3-2021

Isolating The Effect Of Injunctive Norms On Conservation Behavior: New Evidence From A Field Experiment In California

Syon Bhanot
Swarthmore College, sbhanot1@swarthmore.edu

Recommended Citation

Syon Bhanot. (2021). "Isolating The Effect Of Injunctive Norms On Conservation Behavior: New Evidence From A Field Experiment In California". *Organizational Behavior And Human Decision Processes.* Volume 163, 30-42. DOI: 10.1016/j.obhdp.2018.11.002
https://works.swarthmore.edu/fac-economics/462

Isolating the Effect of Injunctive Norms on Conservation Behavior: New Evidence from a Field Experiment in California

Syon Bhanot [1]

Swarthmore College

Forthcoming, Organizational Behavior and Human Decision Processes

November 2018

[1]My thanks to Will Holleran, Ora Chaiken, and the team at WaterSmart Software for making this work possible. Thanks also to Chase Williamson, Sam Wang, Stephanie Kestelman, Rohit Nair, Nerissa Nashin, Haksoo Lee, Tony Trinh, and (especially) Reed Orchinik for research assistance.

Abstract

Social norms messaging campaigns are increasingly used to influence human behavior, with social science research generally finding that they have modest but meaningful effects. One aspect of these campaigns in practice has been the inclusion of injunctive norms messaging, designed to convey a social judgment about one's behavior (often in the form of encouraging or discouraging language, or a visual smiley or frowny face). While some prominent research has provided support for the use of such messaging as a tool for positive behavior change, causal evidence on the effect of injunctive norms messaging as a motivator (as opposed to just one part of a multifaceted messaging campaign) is limited. This paper presents a field experiment on water conservation behavior conducted by an organization in California, involving over 40,000 households, which provides some of the most precise evidence to date regarding the effect of injunctive norms on decision making. I find that not only do injunctive norms encourage conservation behavior, there is also no evidence that they discourage individuals from further attending to norms messaging–regardless of whether the social judgment conveyed is negative or positive. Taken together, this suggests that injunctive norms are a useful tool in "nudge"-style campaigns tackling behavior change.

Keywords: Social Norms, Conservation, Injunctive Norms, Social Judgment, Field Experiment, Behavioral Nudges

Isolating the Effect of Injunctive Norms on Conservation Behavior:
New Evidence from a Field Experiment in California

1 Introduction

There is a rich literature in behavioral economics and social psychology on how social norms can influence behavior. One growing area of research in this domain involves using social norms to influence conservation decisions, with research finding that providing households with information on the energy or water use of neighbors can decrease resource consumption (Allcott, 2011; Brent, Cook, & Olsen, 2015; Ferraro & Price, 2013). Important past work in this area has argued that injunctive norms framing–messaging conveying a social judgment about behavior–is a key element of such informational campaigns (Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007). This work is theoretically important also; by exploring the distinct behavioral impact of injunctive norms, we gain insights regarding how and why social norms impact behavior. That is, if people are influenced by injunctive norms, it implies that individual perceptions regarding what others consider "good" and "bad" behavior are an important part of why social norms are so influential in everyday life.

However, there is very little causal evidence regarding the distinct impact of injunctive norms messaging on behavior, with most experimental work estimating the effects of broader campaigns that incorporate injunctive norms as only one component. Furthermore, as I will discuss, there are some shortcomings in existing research on the effect of injunctive norms that make existing findings not definitive. In this paper, I present experimental evidence from a natural field experiment that addresses many of these shortcomings. Furthermore, the experiment I present allows me to explore the impacts of injunctive norms messaging not only on conservation behavior, but also on the desire for future social norms messages.

This paper builds on a large literature in social science on social norms and interventions that leverage norms for behavior change. Researchers studying social norms often emphasize the distinction between descriptive norms, which simply describe the behavior of others, and injunctive norms, which convey social approval or

disapproval of a given behavior (Cialdini et al., 2006; Jacobson, Mortensen, & Cialdini, 2011; Schultz et al., 2007). In a seminal paper, Schultz et al. (2007) presents the results of a field experiment using door hangers to test the effects of using injunctive and descriptive norms messaging to reduce energy use. The authors find that low energy users increased their energy use when receiving descriptive norm information only (an adverse "boomerang effect" from social norms information), but argue that this effect for low users was eliminated by an additional injunctive norm message (a smiley or frowny face that was hand drawn onto the door hanger). Furthermore, they find that the inclusion of injunctive norms messaging also reduced energy consumption for high users, above and beyond how much descriptive norms reduced use.

While Schultz et al. (2007) offers promising results in support of the independent effect of injunctive norms on behavior, the study had some shortcomings that limit the extent to which its findings can be treated as definitive. First, the injunctive norm message in Schultz et al. (2007) came in the form of a hand-drawn visual, which was absent from the "descriptive norm only" condition. It is possible that the visual cue, independent of the social judgment it conveyed, might have influenced behavior by drawing the reader's eye to the door hanger in a way that the "descriptive-norm-only" door hanger did not. Alternatively, the messaging may have worked because the hand-drawn nature of the visual introduced the idea of being observed by an actual person, which is a similar but not conceptually identical motivator to injunctive norms. Second, the sample size in Schultz et al. (2007) was relatively small (290 households), which makes independent conceptual replications with a more robust sample size an important exercise (Aarts et al., 2015; Camerer et al., 2016; Maniadis, Tufano, & List, 2014). Third, their sample consisted only of households who knew they were being studied, which might have influenced their decisions. Fourth, as with many other studies in this domain, it was impossible to know who actually looked at the informational door hanger, making it difficult to learn more about how the messaging impacted those who were actually attentive to it.

Despite these limitations, Schultz et al. (2007) had significant impacts on both

academic work and in applied settings, with a number of firms and organizations using their findings to develop tools for behavior change in the resource conservation space. In particular, firms like Opower and WaterSmart Software (the partner firm in this study) developed and continue to use social information with injunctive messaging to influence conservation behavior, with recent academic work finding that the messaging developed by these firms has had meaningful effects on the conservation of energy (Allcott, 2011; Allcott & Rogers, 2014) and water (Brent et al., 2015). However, when this research has explored the impact of injunctive norms specifically, the results have been mixed. For example, Allcott (2011) uses a regression discontinuity approach to estimate the impact of the specific smiley faces used in Opower messaging, but fails to find evidence supporting the impact of any one type of injunctive norm on energy use. However, the evidence in Allcott (2011) on injunctive norms also has its limitations. For example, while the regression discontinuity design provides a strong estimate, it is restricted in its causal inference to individuals located around the thresholds. Furthermore, as with most other papers in this literature, Allcott (2011) is unable to distinguish between those who do versus do not actually see the social norms information (as the information is sent in paper form through the mail). Taken together, the body of existing evidence can be summarized as providing imperfect or mixed evidence regarding the independent impact of injunctive norms messaging on behavior.

In this paper, I provide clear, causal evidence on the independent effect of injunctive norms on conservation behavior. Specifically, I present the results from a large-scale, randomized field experiment conducted by WaterSmart Software, a firm in California that works with public utilities to reduce household water use. In the experiment, the firm randomly assigned 45,866 households to one of four different conditions[2]: (1) a control group that received no social information; (2) a "No Drop" treatment group that received social information without a visual cue; (3) a "Drop" treatment group that received social information with a basic visual cue (a plain

[2]In this paper, we only analyze the data for a large subset (41,365) of these households, for reasons outlined in section 3.5.

cartoon water droplet); and (4) an "Injunctive Drop" treatment group that received exactly the same messaging as the "Drop" condition, but with an injunctive message added *into* the visual cue (a smiley, neutral, or frowny face inside the water droplet, depending on the water consumption of the household). Importantly, this experiment took place without the households' knowledge, and was delivered using tracked emails (meaning that whether or not an email was opened was measured). The intervention took place over multiple months, with subjects receiving up to seven emails each with messaging that varied as their water use changed over time.

The nature of this experimental design allows for a variety of contributions to existing literature. First and foremost, the design helps us isolate the causal impact of injunctive norms on behavior (by comparing the Drop and Injunctive Drop conditions). This is of theoretical importance because if injunctive norms messaging "works" here, it suggests that social norms shape behavior, at least in part, because they transmit information about what individuals "should" be doing. Furthermore, I provide evidence not only on the treatment effects of injunctive norms on conservation behavior, but also on individual-level willingness to obtain further social norms information (since I can observe whether emails were opened or not). Second, because of the large sample size in this experiment, I am able to provide evidence that is significantly more statistically robust than previous work with smaller samples. Third, this study represents a "natural field experiment," meaning that subjects were unaware that they were being studied using a randomized experiment (Harrison & List, 2004). This results in findings that are highly generalizable to similar field contexts, where organizations seek to influence real world behavior using injunctive norms messaging.

There are two key findings from the experiment. First, I find clear evidence that the inclusion of an injunctive norm message does reduce water use, a positive and independent effect from injunctive norms messaging that aligns with the findings in Schultz et al. (2007). Notably, there is little evidence to suggest that this effect from injunctive norms is much different for high versus low water users. Second, I find no evidence that injunctive norms messaging discourages individuals from attending to

future social norms messaging. That is, subjects receiving injunctive norms messaging are not significantly less likely to open future email messages motivating them with social norms (and injunctive) messages. Taken together these findings affirm the ability of injunctive norms messaging to encourage conservation behavior in particular, and support the use of such messaging in "nudge"-style campaigns moving forward.

This paper proceeds as follows. Section 2 provides background information on social norms interventions and the effect of injunctive norms on behavior in particular, and then outlines the hypotheses of the experiment. Section 3 outlines the experiment itself. Section 4 offers a description of the empirical methods used for analyzing the data and presents the results. Section 5 discusses the implications of the results and concludes.

2 Background

2.1 Social Norms Messaging

A large and growing body of literature has explored the effects of social norms information on human decision making in a variety of contexts, including savings behavior (Beshears, Choi, Laibson, Madrian, & Milkman, 2015), restaurant orders (Cai, Chen, & Fang, 2009), charitable giving (Frey & Meier, 2004; Krupka & Croson, 2016), and voting behavior (Funk, 2010; Gerber, Green, & Larimer, 2008; Gerber & Rogers, 2009). Among the largest strands of literature in this domain studies interventions that use social norms messaging to influence prosocial, environmentally-friendly behavior, with a series of high-profile papers documenting the robust power of the social norms approach (Allcott, 2011; Allcott & Rogers, 2014; Brent et al., 2015; Cialdini et al., 2006; Ferraro & Price, 2013; Schultz et al., 2007).

