and containment systems; and cooling), “Electrical power” (e.g. scheduling and generation of power), “Operations” (e.g. process and material specific efficiency improvements), and “Combustion systems” (e.g. operation and maintenance of furnaces, ovens, and boilers or possible fuel switching). The positive correlation between the costs and the savings of the projects is evident from the table.

Table 21: Summary of the Treatment Assignment[69]

Audit Range	Follow-up	Crisis	In Cal.	Out of Cal.	Total
Jan 1998 - Aug 1999	Jul 1998 - May 2000	0	284	1314	1598
Sep 1999 - Dec 1999	Mar 2000 - Sep 2000	NA	29	179	208
Jan 2000 - Jun 2000	Jul 2000 - Mar 2001	1	107	389	496
Total			420	1882	2302

[69]Table 21 summarizes how the "treatment" status is defined. In order to keep the treatment and control groups as comparable as possible, it is necessary to look at the firms that chose to be audited under similar conditions, but faced an exogenous change in circumstances while deciding whether or not to implement the recommendations. The starting point of the crises was marked as June 2000. Given that the firms are contacted 6-9 months after the audit we can conclude that the California firms that are visited between January 2000 and June 2000 (inclusive) signed up for the audits without any expectation of the crisis and were contacted for follow up sometime during the crisis. Hence, the decision they made that is recorded in the database must be made sometime during the crisis. On the other hand, firms that signed up and are audited before August 1999 were clearly not affected by the shock during their decision making period *and* were not contacted for follow up sometime during the crisis. The firms that are audited between September and December 1999 might or might not be contacted during the crisis. These observations are dropped because the follow up date is not certain. So to summarize, the firms that were audited between January 1998 and August 1999 are in the "pre-crisis" group, the firms that are audited between September and December 1999 are dropped, and the firms that got audited between January and June 2000 are the "crisis" group. Firms that were audited before January 1998 are not included to prevent any other unobservable factors from convoluting the sample.

Table 22: Comparison of the Covariates Accross the Groups[70]

	Non-Treatment Group	Treatment Group	Difference	P-Value
Payback Time (Years)	1.37	1.37	0.00	0.489
Imp. Cost (1000)	20.79	11.54	9.25	0.300
Ann. Savings (1000)	11.17	9.37	1.81	0.409
Log Ann. Sales	16.77	16.72	0.05	0.311
Energy Intensity	0.90	0.86	0.03	0.350
Plant Size (Mil. Sqf.)	0.95	0.30	0.65	0.257
Log Ann. Energy Exp.	12.58	12.58	-0.00	0.501
No of Employees	165.44	150.32	15.12	0.152
Observations	2094			

[70] Table 22 compares the covariates of the California firms that were in the crisis group to those of the firms that were either pre-crisis group (both in an outside of California) and the firms in the crisis group but outside of California. P-values of the t-test results of their differences are given at the last column. Despite the fact that these factors are controlled for in the ensuing regressions, the overlap in the observables across the groups provides an advantage for their reliability.

Table 23: Groups by Project Types[71]

Project Type	**Non-Treatment Group**		**Treatment Group**		**Total**	
	Obs.	**%**	**Obs.**	**%**	**Obs.**	**%**
Motor sys.	933	46.9	44	41.1	977	46.7
Building and grnds	726	36.5	33	30.8	759	36.2
Therm. sys.	106	5.3	10	9.3	116	5.5
Combust. sys.	75	3.7	15	14.0	90	4.3
Operations	48	2.4	5	4.7	53	2.5
Elec. power	94	4.7	0	0.0	94	4.5
Indust. design	3	0.2	0	0.0	3	0.1
Ancill. costs	2	0.1	0	0.0	2	0.0
Total	1,987	100	107	100.0	2,094	

[71]Tables 23 and 24 compare the breakdown of the project types and standard industrial classification (SIC) across the groups, respectively. Here the "treatment group" refers to the California firms that made the decision while facing the crisis. We see that the share of the projects recommended and the industries represented by the groups are similar. In the regressions the indicator variables for project types and the SIC codes are utilized to control for the time invariant aggregate effect of each project type and industry group. However, the fact that they are well represented across the groups provides further reassurance into the comparability of the firms in the "treatment" group (California firms visited during the "crisis" period) with the other firms in the sample.

