Is System Pressure Reduction a Valuable Water Conservation Tool? Preliminary Evidence From the Irvine Ranch Water District

Policy

Planning

Research

Evaluation

February 14, 2003

A study funded by the Metropolitan Water District of Southern California under their Innovative Conservation Program

By

Anil Bamezai, Ph.D.
(Western Policy Research)

Dale Lessick
(Irvine Ranch Water District)

WESTERN POLICY RESEARCH

171 Pier Avenue
Suite 256
Santa Monica
California 90405

Preface

This report presents findings from an experiment that was run in Irvine Ranch Water District's (IRWD) service area. The purpose of the experiment was to examine the impact of system pressure reduction on residential water consumption, and on the frequency of pressure-related customer service complaints.

Contents

Figures and Tables

Executive Summary

Water utility operators have long known that throttling down system pressure reduces total consumption, and this strategy is sometimes used to deal with short-term water shortages. Water pressure can vary considerably across households if each dwelling is not fitted with its own pressure-regulating valve, especially in areas of varying elevation. And this uneven pressure affects domestic irrigation the most, because the irrigation offshoot is often taken before the pressure-regulating valve if such a valve is present at all. Therefore, there is good reason to believe that reducing pressure will save water. For this reason, system pressure optimization is included as a Potential Best Management Practice (PBMP) in the *Memorandum of Understanding Regarding Urban Water Conservation in California (MOU)*. But, utility-operations staff normally harbor concerns about the impact of pressure reductions on pressure-related customer complaints, although the sensitivity of customer complaints to system pressure has never been systematically studied to our knowledge. This study examines how effective this PBMP is likely to be in practice in terms of both water savings and increased customer complaints. One potential benefit that we do not evaluate in this study, however, is the impact of reduced pressure on lowering water lost to leakage, which may be significant in older cities with a high percentage of unaccounted for water.

The impact of reduced pressure on consumption and customer service calls was examined in Irvine Ranch Water District's (IRWD) service area, mainly because IRWD already had all the key data tracking elements in place. These include: (1) household billing histories; (2) daily system pressure histories by sub-region; and (3) an electronic customer complaint logging system going back in time. In the two selected test neighborhoods of University Park and Racquet Club, pressure was reduced on average by 17.6 and 6 percent respectively during the experimental one-year period. Savings were determined by comparing weather-normalized consumption during the experimental period to two years of pre-experimental (or baseline) billing histories. A similar pre-post comparison was performed for pressure-related customer complaint histories to assess whether such complaints increased during the experimental period.

Our findings demonstrate that reducing system pressure can significantly reduce residential water consumption, especially irrigation, without entailing any significant costs in terms of increased customer complaints. In one of the treatment neighborhoods (University Park) where pressure was reduced significantly (by 17.6 percent), single-family consumption declined by 1.9 percent overall, and by 4.1 percent among those residences with greater-than-average landscapes. We could not detect any significant savings in Racquet Club, however, probably because of the low magnitude of pressure reduction.

Different agencies are likely to realize different levels of savings from pressure reductions, and the preliminary evidence suggests that simple engineering estimates may provide a fairly accurate first-cut as to what can be expected from such measures. The key factors for predicting effectiveness are three; (1) the current level of system pressure, which in turn determines the level of possible reductions; (2) the extent of the practice whereby irrigation offshoots are taken before the residential pressure regulating valve; and (3) the proportion of single-family residential consumption that is accounted for by irrigation.

Although the level of achieved savings may not appear high in terms of percentages, the quantities of water on a service-area wide basis can be quite significant. And given the near-zero cost involved in achieving these savings, cost-effectiveness is virtually guaranteed. These preliminary positive findings suggest that pressure reductions may be a valuable tool for conserving water, and that the basic idea merits further examination. We hope that this report will stimulate others to undertake similar studies elsewhere.

Acknowledgments

We are grateful to IRWD's water operations staff for their assistance in designing this experimental study and for their diligence in implementing it. Special thanks go to Steve Habiger, Steve Cormack, Wayne Wright, Johnnie Johannessen, Tom Roberts, Mike Miller, Andy Willis, Don Geiser, Dave Saindon, Ken Pfister, Mike Purington, and Binh Pham.

1. Introduction

Water utility operators have long known that throttling down system pressure reduces total consumption, and this strategy is sometimes used to deal with short-term water shortages. For this reason, system pressure optimization is included as a Potential Best Management Practice (PBMP) in the *Memorandum of Understanding Regarding Urban Water Conservation in California (MOU)*. This study examines how effective this PBMP is likely to be in practice.