The prior work that is most directly relevant for this paper is Schultz et al. (2007), which presents the results of a field experiment designed to test the effects of using injunctive and descriptive norms messaging to reduce energy use. In the study, the authors randomly assigned 290 recruited households in California to one of two conditions: (1) a "descriptive-norm-only" group receiving a door hanger with

information about the household's energy use and the energy use of the household's neighbors; and (2) a "descriptive-plus-injunctive-information" group, receiving the same descriptive norm information but with an added hand-drawn "frowny" or "smiley" face included (based on whether or not a household's energy use was above or below the average consumption of other homes in their neighborhood). The authors find that low energy users increased their energy use when receiving descriptive norm information only, but find that this effect for low users was eliminated when injunctive norm messaging was added. Furthermore, they find that the inclusion of injunctive norms messaging was also effective for high energy users, as adding an injunctive message resulted in slightly larger decreases in energy use by these users than was achieved by descriptive norms only.

In the decade since the publication of Schultz et al. (2007), private firms like Opower and WaterSmart Software have built large-scale social norm messaging campaigns (incorporating both injunctive and descriptive norms) to encourage conservation. Academic work has explored the causal impact of such efforts, consistently finding that social norms messaging campaigns result in greater reductions in resource consumption (Allcott, 2011; Brent et al., 2015; Ferraro & Price, 2013). In Allcott (2011), for example, the author reports on a series of randomized control trials conducted by Opower, using social norms messaging to influence energy use. Using data from 17 separate Opower projects, the author estimates the causal impact of social norms messaging at roughly 2%, and further finds that the decreases in energy use are especially pronounced for previously high users of energy. Ferraro and Price (2013) find similar results from an experiment conducted in Cobb County, Georgia, estimating that social norms messaging in a paper mailer campaign reduced water use by 2-5%, depending on the "strength" of the social norm language. As in Allcott (2011), the authors also find that the largest effects come from those with the highest ex-ante water use. Finally, Brent et al. (2015) reports on three separate experiments conducted by WaterSmart Software on the efficacy of similar social norms messaging campaigns, finding treatment effects that range from 1-5% reductions in water use. Taken together,

these papers offer evidence that social norms messaging campaigns that use both descriptive and injunctive norms are effective at reducing resource use on the aggregate.

2.2 Injunctive Norms

Most of the existing work on social norms messaging, discussed above, focused primarily on the average treatment effects of the campaigns overall. However, less emphasis was placed on which aspect of the social norms messaging–descriptive or injunctive–was most responsible for the behavior change. Conceptually, there is an important distinction between a descriptive social norm, which simply describes the behavior of others, and an injunctive social norm, which conveys social approval or disapproval of a given behavior (Cialdini et al., 2006; Jacobson et al., 2011; Schultz et al., 2007). That is, a descriptive social norm simply tells you what other individuals *are* doing, whereas an injunctive social norm conveys what others think you *should* be doing. Much of the research on social norms in the domain of conservation behavior has studied social norms messaging campaigns that used both descriptive and injunctive norms simultaneously. Clearly, it is challenging to use such a setup to disentangle the impact of injunctive norms from descriptive ones. In practice this is an important distinction to make–if we are to build effective behavior change campaigns using social norms, we should know if and how injunctive norm messaging supplements or detracts from descriptive norm information.

Overall, evidence regarding the independent impact of injunctive norms on behavior from past research is mixed. One result comes from Cialdini et al. (2006), who ran an experiment at the Petrified Forest National Park in Arizona exploring how messages on park signs influence the likelihood of petrified wood theft. In the study, the authors varied both the general framing of wording on signs (negative framing versus positive framing), and the nature of the norms message (injunctive versus descriptive norms language). The visuals on the signs also varied across the experimental conditions. The authors find that the most efficacious messaging was negatively-worded injunctive messaging (“Please don’t remove the petrified wood from the park”), but

conclude that "norm-based persuasive communications are likely to have their best effects when communicators align descriptive and injunctive normative messages to work in tandem rather than in competition with one another." (Cialdini et al., 2006, p.13)

While Cialdini et al. (2006) does provide evidence regarding the differential impacts of descriptive and injunctive norms, the experimental conditions were not ideally suited to make that distinction, for at least two reasons. First, the fact that the norms conditions manifested in the form of different written messages on the signs makes it hard to prove that the authors' text manipulation was only changing the messaging on the injunctive versus descriptive spectrum.[3] Furthermore, since there was no real control group in their design (treatments were compared to each other), injunctive norms messaging is not (and cannot be) compared directly to an absence of such messaging in their study. Second, the authors varied the visuals on the signs as well, as part of the treatment, which compounds the issue of identifying what exactly drove differences in behavior across conditions. Thus, it is hard to use Cialdini et al. (2006) to make definitive conclusions about the efficacy of injunctive versus descriptive norms.

Of course, Schultz et al. (2007) also provided important evidence on the impact of injunctive norms on behavior, through the use of an experimental condition that added an injunctive norm message (the smiley/frowny face) to a descriptive norm message, which was used alone in a different condition. However, as discussed earlier, their evidence has limitations that influence the extent to which it can be treated as definitive. First, the authors cannot rule out the possibility that the visual cue itself (and not the injunctive norms message it contained) may have caused the behavior change by drawing attention to the door hanger the authors used. Furthermore, the fact that the visual was hand-drawn may have communicated to households that they were being watched by actual people (which is distinct conceptually from injunctive norms as a motivator of behavior change). Second, the study's relatively small sample (290) of

[3]While the authors attempt to justify how individuals interpret norms using a survey with social psychology students, these measures may be affected by the students' knowledge about the intended effects of these norms from their coursework.

households who participated in the experiment knowingly limits the study's generalizability and necessitates replication. Taken together, these limitations suggest that further evidence is needed to better understand if and how injunctive norms contribute to behavior change in the context of large-scale, automated messaging campaigns.

Allcott (2011) also seeks to provide evidence regarding the causal impact of injunctive norms, by investigating experimental data on social norms messaging interventions to reduce energy use by Opower. Specifically, the author uses the fact that Opower varied the injunctive norm message they included on messaging using an algorithm with cutoffs. This allows the author to use a regression-discontinuity approach to estimate the causal impact of specific injunctive norms (by looking at individuals on either side of a given cutoff for a positive versus neutral versus negative injunctive norm message). The author finds no evidence that the nature of the injunctive norm messaging influences energy use, in contrast to the findings of Schultz et al. (2007). However, while the regression-discontinuity approach allows for a relatively compelling causal estimate in this case (with a large sample size), causal inference in the regression-discontinuity approach is limited to individuals located near the thresholds. Therefore, given the conflicting results on injunctive norms messaging in Allcott (2011) relative to previous findings, more evidence is needed in this domain.

Furthermore, as the discussion so far suggests, most work on injunctive norms in the literature on social norms interventions looks at a narrow form of behavior–namely the response on the specific behavior the norms messaging is about (energy or water use, in the conservation space). However, it is plausible that injunctive messaging might influence behavior or attitudes in other ways as well. For example, it is possible that the recipient of the injunctive norm message may react to scrutiny or social judgment by avoiding future "judgmental" information (Golman, Hagmann, & Loewenstein, 2017). In the context of digital messaging campaigns with many emails, this could mean individuals not opening emails they know will contain injunctive norms content, to avoid information about how others are judging them. This can be thought of as an

"ostrich effect" response to injunctive norms, whereby people "stick their head in the sand" via information avoidance when they suspect future information will reflect negatively on them (Eil & Rao, 2011; Karlsson, Loewenstein, & Seppi, 2009; Sicherman, Loewenstein, Seppi, & Utkus, 2016). Alternatively, some evidence suggests that injunctive norms may trigger self-reflection and effortful self-regulation, which could actually increase engagement with future social norms messaging (Jacobson et al., 2011). This is loosely related to what Gherzi, Egan, Stewart, Haisley, and Ayton (2014) refer to as the "meerkat effect," whereby the receipt of information may actually heighten vigilance and attentiveness to further information. Whereas previous studies in this literature have been unable to look into these possible reactions to injunctive norms messaging (since engagement with norms messaging was not observable at the individual level), I am able to offer evidence on this question by exploring whether or not subjects opened the social norms email messages they received.

2.3 Hypotheses

This paper posits two primary hypotheses regarding the impact of injunctive norms on conservation behavior. First, motivated by the findings in Schultz et al. (2007), I hypothesize that injunctive norms messaging, when added to descriptive norms messaging, leads to larger reductions in water use than descriptive norms messaging alone accomplishes. Second, I hypothesize that individuals do not have a strong aversion to injunctive norms messages, in the sense that receipt of injunctive norms messaging will not lower an individual's attentiveness to further norms messages. Admittedly, this second hypothesis relates to a question about which there is little existing evidence, and about which there are conflicting conceptual ideas (as outlined in section 2.2). A plausible alternative hypothesis would be that injunctive norms do influence attentiveness to future messaging, and that this might vary depending on the nature of that messaging. For example, it is plausible that a good social judgment (a smiley face) might trigger a different response, in terms of attentiveness to further messaging, than a bad social judgment (a frowny face). Given the relative shortage of

evidence, therefore, I explore this as a somewhat open question and offer some of the first causal evidence on it.

3 Experiment Overview

3.1 Partners

This experiment was implemented by WaterSmart Software, a California-based firm that works directly with utilities to promote behavior change in water use. WaterSmart's primary means of communication with households is via personalized emails containing Home Water Reports (HWRs), which contain social information on water use and recommendations about efficient household fixture upgrades. Through data sharing agreements with their utility partners, WaterSmart tracks water use and customer engagement over time. For this experiment, WaterSmart partnered with a public utility partner that serves various cities and towns in the greater San Francisco Bay Area.

3.2 Sample

The initial sample selected for inclusion in this experiment was 45,866 single-family homes from 28 different cities and towns in the greater San Francisco Bay Area. Residents of these homes all received water through the public utility partner. Each household in this study was associated with a single, unique water meter, which provided the water use data for this experiment, and a single email address, which received the HWR emails if the household was assigned to a treatment group.[4] Therefore, there is no complication from multiple housing units sharing a single water meter. WaterSmart intended to send a maximum of either six or seven emails to each household in the sample (depending on the exact date of their meter reads).