Table 24: Treatment and Control Groups by SIC Codes

		Non-Treatment Group		Treatment Group		Total	
SIC Code	**Industry**	**No of Obs.**	**Pctg.**	**No of Obs.**	**Pctg.**	**No of Obs.**	**Pctg.**
34	Fabricated Metal Products	320	16.10	25	23.36	345	16.48
30	Rubber & Miscellaneous Plastics Products	295	14.85	10	9.35	305	14.57
35	Industrial Machinery & Equipment	178	8.96	3	2.80	181	8.64
20	Food & Kindred Products	153	7.70	7	6.54	160	7.64
37	Transportation Equipment	146	7.35	11	10.28	157	7.50
33	Primary Metal Industries	135	6.79	21	19.63	156	7.45
36	Electronic & Other Electric Equipment	126	6.34	6	5.61	132	6.30
32	Stone, Clay, & Glass Products	126	6.34	4	3.74	130	6.21
28	Chemical & Allied Products	87	4.38	6	5.61	93	4.44
26	Paper & Allied Products	83	4.18	8	7.48	91	4.35
27	Printing & Publishing	80	4.03	0	0.00	80	3.82
24	Lumber & Wood Products	79	3.98	0	0.00	79	3.77
38	Instruments & Related Products	52	2.62	0	0.00	52	2.48
25	Furniture & Fixtures	31	1.56	0	0.00	31	1.48
39	Miscellaneous Manufacturing Industries	29	1.46	0	0.00	29	1.38
22	Textile Mill Products	28	1.41	0	0.00	28	1.34
23	Apparel & Other Textile Products	21	1.06	5	4.67	26	1.24
29	Petroleum & Coal Products	12	0.60	1	0.93	13	0.62
31	Leather & Leather Products	6	0.30	0	0.00	6	0.29
Total		1,987	100.00	107	100.00	2,094	100.00

Table 25: Regression Results[72]

	Imp. Status	Imp. Status	Imp. Status	Imp. Status	Imp. Status	Imp. Status	Imp. Status
California	-0.121***	-0.237	-0.246	-0.327**	-0.355***	-0.366***	-0.388***
	(0.032)	(0.130)	(0.129)	(0.121)	(0.122)	(0.124)	(0.129)
Crisis	-0.00844	-0.0280	-0.0296	-0.0355	0.0242	0.0123	0.0317
	(0.029)	(0.029)	(0.029)	(0.029)	(0.046)	(0.050)	(0.051)
California x Crisis	0.150**	0.170**	0.169**	0.163**	0.163**	0.165**	0.159**
	(0.063)	(0.064)	(0.064)	(0.062)	(0.062)	(0.062)	(0.064)
Payback Time (Years)			-0.0165	0.00785	0.00780	0.00781	0.00807
			(0.009)	(0.007)	(0.007)	(0.007)	(0.006)
Log Imp. Cost				-0.0546***	-0.0545***	-0.0543***	-0.0548***
				(0.010)	(0.010)	(0.010)	(0.010)
Log Ann. Savings				0.000407	0.00186	0.00177	0.00379
				(0.012)	(0.012)	(0.012)	(0.013)
Montly Trend					-0.00350	-0.00309	-0.00460*
					(0.002)	(0.002)	(0.002)
Constant	0.484***	0.600***	0.634***	0.479	0.559	0.560	0.611*
	(0.014)	(0.127)	(0.127)	(0.282)	(0.286)	(0.286)	(0.295)
Observations	2094	2094	2094	2094	2094	2094	2094
R^2	0.007	0.047	0.050	0.081	0.082	0.082	0.097
Quarter Indicators	No	No	No	No	No	Yes	Yes
SIC Indicators	No	No	No	No	No	No	Yes
State Indicators	No	Yes	Yes	Yes	Yes	Yes	Yes
Firm Characteristics	No	No	No	Yes	Yes	Yes	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.025$, *** $p < 0.005$

References

Abdel-Dayem, M. (1987): "Development of land drainage in Egypt," in *Proc. Symp. 25th International Course on Land Drainage. Twenty-Five Years of Drainage Experience. J. Vos (ed.). Publ*, vol. 42, pp. 195–204.

Adamowski, J., and C. Karapataki (2010): "Comparison of multivariate regression and artificial neural networks for peak urban water-demand forecasting: evaluation of different ANN learning algorithms," *Journal of Hydrologic Engineering*, 15(10), 729–743.

Anderson, S., and R. Newell (2004): "Information programs for technology adoption: the case of energy-efficiency audits," *Resource and Energy Economics*, 26(1), 27–50.