For two reasons, system pressure is often maintained at a higher than optimum level; (1) to reduce customer service calls; and (2) to mask system deficiencies. The frequency of customer service complaints and its relationship with and sensitivity to system pressure, however, has never been scientifically studied. We suspect that exclusive reliance on anecdotal evidence makes system operators overly cautious in this regard, the net result being that water is wasted, especially in irrigation applications where high pressure can lead to misting. A high system pressure is also used to mask problems caused by improper pressure regulation, especially in elevated areas. To assure acceptable pressure at high elevations, main supply line pressure at the base obviously has to be set at a high level. But in such cases it is necessary that households drawing water from the same main line at lower elevations be fitted with pressure regulating valves. Often, this sort of system optimization is not undertaken to an adequate degree leading to water wastage at lower elevations. In the past when water was plentiful and cheap, perhaps not investing in more pressure regulators might have seemed cost-effective, but this may no longer be true.

How much water is system pressure optimization likely to save? Basic laws of fluid mechanics tell us that flow rates through a circular pipe are exponentially related to water pressure. In other words, to double flow rates, pressure must be quadrupled all else being equal. Or, to put it conversely, decreasing pressure, say, by 10 percent will likely reduce flow rates by approximately 5 percent ($(1-\sqrt{(1-0.1)}) = 0.051$). Does this mean consumption will also decrease by 5 percent? Well, not necessarily. For indoor fixtures, such as toilets, dishwashers, and laundry machines, lower flow rates only increase fill times without affecting consumption. Other end uses, such as showers and faucets may respond to a greater degree. But, for two reasons, we expect (automatic) irrigation to respond the most to system pressure reduction; (1) the irrigation offshoot is generally taken before the residence's pressure regulating valve if such a valve is present at all; and (2) residential customers are generally thought to over-irrigate and therefore unlikely to notice any significant change in their landscapes as a result of pressure reduction. Assuming that irrigation represents about a third of total residential consumption, we expect pressure reductions in the neighborhood of 10 percent to reduce total consumption by 1 to 2 percent. This hypothetical estimate represents a large amount of water when computed on a service area wide basis, especially considering the near-zero cost of obtaining it.

One aspect that we do not evaluate in this study is the impact of reduced pressure on lowering water lost to leakage, which may be significant in older cities with a high percentage of unaccounted for water.

The balance of this report is organized as follows. Section 2 describes the study design. Section 3 presents the key findings. Section 4 concludes. Technical issues about the underlying statistical models are described in Appendix A.

2. Study Setting and Methods

This study was implemented in Irvine Ranch Water District's (IRWD) service area mainly because IRWD already had all the key data tracking elements in place. These include: (1) household billing histories; (2) daily system pressure histories by sub-region; and (3) an electronic customer complaint logging system going back in time. There was no way to track reduction in system leakage by sub-region, however. Therefore, this potential positive benefit of reducing system pressure remains unmeasured.

Two neighborhoods of IRWD (University Park and Racquet Club) were selected for pressure reduction. We call them the "treatment" group. These two treatment neighborhoods include many residential customers, and historically have been maintained at a pressure ranging between 70 and 85 pounds per square inch (Figure 1). During the experimental period (August, 2001—August, 2002), pressure was reduced in two steps staggered roughly six months apart. During the second half of the experimental period, pressure was kept as close as possible between 60 and 65 pounds per square inch. IRWD considers a pressure of 60 pounds per square inch to be the minimum acceptable level. The main supply line to each of these treatment neighborhoods is fitted with pressure-regulating valves that can be remotely tracked and controlled through IRWD's SCADA system. Maintaining even pressure, even with a computer-controlled system, however, was not easy because of large daily fluctuations in system pressure.

An additional three comparison neighborhoods (Culverdale, Laurelwood, and Peppertree) were selected to serve as the "comparison" group, that is, the group where pressure was not altered in any deliberate way. Including a comparison group improves the credibility of any evaluation. These three comparison neighborhoods were also suitably isolated from the larger system through their own pressure-regulating valves, but these valves were not hooked to the SCADA system. So detailed pressure histories going back in time were not available for the comparison neighborhoods.

After the experiment was run for a full year, we obtained customer billing histories and logs of pressure-related customer service complaints for a full three years (two pre-experimental or baseline years, one experimental year). The billing histories were matched with daily weather data drawn from CIMIS's database, and then statistically weather normalized to estimate water savings. A before-and-after comparison of pressure-related complaints was also conducted to evaluate whether such complaints increased in the treatment neighborhoods.

The next section summarizes the study's findings.