The five towns that supplied the largest number of subjects to the sample were Oakland, Richmond, San Leandro, Berkeley, and San Ramon. Prior to the start of this

[4]Because each household was associated with one subject and their email address, I use "subjects" and "households" somewhat interchangeably in the paper.

experiment, WaterSmart was already sending HWRs to other households in this region through their partnership with the public utility. However, this study targeted households who were being added to the existing base of customers receiving HWRs, so all households in this study were receiving HWR emails for the first time.

3.3 Study Design

All 45,866 households described were randomly assigned to one of four conditions: Control; No Drop; Drop; and Injunctive Drop. The Control group (10% of the sample) received no emails. The No Drop group (40% of the sample) received standard HWR emails, with peer information and other messaging, but no visual cue or injunctive norms messaging. The Drop group (40% of the sample) received the same HWR emails as the No Drop group, but with an additional visual cue at the top of the email (a water droplet). Finally, the Injunctive Drop group (10% of the sample) received the same HWR emails as the Drop group, but with the addition of an injunctive norm message inside the water droplet visual. Specifically, in the Injunctive Drop condition the droplet visual contained a smiley, neutral or frowny face (visible in Figure 2) to represent the household's water use performance relative to similar homes. The precise injunctive visual used in each HWR email in this condition was determined internally by WaterSmart, using an algorithm that took household demographic characteristics into account (to determine household water consumption relative to need). The Drop and No Drop groups were larger in size than the Control and Injunctive Drop conditions because of internal priorities at WaterSmart. An example of the visual differences between the treatments is provided in Figure 1. Note that the content below the drop visual (or absence of it) is the same across treatments, though both the text and the bar graph visuals naturally varied across households and over time (but in a manner that was unrelated to the treatment assignments in this experiment).

3.4 Demographics, Randomization, and Balance Check

Randomization was done by WaterSmart, using a simple randomization for the 45,866 households in the sample. The firm has a track record of experimentation and

conducting randomized evaluations with academic partners (Bhanot, 2017; Brent et al., 2015). Table 1 shows the mean values for various observable characteristics by condition, and overall, for the full sample of 45,866 households. To test the balance of the sample on the observed demographic characteristics, I run a set of regressions of each treatment condition on the various demographic characteristics, excluding data from all other conditions except the control. This allows for a set of pairwise balance checks for each treatment condition against the control. I then use F-tests to test for joint significance and balance. Table 2 shows the results of these analyses. In addition, Table 2 shows the results of a similar test for balance between the treatment conditions only, excluding all control observations; this is added because the key specifications in this paper involve comparisons between treatments.

When comparing each condition to the control in columns 1-3 in Table 2, we see signs of small but statistically significant imbalances in demographic characteristics between each individual treatment and the control, apparently stemming from the control group in particular. These differences seemingly occurred by chance, but necessitate the use of demographic controls in the analysis, particularly for analysis using the control group. Importantly however, there does not appear to be any imbalance between the households assigned to the three treatment conditions, as visible in columns 4-5 of Table 2. This is important, since between-treatment differences are the centerpiece of this paper's analysis.

The experiment proceeded based on this randomization, and a few details of implementation are worth emphasizing. First, to determine the nature of the information included in the HWR emails, WaterSmart used an internal algorithm to classify households based on water use need (using occupancy and lot size, among other variables). This classification was then used as the basis of comparison to populate the HWRs' social information content (including the nature of the injunctive norm for those in the Injunctive Drop condition, as mentioned earlier). Second, of the 41,280 subjects assigned to one of the three treatment groups from the full sample, most (34,327, or 83.16%) received six or seven experimental HWR emails as planned. Online

Appendix Table A.1 outlines the number of subjects in each treatment group, and the distribution of the number of emails received by subjects in each group. Third, because water use is highly seasonal (higher in the summer than in the winter), I use month-in-sample fixed effects to control for trends in water use across the sample over time. Furthermore, I control for each household's water use in the matching calendar month pre-experiment when analyzing post-experiment observations (as described in section 4), to control for any household-specific trends in water use by calendar month, which recur from year to year.

3.5 Data Restrictions

The dataset including the full 45,866 households had to be trimmed prior to analysis, which impacts how the results should be viewed. There were two sets of data restrictions — one set that involved removing data across conditions (treatments and control) and one set that involved removing data from treatment households only. Importantly, the first set of restrictions should not differentially impact the control and treatment conditions in terms of balance, but the second set may cause some imbalances. This ultimately means that we should not infer too much from the specific estimates from the analysis involving the control group, because of selection issues (especially in light of the imbalances found in the randomization check in section 3.4). However, this is not critical to the central findings of the paper, which relate to how the treatments compare to one another. In this section I will outline these two sets of data restrictions.

The first set of data restrictions applied to households across conditions. First, I removed 67 households that had at least one reading of over 5,000 gallons per day of water use in the period analyzed in the data (December 2013 onward). This was based on advice from WaterSmart, who identified these homes as outliers who likely had a major water leak in their high water use period or periods. Second, I focused my analysis on only the subset of remaining households that had a water meter read between November 17, 2014 and January 21, 2015. To understand the reason for this, note that while the experiment began in late 2014 and carried on through the latter half of 2015, households assigned to the treatment conditions began receiving HWRs at different times for a variety of scheduling and logistical reasons. In other words, there was no clear and universal "start date" for the experiment. This complicates analysis that involves the control group, because it is not immediately obvious when control households "would have" received their first mailer had they been treated. However, for households who got a water meter read in the window of dates noted above, I can directly link their water meter read dates with the dates they received (or "would have received," for the Control group) their first HWR (between December 9, 2014 and January 30, 2015). The exact process I used to do this is described in the Online Appendix. This process removed 706 households from the analysis, across conditions, leaving 45,093 households from the original sample.

A second set of data restrictions was then applied to these 45,093 households, and was specific to households in the three treatment groups only. First, 396 treatment households were dropped because they never received an HWR email; this occurred because of logistical issues (families moving before the experiment began, issues with email addresses, etc). Next, I removed 1,865 treatment households that did not receive their first HWR in the window of time I focused on in this study (December 9, 2014 to January 30, 2015). That is, some individuals in the sample received their first HWR much earlier in 2014 or much later in 2015. Finally, 1,467 treatment households were removed because of misalignments between the water use measured in their meter reads and the water use feedback provided in their subsequent HWR. In other words, for these

households, the number of gallons per day from their last water meter read was not reflected in their next HWR at least once, because of timing issues related to exactly when their reads were done and when the emails went out. In all of these instances, water use rates from an earlier read were used in HWRs, but because of this asymmetry in reads and HWR content, I exclude these anomalous households from analysis.

After these trims of the dataset, 41,365 households remained from the initial 45,866, and these households comprise the sample that I empirically investigate here. Importantly, the second set of data restriction processes described above involved removing only households from the treatment groups because of anomalies related to the content or delivery of HWRs specifically. Though there is no strong ex-ante reason to believe that this would cause major imbalances between Control households and those in the treatment conditions, these adjustments may influence the analysis involving comparisons between the Control and the treatment conditions. However, it would not affect comparisons across treatments, which is the primary focus of this paper.

4 Empirical Methods and Results

In this section, I outline the empirical approaches used to analyze the experimental data, and present the results. Because I address various questions in this paper, I present these methods and findings by research question in the subsections below.

4.1 What is the unique effect of injunctive norms messaging on water use behavior?

The central question in this paper is whether and how injunctive norms affect prosocial behavior, in this case water consumption. To estimate this, we must exploit the differences between the treatments, rather than their differences relative to a control group that did not receive emails at all. In particular, recall that the No Drop condition provides only social information, the Drop condition adds an additional visual cue to the No Drop condition, and the Injunctive Drop condition provides an injunctive norm message within the visual cue. Therefore by comparing these conditions we can isolate both the causal impact of including a visual cue on water use (by comparing the No

Drop and Drop treatments) and, most importantly, the causal impact of injunctive norms messaging on water use (by comparing the Drop and Injunctive Drop treatments).

The general specification I use to estimate these causal impacts is as follows. First, I restrict my analysis to only post-treatment observations for each household, and regress water use for all post-treatment observations (measured in gallons per day) on dummy variables for assignment to each experimental condition (omitting one condition). These dummy variables for a regression involving n+1 conditions are captured by the $\sum_1^n[\beta_n(T_n)_i]$ term below. I then add controls for the following: 1) the water use for a given household in the matching calendar month pre-treatment (drawn from December 2013 to November 2014 water use data and denoted $GPDpre_{im}$ below); 2) various demographic characteristics at the household level (number of occupants, home size, number of floors, lot size, irrigable area, number of bedrooms, and number of bathrooms), denoted λ_i below; 3) wave fixed effects (denoted γ_w below); and 4) month-in-sample fixed effects (effectively a set of dummy variables for the month and year combination for a given meter read observation, to control for general water use trends over time in the sample), denoted δ_m below. This empirical approach generally follows that used by List, Metcalfe, Price, and Rundhammer (2017) and Allcott and Rogers (2014), who perform similar analyses on similar data. The full specification with all controls is given below:

$$GPD_{imw} = \beta_0 + \sum_1^n[\beta_n(T_n)_i] + GPDpre_{im} + \lambda_i + \gamma_w + \delta_m + \epsilon$$

Returning to the central research question, to isolate the causal impact of the injunctive norms messaging I run the specification above excluding data from the control group, thereby comparing the treatments directly to one another. The results are reported in columns 1-3 in Table 3.[5] In the preferred specification in column 3, the effect of the Injunctive Drop condition is significantly larger than that of the No Drop

[5]Note that all analyses in Table 3 were also conducted without the data restrictions outlined in section 3.5. These results are qualitatively similar to those in Table 3, and are visible in Online Appendix Table A.3.

condition, which is omitted from the regression (an effect size of 2.5 gallons per day, which is significant at a 1% level). Furthermore, the p-values from Wald tests comparing the regression coefficients for the Injunctive Drop condition and the Drop condition (thereby directly testing the effect of the injunctive norm messaging) are consistently significant at the 5% or 10% level across specifications. This is the critical empirical result, as it implies that the reason the Injunctive Drop condition performs best has to do with the social judgment in the visual, rather than the visual cue of the droplet alone (indeed, the coefficient on "Drop" is not significant here).