Auffhammer, M., and R. Steinhauser (2012): "Forecasting the path of US CO2 emissions using state-level information," *Review of Economics and Statistics*, 94(1), 172–185.

Babin, F. G., C. E. Willis, and P. G. Allen (1982): "Estimation of substitution possibilities between water and other production inputs," *American Journal of Agricultural Economics*, 64(1), 148–151.

Ben-Akiva, M., and D. Bolduc (1996): *Multinomial probit with a logit kernel and a general parametric specification of the covariance structure*. D epartement d' economique, Universit e laval with Department of Civil and Environmental Engineering, Massachusetts Institute of Technology.

Berndt, E. K., B. H. Hall, R. E. Hall, and J. Hausman (1974): "Estimation and inference in nonlinear structural models," *Annals of Economic and Social Measurement*, 3(4), 653–665.

Bhatti, A. (1987): "A review of planning strategies of salinity control and reclamation projects in Pakistan," in *Proceedings, Symposium 25th International Course on Land Drainage. ILRI publ*, vol. 42.

Boland, J. J. (1997): "Assessing urban water use and the role of water conservation measures under climate uncertainty," *Climatic Change*, 37(1), 157–176.

Borenstein, S. (2002): "The trouble with electricity markets: understanding California's restructuring disaster," *The Journal of Economic Perspectives*, 16(1), 191–211.

Brownell, H., and S. D. Eaton (1975): "The Colorado River Salinity Problem with Mexico," *American Journal of International Law*, pp. 255–271.

Calloway, J. A., A. K. Schwartz, and R. G. Thompson (1974): "Industrial economic model of water use and waste treatment for ammonia," *Water Resources Research*, 10(4), 650–658.

Cook, Z., S. M. Urban, M. A. Maupin, R. Pratt, and J. Church (2001): "Domestic, Commercial, Municipal and Industrial Water Demand Assessment and Forecast in Ada and Canyon Counties, Idaho," Idaho Department of Water Resources.

Datta, K., V. Sharma, and D. Sharma (1998): "Estimation of a production function for wheat under saline conditions," *Agricultural Water Management*, 36(1), 85–94.

de Bono, A., D. Del Pietro, G. Giuliani, A. Harayama, G. Le Sourd, and R. Diana (2004): *Freshwater in Europe: facts, figures and maps*. United Nations Environment Programme.

De La Torre, A. (1987): "Drainage and land reclamation in Peru," *Publication. International Inst. for Land Reclamation and Improvement (Netherlands)*.

De Rooy, J. (1974): "Price responsiveness of the industrial demand for water," *Water Resources Research*, 10(3), 403–406.

DeCanio, S., and W. Watkins (1998): "Investment in energy efficiency: do the characteristics of firms matter?," *Review of Economics and Statistics*, 80(1), 95–107.

Dinar, A., J. Letey, and K. C. Knapp (1985): "Economic evaluation of salinity, drainage and non-uniformity of infiltrated irrigation water," *Agricultural water management*, 10(3), 221–233.

Dinar, A., J. Rhoades, P. Nash, and B. Waggoner (1991): "Production functions relating crop yield, water quality and quantity, soil salinity and drainage volume," *Agricultural Water Management*, 19(1), 51–66.

Dziegielewski, B., and J. J. Boland (1989): "FORECASTING URBAN WATER USE: THE IWR-MAIN MODEL1," *JAWRA Journal of the American Water Resources Association*, 25(1), 101–109.

El-Ashry, M. T., J. van Schilfgaarde, and S. Schiffman (1985): "Salinity pollution from irrigated agriculture," *Journal of Soil and Water Conservation*, 40(1), 48–52.

Francois, L. (1982): "Narrow row cotton (Gossypium hirsutum L.) under saline conditions," *Irrigation Science*, 3(3), 149–156.

Greene, W. H. (2008): *Econometric analysis*. Pearson Education, Inc.

Hajivassiliou, V. A., and P. A. Ruud (1994): "Chapter 40 Classical estimation methods for {LDV} models using simulation," in *Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden, vol. 4, pp. 2383 – 2441. Elsevier.

Hanak, E. (2011): *Managing California's water: From conflict to reconciliation*. Public Policy Instit. of CA.

Hanson, B., S. R. Grattan, and A. Fulton (1999): *Agricultural salinity and drainage*. University of California Irrigation Program, University of California, Davis.

