Title
Essays in Behavioral and Environmental Economics

Author
Gorenshteyn, Dina

Publication Date
2019

Peer reviewed|Thesis/dissertation

Essays in Behavioral and Environmental Economics

by

Dina Gorenshteyn

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David L. Sunding, Chair
Professor Meredith Fowlie
Professor Lucas Davis

Summer 2019

Essays in Behavioral and Environmental Economics

Abstract

Essays in Behavioral and Environmental Economics

by

Dina Gorenshteyn

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor David L. Sunding, Chair

Financial incentives are traditionally relied upon by economists to influence behavior, however there is increasing evidence that non-market-based solutions may be preferable. This outcome particularly emerges when dealing with inattentive individuals or in situations where financial approaches face prohibitive legislation, budgetary restrictions, or political push-back. Understanding the impact of using non-market-based behavioral instruments, and how they compare to the use of market-based financial incentives, is crucial for governments to design effective public policy, as well as for private institutions to maximize profit. In this dissertation, I consider the setting of drought-prone California and its water consumers. The first two chapters contribute to the behavioral and environmental economics literature by exploring the impact of using non-market-based techniques – namely the use of information provision and moral suasion – in encouraging water conservation. The last chapter highlights the importance of finding effective approaches to water conservation by quantifying the impact of drought on economic outcomes.

In the first chapter, I use a novel natural experiment and a randomized field experiment to investigate how information and financial incentives compare in influencing the behavior of inattentive customers. Specifically, I study how these two tools compare in encouraging water customers to fix in-home leaks. I find that information is a powerful, low-cost tool in swaying behavior, and can be even more effective than financial incentives. Importantly, the impact of a financial incentive depends on a customer's typical bill-to-bill charge variance, such that high-variance customers are less likely to react to a financial incentive and respond more to a clear informational signal. Financial incentives under one standard deviation of customer bill-to-bill variation likely go unnoticed. In the observed setting of customers with in-home leaks, the average customer has a 30% relative standard deviation of the month-to-month bill and only responds to a bill increase of 50% or more. Further, the impact of financial incentives vary by customer income and are considerably less effective on customers with automatic bill payment. I also find that while sending information by mail is effective, delivery through email or text may be preferable in time sensitive situations and for encouraging the use of online resources.

In the second chapter, I assess the impacts of using moral suasion via public appeals to encourage behavior change. As a case study, I analyze the effect of Governor Jerry Brown's public pleas for water conservation in the face of California's record-breaking drought. Using

high frequency hourly consumption data at the household-level for the years 2012-2015, I conduct an event study to understand the level of short-term water conservation associated with these appeals. I find statistically and economically significant decreases in water consumption in the single-family residential sector of San Francisco in the two weeks following a well-publicized public appeals announcement. These short-term decreases range from 1.9 - 4.6% of total single-family residential water demand.

In the third chapter, co-authored with David Sunding and Maximilian Auffhammer, we evaluate the effect of the drought on economic outcomes. In this study we use *ex post* impact assessment methods to measure the effect of drought on farm employment and harvested acreage for the fifth largest economy in the world – California. We find evidence of a statistically and economically significant relationship between surface water imports and both employment and harvested area. We also present evidence that the effects of drought are smaller in areas with better access to local water supplies, especially groundwater, and have declined over time. The latter observation is consistent with observed shifts in land allocation toward perennial crops and with increased reliance on groundwater extraction, particularly in dry years. These trends may not be sustainable in light of the State's recent efforts to curb groundwater overdraft. Our results suggest that absent other interventions, the future effects of drought on economic outcomes in California agriculture could be even larger than those observed in the recent past.

To *Mom and Dad*
for instilling in me curiosity and steadfastness,
and for engulfing me in warmth.

To *Edward*
for a lifetime of mentorship and friendship.

To *Sveta and Yura*
for unwavering support.

To *Rita and Ilya*
for making it all possible.

Contents

Acknowledgments

I would like to thank my dissertation chair, David Sunding, for years of constant support, sage guidance, and endless opportunities that have allowed me to thrive and become the economist I am today. I would also like to thank my other committee members: I thank Meredith Fowlie for her clear-cut and invaluable feedback, and Lucas Davis for taking the time to carefully think through the details with me and for providing indispensable advice. I am also grateful to the San Francisco Public Utilities Commission for trusting me with their data and supporting my research.

I thank Andreas for inspiring and arming me to take on this challenge. I thank Ellen, Max, Steve, Jim, Sofia, and Severin, as well as the participants of the Environmental and Resource Economics Seminar for helpful suggestions and comments. I thank Carmen and Diana for making the logistical parts of grad school a breeze, and for always lending an ear when the rest gets tough.