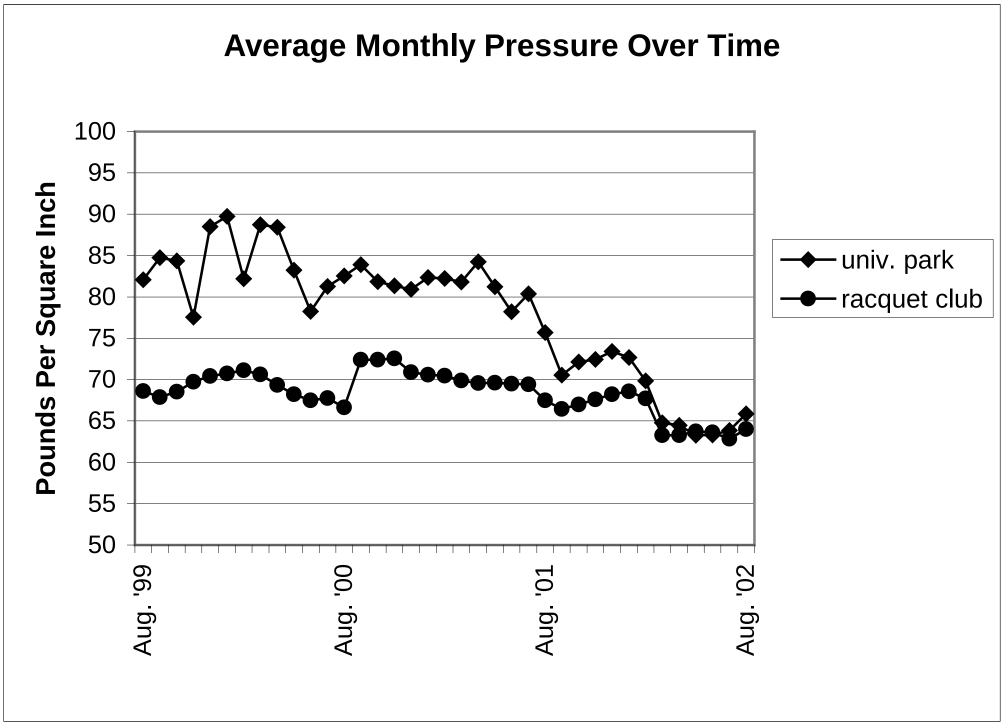

Figure 1 Average monthly pressure over time in the treatment neighborhoods

3. Study Findings

Customer Service Requests

A key component of our study was to track the impact of pressure reduction on customer service requests. These data were collated for a three-year period corresponding to the water savings analyses (two pre-experimental or baseline years, one experimental year).

IRWD's computerized system catalogs pressure-related customer service requests in three ways; (1) those pertaining to low pressure; (2) those pertaining to high pressure; and (3) finally those that pertain to fluctuating pressure. Tables 1 and 2 provide a frequency count of these customer service requests by type of residence (single-family vs. condominium), treatment status (control vs. treatment), and time period (baseline vs. intervention).

While comparing differences between the baseline and intervention periods, keep in mind that the former period includes two years, while the latter includes only one year. So, for example, there were 5 low-pressure complaints from single-family residences located in the control villages during the pre-intervention period, and 3 during the post-intervention period, respectively (Table 1). On a per year basis, this amounts to a very insignificant change. In both the treatment villages, however, a marginal up-tick in low-pressure complaints can be observed among single-family residences.

In the case of condominiums, lowering pressure appears to have produced only benefits because the frequency of high-pressure complaints in University Park decreased without any appreciable increase in the frequency of low-pressure complaints.

The causes underlying low-pressure complaints generally included malfunctioning pressure regulating valves or valves with angle stops that were not fully open. In other words, complaints that required a one time fix, not constant management.

Overall, the data presented in Tables 1 and 2 do not indicate that pressure reduction is likely to lead to a flurry of complaints, in turn fueling perceptions of diminished customer service. Hence, there seems to be no reason why water agencies should shy away from optimizing pressure in their supply systems to promote the efficient utilization of scarce water resources.

Table 1 Number of customer service requests from single-family residences

Sub-region	Time Period	Type of Complaint: Low Pressure	High Pressure	Fluctuating Pressure
Control villages	Baseline	5	1	2
	Intervention	3	0	1
Treatment village (Racquet Club)	Baseline	3	0	0
	Intervention	3	0	0
Treatment village (University Park)	Baseline	1	1	0
	Intervention	2	0	0

Table 2 Number of customer service requests from condominiums

Sub-region	Time Period	Type of Complaint: Low Pressure	High Pressure	Fluctuating Pressure
Control villages	Baseline	3	0	0
	Intervention	0	0	0
Treatment village (Racquet Club)	Baseline	0	0	0
	Intervention	0	0	0
Treatment village (University Park)	Baseline	8	6	1
	Intervention	5	0	0

Water Savings

Table 3 shows the average daily pressure in the two treatment neighborhoods during the baseline and intervention periods. These averages are derived from the same data that underlie Figure 1. Between the baseline and intervention periods, University Park and Racquet Club experienced pressure reductions of roughly 17.6 and 6 percent, respectively. Because historically Racquet Club had lower pressure than University Park, the scope for further reduction was obviously lower (water operations staff at IRWD did not think it wise to lower pressure below 60 pounds per square inch.)