Hassett, K. A., and G. E. Metcalf (1992): "Energy tax credits and residential conservation investment," *Working Paper No 4020, National Bureau of Economic Research Working Paper Series*.

Hausman, J. A. (1979): "Individual discount rates and the purchase and utilization of energy-using durables," *The Bell Journal of Economics*, pp. 33–54.

Hoffman, G. J. (2010): "Salt Tolerance of Crops in the Southern Sacramento-San Joaquin Delta," *Report for the California Environmental Protection Agency*.

Jacobsen, T., and R. M. Adams (1958): "Salt and silt in ancient Mesopotamian agriculture," *Science*, 128(3334), 1251–1258.

Jaffe, A., and R. Stavins (1994): "The energy-efficiency gap What does it mean?," *Energy Policy*, 22(10), 804–810.

Kenny, J. F., N. L. Barber, S. S. Hutson, K. S. Linsey, J. K. Lovelace, and M. A. Maupin (2014): *Estimated use of water in the United States in 2010*. US Geological Survey Reston, VA.

Knapp, K. C., B. K. Stevens, J. Letey, and J. Oster (1990): "A dynamic optimization model for irrigation investment and management under limited drainage conditions," *Water Resources Research*, 26(7), 1335–1343.

Kurukulasuriya, P., R. Mendelsohn, et al. (2007): *Crop Selection*, vol. 4307. World Bank Publications.

Maas, E., and G. Hoffman (1977): "Crop salt tolerance: evaluation of existing data," *J. Irrig. Drain*, 504(103), 115–134.

MacEwan, D., and R. E. Howitt (2012): "Behavioral Salinity Response: Estimating Salinity Policies from Remote Sensed Micro-Data," in *2012 Conference (56th), February 7-10, 2012, Freemantle, Australia*, no. 124331 in 1. Australian Agricultural and Resource Economics Society.

McCracken, M. W., and K. D. West (2002): "Inference about predictive ability," *A Companion to Economic Forecasting*, pp. 299–321.

Morales, M. A., J. M. Martin, and J. P. Heaney (2009): "Methods for Estimating Commercial, Industrial and Institutional Water Use," in *FSAWWA Water Conference*.

Oosterbaan, R. (1988): "Effectiveness and Social/Environmental Impacts of Irrigation Projects: A Criticial Review," *ILRI Annual Report*, 1988, 18–34.

Pannell, D. J. (2001): "Dryland salinity: economic, scientific, social and policy dimensions," *Australian Journal of Agricultural and Resource Economics*, 45(4), 517–546.

Pitman, M. G., and A. Läuchli (2002): "Global impact of salinity and agricultural ecosystems," in *Salinity: environment-plants-molecules*, pp. 3–20. Springer.

Reiss, P., and M. White (2008): "What changes energy consumption? Prices and public pressures," *The RAND Journal of Economics*, 39(3), 636–663.

Renzetti, S. (1992): "Estimating the structure of industrial water demands: the case of Canadian manufacturing," *Land Economics*, pp. 396–404.

Revelt, D., and K. Train (1998): "Mixed logit with repeated choices: households' choices of appliance efficiency level," *Review of economics and statistics*, 80(4), 647–657.

Reynaud, A. (2003): "An econometric estimation of industrial water demand in France," *Environmental and Resource Economics*, 25(2), 213–232.

Sanghi, A., and R. Mendelsohn (2008): "The impacts of global warming on farmers in Brazil and India," *Global Environmental Change*, 18(4), 655–665.

Sanstad, A., C. Blumstein, and S. Stoft (1995): "How high are option values in energy-efficiency investments?," *Energy Policy*, 23(9), 739–743.

Schneider, M. L., and E. E. Whitlatch (1991): "User-specific water demand elasticities," *Journal of Water Resources Planning and Management*, 117(1), 52–73.

Seo, S. N., and R. Mendelsohn (2008): "An analysis of crop choice: Adapting to climate change in South American farms," *Ecological Economics*, 67(1), 109–116.

Snow, L. A. (2010): "Highlights of the BDCP," *Bay Delta Conservation Plan*.

Tonn, B., and M. Martin (2000): "Industrial energy efficiency decision making* 1," *Energy policy*, 28(12), 831–843.

Train, K. (1985): "Discount rates in consumers' energy-related decisions: a review of the literature," *Energy*, 10(12), 1243–1253.

