

The Effect of Social and Consumption Analytics on Residential Water Demand

Mehdi Nemati* Steven Buck† Hilary Soldati‡

January 6, 2017

2017 ASSA Annual Meeting–Chicago

AERE section on Behavioral Interventions and Water Conservation

Working Paper¹

Abstract

In this paper, the effects of Dropcountr on water usage were examined using household-level panel data for the City of Folsom, California, from January-2013 to September-2016. Results suggest that the introduction of the Dropcountr services for the population of households participating in Dropcountr causes an aggregate treatment effect of 7% reduction in water usage with a significant variation in the effect across households dependent on baseline consumption quintile. In response to the Dropcountr services, households in the highest quintile of baseline consumption reduce water usage by an estimated 13%.

JEL Codes: Q250, D12

Keywords: Automated meters, Non-price conservation, California water, Urban water demand

*University of Kentucky, email: mehdi.nemati@uky.edu

†University of Kentucky, email: steven.buck@uky.edu

‡University of California, Berkeley, email: soldati@berkeley.edu

¹Copyright 2016 by Mehdi Nemati, Steven Buck, and Hilary Soldati. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

1 Introduction

While many arid regions already struggle to balance supply and demand of water resources, climate change will not only exacerbate many of these existing tensions, but will also introduce new conflicts. Mechanisms which not only reduce water consumption, but do so in a cost effective manner, are invaluable to meet these current and future water resource challenges. Although limited academic evidence is available as to whether social comparison programs may offer such benefits in the residential water sector, initial research suggests meaningful potential. Relying on price adjustments to reduce household water demand results in uncertainty in revenue forecasting for utilities and stirs political rancor due to equity concerns for this basic good.

Moreover, debate persists in the academic literature as to the significance of and the type of (average versus marginal) prices effects on consumption decisions. Generating frequent and highly granular micro-level household data through partnerships between a digital social comparison product and water service providers will improve academic and policy-maker information around decision-making over residential water demand management programming. Well-designed experiments and partnerships have the potential to reduce consumption, while also providing more precise estimates about how various price and non-price management tools, as well as household characteristics, determine water consumption. Additionally, such information could be leveraged not only to direct more effective and efficient water management strategies, but also to enable improved forecasting of future water demand, which is necessary in determining optimal state and regional regulatory and infrastructure choices. Hence, this type of research is important in developing solutions to water resource challenges that are impactful, cost-effective, and efficient.

This paper will contribute to a substantial body of similar research in the energy sector and growing, but less developed work, in social comparison programs for the water sector. Experimental designs in numerous markets with *Opower*, an information sharing and social comparison tool used in residential electricity management programming, have allowed for a multitude of research questions to be explored with respect to residential energy

consumption. In general, these findings show an economically and statistically significant average treatment effect, with evidence of heterogeneous impacts and advantages over other programs in reducing energy consumption in a cost-effective manner (Allcott, 2011, 2012; Ferraro et al., 2011). Limited academic analysis has been generated in the water sector, however; the authors are aware of only two analyses published in peer-reviewed academic journals, which examined the effect of *WaterSmart* services in three California utilities and isolated program in Cobb County Georgia (Brent et al., 2015; Ferraro and Price, 2013).

This paper examines the effect of a social-norm-based conservation program on households' water usage. The program under study is administered by Dropcountr (DC), which is a mobile and web application that provides information to water utilities and their customers. Their program provides information on (i) current water usage, (ii) a comparison to the previous usage, (iii) comparison to similar nearby households, and (iv) the efficient budget for households. In addition, the web interface that also provides tips about where households can save water and connects them to the existing water utility rebate programs on water saving appliances. DC also monitors households' hourly water usage data to identify possible leaks in their water system. They use unexpected boosts in water consumption as a signal of a leak in the households' water system and sends an email message or phone alert to the customer. Hence, DC is designed to motivate households to reduce their water use by changing their behavior, adopting water efficient technologies or finding leaks.

For the City of Folsom, the data used for this analysis includes two years of historical consumption, along with twenty months of data under the DC pilot program, spanning January 2013 through September 2016. This program was designed by DC as an opt-in program, therefore analysis of a treatment effect is challenged by this non-experimental design. However, various statistical tools will be explored to minimize the challenges of interpreting results.