These results suggest that the inclusion of injunctive messaging is a crucial component of the efficacy of this social information campaign. Indeed, unlike previous work in this domain, this study provides causal evidence on the role of injunctive messaging specifically, as the only difference between the Injunctive Drop and Drop conditions was the inclusion of the smiley/neutral/frowny face in the drop visual included in both conditions.

The analysis in Table 3 focuses on the average treatment effects of each condition, but it is also plausible that the injunctive norm messaging might influence behavior differently for those using less versus more water (as this would shape whether they see regular social approval or disapproval in the injunctive messaging campaign). I explore this question by computing disaggregated average treatment effects of the injunctive norm messaging (by comparing water use in the Injunctive Drop and Drop treatment groups), by decile of pre-experiment water use.[6] The pre-experiment water use measure I use is the mean water use for a given household from December 2013 to November 2014, and deciles were constructed for this variable within each of five possible categorical "buckets" for the number of occupants in a household (1, 2, 3, 4, and 5 or more persons). This allows me to control somewhat for differential water use needs by family size. The regression specification I use is precisely the same as that used in column 3 in Table 3 (with all controls), however I restrict analysis to only the Drop and

[6]The use of deciles is admittedly somewhat arbitrary, but I argue it is reasonable given the sample size and the use of deciles in similar analysis in related past work (including Allcott (2011)).

Injunctive Drop treatment groups here, and use the Drop treatment group as the omitted group in the regression. This means that the coefficients on the Injunctive Drop dummy variable in the regressions correspond to the average treatment effect estimates of the injunctive norm messaging on water use for each of the ten deciles of pre-experiment water use.

The results of this analysis are displayed in Figure 3, which shows the average treatment effect estimates drawn from the ten regression coefficients in this analysis, along with the corresponding standard errors. Because the use of deciles "splits" the data into many subgroups, it is important not to over-infer based on these estimates; instead, they should be viewed as providing clues regarding how the intensity of prior water use correlates with response to injunctive norms. The results suggest that those who use very little water ex-ante seem to respond to injunctive norms messaging by slightly lowering their use (see deciles 1-3), as do those with relatively high (but not extremely high) water usage (deciles 8-9). The results are mixed for those in the middle of the distribution (deciles 4-7). Interestingly, those who use the most water ex-ante (decile 10) seem to display an adverse reaction to injunctive norms messaging, as captured by the significant increase in water use for decile 10 households in the Injunctive Drop treatment relative to those in the Drop treatment. Note that whereas Schultz et al. (2007) only disaggregate their analysis using a median split of the data, the larger dataset here allows for more detailed analysis–and an interesting pattern emerges. As with Schultz et al. (2007), I find that the injunctive norms messaging seems to be having a positive impact on low water users, but I also find suggestive evidence that the story is not all rosy at the upper end of the water use distribution. In particular, injunctive norms may be backfiring here for the highest water users in the population, causing them to use more water than they would have if they had simply seen a drop visual without the added social judgment.

Overall, the results presented here provide evidence that, on average, the inclusion of injunctive norms in resource conservation messaging campaigns using social norms has a beneficial impact on the efficacy of these campaigns. However, the disaggregated

analysis using pre-experiment water use suggests some heterogeneity in response to injunctive norms messaging. However, given the fine slicing of the data here, more research is needed before definitive conclusions can be reached regarding the effects of injunctive norms on different subgroups.

4.2 How much of the overall treatment effect of the mailers comes from injunctive norms?

To benchmark the average treatment effect sizes for injunctive norms messaging observed in section 4.1, I also conducted analysis including the control group that did not receive HWRs. Doing so allows me to estimate the magnitude of the causal effect of the injunctive norms messaging in isolation, as a percentage of the "total" effect of all aspects of the HWR with the most content (namely the Injunctive Drop HWR email, which had descriptive norms content, tips about water conservation, injunctive norms messaging, a visual cue, etc). Note that this analysis also serves to replicate previous findings on the average treatment effects of social information campaigns for conservation (Allcott, 2011; Brent et al., 2015; Ferraro & Price, 2013). This is important for two reasons. First, past studies have explored paper mailers and not email messaging campaigns, so the average treatment effects estimated in this paper add to a relatively small literature regarding the effectiveness of digital campaigns for behavior change. Second, a great deal of recent research has emphasized the critical importance of independent replications in their own right (Aarts et al., 2015; Camerer et al., 2016; Maniadis et al., 2014).

Columns 4-6 in Table 3 display the average treatment effect estimates for each treatment group, relative to the control group, using the same three regression specifications as used in columns 1-3. In all specifications, there is a statistically significant treatment effect on the amount of water used for all treatment conditions, with all conditions reducing water use relative to the control. In the preferred specification in column 6, with a full set of controls and fixed effects, the effect estimates range from 5.7-8.3 gallons per day (roughly 3-5% of mean use). These effect sizes are

consistent with those in previous work (Allcott, 2011; Brent et al., 2015; Ferraro & Price, 2013), and show that social information provision was effective in encouraging conservation behavior in this context. Note that the precision of these estimates should be treated with caution, given the demographic imbalances in the control group specifically (see section 3.4) and the data restriction processes (see section 3.5).

Nevertheless, these estimates provide a benchmark for evaluating the relative importance of injunctive norms messaging. Specifically, if we take the point estimates in Table 3 at face value, we obtain a rough estimate of the proportion of the efficacy of the Injunctive Drop HWR emails that can be attributed to each aspect of the messaging that the Injunctive Drop HWR emails contained (the general social information and other HWR content, the visual cue, and the injunctive norm message). Specifically, assuming an average treatment effect of 8.256 gallons per day (column 6), the results suggest that the injunctive norm messaging is responsible for roughly 21% of the overall treatment effect ((8.256-6.525)/8.256), the visual cue for roughly 10% ((6.525-5.735)/8.256), and the rest of the social information and HWR content for the remaining 69% (5.735/8.256).

4.3 Is there evidence that injunctive norms messaging discourages subjects from attending to future social norms messaging?

One potential side effect of injunctive norms messaging is that it may influence subjects' likelihood of attending to future social norms messaging. As discussed in section 2.2, the predicted effect here could go either way. For example, being told that your behavior is met by social disapproval may discourage you from learning about others' attitudes about your behavior in the future (an "ostrich effect"), or it may trigger effortful self-regulation or a "meerkat" effect (Gherzi et al., 2014; Jacobson et al., 2011; Karlsson et al., 2009). The same ambiguity exists for individuals who see that their behavior is socially approved–that is, these individuals may not see a need to pay attention to future injunctive messaging because their behavior has already been "approved of" by society, or it may heighten their desire for further social approval.

To investigate this question, I use multiple empirical approaches. First, I focus on the treatment groups only, and restrict my attention to the emailed HWR that immediately followed the first emailed HWR that a given subject actually opened. Note that the first email a given subject opened is not necessarily the first email a given subject received–it is possible that they did not open the first email they received from WaterSmart. But because I can see in the data whether or not a given email was opened, I know which email was the first one that was actually opened, meaning I know when a person first saw the injunctive norm messaging (or not, for those conditions without an injunctive norms message).[7] Then, by using whether or not subjects opened the email immediately following their first visual engagement with an HWR email as an outcome variable (the dummy variable $EmailOpened_{imw}$ below), I can test for the effect of the email's content on the desire to engage with further social norms information. The regression specification I run to conduct this analysis is as follows:

$$EmailOpened_{imw} = \beta_0 + \beta_1(T_{Drop})_i + \beta_2(T_{InjunctiveDrop})_i + \lambda_i + \gamma_w + \delta_m + \epsilon$$

Note that I add controls for: 1) various demographic characteristics at the household level (number of occupants, home size, number of floors, lot size, irrigable area, number of bedrooms, and number of bathrooms), denoted λ_i; 2) wave fixed effects (denoted γ_w); and 3) month-in-sample fixed effects, denoted δ_m. The results of this analysis are visible in Table 4.

These results provide compelling evidence that both the visual cue and, more importantly, the injunctive norms messaging have, on average, little to no effect on the likelihood that a subject opens the HWR email that follows the first HWR email they open. In particular, as shown in column 1 in Table 4, 65.7% of subjects in the No Drop condition open the HWR email immediately following the first one they open, and neither the Drop condition nor the Injunctive Drop condition significantly reduce that figure. It is worth noting that the negative coefficient on the Injunctive Drop condition is, in the regression without controls, significant at the 10% level, but the magnitude of

[7] Critically, there was no way for subjects to see the content of an email (and thus the treatment) without opening it, so I would not expect any selection into email opening by condition.

the coefficient decreases with controls and is no longer distinguishable from zero. When restricting this analysis to the Drop and Injunctive Drop conditions only (in columns 3 and 4), to isolate the average effect of injunctive norms messaging in particular on the likelihood of opening the next email, I again find no evidence of injunctive norms significantly changing individual willingness to engage with future content. These results suggest that injunctive norms do not result in significantly more or less attention being paid to future messaging, at least as measured by subjects' willingness to engage with the next piece of messaging they are sent.

The analyses in Table 4 focused on the average effects of injunctive norms across the sample, however. To explore the impact of the specific injunctive messaging types used here (smiley vs. neutral vs. frowny face), I exploit the fact that I not only know when an individual actually opened an email for the first time, but I also know what content they observed when they opened that email. In particular, I am able to restrict attention to only those subjects who, when they first opened an email, saw a frowny/neutral/smiley face in the Injunctive Drop treatment and those in the Drop treatment who "would have seen" a frowny/neutral/smiley face (had they been assigned to the Injunctive Drop condition). By directly comparing those in the Injunctive Drop and Drop conditions who saw/would have seen a smiley face when they first opened an email, I am able to obtain a causal estimate of the effect of a positive injunctive norm message (conveying social approval) on future engagement with HWR emails (again using their likelihood of opening their next email). The regression specification I use here is identical to that used in Table 4, column 4. I then repeat this analysis for the neutral and frowny faces, to estimate the treatment effects of each injunctive message type.

The results of this analysis are visible in Figure 4, which plots the average treatment effects of each injunctive message type on the likelihood that subjects opened the next HWR email. This figure provides limited evidence that the specific injunctive norm messages have an impact on the likelihood of subjects to open ensuing mailers. If anything, there is some evidence that the inclusion of the smiley face visual for low

water users encourages a slight decrease in the probability of opening the next email, but the coefficient in the regression is not quite significant at the 10% level ($p = 0.103$). Meanwhile, the coefficient for the frowny face injunctive message is positive but indistinguishable from zero ($p = 0.536$), which is strong evidence against the existence of any "ostrich effect" from including an injunctive message of social disapproval for high water users. That is, adding a frowny face to the HWR for high water users did not encourage these subjects to shy away from opening the next HWR email they received.