Finally, I thank my family and friends for their tireless love, support and encouragement – over thousands of cups of tea – without which this dissertation would not have been possible and my sanity would not have remained intact.

Chapter 1

Words Speak Louder Than Money – Influencing Behavior with Information and Financial Incentives

1.1 Introduction

The interest in influencing behavior is ubiquitous. Firms and governments alike seek the best tools to encourage actions such as purchasing a new product, adopting a new technology, conserving natural resources, or developing healthier habits. Financial incentives have been a well-documented approach in effectively promoting all types of behavior,[1] however they may not always be the best instrument, especially when dealing with recurring customers who have become inattentive to their bills. Inattention is rational as it is costly in terms of cognitive resources and time (Gabaix et al., 2006; Sallee, 2014), and is further heightened when customers are dealing with complex pricing schemes (Borenstein, 2009; Ito, 2014) and are on automatic bill payment (Finkelstein, 2009; Sexton, 2015).

Given customer inattention, those interested in influencing behavior may look to alternative approaches. One such approach is providing customers with information about the desired action. Information can come in many forms. Customers could be given technical information about the product use or desired action, including general statistics across the population or customized to the particular customer. The information could further include pro-social/moral messaging such as how customer behavior affects the environment. Moreover, the information could contain social comparisons about how the customer's behavior stacks up against his peers.

The literature has shown that providing information can be a powerful tool. This result is found across many areas of economics. It has been demonstrated in studies of education attainment that giving low-income families more information about schooling will lead to better schooling choices (Hoxby and Turner, 2013; Hastings and Weinstein, 2008). Similarly,

[1] Including, but not limited to, motivating exercise (Charness and Gneezy, 2009), inducing smoking cessation (Volpp et al., 2009), spurring charitable giving (Landry et al., 2005), improving academic refereeing (Chetty et al., 2014), and encouraging energy conservation (Reiss and White, 2008).

in the health sector, increased information provisions can improve Medical Part D drug plan choices (Kling et al., 2012) and physician performance (Handel et al., 2014). In the financial sector, information has been shown to help individuals better assess the costs of payday loans (Bertrand and Morse, 2011). There have also been various studies on the effect of information on energy and water consumption (Allcott, 2011; Ferraro and Price, 2011; Jessoe and Rapson, 2014; Ito et al., 2018; Reiss and White, 2008; Davis and Metcalf, 2016; Gosnell et al., 2016).

While the aforementioned studies have laid a strong foundation towards the understanding of how information affects behavior, there is still room to contribute to the literature. Some of the studies focused on a particular subset of the population such as low-income families, seniors, or male airplane pilots, yielding results that may not be representative of all demographics (Hoxby and Turner, 2013; Hastings and Weinstein, 2008; Kling et al., 2012; Gosnell et al., 2016). Other studies focused on all demographics, but rely on recruited subjects, which raise concerns about external validity associated with selection bias, the Hawthorne Effect, and priming (Jessoe and Rapson, 2014; Ito et al., 2018). A subset of studies are representative of all demographics and do not rely on recruited subjects, however these studies tend to focus only on informational treatments and do not have a direct comparison of outcomes to financial treatments (Allcott, 2011; Ferraro and Price, 2011).

The objective of this paper is to compare the effects of informational notifications and financial incentives in influencing the behavior of inattentive customers. The setting used for the analysis is a natural experiment focused on single-family residential (SFR) water consumers, serviced by San Francisco Water Power Sewer (hereinafter "SF Water"), that experience in-home leaks such as constantly running toilets and leaky faucets. An average leak costs a customer \$320 and wastes about 12 thousand gallons of water.[2] There are three ways a customer becomes aware of a leak: (1) by seeing or hearing the problem, (2) getting a notification from the water utility with a leak alert, or (3) receiving an unusually high bill.

To understand how an informational notification (i.e. leak alert from the water utility) compares to a financial incentive (i.e. an unusually high bill) in encouraging customers to fix leaks, I exploit the exogenous variation in leak timing (e.g. day-of-week and day-of-month that the customer starts leaking). Leak timing randomly affects whether a customer receives an informational notification from the water utility or a bill statement with a bill spike first. This allows me to directly compare how customers respond to the two treatments. I am able to further test how outcomes vary for different levels of financial incentives since the timing of the leak relative to the customers bill cycle completion date, as well as the physical size of the leak, randomly affect the size of the bill spike. Finally, I conduct a randomized field experiment to test whether sending information through multiple contact points, including electronic methods, is more effective than sending just a paper mailer. I follow-up with a survey to collect qualitative evidence in support of the empirical findings.