To preview results, this initial research suggests that DC has a statistically and economically significant conserving effect on water consumption at the household level for customers who enrolled in the service. The estimated aggregate treatment effect in the City of Folsom

is a 7% reduction in average monthly consumption for the enrolled households. There appears to be a stronger effect for those households identified as high water consumers in the baseline period during the summer months. This paper also finds evidence of a “*boomerang*” effect for those households in the lower portion of baseline distribution, explained in the Analysis section below. These results are particular to the City of Folsom with an opt-in program design. The precise magnitude of a DC effect on household water consumption will vary both by location, experimental design, and by time-specific conditions, such as weather conditions and variations in other determinants of water consumption that correlate with time and location.

This paper proceeds as follows: Section 2 discusses relevant academic literature; Section 3 offers an overview of the DC business model and description of services; Section 4 describes data and method; analysis of the program is presented in Section 5; the paper concludes with discussion, summary, and policy implications in Section 6.

2 Relevant Literature

This paper has relevance to existing literature in two particular areas: estimating the effect of social comparison on consumption decisions, in general, and understanding the determinants of residential water demand, in specific. Price response in household water consumption has been studied extensively in the academic literature. Debate persists in how decision-makers are affected by both the qualitative aspects of price (block rates versus uniform pricing and average versus marginal) and the quantitative changes (estimating elasticities) ([Dalhuisen et al., 2003](#); [Ito, 2014](#); [Olmstead et al., 2003, 2007](#); [Olmstead and Stavins, 2007](#)). However, price instruments to reduce residential demand are considered a political liability, complicate revenue estimation for utilities, and inspire concerns over the impacts to lower income households ([Agthe and Billings, 1987](#)). Additionally, it is widely understood that other factors determine residential water demand, such as: income, household size, lot size, landscaping, and weather. [Buck et al. \(2015\)](#) uses a data-driven process to identify model performance

in predicting residential water demand, which reveals that price is not necessarily the most important determinant.

Consistent with this, utilities often employ non-price demand side management (DSM) strategies to influence household water consumption. [Renwick and Green \(2000\)](#), estimate the effects of six different categories of non-price DSM policies, which include information and rebate opportunities. Not surprisingly, they find that mandatory policies result in larger demand reductions, relative to voluntary programs. They also identify areas where more research is needed, including the effect of household characteristics and of multiple, simultaneous policy tools on aggregate demand. Services such as DC, which not only have the technological flexibility to vary signals, are able to amass frequent, granular data that can be used to fill knowledge gaps. Additionally, recent research has estimated household willingness-to-pay to avoid water service disruptions for some California utilities ([Buck et al., 2015, 2016](#)). These estimates may help utilities evaluate the conservation benefits that are possible through various categories of messaging, including social norms, information, and prosocial language.

Social comparison of household consumption first began in the residential electricity sector. The leading figure in this movement has been *Opower*, which partners with utilities to create content with the objective of reducing household electricity demand and improving efficiency and conservation. A growing collection of research in this field has provided estimates on program effectiveness, as well as evaluating persistence of treatment and examining site selection bias ([Allcott, 2011](#); [Allcott and Rogers, 2012](#); [Ayres et al., 2012](#)). These analyses estimate treatment effects in the range of 1.2 -3.3%, which varies according to location and program implementation, but appears to persist over time. Research on heterogeneous effects suggests that targeted content, that considers subpopulation attributes, improves messaging response ([Costa and Kahn, 2013](#)). [Allcott \(2012\)](#) identifies a problem in site and population selection bias, where program evaluation of early-adopting utilities overstate the treatment effect relative to implementation across less environmentally progressive regions and populations.

This business model of combining social, behavior, and data science to impact household decision-making is being replicated in the water sector. *WaterSmart* Software has been building partnerships with water utilities in California, as well as other states, for the past several years. In one analysis, this service has been shown to cause a 5% reduction in average consumption for two California markets, with no statistically significant effect in a third (Brent et al., 2015). A 2007 randomized experiment in Cobb County Georgia found strong evidence that social comparison messages had a substantially larger impact than prosocial content and technical recommendations (Ferraro and Price, 2013). They find an estimated 4.8% effect when treatment combines social comparison, prosocial messaging, and technical suggestions. Both the *WaterSmart* program and Georgia study find significant heterogeneity in treatment effect across household types, while only the *WaterSmart* analysis observes stable persistence in treatment effect over time. DC differs from both of these programs for their emphasis on leveraging digital communication platforms, rather than paper reports, which allows for greater flexibility in message content, more frequent and varied content, and the option to survey customer feedback.