Table 3 Average daily pressure over time (psi)

Treatment Neighborhood	Average pressure: Baseline Period	Intervention Period	Pressure change
University Park	82.9	68.3	-17.6%
Racquet Club	69.8	65.6	-6.0%

Using information contained in IRWD's billing system, we then identified single-family dwelling units that had been occupied by the same household during the three-year analysis period (August, 1999—August, 2002). Water savings were derived only from accounts meeting this criterion, so as to maintain as best as possible an apples-to-apples comparison over time. For these select accounts, the billing histories were next matched with daily weather data (obtained from CIMIS). Finally, statistical models were used to weather-normalize these billing histories, and to estimate water savings. Appendix A describes the statistical methodology in greater detail.

Two water savings estimates were derived. The first was based upon data from all single-family accounts where occupants met the three-year residency criterion. The second was based on a subset of the above group that in addition also had large landscapes. Large landscape accounts were identified as those with outdoor variances. IRWD has a water-budget based pricing system, where deviations from the allocated budget carries increasingly stiff financial penalties. Every household receives a standard indoor and outdoor water-use allocation—of which the latter is pegged to the daily evapotranspiration rate. But if households can demonstrate either a larger-than-average family size, or larger-than-average landscape, they can apply for a variance to have these allocations adjusted upward.

Table 4 shows our estimates of water saved by single-family residences in the treatment neighborhoods of University Park and Racquet Club. We were unable to detect significant water savings in Racquet Club, both, overall, or among the subset that had requested outdoor variances. Given that pressure reductions in Racquet Club were of a relatively small magnitude, perhaps these findings are not all that surprising. On the other hand, water savings in University Park were both statistically significant and generally in line with our *a priori* expectations. Here, annual consumption declined 1.9 percent overall, and 4.1 percent among the subset with outdoor variances.

Table 4 Annual reduction in consumption (single family residences)

	All single-family accounts		Outdoor variance accounts	
Neighborhood	**Annual reduction**	**Number of accounts**	**Annual reduction**	**Number of accounts**
University Park	1.9%	304	4.1%	68
Racquet Club	≈0.0%	167	≈0.0%	33

NOTE: University Park savings are statistically significant. Racquet Club savings are not statistically distinguishable from zero.

Examining savings for the larger-than-average landscape accounts was useful for two reasons. First, it helped to test our hypothesis that pressure reduction largely works because of its ability to reduce irrigation demand. Second, the savings estimates are more reliable in a statistical sense because they are based upon accounts that are inherently more comparable to one another since all have an outdoor variance.

How credible is the savings estimate of 4.1 percent for outdoor-variance accounts of University Park? Table 3 shows that, on average, pressure was reduced by 17.6 percent in University Park. Assuming outdoor consumption is approximately 45 percent of total consumption among the outdoor-variance accounts (an assumption for which there is both anecdotal and empirical support [Hunt et al., 2001]) our engineering savings estimate works out to almost exactly 4.1 percent ($(1-\sqrt{(1-0.176)})\times 0.45 = 0.041$). This closeness between the engineering and statistical estimates suggests that water agencies may be in a good position to estimate likely savings from this potential best management practice.

We also attempted to estimate water savings achieved in Condominiums as a result of this experiment, but were unable to detect any. This was in line with our expectation, but we also believe that our statistical analysis for Condominiums is less reliable. We had no information about the occupancy rate of Condominiums over time, nor any information about occupant characteristics. Thus it is less clear that one is comparing apples and apples across the treatment and comparison groups, or over time. Although we did not have detailed occupant characteristics for single-family dwellings either, at least we were able to identify those units where the identity of the residents had remained the same over the three-year analysis period. And then by focusing only on outdoor-variance accounts, we were able to tighten the statistical comparability even more.

4. Conclusions

This study demonstrates that reducing system pressure can significantly reduce residential water consumption, especially irrigation, without entailing any significant costs in terms of increased customer complaints. In one of the treatment neighborhoods (University Park) where pressure was reduced significantly (by 17.6 percent), single-family consumption declined by 1.9 percent overall, and by 4.1 percent among those residences with greater-than-average landscapes. Different agencies are of course likely to realize different levels of savings from pressure reductions, and the preliminary evidence suggests that simple engineering estimates may provide a fairly accurate first-cut as to what can be expected from such measures. The key factors for predicting effectiveness are three; (1) the current level of system pressure, which in turn determines the level of possible reductions; (2) the extent of the practice whereby irrigation offshoots are taken before the residential pressure regulating valve; and (3) the proportion of single-family residential consumption that is accounted for by irrigation.