Unlike most of the previous studies, this analysis is able to compare the effects of financial and informational treatments, while producing results that are representative of the overall population and avoid many issues of external validity. The results are representative of the

[2]These figures are based on the study period data assuming no leak alert notifications. The cost is based on FY17-18 single-family residential rates, including water and sewage fees.

overall population as all demographics are consumers of water, and are vulnerable to leaks.[3] Further, the study steers clear of selection bias since it relies on a naturally occurring, non-recruited set of customers that happen to spring leaks. Additionally, the study is not tainted with the Hawthorne Effect. Customers in the natural experiment are not explicitly aware that they are part of a study, so they should not exhibit any behavior change in response to being observed. Moreover, the study does not suffer from priming. Since the customers in the study do not know that they will be receiving a notification or an unusually high bill, they are not primed to look for these things, so the results capture attention towards unsuspecting treatments. Lastly, since the financial treatment in the study is administered as part of the monthly bill – which is the case in many settings – and is not stated separately, the study is able to assess the effect of financial incentives that appear as part of a recurring bill statement.

I find that informational notifications are in fact successful in encouraging behavior, and are actually more effective than the average bill increase. This is consistent with the informational notifications being more salient than the bill increases. Since customer bills have natural month-to-month consumption variance, even a substantial bill increase may be difficult to spot. The average customer has a 30% relative standard deviation of the month-to-month bill and responds to a bill increase of 50% or more.[4] These findings suggest that customer attention is only captured when the bill increase surpasses one standard deviation of the month-to-month bill variation, an important consideration when deciding on financial incentive size.

Interestingly, two-thirds of customers fix leaks with just a physical clue or an informational notification (i.e. no explicit financial incentive). The remaining customers end up receiving a financial incentive through a bill increase. I find that the larger the financial incentive, the more quickly these customers fix leaks. As expected, customers with lower incomes are substantially more responsive to financial incentives than those with higher incomes. Moreover, customers that are on automatic bill payment do in fact respond to financial incentives, but are much less sensitive than their counterparts, responding only at fairly large bill increases.

The data also shows that while sending information by mail is effective, delivery through text or email may be preferable in time sensitive situations and for encouraging the use of online resources. In practice, sending customers informational notifications via mail saves an average customer \$115 in avoided water costs, while costing the water utility only \$2.74 per leak for printing and mailing. Sending electronic notifications are estimated to save an average customer an additional \$25, while costing the water utility an average of \$3.22 per leak for the system maintenance and communication fees.

These findings are statistically and economically significant as water utilities represent large, ubiquitous markets where even small changes in behavior on the consumer-level have substantial monetary and environmental impacts on the aggregate. In San Francisco alone, informational notifications save single-family residential customers 27 million gallons of water

---

[3]While those prone to leak are not completely random, they are still generally representative.

[4]The relative standard deviation is defined as one standard deviation of the past 6 monthly bill charges divided by the most recent pre-leak bill charge. Bill increases are measured as the percent increase relative to the most recent pre-leak bill charge.

and \$0.7 million annually, with savings expected to considerably increase in the coming years. Potential state-wide impacts of information are noteworthy as single-family residential water demand represents only 20% of SF Water retail demand, and SF Water is just one of hundreds of water utilities in California. Further, using information to influence water conservation via fixing leaks is a sustainable form of conservation, is positively received by customers[5], and is among the most cost-effective water conservation strategies utilized by SF Water. These finding are particularly important as it is expected that droughts in the western United States will increase in frequency and intensity due to climate change, making effective water conservation strategies all the more essential.[6]

The paper proceeds as follows. Section 1.2 discusses the background details regarding leak detection and leak notification. Section 1.3 lays out the testable hypotheses and outlines the research design. Section 1.4 describes the data. Section 1.5 presents the empirical analysis and results. Section 1.6 reports the implied savings, and Section 1.7 provides concluding thoughts.

1.2 Background

This study focuses on single-family residential water customers serviced by SF Water that are identified as likely experiencing a leak during the year and a half study period of September 2017 through February 2019. Residential in-home leaks can waste a great deal of water and cost customers a lot of money. In-home leaks are typically caused by leaky toilets, but in many cases are due to other sources such as leaky or left on irrigation systems, leaky faucets, or burst pipes. On average, there are roughly 100 new incidents of leaks in single-family homes serviced by SF Water every week. At this rate, almost 5% of all single-family residential customers experience a leak annually. This is a conservative estimate as the water utility only detects fairly large leaks of at least one cubic foot per hour (cf/hr), where one cubic foot is equivalent to 7.5 gallons. As a result, there are likely many more unaccounted for incidences of smaller leaks.