We analyze the effect of enrollment in DC service on average monthly water consumption in the City of Folsom. We provide evidence that DC effect in the water sector is comparable and even larger than Opower’s effect in the energy sector. In addition, we examine the heterogeneity in the treatment effect by baseline water usage. Understanding variation in treatment effects of DC helps target subgroups in a cost effective manner. Also, this result helps researchers to understand generalizability of the treatment effects to different populations and places (Heckman et al., 1997; Djebbari and Smith, 2008; Ferraro and Miranda, 2013; Manski, 2004; Imai et al., 2013; Ferraro and Miranda, 2013).

3 Overview of Dropcountr Services

DC users have anytime access to water usage and other information via their mobile devices (iOS and Android) or by logging into their account on the web. In addition, DC sends users

a monthly email summary of their water use, including contextual comparisons and water utility announcements. While DC can and does work with utilities who read their meters monthly or bi-monthly, DC is especially well suited for utilities who have migrated to smart metering.

Users who have downloaded the mobile application receive “push” notifications to their mobile devices. These notifications can alert households when they may be approaching the next tier for a block-pricing utility, indication of leaks, rebate opportunities or other tips. The web platform allows customers to access their DC account, where they can explore their monthly report in more detail and access similar information that may be generated through the mobile alerts. Additionally, DC will produce and mail paper water use reports for utilities that request it.

The “Your Water” interface on both mobile and web apps includes four main features: summary statistics of usage, which includes reference to an individualized “goal”; comparison of usage to “similar” and “efficient” households; and conservation tips tailored to their account characteristics. An example of this interface may be found in [Figure 1](#).

The top portion provides statistics on monthly and average daily consumption, along with a graphical representation of their historical consumption over the previous 12-month period. In addition, this portion of the report evaluates the households’ performance in achieving their “goal” water usage. A goal is in effect an account-specific value, and represents the amount of water required by the account each month of the year. The goal is the sum of an indoor budget, primarily determined by household occupancy, and an outdoor budget, which based on parcel size, irrigable area and local weather and other climate factors such as local evapotranspiration constants. The industry standard and baseline assumption is that 50% of parcel area is irrigated; households may update this irrigation profile, along with other household features, in their DC account.

The social comparison component informs customers how their usage compares to “similar” or “nearby” households and “efficient” households. A “similar/nearby” household lies within a specified radius of the given account and is comparable in features, such as lot size

and household occupancy. Households with consumption below a certain percentile of the distribution are labeled “efficient” by DC. The “Relevant water saving tips” portion of the report encourages water savings by suggesting two conservation tips per report, out of over 100 recommendations, which are tailored to that particular household’s profile and past use. Finally, customers are encouraged to log into their online account, where they may explore their report in greater detail and receive further conservation information. [Table 1](#) indicates types and number of messages sent by DC to the enrolled customers in the city of Folsom.

4 Data and Empirical Strategy

4.1 Data

In mid-December of 2014, all account holders in the City of Folsom service area were offered the option of participating in the DC pilot program on a “first come, first served” basis. Offer of service came as a paper advertisement, on city letterhead, with a monthly bill and included a market insert that illustrated the look and style of the DC web and mobile platforms. The utility contracted for a maximum of 5,000 accounts, with current enrollment just over 3,350 accounts. Progression of DC enrollment over the treatment period in the City of Folsom service area is presented in [Figure 2](#).

For this analysis, households who participated in the DC service offer at any point during the study period will be referred to as “treated” households, while those who do not are “control” households. The first full month after which a household has received their first DC report is considered the first treatment month. Therefore, in the City of Folsom since enrollment began in December 2014, the first reports were generated in January 2015, makes the first possible treatment month. This approach is consistent in defining treatment for both *Opower* and *WaterSmart* program analysis. Using this definition of treatment, rate of enrollment is represented in [Figure 2](#).