I next repeat similar analysis to that done above (in Table 4 and Figure 4), but instead of using the opening of the HWR email immediately following the first opened HWR email as the outcome measure, I instead use the percentage of emails a given subject opened (from amongst those received) after they first opened an HWR email. This serves as another test of how the injunctive norm messaging in the first opened HWR email influences ensuing behavior.

I begin with a regression analysis that mirrors that used for the analysis in Table 4, but with this different outcome variable. The results, visible in Table 5, suggest that there is no relationship between the inclusion of injunctive norms messaging and the percentage of HWR emails opened by subjects, with all coefficients on treatment variables being both very small in magnitude and not statistically significant. This is consistent with the findings in Table 4, and supports the conclusion that the injunctive norms messaging is not changing the individual-level likelihood of engaging with HWR emails.

Finally, I replicate the process used to generate Figure 4, but again using the new outcome variable, to determine the effect of the specific injunctive message seen in the first HWR email opened on the percentage of future social norms email messages opened. The average treatment effects are depicted in Figure 5. These results also support the conclusion that the specific injunctive norms messages seen on the first opened HWR (a smiley, neutral, or frowny face) did not have an impact on the percentage of ensuing HWR emails that a given subject opened. Again, this is consistent with a conclusion that while the injunctive norms messaging may have influenced water

use behavior, observing the social judgment in the injunctive norms message did not have an effect on the willingness of subjects to open the HWR emails that followed.

5 Discussion and Conclusions

In this paper, I provide the most statistically-robust evidence to date regarding the impact of injunctive norms messaging on decision making in the context of social norms messaging campaigns (which are increasingly used by organizations to change human behavior). Specifically, the large-scale field experiment that I report on used treatment conditions that allow me to isolate the effect of injunctive norms messaging specifically. There are two primary findings from my analysis. First, I find compelling evidence that the inclusion of injunctive norms messages that convey a social judgment has, on average, a beneficial impact, leading to reductions in water use that are larger than those from an email campaign omitting such messages. Second, I find no evidence that the inclusion of injunctive norms messaging changes the likelihood that subjects engage with future social norms messaging emails (as measured by the opening of these emails). Furthermore, I find very little evidence that the specific types of injunctive norms messages used (smiley, neutral, or frowny faces) influence how willing subjects are to open future messages. If anything, there is weak evidence that positive social judgments encourage high achievers to ignore future messaging (see Figures 4 and 5), but these findings were not statistically significant.

There are important theoretical and conceptual implications of this work. First, these findings suggest that a non-trivial part of the effectiveness of social norms messaging campaigns is the social judgment that is conveyed by these efforts. Researchers studying social norms have long emphasized the importance of collective perceptions of what others "ought to do" as playing a key role in establishing norms (Bicchieri & Chavez, 2010). The results in this paper support that interpretation, as they suggest that the impact of norms messaging is reinforced when the messaging comes with a clear social judgment about behavior.

Second, these results also offer important evidence regarding how messaging

containing social judgments influence people's willingness to engage with information. In light of the ongoing "information revolution" for both individuals and organizations, researchers in the behavioral sciences have become increasingly interested in the demand for information at both the individual and organizational levels (Ganguly & Tasoff, 2016; Gherzi et al., 2014; Golman et al., 2017; Golman & Loewenstein, 2016; Karlsson et al., 2009). The results presented here suggest that when people receive information about their behavior along with social judgments, their demand for further information is not significantly affected by the inclusion of a social judgment. In some sense, this is not surprising–as past research has suggested, information (both useful and useless) can contribute to individual utility in both positive and negative ways (Eliaz & Schotter, 2007; Golman et al., 2017). However, there is little real world evidence on the demand for information from social norms campaigns, which are increasingly used by organizations like WaterSmart to shape behavior on a large scale. By providing causal evidence on this question, this paper makes a contribution to this burgeoning literature.

Finally, the practical takeaway for organizations and individuals looking to shape behavior is simple. Based on the findings here, using injunctive norms messaging to change behavior appears to be an effective tool for behavior change. And in particular, while further research is required to better understand whether or not this is true in other domains, the results in this paper suggest that injunctive norms should be used with confidence by organizations seeking to influence resource conservation behavior specifically.

6 References

Aarts, A. A., Anderson, J. E., Anderson, C. J., Attridge, P. R., Attwood, A., Axt, J., ... Zuni, K. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251). doi: 10.1126/science.aac4716

Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*. doi: 10.1016/j.jpubeco.2011.03.003

Allcott, H., & Rogers, T. (2014). *The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation.* doi: 10.1257/aer.104.10.3003

Beshears, J., Choi, J. J., Laibson, D., Madrian, B. C., & Milkman, K. L. (2015). The Effect of Providing Peer Information on Retirement Savings Decisions. *Journal of Finance*. doi: 10.1111/jofi.12258

Bhanot, S. P. (2017). Rank and response: A field experiment on peer information and water use behavior. *Journal of Economic Psychology*. doi: 10.1016/j.joep.2017.06.011

Bicchieri, C., & Chavez, A. (2010). Behaving as expected: Public information and fairness norms. *Journal of Behavioral Decision Making*, *23*(2), 161–178. doi: 10.1002/bdm.648

Brent, D. A., Cook, J. H., & Olsen, S. (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*. doi: 10.1086/683427

Cai, H., Chen, Y., & Fang, H. (2009). Observational learning: Evidence from a randomized natural field experiment. *American Economic Review*. doi: 10.1257/aer.99.3.864

Camerer, C. F., Dreber, A., Forsell, E., Ho, T. H., Huber, J., Johannesson, M., ... Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, *351*(6280), 1433–1436. doi: 10.1126/science.aaf0918

Cialdini, R. B., Demaine, L. J., Sagarin, B. J., Barrett, D. W., Rhoads, K., & Winter,

P. L. (2006). Managing social norms for persuasive impact. *Social Influence*. doi: 10.1080/15534510500181459

Eil, D., & Rao, J. M. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, *3*(2), 114–138. doi: 10.1257/mic.3.2.114

Eliaz, K., & Schotter, A. (2007). Experimental testing of intrinsic preferences for noninstrumental information. In *American economic review* (Vol. 97, pp. 166–169). doi: 10.1257/aer.97.2.166

Ferraro, P. J., & Price, M. K. (2013). Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment. *Review of Economics and Statistics*. doi: 10.1162/REST{_}a{_}00344

Frey, B. S., & Meier, S. (2004). Pro-social behavior in a natural setting. *Journal of Economic Behavior and Organization*. doi: 10.1016/j.jebo.2003.10.001

Funk, P. (2010). Social Incentives and Voter Turnout: Evidence from the Swiss Mail Ballot System. *Journal of the European Economic Association*. doi: 10.1111/j.1542-4774.2010.tb00548.x

Ganguly, A., & Tasoff, J. (2016). Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future. *Management Science*(September 2017), mnsc.2016.2550. doi: 10.1287/mnsc.2016.2550

Gerber, A. S., Green, D. P., & Larimer, C. W. (2008). Social pressure and voter turnout: evidence from a large-scale field experiment. *American Political Science Review*, *102*(1), 33–48.

Gerber, A. S., & Rogers, T. (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *Journal of Politics*. doi: 10.1017/S0022381608090117

Gherzi, S., Egan, D., Stewart, N., Haisley, E., & Ayton, P. (2014). The meerkat effect: Personality and market returns affect investors' portfolio monitoring behaviour. *Journal of Economic Behavior and Organization*. doi: 10.1016/j.jebo.2014.07.013

Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information Avoidance. *Journal*

of Economic Literature, *55*(1), 96–135. Retrieved from http://pubs.aeaweb.org/doi/10.1257/jel.20151245 doi: 10.1257/jel.20151245

Golman, R., & Loewenstein, G. (2016). Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information. *Decision*. Retrieved from http://dx.doi.org/10.1037/dec0000068%5Cnhttp://doi.apa.org/getdoi.cfm?doi=10.1037/dec0000068 doi: 10.1037/dec0000068

Harrison, G. W., & List, J. A. (2004). Field experiments. *Journal of Economic Literature*, *42*(4), 1009–1055.

Jacobson, R. P., Mortensen, C. R., & Cialdini, R. B. (2011). Bodies Obliged and Unbound: Differentiated Response Tendencies for Injunctive and Descriptive Social Norms. *Journal of Personality and Social Psychology*. doi: 10.1037/a0021470

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*. doi: 10.1007/s11166-009-9060-6

Krupka, E. L., & Croson, R. T. (2016). The differential impact of social norms cues on charitable contributions. *Journal of Economic Behavior and Organization*. doi: 10.1016/j.jebo.2016.05.005

List, J., Metcalfe, R., Price, M., & Rundhammer, F. (2017). Harnessing Policy Complementarities to Conserve Energy: Evidence from a Natural Field Experiment. *NBER Working Paper No. 23355*. Retrieved from http://www.nber.org/papers/w23355.pdf%5Cnhttp://www.nber.org/data-appendix/w23355/HarnessingPolicyComplementarities-Appendix.pdf doi: 10.3386/w23355

Maniadis, Z., Tufano, F., & List, J. a. (2014). One Swallow Doesn' t Make a Summer: New Evidence on Anchoring Effects. *American Economic Review*, *104*(1), 277–290. Retrieved from http://pubs.aeaweb.org/doi/abs/10.1257/aer.104.1.277 doi:

10.1257/aer.104.1.277

Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). Destructive, and Reconstructive Power of Social Norms. *Psychological Science*. doi: 10.1111/j.1467-9280.2007.01917.x

Sicherman, N., Loewenstein, G., Seppi, D. J., & Utkus, S. P. (2016). Financial Attention. *Review of Financial Studies*, *29*(4), 863–897. doi: 10.1093/rfs/hhv073

Table 1

Demographics by Experimental Condition

	(1) Control	(2) No Drop	(3) Drop	(4) Injunctive Drop	(5) All
Number of occupants	3.158	3.276	3.270	3.246	3.259
	(1.100)	(1.230)	(1.209)	(1.212)	(1.208)
Home size (sqft.)	1837.1	1822.4	1815.0	1809.9	1819.7
	(794.6)	(789.5)	(779.3)	(816.4)	(788.7)
Number of floors	1.416	1.405	1.408	1.397	1.406
	(0.511)	(0.512)	(0.516)	(0.508)	(0.513)
Lot size (sqft.)	8127.9	8098.8	8016.5	7972.5	8056.0
	(7316.4)	(7289.9)	(7187.9)	(7109.5)	(7233.8)
Irrigable area	4430.4	4345.9	4297.1	4308.5	4331.1
	(4407.2)	(4280.5)	(4224.8)	(4260.2)	(4269.2)
Number of bedrooms	3.265	3.248	3.237	3.249	3.245
	(0.973)	(1.120)	(1.016)	(1.021)	(1.055)
Number of bathrooms	2.221	2.211	2.207	2.206	2.210
	(0.980)	(0.968)	(0.989)	(1.084)	(0.990)
Pretreatment Mean Water Use	260.4	257.9	255.5	256.4	257.0
	(170.2)	(166.0)	(164.3)	(162.6)	(165.4)
N	4586	18347	18346	4587	45866

Notes. Standard deviations are in parentheses. The reported values for the variables lot size and irrigable area reflect the data following the removal of the largest 0.5% of observations for each of those variables from the sample, as these variables have extremely large outliers that differentially skew the means and standard deviations across conditions. Pretreatment mean water use is measured in gallons per day, and is computed using mean water use based on reads from December 2013 to November 2014.