Although the level of achieved savings may not appear high in terms of percentages, the quantities of water on a service-area wide basis can be quite significant. And given the near-zero cost involved in achieving these savings, cost-effectiveness is virtually guaranteed. These preliminary positive findings suggest that pressure reductions may be a valuable tool for conserving water, and that the basic idea merits further examination. We hope that this report will stimulate others to undertake similar studies elsewhere.

Appendix A—Model Specification and Estimation

Conceptual Model

A logical way of modeling staggered billing data is to conceive the model at a daily level and then scale it up to the meter-read level. Equation (1) expresses logarithmically transformed daily consumption (U_{it}) for customer (i) at time (t) as a function of the daily weather index (W_t), say, the evapotranspiration rate, customer characteristics (X_i), daily intercept terms (α_t) and random error (ε_{it}). This model is very flexible insofar the intercept terms and weather coefficients are conceptually allowed to vary on a daily basis. Intercept terms are necessary to capture indoor use. Also, intervening human factors make irrigation's relationship with weather somewhat sticky. To some extent irrigation decisions are based upon experience and "gut feel." A weather index alone is therefore unlikely to capture the full variation in consumption by time of year.

$$Ln(U_{it}) = \alpha_t + \beta_t W_t + \eta X_i + \varepsilon_{it} \qquad (1)$$

$$\text{where} \quad \varepsilon_{it} \sim N(0,\sigma^2)$$

Daily consumption is logarithmically transformed because water consumption is generally distributed with a long right-hand tail. And usually, even accounting for customer heterogeneity and seasonality is not sufficient for normalizing model error. A couple of explanations can be offered for skewed model error. First, the most seasonal component of consumption—irrigation—is a discrete event, even when scheduled according to scientific principles. A landscape is supposed to be irrigated when daily evapotranspiration has depleted the soil water content below a certain threshold (Snyder and Sheradin, 1992). When daily evapotranspiration is low and uncertain, or rainfall is received periodically, average daily consumption may exhibit a rightward skew. Second, landscape professionals often set irrigation schedules by varying a preset baseline schedule in proportion to changes in the evapotranspiration rate. Errors are therefore proportionally magnified or diminished.

Averaging consumption across the (N) days included in a read taken at time (T) yields the meter read-level model (Equation 2). Throughout, summation operators are subscripted backward in time because meter read-dates signal the end of a consumption period. If consumption days (N) vary markedly across reads, averaging insures error homoscedasticity at the meter-read level when daily error is homoscedastic. Of course, in spite of averaging, meter read-level error will be heteroscedastic if daily error itself is heteroscedastic, in which case (2) should be estimated using generalized least squares. Autocorrelation is a different matter, however. Because of error averaging, autocorrelation at the meter-read level should be low to nonexistent even if daily error is highly autocorrelated. It can be mathematically shown that if daily autocorrelation is as high as 0.9, even then observed autocorrelation will only be 0.092 for 30-day cycle reads, and 0.025 for 61-day cycle reads (Bamezai, 1997).

$$\frac{1}{N}\sum_{t=T}^{T-N} Ln(U_{it}) = \sum_{t=T}^{T-N}\alpha_t \frac{1}{N} + \sum_{t=T}^{T-N}\beta_t \frac{W_t}{N} + \eta X_i + \frac{1}{N}\sum_{t=T}^{T-N}\varepsilon_{it} \qquad (2)$$

where $\frac{1}{N}\sum_{t=T}^{T-N}\varepsilon_{it} \sim N(0,\sigma^2)$

Estimation of (2) as it stands requires the creation of at least 365 daily indicator variables (equal to *1/N* for days included in the read) for capturing the daily intercepts and another 365 interactions of these indicators with the daily weather index to capture the daily weather response. For days not included in a specific meter read the corresponding daily indicators and their interactions take on the value of zero. Such an enormous estimation exercise is unlikely to succeed not only because of the immense computing resources required but also because of multicollinearity among many of the daily indicator variables. Meter reads must be available for every day in the year to provide the variation necessary for estimating these daily parameters, but read-dates are often clustered by design. Thus, for estimation purposes, it is necessary to impose some simplifying restrictions on these daily parameters.

An option is to assume that the daily intercepts (α_t) and the weather response coefficients (β_t) are equal for all days in a given month. Doing so reduces the estimation problem down to 12 monthly intercepts, 12 weather coefficients, and other customer characteristics included in the model. It is not necessary to place the same restrictions on (α_t) and (β_t). For example, the daily intercept terms (α_t) may be fit with piece-wise linear or cubic splines (Suits et al., 1978; Robb, 1980), while the weather coefficients (β_t) may be assumed constant for either all days in a month or all days in a season. The daily intercepts can also be captured using Fourier harmonics (Bamezai, 1996).