[Table 2](#) present summary statistics for the number of households and also a number of

observations before and after treatment in each group. The treatment group includes 3,089 households and the control group includes 15,986 households. A number of observations before and during the treatment period are approximately the same in both groups.

[Table 3](#) provide the basic double difference result in levels ($[\text{consumption in treatment post policy} - \text{consumption in control post policy}] - [\text{consumption in treatment pre-policy} - \text{consumption in control pre-policy}]$) for the City of Folsom. The result in [Table 3](#) indicates that water consumption in treated households was reduced on average by 0.9 CCFs (748 gallons) per month due to the DC service which is equivalent to 5.56% of average monthly usage.

Further investigation of the pre-trends between control and treatment groups is analyzed using graphical analysis. [Figure 3](#) present average water consumption by month in the treatment and control groups with a vertical dash line which indicates the treatment start time. This graph illustrates that despite differences in average consumption across the treated and control groups before treatment, there exists a visually distinct increase in this difference in average monthly consumption between treated and control households following the introduction of DC service (indicated by the vertical dashed line). In other words, we observe graphical evidence that there is a larger difference in average water usage after treatment between those households that enrolled in DC and those that did not enroll.

We observe this difference more clearly by plotting the difference in average monthly consumption as a percent difference between the two groups across the sample time horizon. [Figure 4](#) illustrates how this percent difference changes across the sample period. Reflecting the pattern observed in [Figure 3](#), we see that there is a significant increase in the difference in average monthly consumption as a percent between the pre-period, prior to the availability of DC services, and the post-period, with households under DC treatment. The dashed pink line represents the average percent difference in the pre- and post-periods. In the pre-period, we observe that households who become DC enrolled consume approximately 2% less water per month on average. Whereas, in the post-period, households who are enrolled in DC consumed about 8% less water per month on average.

Figure 5, Figure 6, and Figure 7 illustrate how difference between treated and control groups in the City of Folsom changes across the households with different baseline consumption. For the purpose of these figures, quantiles of consumption are defined based on the average baseline summer usage. Quantiles threshold in CCFs are: 13.11 and lower as first quantile, between 13.11 and 20.26 as second, between 20.26 and 27.93 as third, between 27.93 and 39.33 as fourth, and higher than 39.33 as fifth quantile. These figures illustrate that there are larger increases in the difference in average monthly consumption between treated and control households following the introduction of DC service for the higher quantiles.

In addition to this graphical evidence of parallel trends, various fixed effects are employed to account for both seasonal, annual, and household invariant factors that may determine consumption. Given the extensive amount of baseline data and number of observations, these fixed effects are able to explain a large amount of variation that could otherwise bias results.

For regression analysis purpose, we organize two panel datasets of household-level monthly water consumption in the City of Folsom water utility service areas. City of Folsom panel begins January 2013 and ends in August 2016, this period includes the start date of the DC service (December 2014). The regression results measure the effect of DC taking into account household characteristics that also affect consumption (e.g. lot size) as well any seasonal or year-specific effects on consumption. The average effect of DC enrollment on water consumption is estimated by defining two groups; households who enrolled in DC (treated households) and households who did not enroll in DC (control households).

In the difference-in-differences regression, the outcome of interest is the log of the households' monthly water consumption. Equation 1 indicates the preferred specification. which q_{hmy} is the water consumption in the household h at month m and year y . The variable of interest is DC which denotes whether a household observation is in the treatment group during the post period in which DC was active. We include household-calendar month fixed effects (γ_{hm}) which controls for two types of variables. First, these control for time-constant variables specific to a household, e.g., number of toilets which is fixed for the vast majority

of households in our sample. Second, the household-calendar month fixed effects control for calendar month specific water-use factors specific to each household, e.g., household X has an outdoor irrigation system set to medium irrigation every May and set to high every July. We also include calendar-month year fixed effects (μ_{my}), which control for consumption factors which are common to all household within a given calendar month for a specific year, e.g., an unseasonably warm October in 2015 or time-specific regulation such as the 2015-16 water restrictions administered by the California Water Resources Control Board, in our preferred specification; ϵ_{hmy} captures all unobservables which affect the dependent variable.

$$\log(q_{hmy}) = \alpha_1 \cdot \text{Dropcountr} + \gamma_{hm} + \mu_{my} + \epsilon_{hmy} \quad (1)$$

5 Result

We begin with a presentation of the average change in the quantity of water consumption due to the DC service, then we explore heterogeneity of this effect across baseline consumption quantiles.