Table 2

Randomization Checks

	Treatments vs. Control			Treatments Only	
	(1) No Drop	(2) Drop	(3) Injunctive Drop	(4) Drop	(5) Injunctive Drop
Number of occupants	0.0186***	0.0187***	0.0284***	-0.000891	-0.00286
	(0.00253)	(0.00258)	(0.00537)	(0.00245)	(0.00247)
Home size (sqft.)	-0.00000228	-0.00000370	-0.0000122	-0.00000550	-0.00000709
	(0.00000625)	(0.00000636)	(0.0000126)	(0.00000615)	(0.00000622)
Number of floors	-0.0163**	-0.0135**	-0.0345***	0.00532	-0.00347
	(0.00634)	(0.00640)	(0.0126)	(0.00625)	(0.00634)
Lot size (sqft.)	5.49e-08	-0.000000926	-0.000000480	-0.000000337	-0.000000108
	(0.000000259)	(0.000000873)	(0.00000141)	(0.000000322)	(0.000000262)
Irrigable area	-7.84e-08	0.00000132	0.000000687	0.000000482	0.000000154
	(0.000000370)	(0.00000125)	(0.00000201)	(0.000000459)	(0.000000375)
Number of bedrooms	-0.00643**	-0.00982***	-0.00951	-0.00189	0.00275
	(0.00327)	(0.00370)	(0.00783)	(0.00333)	(0.00328)
Number of bathrooms	0.00148	0.00375	0.00892	0.00347	0.00443
	(0.00491)	(0.00458)	(0.00892)	(0.00464)	(0.00474)
Pretreatment Mean Water Use	-0.0000353*	-0.0000404**	-0.0000640*	-0.0000130	-0.00000606
	(0.0000183)	(0.0000185)	(0.0000367)	(0.0000182)	(0.0000185)
Constant	0.794***	0.802***	0.512***	0.508***	0.211***
	(0.0110)	(0.0113)	(0.0224)	(0.0110)	(0.0111)
Observations	20944	20960	8400	33552	20992
R^2	0.003	0.003	0.004	0.000	0.000
F-test p-value	0.00	0.00	0.00	0.66	0.72

Notes. Column 1-3 show pairwise balance checks for each treatment condition against the control, and 4-5 show pairwise balance checks for the Drop and Injunctive Drop conditions against the No Drop condition. Pretreatment mean water use is measured in gallons per day, and is computed using mean water use based on reads from December 2013 to November 2014. F-tests were used to test for joint significance and balance. Standard errors are in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Table 3

Average Treatment Effects on Water Use

	DV: Gallons per Day (GPD)					
	Excluding Control			All Conditions		
	(1)	(2)	(3)	(4)	(5)	(6)
Injunctive Drop	-3.519***	-2.580***	-2.511***	-9.694***	-8.328***	-8.256***
	(1.043)	(0.923)	(0.914)	(1.336)	(1.159)	(1.146)
Drop	-1.438**	-0.777	-0.798	-7.614***	-6.510***	-6.525***
	(0.675)	(0.588)	(0.582)	(1.074)	(0.916)	(0.905)
No Drop				-6.176***	-5.737***	-5.735***
				(1.073)	(0.910)	(0.898)
Pre-Exp. Cal. Month GPD		0.409***	0.413***		0.412***	0.414***
		(0.00527)	(0.00549)		(0.00517)	(0.00539)
Observations	228461	193749	193749	256031	216935	216935
R^2	0.000	0.356	0.369	0.000	0.360	0.374
Mean GPD for Omitted Group	198.84	198.84	198.84	205.01	205.01	205.01
Unique Households in Sample	36787	33669	33669	41240	37724	37724
Demographic Controls	No	Yes	Yes	No	Yes	Yes
Month-In-Sample Fixed Effects	No	No	Yes	No	No	Yes
Wave Fixed Effects	No	No	Yes	No	No	Yes
Wald Statistic p-value	0.0462	0.0519	0.0619	0.0462	0.0500	0.0592

Notes. This table shows the main results from this experiment, in the form of average treatment effects on water use (in gallons per day, or GPD), using linear regression models. Specifications 1-3 show the average treatment effects of the Injunctive Drop and Drop conditions relative to the No Drop condition, and 4-6 show the average treatment effects of the treatment conditions (Injunctive Drop, Drop, and No Drop) relative to the Control. Regressions with and without controls are included–the controls are: 1) demographic controls at the household level, which consists of home size, lot size, irrigable area, and the number of occupants, floors, bedrooms and bathrooms; 2) controls for pre-treatment mean water usage for a given observation (from water use for the household in the matching calendar month in the pretreatment period); 3) month-in-sample fixed effects (uniquely identifying each month-year combination); and 4) wave fixed effects. A p-value for a Wald test is also reported, which relates to the null hypothesis that the coefficients for the Injunctive Drop and Drop conditions are equal. Robust standard errors are in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Table 4

ATEs on the Likelihood of Opening the Next HWR Email

	DV: Probability of Opening Next HWR Email			
	All Treatments		Excluding No Drop	
	(1)	(2)	(3)	(4)
Injunctive Drop	-0.0177*	-0.0122	-0.0135	-0.00915
	(0.00929)	(0.00933)	(0.00929)	(0.00933)
Drop	-0.00422	-0.00315		
	(0.00586)	(0.00593)		
Observations	29673	27518	16538	15368
R^2	0.000	0.053	0.000	0.054
Mean Open Rate for Omitted Group	0.657	0.657	0.653	0.653
Demographic Controls	No	Yes	No	Yes
Month-In-Sample Fixed Effects	No	Yes	No	Yes
Wave Fixed Effects	No	Yes	No	Yes

Notes. This table shows the average treatment effects on the likelihood of opening the next social norms message following the first HWR email open, using linear regression models. Specifications 1-2 show the average treatment effects of the Injunctive Drop and Drop conditions relative to the No Drop condition, and 3-4 show the average treatment effects of the Injunctive Drop condition relative to the Drop condition. Regressions with and without controls are shown–the controls are: 1) demographic controls at the household level, which consists of home size, lot size, irrigable area, and the number of occupants, floors, bedrooms and bathrooms; 2) month-in-sample fixed effects; and 3) wave fixed effects. Robust standard errors are in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Table 5

ATEs on the Percentage of HWR Emails Opened (After the First HWR Email Open)

	DV: Percentage of Emails Opened			
	All Treatments		Excluding No Drop	
	(1)	(2)	(3)	(4)
Injunctive Drop	-0.00439	-0.00579	-0.00456	-0.00681
	(0.00536)	(0.00549)	(0.00536)	(0.00549)
Drop	0.000168	0.00104		
	(0.00343)	(0.00353)		
Observations	31279	28959	17430	16166
R^2	0.000	0.012	0.000	0.013
Mean Open Rate for Omitted Group	0.490	0.490	0.490	0.490
Demographic Controls	No	Yes	No	Yes
Wave Fixed Effects	No	Yes	No	Yes

Notes. This table shows the average treatment effects on the percentage of HWR emails opened following the first HWR email open, using linear regression models. Specifications 1-2 show the average treatment effects of the Injunctive Drop and Drop conditions relative to the No Drop condition, and 3-4 show the average treatment effects of the Injunctive Drop condition relative to the Drop condition. Regressions with and without controls are shown–the controls are: 1) demographic controls at the household level, which consists of home size, lot size, irrigable area, and the number of occupants, floors, bedrooms and bathrooms; and 2) wave fixed effects. Robust standard errors are in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Figure 1. No Drop, Drop and Injunctive Drop Visuals Provided by WaterSmart