Because monthly restrictions are perhaps the most obvious choice with billing data that follow a 30-day cycle, the implication of these restrictions is developed in greater detail. Equation (3) shows what these restrictions imply for meter reads that span a total of (*N*) days, with (*m*) days falling in one month and (*n*) days in the next.

$$\frac{1}{N}\sum_{t=T}^{T-N} Ln(U_{it}) = \alpha_m \frac{m}{N} + \alpha_n \frac{n}{N} + \beta_m \sum_{t=T}^{T-m}\frac{W_t}{N} + \beta_n \sum_{T-m-1}^{T-m-n}\frac{W_t}{N} + \eta X_i + \frac{1}{N}\sum_{t=T}^{T-N}\varepsilon_{it} \qquad (3)$$

To estimate (3) it is necessary to allocate the total number of days covered by a meter read to each month. In other words, 12 monthly variables must be created of which 2 take on the values (*m/N*) and (*n/N*) for any given read, the rest being zero. Similarly, the daily weather index during a read interval must also be split into month-specific aggregates. Once again 12 weather variables are required of which only at most 2 take on a nonzero value for any given read. Meter reads taken bimonthly can be handled just as easily in the above framework, the only difference being that such reads are likely to span 3 instead of 2 months.

Construction of the dependent variable in (3), however, still poses a minor problem. The dependent variable is equal to the sum of logarithmically transformed daily consumption. But billing histories yield only the sum of daily untransformed consumption which after a logarithmic transformation does not equal the desired dependent variable (Equation 4).

$$\frac{1}{N}\sum_{t=T}^{T-N} Ln(U_{it}) \neq Ln(\frac{1}{N}\sum_{t=T}^{T-N} U_{it}) \tag{4}$$

The above inequality, however, can easily be resolved by leaning on well-known properties of a lognormal distribution.

If $$Ln(U_{it}) \sim N(\mu_t, \sigma^2)$$

then $$E(\frac{1}{N}\sum_{t=T}^{T-N} Ln(U_{it})) = \frac{1}{N}\sum_{t=T}^{T-N} \mu_t \tag{5}$$

Similarly

$$Ln(E(\frac{1}{N}\sum_{t=T}^{T-N} U_{it})) = Ln(\frac{1}{N}\sum_{t=T}^{T-N} e^{\mu_t + \frac{\sigma^2}{2}}) = \frac{1}{N}\sum_{t=T}^{T-N} \mu_t + \frac{\sigma^2}{2} + Ln(\frac{1}{N}\sum_{t=T}^{T-N} e^{\varepsilon_t}) \tag{6}$$

where $$\varepsilon_t = \mu_t - \frac{1}{N}\sum_{t=T}^{T-N} \mu_t$$

Under most plausible scenarios of the rate of change in average daily consumption (μ_t) over the course of 30 or 61 days, the last term in (6) converges to zero. In other words, the two quantities cited in (4) differ approximately by a constant (that is, half of the daily variance), hence are readily substitutable.

Approximating Nonlinearity and Reducing Measurement Error

If data and model diagnostics indicate that the weather index (say, the evapotranspiration rate) should either be logarithmically transformed, or that higher powers should be included as well, the framework developed in (1) through (6) can easily include such possibilities. One such case is discussed below for illustration.

Assume daily consumption is a quadratic function of weather instead of a linear function (Equation 7).

$$Ln(U_{it}) = \theta_t + \omega_t W_t + \psi_t W_t^2 + \eta X_i + \varepsilon_{it} \qquad (7)$$
$$\text{where} \quad \varepsilon_{it} \sim N(0, \sigma^2)$$

Under the assumption of monthly restrictions, estimation of (7) now requires 12 additional variables to capture the weather index's second power. But by applying a linear approximation to (7) both the computational burden and the impact of measurement error can be minimized. The daily weather index is first reexpressed in terms of deviations from the daily mean, but then higher powers of the deviations are dropped (Equation 8).

$$Ln(U_{it}) = \theta_t + \omega_t(\overline{W_t} + \Delta W_t) + \psi_t(\overline{W_t} + \Delta W_t)^2 + \eta X_i + \varepsilon_{it} \qquad (8)$$
$$\Rightarrow Ln(U_{it}) \approx \alpha_t + \beta_t \Delta W_t + \eta X_i + \varepsilon_{it}$$
$$\text{where} \quad \alpha_t = \theta_t + \omega_t \overline{W_t} + \psi_t \overline{W_t}^2$$
$$\beta_t = \omega_t + 2\psi_t \overline{W_t}$$

After the linear approximation the essential structure of (8) is identical to (1), except that by working with daily deviations in the weather index, an approximate nonlinear weather specification is implicitly assumed without any increase in the computational burden. Bamezai (1997) demonstrates the validity of this approximation. Two additional benefits also accrue from the above approximation. First, the daily intercepts (or monthly if so constrained) provide a direct measure of average consumption on a particular day (or month) in a "normal" weather year because the differenced weather index is centered at the mean by construction. Second, a systematic time bias in the weather index's mean caused by lack of information about plant material by customer is likely to influence the deviations significantly less. A differenced weather specification (8) therefore simultaneously minimizes the impact of systematic measurement error while capturing an approximate nonlinear weather response. Even if weather response is linear, a differenced weather index is preferable to an undifferenced index: either index will yield identical results in the absence of measurement error, but the former is likely to be more accurate in the presence of measurement error.