5.1 Average effect of Dropcountr

Results for the difference-in-differences specifications are presented in [Table 4](#). Log of monthly water consumption in households is the dependent variable in all of the specifications. Standard errors for all of the specifications are reported in the parenthesis and are clustered at the level of the households to account for within-household serial correlation in the error term. First column of both tables has households-by-month fixed effects, month-by-year fixed effects, and DC effect. The DC effect is defined by an interaction between post-period and treatment households.

The point estimate of average treatment effect (DC effect) using column (1) specification

indicates that DC service suggests 5% reduction in monthly water consumption, on average. This result is both statistically and economically significant, meaning we can reject the hypothesis that there is no effect of DC enrollment on average monthly water consumption. The change in average gallons per day is an estimated 24 fewer gallons for the average enrolled household. To put these reductions in perspective: the average shower uses 16-40 gallons (depending on shower head efficiency), clothes washing machines require 25-40 gallons per wash, while dishwashers use 6-16 gallons per load. In addition, the estimates reported here are consistent with those found for *WaterSmart* Software of a 4.9-5.1% average treatment effect for two experimental designs (where no effect was found for a third utility) ([Brent et al., 2015](#)).

Notably, although the previous graphs suggest that all households reduced consumption in the post-period, the controls in our regression analysis allow identification of DC's effect on household consumption that takes this general reduction into account. Thus, we find that DC treated households reduced consumption during the post-period more than households who did not enroll in DC. Taking into account baseline differences and controlling for consumption factors as described in the discussion of the econometric model presented in [Equation 1](#).

5.2 Investigating Heterogeneity

In this section, we move beyond estimation of average treatment effects and we consider estimating heterogeneity of household's responses to DC. Understanding heterogeneity of treatment effect will allow to target households that are more responsive which will be a cost-effective strategy ([Heckman et al., 1997](#); [Djebbari and Smith, 2008](#); [Ferraro and Miranda, 2013](#)). Also, investigating treatment effect by subgroups helps researchers understand generalizability of the result of this study to other populations and places ([Manski, 2004](#); [Imai et al., 2013](#); [Ferraro and Miranda, 2013](#)).

We explore heterogeneity of treatment effect by average summer baseline (pre-) period water consumption. For each household, we calculate the mean summer pre-treatment wa-

ter consumption. Next, we create dummy variables for whether that mean summer pre-treatment water consumption is in the first, second, third, fourth, or fifth quantile of the whole sample summer pre-treatment consumption (i.e. Q.1, Q.2, etc.). Next, we interact these dummies with treatment household and time dummy indicators. We defined baseline consumption quantiles as 20% and lower, between 20% and 40%, between 40% and 60%, between 60% and 80%, and higher than 80% percentiles. Quantiles threshold in CCFs are: 13.11 and lower as first quantile, between 13.11 and 20.26 as second, between 20.26 and 27.93 as third, between 27.93 and 39.33 as fourth, and higher than 39.33 as fifth quantile.

Results for this specification are reported in column (2) of [Table 4](#). The control variables in this regression correspond to Columns 1 in same table. We find that the DC effect is monotonically increasing in baseline consumption level—the largest effect is observed for the group with highest baseline consumption. These results are consistent with the average effect for all households that is estimated and presented in the first column of the same table. Preliminary analysis suggests that households in the highest quintile of baseline consumption reduce consumption by an estimated 13% in response to the DC service. However, there appears to be a 7.2% increase in usage in average monthly consumption for those households in the lower quartile of baseline consumption. This response is referred to as a “boomerang effect”, where customers who learn that they are actually using less than their neighbors or other households like their own increase their demand ([Clee and Wicklund, 1980](#)). It should be noted that the analyses on both *Opower* and *WaterSmart* do not find evidence of a boomerang effect in any of the studied markets. The techniques employed here take a rather coarse approach to segmenting the population. Continued work on this project will explore a potential boomerang effect in greater detail.