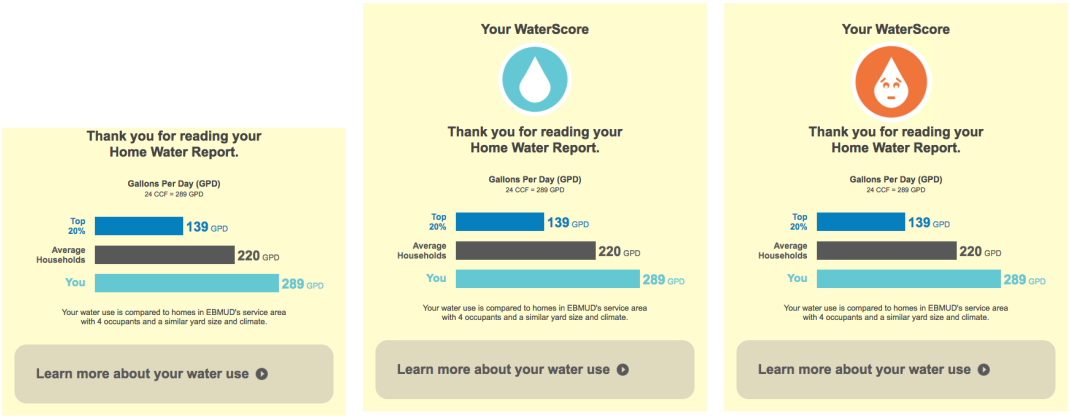

(a) Treatment 1. No Drop (b) Treatment 2. Drop (c) Treatment 3. Injunctive Drop

Figure 2. Other Types of Injunctive Faces Provided by WaterSmart

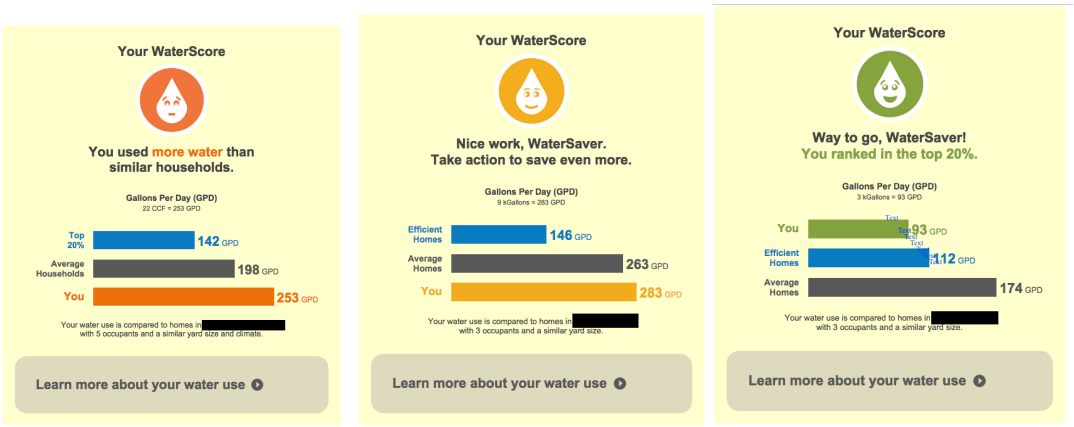

(a) Frowny Visual (b) Neutral Visual (c) Smiley Visual

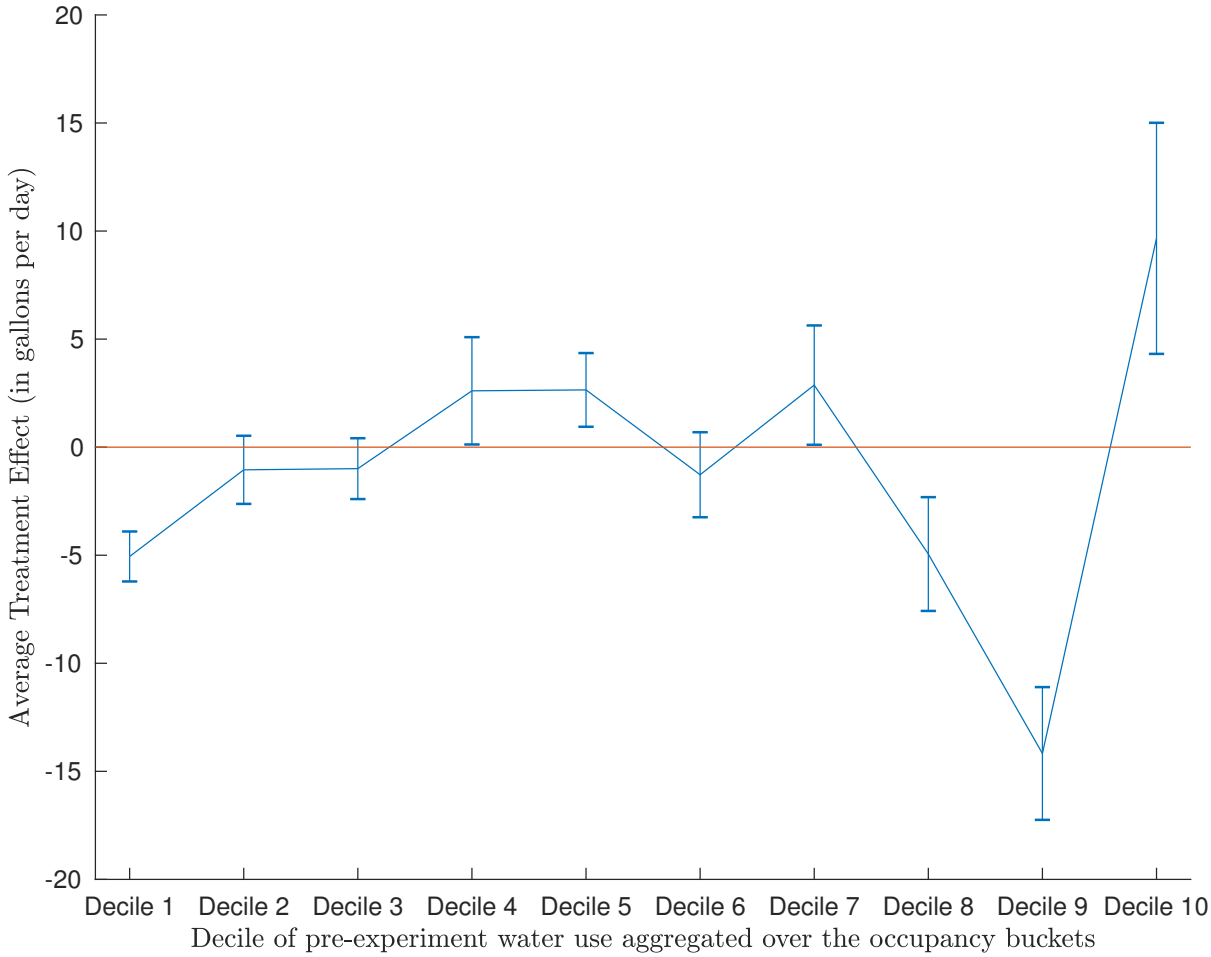

Figure 3. Disaggregated ATEs by Decile of Mean Pre-Treatment Water Use. The error bars mark ±1 robust standard error. The deciles of pre-treatment mean water usage were constructed within each of five categorical buckets for the number of occupants in a given home (1, 2, 3, 4, and 5 or more persons), then aggregated. Decile 1 indicates the bottom decile (i.e., the lowest pre-experiment household water usage decile in each occupancy bucket), and Decile 10 indicates the top decile (i.e., the highest pre-experiment household water usage decile in each occupancy bucket). The coefficients for deciles 1 and 9 are statistically significant at the 1% level, while the coefficients for deciles 8 and 10 are statistically significant at the 10% level.

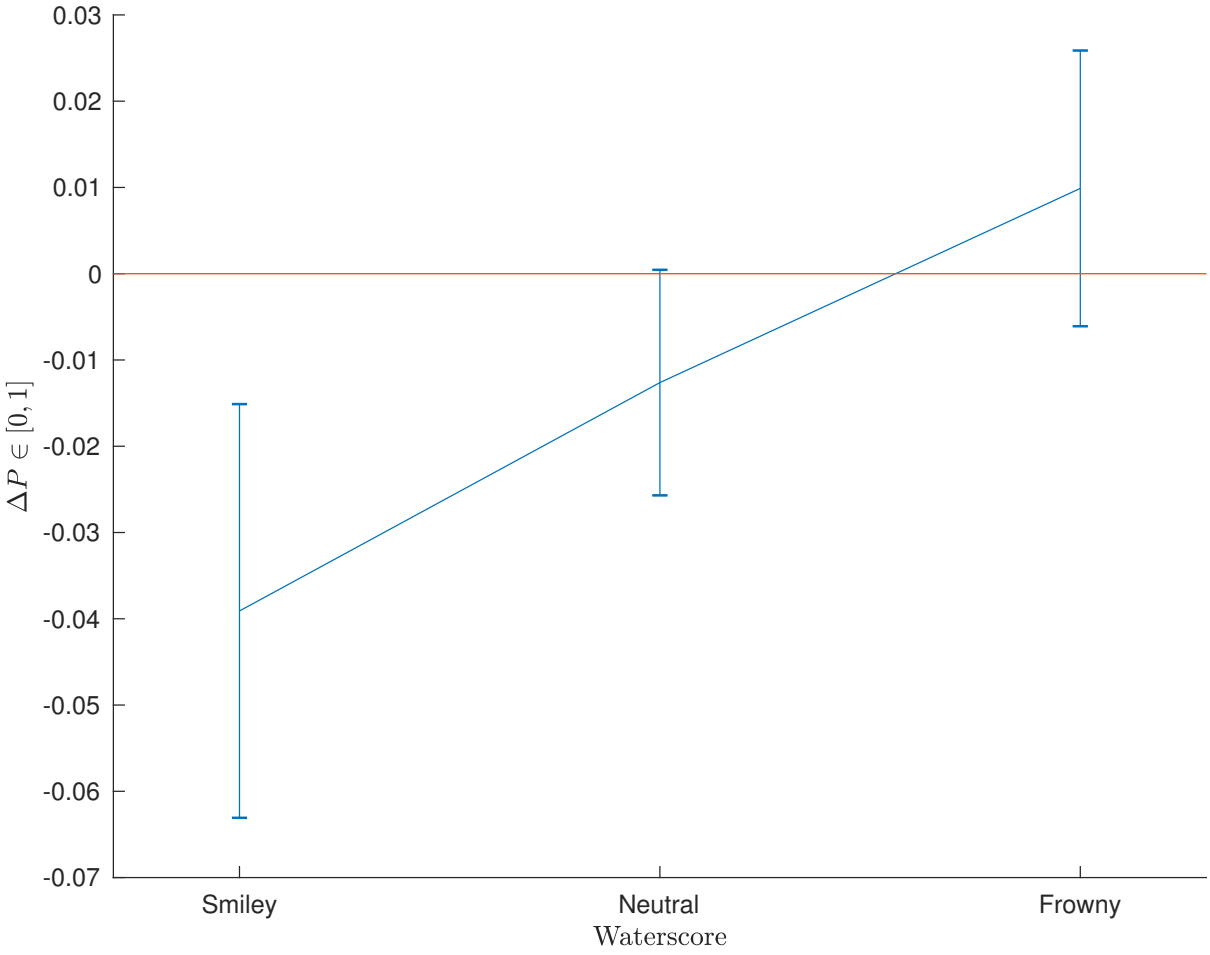

Figure 4. ATEs on the Likelihood of Opening the Next HWR Email, by Injunctive Norm Type. This figure shows the average treatment effects on the likelihood of opening the next HWR email following the first HWR email open, disaggregated by the type of injunctive norm seen on the first HWR email open, using linear regression models. The average treatment effects of the Injunctive Drop condition relative to the Drop condition are estimated using specifications with a full set of controls (i.e., that used in Table 4, column 4). The controls are: 1) demographic controls at the household level, which consists of home size, lot size, irrigable area, and the number of occupants, floors, bedrooms and bathrooms; 2) month-in-sample fixed effects; and 3) wave fixed effects. The error bars mark ±1 robust standard error, and the p-values from hypotheses tests comparing each coefficient to zero are 0.103, 0.334, and 0.536 for the smiley, neutral, and frowny messages, respectively.