Train, K. E. (2009): *Discrete choice methods with simulation*. Cambridge university press.

Vogel, N. (2000): "How states consumers lost with electricity deregulation," *The Los Angeles Times*, 9.

Weare, C. (2003): *The California electricity crisis: causes and policy options*. Public Policy Instit. of CA.

Wooldridge, J. M. (2002): *Introductory Econometrics: A Modern Approach*. South-Western College.

Yaron, D., and E. Bresler (1970): "A model for the economic evaluation of water quality in irrigation," *Australian Journal of Agricultural and Resource Economics*, 14(1), 53–62.

Yaron, D., E. Bresler, H. Bielorai, and B. Harpinist (1980): "A model for optimal irrigation scheduling with saline water," *Water Resources Research*, 16(2), 257–262.

Ziegler, J. A., and S. E. Bell (1984): "Estimating Demand for Intake Water by Self-Supplied Firms," *Water Resources Research*, 20(1), 4–8.

Appendices

A

Table A1: Salinity Tolerance Information of the Crops

Crop	Class	Threshold	Decline	Rating	Source
Alfalfa	F	2	7.3	MS	2
Almond	D	1.5	19	S	2
Apple	D	1.33	19.08	S	4
Apricot	D	1.6	24	S	2
Asparagus	T	4.09	2	T	2
Barley	G	8	5.5	T	2
Bean	T	1	19	S	2
Beet	T	4	9	MT	1
Berry	T	1.5	22	S	3
Bok Choy	T	1.8	9.69	MS	7
Broccoli	T	2.8	9.19	MS	1
Cabbage	T	1.8	9.69	MS	1
Carrot	T	1	14	S	1
Celery	T	1.8	6.2	MS	1
Cherry	D	1.5	19	S	5
Clover	B	1.5	12	MS	1
Collard	T	1.8	9.69	MS	7
Corn	F	1.7	12	MS	2
Cotton	T	7.7	5.2	T	1
Cucumber	T	2.5	13	MS	1
Dry Bean	F	1	19	S	2
Fig	D	4.52	7.56	MT	4
Forage	B	4.5	2.6	MT	1
Grape	V	1.5	9.6	MS	1
Grass Seed	T	4.59	7.6	MT	1
Herb, Spice	T	-	-	-	9
Kale	T	1.8	9.69	MS	7
Kiwi	D	1.5	22	S	6
Lettuce Leaf	T	1.3	13	MS	1
Melon	T	1	8.4	MS	1

Table A1: continued

Crop	Class	Threshold	Decline	Rating	Source
Mustard	F	1.8	9.69	MS	7
Oat	G	7.46	5.48	T	1
Olive	D	4	12	MT	1
Onion	T	1.2	16	S	1
Ot-Flower Seed	T	-	-	-	9
Parsley	T	1	14	S	8
Pastureland	B	7.5	6.9	T	1
Peach	D	1.7	21	S	1
Pear	D	1.33	19.08	S	4
Peas	T	3.4	10.6	MS	1
Pecan	D	1.84	10.17	MS	4
Pepper	T	1.5	14	MS	1
Pistachio	D	1.84	10.17	MS	4
Potato	T	1.7	12	MS	1
Pumpkin	T	1	8.4	MS	8
Rice	F	1.9	9.1	MS	1
Rye	G	11.4	10.8	T	1
Ryegrass	B	5.6	7.6	MT	1
Safflower	G	4.52	7.56	MT	4
Sorghum	G	6.8	16	MT	1
Spinach	T	2	7.6	MS	1
Squash	T	3.2	16	MS	1
Sudangrass	F	2.9	4.3	MT	3
Sunflower	F	4.8	5	MT	1
Sweet Basil	T	-	-	-	9
Tomato	T	2.5	9.9	MS	1
Triticale	G	6.1	2.5	T	1
Turf	T	4.52	7.56	MT	4
Turnip	T	0.9	9	MS	1
Vegetable	T	-	-	-	9
Walnut	D	1.5	20	S	1
Watermelon	T	1	8.4	MS	8
Wheat	G	3.5	4	MT	1