The coefficient -0.047 in column (1) of [Table 4](#) summarizes average percent reduction across all households. This is different than the aggregate reduction in consumption resulting from DC because it does not take into account the fact that households with high levels of baseline use experienced larger percentage reductions than households with lower baseline use. As a consequence, the average percentage reduction capture by the coefficient in column (1) is less than the aggregate effect of DC. In summary, in terms of overall impact of DC,

the object of interest is the aggregate treatment effect, which we estimate to be -6.98% for the population of households participating in DC. Assuming all of Folsom participated in DC and had a similar response, then the aggregate reduction in water consumption for the single family residential sector in Folsom would be -6.94%. This is slightly lower than the aggregate effect for participating households because the composition of households in terms of baseline use is shifted towards higher end users for the overall population in Folsom ².

Table 5 summarizes reductions in water usage in levels rather than percentage reductions (also taking into account timing of enrollment) due to Dropcountr for households who participated in the program. Total reduction in water consumption due to Dropcountr is 37.5 million gallons from January-2015 to September-2016 (inclusive). Dropcountr caused 9.5 and 23 million gallons reduction in consumption for households in quintile 4 and 5, respectively.

6 Conclusions and Policy Implications

This study provides insight how a social-norm-based conservation programs effect on water usage. Specifically, the effect of DC on water usage was examined by using household-level panel data and adopting a difference-in-differences approach. Results suggest that the introduction of the DC services for the population of households participating in DC causes aggregate treatment effect of 7% reduction in water usage. In addition, analysis suggests that in response to the DC service households in the highest quintile of baseline consumption reduce water usage by an estimated 13% in the City of Folsom –at the margin, this is a large

² Aggregate treatment effect for the population of households participating in DC is calculated using following equation:

$$\text{Aggregate Effect} = \frac{\sum_{i=1}^5 \bar{q}_i * \beta_i * (NHH_i)}{\sum_{i=1}^5 \bar{q}_i * (NHH_i)} \quad (2)$$

where: *Aggregate Effect* is aggregate treatment effect for the population of households participating in DC, \bar{q}_i indicates average usage in 2013 for households who eventually enrolled in DC, β_i indicates estimated coefficient for the quintile i from Table 4, and NHH_i indicates number of enrolled households in quintile i .

effect.

These are also evidence that not all of the households react alike to DC. The results hold as a general rule, those in the higher quantiles of the baseline water usage had the largest responses. This result is comparable with the existing literature ([Allcott, 2011](#); [Ferraro and Miranda, 2013](#); [Brent et al., 2015](#)). Such a result indicates the effectiveness of sub-group targeting in social-norm-based conservation programs towards baseline users with higher consumption.

Future analyses that we aim to investigate are (i) the persistence of these effects, (ii) whether there are other subgroups to target besides high baseline users (e.g., high income households or those with large lot sizes), (iii) the channels through which the DC program acts (e.g., consumption feedback, social comparison, household budget, etc.), and (iv) whether the program's effect can be magnified when coupled with other conservation programs (e.g. daily water readings, lawn replacements, media messaging, etc.).

References

- Agthe, D. E. and R. B. Billings (1987). Equity, price elasticity, and household income under increasing block rates for water. *American Journal of Economics and Sociology* 46(3), 273–286.
- Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics* 95(9), 1082–1095.
- Allcott, H. (2012). Site selection bias in program evaluation. Technical report, National Bureau of Economic Research.
- Allcott, H. and T. T. Rogers (2012). How long do treatment effects last? persistence and durability of a descriptive norms intervention’s effect on energy conservation.
- Ayres, I., S. Raseman, and A. Shih (2012). Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *Journal of Law, Economics, and Organization*, ews020.
- Brent, D. A., J. H. Cook, and S. Olsen (2015). Social comparisons, household water use, and participation in utility conservation programs: Evidence from three randomized trials. *Journal of the Association of Environmental and Resource Economists* 2(4), 597–627.
- Buck, S., M. Auffhammer, S. Hamilton, D. Sunding, et al. (2015). Measuring the welfare losses from urban water supply disruptions. Technical report.
- Buck, S., M. Nemati, D. Sunding, et al. (2016). The welfare consequences of the 2015 california drought mandate: Evidence from new results on monthly water demand. In *2016 Annual Meeting, July 31-August 2, 2016, Boston, Massachusetts*, Number 236049. Agricultural and Applied Economics Association.
- Buck, S., H. Soldati, D. L. Sunding, et al. (2015). Forecasting urban water demand in california: Rethinking model evaluation. In *2015 AAEA & WAEA Joint Annual Meeting, July 26-28, San Francisco, California*, Number 205737. Agricultural and Applied Economics Association & Western Agricultural Economics Association.