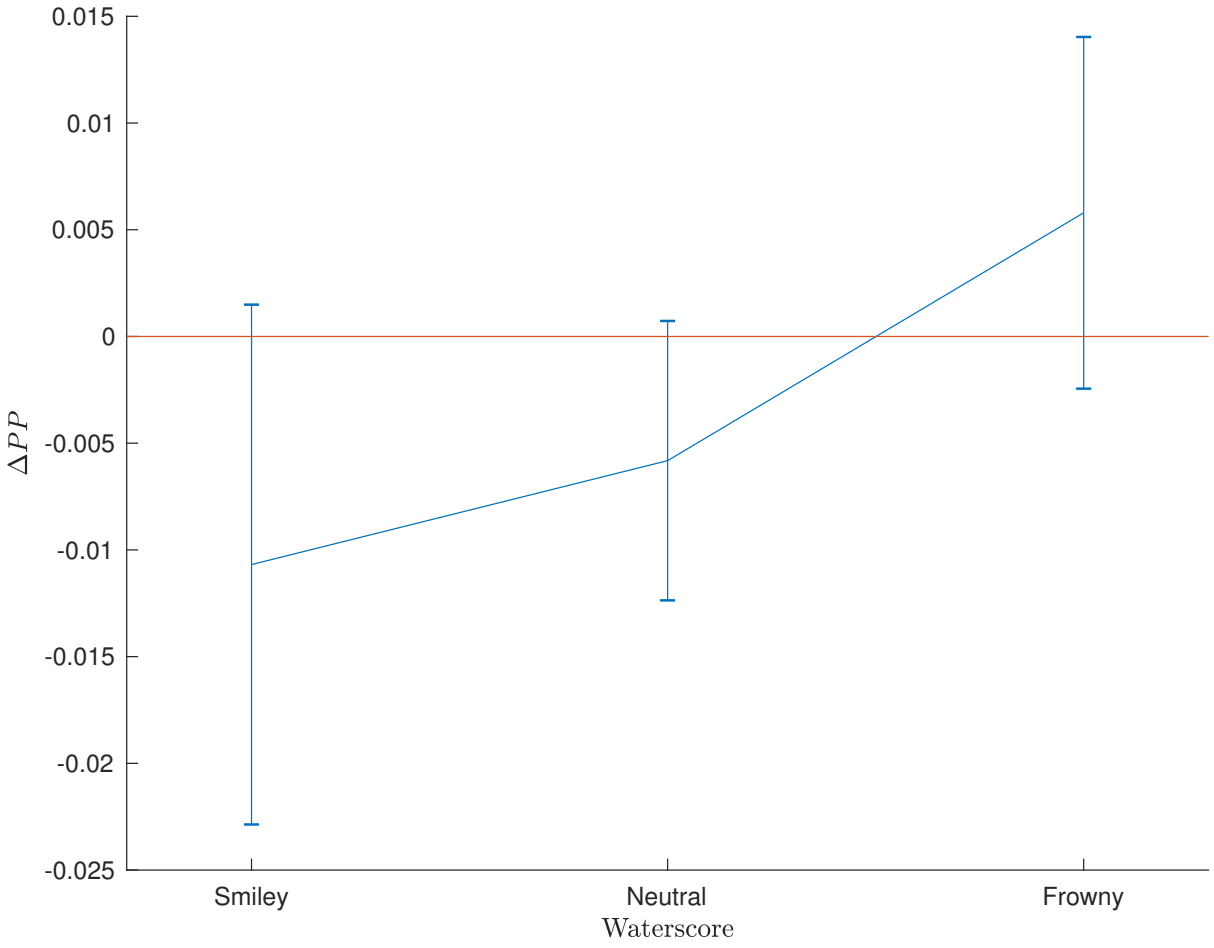

Figure 5. ATEs on the Percentage of HWR Emails Opened (After the First HWR Email Open), by Injunctive Norm Type. This figure shows the average treatment effects on the percentage of HWR emails opened following the first HWR email open, disaggregated by the type of injunctive norm seen on the first HWR email open, using linear regression models. The average treatment effects of the Injunctive Drop condition relative to the Drop condition are estimated using specifications with a full set of controls (i.e., that used in Table 4, column 4). The controls are: 1) demographic controls at the household level, which consists of home size, lot size, irrigable area, and the number of occupants, floors, bedrooms and bathrooms; 2) month-in-sample fixed effects; and 3) wave fixed effects. The error bars mark ±1 robust standard error, and the p-values from hypotheses tests comparing each coefficient to zero are 0.38, 0.374, and 0.482 for the smiley, neutral, and frowny messages, respectively.

Online Appendix:

Isolating the Effect of Injunctive Norms on Conservation Behavior: New Evidence from a Field Experiment in California

Table A.1

The Distribution of the Number of Emails Received by Subjects in Each Treatment Group

	Number of emails received								Total
	0	1	2	3	4	5	6	7	
No Drop Group	346 (1.89%)	288 (1.57%)	319 (1.74%)	378 (2.06%)	546 (2.98%)	1241 (6.76%)	12506 (68.16%)	2723 (14.84%)	18347 (100%)
Drop Group	361 (1.97%)	278 (1.52%)	319 (1.74%)	352 (1.92%)	522 (2.85%)	1256 (6.85%)	12483 (68.04%)	2775 (15.13%)	18346 (100%)
Injunctive Drop Group	87 (1.90%)	72 (1.57%)	81 (1.77%)	94 (2.05%)	134 (2.92%)	279 (6.08%)	3182 (69.37%)	658 (14.34%)	4587 (100%)
Total	794 (1.92%)	638 (1.55%)	719 (1.74%)	824 (2.00%)	1202 (2.91%)	2776 (6.72%)	28171 (68.24%)	6156 (14.91%)	41280 (100%)

Description for Online Appendix Table A.2

As described in the paper, in section 3.5, I performed a large trim of the dataset motivated by the need to identify appropriate control households for the treatment households, who started receiving HWRs on different dates. Note that this is not strictly necessary for analysis that focuses only on treated households, which comprises the bulk of the analysis in the paper--it is only influential for analysis of treatment effects relative to not receiving emails. The challenge here was to determine when Control households "would have" received their first HWR had they been treated. To do this, I focused on treatment subjects only and compared the dates of water meter reads in the main window of HWR email initiation (November 17, 2014 to January 21, 2015) to the dates that these subjects actually received their first HWR by email (which ranged for this group from December 9, 2014 to January 30, 2015). There was a clear pattern, whereby nine distinct ranges of dates for water meter reads were associated with receiving a first HWR email on a specific, later date.[1] Using this information, I linked each subject (in all conditions, including the Control) to one of nine "waves" based on their meter read dates falling in each of these nine ranges. This effectively grouped subjects in all conditions into nine distinct subgroups, who received, or "would have" received in the case of the Control, their first HWR email on each of nine specific dates. Details regarding these wave assignments and dates are included in Online Appendix Table A.2. I then dropped from analysis the 706 households that did not have a meter read date in the nine "wave" windows.

[1]For a few meter read dates there was a bit more ambiguity about when first emails were sent than for others, and in these cases the most common "date of first email" for a given meter read date was used to determine when Control subjects "would have" received their first email had they been treated.

Table A.2

Wave Date Range Description

Wave Date Range	Associated First HWR Email Date	Control Households	Treatment Households
November 17, 2014-November 24, 2014	December 9, 2014	722	6250
November 25, 2014-December 1, 2014	December 12, 2014	406	3513
December 2, 2014-December 5, 2014	December 18, 2014	463	3865
December 6, 2014-December 12, 2014	December 23, 2014	459	3742
December 13, 2014-December 22, 2014	January 2, 2015	683	6318
December 23, 2014-December 24, 2014	January 6, 2015	164	1502
December 26, 2014-December 31, 2014	January 14, 2015	430	3775
January 2, 2015-January 6, 2015 & January 8, 2015-January 13, 2015	January 23, 2015	736	6356
January 7, 2015 & January 14, 2015-January 21, 2015	January 30, 2015	441	4087

Table A.3

Average Treatment Effects on Water Use

	DV: Gallons per Day (GPD)					
	Excluding Control			All Conditions		
	(1)	(2)	(3)	(4)	(5)	(6)
Injunctive Drop	-3.790***	-2.468**	-2.425**	-10.85***	-10.52***	-10.58***
	(1.211)	(1.121)	(1.114)	(1.877)	(1.942)	(1.934)
Drop	-1.356	-0.313	-0.333	-8.416***	-8.342***	-8.455***
	(0.867)	(0.786)	(0.783)	(1.675)	(1.769)	(1.759)
No Drop				-7.060***	-8.030***	-8.132***
				(1.722)	(1.795)	(1.790)
Pre-Exp. Cal. Month GPD		0.305***	0.303***		0.316***	0.314***
		(0.0357)	(0.0368)		(0.0338)	(0.0349)
Observations	247135	209136	209136	274949	232538	232538
R^2	0.000	0.198	0.207	0.000	0.185	0.193
Mean GPD for Omitted Group	200.11	200.11	200.11	207.17	207.17	207.17
Unique Households in Sample	40020	36589	36589	44511	40679	40679
Demographic Controls	No	Yes	Yes	No	Yes	Yes
Month-In-Sample Fixed Effects	No	No	Yes	No	No	Yes
Wave Fixed Effects	No	No	Yes	No	No	Yes
Wald Statistic p-value	0.0333	0.0487	0.0541	0.0333	0.0460	0.0509

Notes. This table shows analyses that precisely mirror those in Table 3 in the main paper, but without any of the data restrictions used (as described in Section 3.5 in the paper). Note that households without post-treatment data are still excluded here, since they did not have any outcome data to be analyzed. The results are qualitatively similar to those in Table 3 in the main paper. Robust standard errors are in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$