While almost all leaks are eventually resolved, the speed at which they are addressed affects the amount of water and money wasted. The smallest leak that the water utility detects costs a customer about \$5 a day.[7] Absent any informational notifications, the average customer takes about a month to fix a leak, wasting roughly 12 thousand gallons of water, and paying \$320 extra in water bills. Those that take longer may end up paying several thousands of dollars. For reference, the average single-family residential customer pays roughly \$90 on

[5] Survey results show all but one respondent requested future notifications regarding leaks.

[6] As described by the California Department of Water Resources report: "California's Most Significant Drought: Comparing Historical and Recent Conditions", 2015.

[7] Cost is based on SF Water FY17-18 water (\$8.81/ccf) and wastewater (\$12.40/ccf) variable charge, where ccf is hundred cubic feet (748 gallons). The water variable cost is the second tier of potable water rates where the first tier is for consumption of the first 4 ccf and the second tier is for all consumption above 4 ccf. Since average customers consume about 5 ccf without a leak, it is likely the water spent on a leak would fall in the second tier. The wastewater variable cost is not tiered. With a flow factor of 90% to wastewater, one ccf costs \$8.81 + .90 * \$12.40 = \$19.97. One day of the smallest leak of 1cf/hr equals .24 ccf, costing \$19.97 * .24 = \$4.79.

water per month.[8] In aggregate, leaks represent roughly 1.2% of total single-family water demand, costing $1.64 million to customers annually. Again, this is a conservative estimate as it does not include leaks smaller than 7.5 gallons per hour.

1.2.1 Detecting Leaks

SF Water can detect the start and end of a leak. The water utility tracks household level hourly water consumption through their smart meter technology. If a household meter shows nonstop water usage of at least 1 cubic foot (7.5 gallons) of water per hour for 72 consecutive hours, the water utility flags the household as likely experiencing a leak. Such nonstop water usage usually indicates a leak, however in some cases may be caused by an *unintended behavior* such as inadvertently leaving the irrigation on, or by an *intended behavior* such as running a medical device or filling a swimming pool.

The first day of the flagged 72 hours is considered the first day of the leak. The day a stop in the continuous use is detected is considered to be the last day of the leak. The water utility monitors the water account for an additional 28 days after the continuous use stops to officially consider the leak fixed.

1.2.2 Leak Alert Notification

Customers are responsible for fixing their in-home leaks and for any associated cost of water loss from inaction. Traditionally customers become aware of a leak by either physical clues (seeing or hearing the issue), or through a monetary clue by receiving an unusually high bill. Since SF Water is now able to detect leaks just a few days from the start, the water utility sends a notification to the customer with an alert about the potential leak. Notifications are sent by mail, robo-call (i.e. interactive voice response), email, and SMS text via mobile phone. All customers receive a mailed notification. Customers with additional contact information on file with the water utility are sent the other modes of notifications in addition to the mailer. Figure 1.1 shows an example of the interior messaging of the mailed letter notification. The other modes of communication have a very similar message. More details about the notification messaging and the exterior of the notification are found in Appendix A.1.

Once a household is flagged by the water utility as likely having a leak, they enter a queue to be sent notifications. Once the household's turn is reached in the queue, typically sometime that day, or the following business day, an email, text, and robo-call are sent out to the customer within a 30 second window.[9] All mailers for the week are generated on Sunday night and physically mailed on Mondays. As a result, all customers identified as having a potential leak for the week are mailed a notification on the following Monday, and likely receive the mailer on Wednesday.

---

[8] This figure is based on pre-study period 2016 billing data for all single-family residential customers serviced by SF Water. Water bills include water and wastewater charges.

[9] The email, text, and robo-call are sent throughout the day Monday - Friday, 10 AM - 4 PM. Any customers not reached in the queue by the end of the day are likely notified the next morning.