Weather Index Construction

For the analyses that follow, weather variation is captured through a rainfall adjusted evapotranspiration-rate index (Equation 9). The evapotranspiration rate measures a plant's total water demand. It is necessary to subtract effective rainfall from the evapotranspiration rate to accurately predict net irrigation demand. The daily evapotranspiration and rainfall data are obtained from CIMIS's Irvine station.

$$W_t = \max[0, (ET_t^R K_t^C - P_t u)] \qquad (9)$$

where

W	daily weather index (inches)
ET^R	daily reference evapotranspiration rate (inches)
K^C	monthly crop coefficient
P	daily precipitation (inches)
u	effective proportion of precipitation
SR	surplus effective rainfall carryover

CIMIS's ET^R represents the water demand of 4- to 6-inch-tall, cool-season grass transpiring at its maximum rate. In reality, plant height, plant roughness, plant age, ground shading, and other factors, all influence actual evapotranspiration needs of a plant (Snyder, 1993). To stay consistent with how IRWD calculates outdoor allocations, monthly crop coefficients for cool season turf are used to correct the reference evapotranspiration rate. Meyer and Gibeault (1987) originally estimated these crop coefficients. Half of daily rainfall is assumed to be effective as per CIMIS's recommendation, but when effective rainfall exceeds total evapotranspiration demand, net evapotranspiration demand is floored at zero. As mentioned earlier, the science underlying irrigation is essentially a stock and flow problem (Snyder and Sheradin, 1992). Soil moisture content (stock) must be maintained within a certain threshold. Evapotranspiration (flow) reduces the stock on a daily basis, effective rainfall adds to it intermittently, with irrigation acting as the balancing lever. A weather index constructed using a stock and flow framework is likely to be a better predictor of irrigation demand—the most weather-sensitive portion of total demand.

Model Results

Water reductions were estimated by analyzing billing histories controlling for weather and other unobserved time-invariant differences (fixed effects) across households. All households analyzed here follow a 30-day billing cycle. The model relates the logarithm of average daily consumption to a vector of covariates including monthly time indicators and weather deviations constructed as per the conceptual framework described earlier. Weather variables were constructed from daily deviations in the weather index to capture a nonlinear weather response. Weather effects for the months of July through October were pooled since the month-specific weather effects were not significantly different from one another. The weather index is not logarithmically transformed since it does not exhibit a rightward skew. Upon examination, model errors revealed significant heteroscedasticity. The final models therefore incorporate a heteroscedasticity correction based upon methods discussed by Carroll and Ruppert (1988).

Table 5 displays two models; the first is based upon all households that met the three-year residency criterion; the second only upon the subset that in addition also had outdoor variances. These models exhibit strong predictive power as indicated by the highly significant time dummies and weather coefficients, as well as a relatively high adjusted-R Square. The time dummies behave as expected, indicating minimum usage in the month of February and maximum in the months of July and August.

In a non-randomized evaluation design of the sort used here, the behavior of the comparison group over time can crucially affect estimated savings. In the full-sample model, the results show that consumption in the comparison group increased approximately 1 percent during the intervention period relative to the baseline period. The way our models are set up, this rise is added to the change in consumption of the treatment group between the intervention and baseline periods under the conjecture that treatment-group consumption would also have risen by 1 percent in the absence of pressure reductions. Some may quibble with this line of reasoning, however, and instead argue that the comparison group's 1 percent increase should be ignored, in which case the savings estimate for University Park would drop to 0.9 percent and Racquet Club would be found to have dissaved after the pressure reduction. Since the latter conclusion is highly counterintuitive, we feel that information about the rise in consumption in the comparison neighborhoods should not be ignored.

These issues do not plague the second model, however, which is based only upon the outdoor-variance accounts, because here the comparison group shows no significant change in consumption between the baseline and intervention periods. The second model is thus statistically more solid and offers stronger evidence supporting our hypothesis that pressure reductions reduce water consumption.