- Clee, M. A. and R. A. Wicklund (1980). Consumer behavior and psychological reactance. *Journal of Consumer Research* 6(4), 389–405.
- Costa, D. L. and M. E. Kahn (2013). Energy conservation “nudges” and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association* 11(3), 680–702.
- Dalhuisen, J. M., R. J. Florax, H. L. De Groot, and P. Nijkamp (2003). Price and income elasticities of residential water demand: a meta-analysis. *Land economics* 79(2), 292–308.
- Djebbari, H. and J. Smith (2008). Heterogeneous impacts in progress. *Journal of Econometrics* 145(1), 64–80.
- Ferraro, P. J. and J. J. Miranda (2013). Heterogeneous treatment effects and mechanisms in information-based environmental policies: Evidence from a large-scale field experiment. *Resource and Energy Economics* 35(3), 356–379.
- Ferraro, P. J., J. J. Miranda, and M. K. Price (2011). The persistence of treatment effects with norm-based policy instruments: evidence from a randomized environmental policy experiment. *The American Economic Review* 101(3), 318–322.
- Ferraro, P. J. and M. K. Price (2013). Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. *Review of Economics and Statistics* 95(1), 64–73.
- Heckman, J. J., J. Smith, and N. Clements (1997). Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies* 64(4), 487–535.
- Imai, K., M. Ratkovic, et al. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics* 7(1), 443–470.
- Ito, K. (2014). Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *The American Economic Review* 104(2), 537–563.

- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica* 72(4), 1221–1246.
- Olmstead, S. M., W. M. Hanemann, and R. N. Stavins (2003). Does price structure matter? household water demand under increasing-block and uniform prices. *New Haven: School of Forestry and Environmental Studies, Yale University, Working Paper*.
- Olmstead, S. M., W. M. Hanemann, and R. N. Stavins (2007). Water demand under alternative price structures. *Journal of Environmental Economics and Management* 54(2), 181–198.
- Olmstead, S. M. and R. N. Stavins (2007). Managing water demand: Price vs. non-price conservation programs. *A pioneer institute white paper* 39.
- Renwick, M. E. and R. D. Green (2000). Do residential water demand side management policies measure up? an analysis of eight california water agencies. *Journal of Environmental Economics and Management* 40(1), 37–55.

Figures and Tables



Figure 1: Dropcountr Home Water Use Report sample




Figure 2: Progression of DC enrollment over the treatment period in City of Folsom, CA. Total number of enrolled households by end of September 2016, was 3,350.




Figure 3: Average monthly consumption by Dropcountr enrollment status in City of Folsom, CA. Vertical dashed line indicates start of treatment period (January- 2015)




Figure 4: Difference is average monthly consumption, as a percent, across time by Dropcountry enrollment status in City of Folsom, CA. Vertical dashed line indicates start of treatment period. Pink dash represents the average percent difference in household consumption for the pre- and post-periods. Average percent difference in household consumption for the pre-periods is 2.25 and for post-periods is 8.63.




Figure 5: Average monthly consumption by Dropcountr enrollment status and baseline consumption in quantiles **one** and **two** in City of Folsom, CA. Vertical dashed line indicates start of treatment period (January- 2015). Quantiles of consumption are defined based on the average baseline summer usage. Quantiles threshold in CCFs are: 13.11 and lower as first quantile, between 13.11 and 20.26 as second, between 20.26 and 27.93 as third, between 27.93 and 39.33 as fourth, and higher than 39.33 as fifth quantile.




Figure 6: Average monthly consumption by Dropcountr enrollment status and baseline consumption in quantiles **three** and **four** in City of Folsom, CA. Vertical dashed line indicates start of treatment period (January- 2015). Quantiles of consumption are defined based on the average baseline summer usage. Quantiles threshold in CCFs are: 13.11 and lower as first quantile, between 13.11 and 20.26 as second, between 20.26 and 27.93 as third, between 27.93 and 39.33 as fourth, and higher than 39.33 as fifth quantile.