B

Table B2: Agencies and retailers included in the study

Member Agency	Retailer	Years Available
Anaheim	Anaheim	2001-2003, 2007-2008
Beverly Hills	Beverly Hills	2000 - 2003
Burbank	Burbank	1994, 1996-1998, 2000-2008
Calleguas MWD	Westlake	1996-1998, 2000-2010
	Camarillo	1998, 2000-2010
	Camrosa WD	1994-1995, 2000-2008
	Oxnard	2003-2008
	Simi Valley	2003-2004, 2007-2010
	Thousand Oaks	2000-2002, 2004-2007
Central Basin MWD	East Los Angeles	1996-1998, 2000-2010
	Cerritos	1996-1998, 2000-2004, 2006-2010
	Downey	2006-2010
	Lakewood	2000-2003, 2005-2010
	Orchards Dale WD	1996-1997, 2000-2010
	Paramount	2005-2010
	Pico Riviera	2001, 2008-2010
	Pico WD	1994, 1996-1998, 2000-2003, 2005-2007
	Whittier/La Mirada	1995-1998, 2000-2010
	Vernon	2002-2010
Compton	Compton	2008-2010
Eastern MWD	Eastern MWD	2004-2010
	Rancho California WD	1997-1998, 2000-2010
Foothill MWD	La Canada ID	2002-2010
Fullerton	Fullerton	1994, 1996-1997, 2000-2010

[72]Sources: 1: Hanson, Grattan, and Fulton (1999). 2: Hoffman (2010). 3: Maas and Hoffman (1977). 4: Hoffman (2010). Only tolerance group information was available. Average value for the threshold and slope of the corresponding tolerance group was used. 5: Value for almond was used. 6: Value for berry was used. 7: Value for cabbage was used. 8: Value for carrot was used. 9: No information was available.

Table B2: continued

Member Agency	**Retailer**	**Years Available**
Glendale	Glendale	1998, 2004-2010
IEUA	Ontario	1997-1998, 2000-2003, 2005-2009
	Upland	2000, 2002-2008, 2010
Las Virgenes MWD	Las Virgenes MWD	1994, 1996-1997, 2000-2010
Long Beach	Long Beach	1996-1998, 2000-2010
Los Angeles	Los Angeles	1996-1998, 2000-2010
MWDOC	Buena Park	1994, 1996-1997, 2000-2010
	Fountain Valley	1994, 1996-1997, 2000-2010
	Garden Grove	2006-2010
	Huntington Beach	2001-2010
	Mesa Consolidated WD	1996, 1998, 2000, 2006, 2008-2010
	Westminster	1996-1997, 2000-2010
	Yorba Linda WD	1994, 1996-1998, 2000-2010
Pasadena	Pasadena	2007-2010
San Diego CWA	Carlsbad MWD	2001-2003, 2005-2010
	City San Diego	2001-2004, 2006-2010
	Escondido	1998, 2000-2010
	Fallbrook PUD	2000-2001, 2003-2010
	Helix	2001-2010
	Oceanside	2000-2010
	Olivenhain MWD	2000-2010
	Otay, Padre Dam MWD Eastern	1998, 2000-2003, 2007
	Poway	2000-2001, 2005-2010
	Rainbow MWD	2003-2005, 2008, 2010
	Ramona MWD	1998, 2000, 2001-2010
	Rincon del Diablo MWD	2001-2004, 2006-2010
	San Dieguito WD	2004-2005

Table B2: continued

Member Agency	Retailer	Years Available
	Santa Fe ID	1997, 2000-2010
	Sweetwater Authority	2000-2010
	Vallecitos WD	2000-2010
	Valley Center MWD	2001-2007
	Vista ID	2000-2010
San Fernando	San Fernando	2000-2001, 2003-2005, 2007, 2009-2010
Santa Ana	Santa Ana	2001-2003, 2009-2010
Santa Monica	Santa Monica	1994, 1996-1997, 2000-2010
Three Valleys MWD	Covina	2001-2009
	Pomona	2000-2010
	Walnut Valley WD	2000-2010
Torrance	Torrance	1994-1995, 2000-2001
Upper San Gabriel Valley MWD	Alhambra	2000-2007, 2009-2010
	Arcadia	2008-2010
	Azusa	2001-2010
	Monrovia	2006-2010
West Basin MWD	Hermosa Redango	1996-1998, 2000-2003, 2005-2010
	El Segundo	2004-2005, 2007-2008
Western MWD	Corona, Elsinore Valley MWD	1996-1998, 2001-2003, 2005-2006, 2008-2010
	Jurupa CSD	2000-2003, 2007-2010
	Norco	2001-2010
	Western MWD	2000-2010