We also attempted to estimate a pressure elasticity of demand by omitting the indicator variables and instead including in the model average pressure corresponding to each meter read. For this approach, only data from the treatment neighborhoods could be used because daily pressure histories were not available from the comparison neighborhoods. Estimated elasticities ranged between 0.1 and 0.2, but were significant only at the 10 percent level. Why we could not estimate a pressure-elasticity with greater precision when we could successfully estimate savings using an indicator variables approach remains unclear. We suspect that pressure data may be contaminated by some amount of measurement error, and that in any event they exhibit a high level of daily variability, which degrades the precision of the pressure-elasticity estimate. We also could not statistically detect differences in achieved savings between the first and second halves of the intervention period. Hopefully, future similar studies will overcome these precision problems by working with larger sample sizes.

Table 5 Estimated Fixed Effects Water Savings Model

Covariate	Full sample Coefficient (Std. Err.)	Outdoor variance only Coefficient (Std. Err.)
February indicator	-0.164**	-0.128**
	(0.016)	(0.035)
March indicator	-0.020	0.006
	(0.014)	(0.031)
April indicator	0.178**	0.253**
	(0.015)	(0.031)
May indicator	0.249**	0.323**
	(0.014)	(0.031)
June indicator	0.355**	0.466**
	(0.015)	(0.032)
July indicator	0.440**	0.548**
	(0.015)	(0.031)
August indicator	0.442**	0.549**
	(0.015)	(0.031)
September indicator	0.365**	0.450**
	(0.015)	(0.030)
October indicator	0.243**	0.343**
	(0.015)	(0.031)
November indicator	0.061**	0.128**
	(0.015)	(0.030)
December indicator	-0.025	0.040
	(0.017)	(0.035)
January weather deviation	-0.471	3.426
	(1.285)	(2.792)
February weather deviation	4.579**	5.237**
	(0.273)	(0.596)
March weather deviation	3.921**	6.632**
	(0.414)	(0.965)
April weather deviation	1.061**	0.253
	(0.308)	(0.666)
May weather deviation	1.552**	3.192**
	(0.411)	(0.858)
June weather deviation	2.450**	1.680**
	(0.400)	(0.820)
July through October weather deviation	2.364**	2.552**
	(0.124)	(0.245)
November weather deviation	3.565**	4.850**
	(0.421)	(0.866)
December weather deviation	6.125**	6.245**
	(0.286)	(0.575)
Post-intervention period indicator	0.012**	0.013
	(0.004)	(0.008)
Racquet Club indicator ×	0.011	0.001
Post-intervention period indicator	(0.007)	(0.005)
University Park indicator ×	-0.020**	-0.042**

Covariate	Full sample Coefficient (Std. Err.)	Outdoor variance only Coefficient (Std. Err.)
Post-intervention indicator	(0.006)	(0.012)
Constant	0.000	0.000
	(0.001)	(0.002)
Adjusted R-square	0.368	0.444
Number of single-family accounts used in the model	Comparison-557 University Park-304 Racquet Club-167	Comparison-120 University Park-68 Racquet Club-33

**Significant at 5 percent level or less.

References

Bamezai, A., On weather normalizing customer-level billing data, *Water Resources Research,* 33(5), 1171-1178, 1997.

Bamezai, A., Evaluation of water conservation programs in thorny nonexperimental settings, *Water Resources Research*, 32(4), 1083-1090, 1996.

Hunt, T., Lessick, D., Berg. J., Wiedmann, J., Ash, T., Pagano, D., Marian, M., and A. Bamezai, *Residential weather based irrigation scheduling: Evidence from the Irvine "ET Controller" study*, Irvine Ranch Water District, Irvine, 2001.

Carroll, R. J. and D. Ruppert, *Transformation and Weighting in Regression*. New York: Chapman and Hall, 1988.

Meyer, J. L., and V. A. Gibeault, Turfgrass performance when under-irrigated*, Applied Agricultural Research*, 2(2), 117-119, 1987.

Pagano, D., Barry, J., and Western Policy Research, *Efficient Turfgrass Management: Findings from the Irvine Spectrum water Conservation Study*, Metropolitan Water District of Southern California, Los Angeles, 1997.

Ramsey, J. B., Tests for specification errors in classical linear least squares regression analysis, *Journal Royal Statistical Society, Series B*, 31(2), 350-371, 1969.

Robb, A. L., Accounting for seasonality with spline functions, *Review of Economics and Statistics*, 62(2), 321-323, 1980.

Snyder, R. L., Irrigating up crops efficiently with sprinklers, *drought tips number 92-52*, Department of Water Resources, Sacramento, California, 1993.

Snyder, R. L., and H. L. Sheradin, Water balance irrigation scheduling using CIMIS ETo, d*rought tips number 92-20*, Department of Water Resources, Sacramento, California, 1992.

Suits, D. B., A. Mason and L. Chan, Spline functions fitted by standard regression methods, *Review of Economics and Statistics*, 60(1), 132-139, 1978.