Figure 7: Average monthly consumption by Dropcountr enrollment status and baseline consumption in quantiles **four** and **five** in City of Folsom, CA. Vertical dashed line indicates start of treatment period (January- 2015). Quantiles of consumption are defined based on the average baseline summer usage. Quantiles threshold in CCFs are: 13.11 and lower as first quantile, between 13.11 and 20.26 as second, between 20.26 and 27.93 as third, between 27.93 and 39.33 as fourth, and higher than 39.33 as fifth quantile.

Table 1: Summary of message types sent by Dropcountr to the enrolled customers in City of Folsom

Message Types	Number of times sent
Utility admin message	50,219
Monthly report email	36,356
Unsolicited monthly report email	3,446
Leak alert	2,541
New user tips	767
Total	93,329

Table 2: Summary statistics of Summary Statistics of Data Availability for Analysis in City of Folsom, CA. Monthly consumption values in CCFs for baseline period: January 2013 through December 2014.

	All accounts	Control group	Treatment group
Number of accounts	19,075	15,986	3,089
Pre-period observations	437,327	365,515	71,812
Treatment period observations	330,172	272,653	57,519
Baseline:			
Average	20.13	20.11	20.24
25th percentile	7.78	7.65	8.05
Baseline median	15.00	14.88	15.70
75th percentile	26.51	26.36	27.00

Table 3: Average Monthly Water Consumption of Treated and Control Group (CCFs) in City of Folsom, CA

	(1)	(2)	(3)	(4)
	Control Households	Treated Households	Difference (levels)	Difference (%)
Pre-period	20.11	20.24	0.13	0.66
Post-period	15.62	14.85	-0.76	-4.90
Double Difference	-4.49	-5.39	-0.90	-5.56

Notes: Households that never enrolled in Dropcountr consumed on average 20.11 CCF of water pre-period; this number reduced to 15.62 CCF in post-period. However, households that eventually enrolled in Dropcountr consumed 20.24 CCF of water pre-period and 14.85 CCF in post-period. Comparing two groups indicates that Dropcountr reduced water consumption in treatment group by 0.9 CCF per month. In percentage terms, Dropcountr reduced water consumption in treatment group by 5.56%.

Table 4: Average treatment effect of opt-in Dropcountr enrollment and Heterogeneous effects by consumption quantile in City of Folsom, CA

	(1)	(2)
Dropcountr Average Effect	-0.047*** (0.003)	-
Dropcountr Effect in Quantile 1		0.072*** (0.009)
Dropcountr Effect in Quantile 2		-0.018** (0.007)
Dropcountr Effect in Quantile 3		-0.044*** (0.006)
Dropcountr Effect in Quantile 4		-0.060*** (0.006)
Dropcountr Effect in Quantile 5		-0.129*** (0.007)
Dropcountr Aggregate Effect		-6.9%
Month by Year Effects	Yes	Yes
Household by Month Fixed Effects	Yes	Yes
Observations	767,000	767,000
R-squared	0.185	0.185

Notes: Quantiles of consumption are defined based on the average baseline summer usage. Quantiles threshold in CCFs are: 13.11 and lower as first quantile, between 13.11 and 20.26 as second, between 20.26 and 27.93 as third, between 27.93 and 39.33 as fourth, and higher than 39.33 as fifth quantile. Dropcountr aggregate effect is calculated for the population of households participating in Dropcountr. *p<0.1; **p<0.05; ***p<0.01

Table 5: Cumulative water savings in absolute terms for households who participate in Dropcountr program by September-2016 (All of the consumption numbers are in thousand gallons) in the City of Folsom

Quintile	Number of households	Consumption after enrollment	Dropcountr Effect	Consumption if not enrolled	savings
1	333	25,439	0.072	23,730	-1,709
2	586	67,815	-0.018	69,058	1,243
3	779	118,971	-0.044	124,446	5,476
4	767	149,931	-0.060	159,501	9,570
5	535	154,607	-0.129	177,505	22,898
Total	3,000	516,762	-	554,240	37,478

Notes: All of the consumption numbers are in thousand gallons. Overall, we estimate Dropcountr reduced aggregate water consumption by 37.5 million gallons for program participants between enrollment up to September 2016.