Figure 1.1: Notification Messaging (Interior)

URGENT: Potential Water Leak at 546 Clipper.
Courtesy Notice # 1

Dear SFPUC Customer,

Our data shows **nonstop water usage at your home between 12/15–12/17/2017** of at least 60 gallons per hour. This may mean you have a **plumbing leak!**

Please log onto **MyAccount.sfwater.org/find** to review your daily water use and check for unusual increases. If you are not currently registered for MyAccount it only takes a few minutes. For tips on how to detect and fix a leak or for free assistance available through the SFPUC, please visit **sfwater.org/homeleaks/find**. You can also obtain a copy of our Leak Guide at the first floor Customer Service Counter at 525 Golden Gate Avenue.

Remember, while leaks can happen to anyone, it is your responsibility to resolve plumbing leaks in your home in a timely manner. Tenants receiving this notice may want to contact the property owner for more direction on leak repairs. For further questions, call **(415) 551-3000** 8AM-5PM, Monday to Friday or email **customerservice@sfwater.org**.

Thank you,
SFPUC Customer Service

Since leaks are flagged by the system 72 hours into the leak, customers receive their mailer anywhere from 7 to 13 days after the beginning of their leak. If customers have additional contact points on file, they may receive some type of notification (i.e. email, text, robo-call) as soon as 3 days into the leak. Once the initial round of notifications are sent out, households are notified again 14 days later, and then a third time 8 weeks after the second round of notifications.

1.3 Research Design

The research design consists of two distinct experiments. The first is a natural experiment resulting from the random timing of leaks. This experiment allows me to analyze how customers respond to information and to financial incentives. The second is a randomized field experiment conducted in collaboration with SF Water. The field experiment allows me to analyze how to best deliver information. Before describing the details of each experiment, I outline a set of testable hypotheses.

Testable Hypotheses

Hypothesis 1. People respond to informational notifications (no explicit financial incentive needed).

Hypothesis 2. Information versus financial incentive.
(a) The average bill increase is not as effective as an informational notification.
(b) A bill increase is at least as effective as an informational notification above a certain percent increase threshold.

Hypothesis 3. For customers that receive both an informational notification and a bill increase, the greater the financial incentive the more effective it is in influencing behavior.

Hypothesis 4. Informational notifications delivered through multiple contact points, including electronic methods, are more effective than just a paper mailer.

1.3.1 Natural Experiment Design

The random timing of a leak determines the amount of days until a customer receives a notification, the amount of days until a customer receives an unusually high bill, and the size of the bill increase. Below I outline how exactly the leak timing affects each of these outcomes, and how this set-up allows me to test the aforementioned hypotheses. It's worth noting that every customer eventually receives both a bill and a notification.

The random *day-of-week* that a customer springs a leak affects the amount of days until the arrival of the notification. This occurs because, as discussed before, all mailer notifications for leaks during the week are batched on the Sunday of the week and reach the customer the following Wednesday. As a result, customers receive their notifications anywhere from three to thirteen days into a leak. This exogenous variation in days to receive a notification allows me to test how receiving a notification one day earlier affects the total amount of days leaking. In other words, I can test if an informational notification is effective in influencing behavior (i.e. *Hypothesis 1*).

The random *day-of-month* that a customer springs a leak relative to the random day-of-month that the customer's bill cycle ends affects the amount of days until a customer receives an unusually high bill. Since the days until the notification and the days until the bill arrival is random, I am able to test whether a notification or an average financial incentive is more effective given that they arrive on the same amount of days into a leak (i.e. *Hypothesis 2a*). Moreover, the random amount of days until a customer receives the bill affects the size of the bill spike. Variation in the bill spike size is also affected by the pseudo-random size of the leak.[10] This set-up allows me to test how different sizes of financial incentives affect days leaking (i.e. *Hypothesis 3*) and help pinpoint the percent threshold increase at which customers notice something is unusual with the bill and start reacting as they would had they received an informational notification (i.e. *Hypothesis 2b*).

I test whether a blast of notifications is more effective than just a paper mailer (i.e. *Hypothesis 4*) using a randomized field experiment, described below.

1.3.2 Randomized Field Experiment Design

In partnership with SF Water, I designed and implemented a randomized field experiment that varied whether customers with leaks received a paper mailer or a multi-pronged blast of notifications that included a mailer along with electronic notifications, specifically robo-calls, emails, and/or texts.[11] All 1900 single-family residential households suspected of in-home leaks during the study period of September 2017 to January 2018 were included in the experiment.

Experimental Groups

Households were randomly assigned to one of two groups:

Control Group (CG): Households in this group were sent a mailer notification regarding their potential leak. While these customers may have additional contact information on file, no other types of notifications were sent.

- Customers in the Control Group receive their mailer anywhere from 7 to 13 days after the beginning of their leak.

Treatment Group (TG): Households in this group were sent a blast of notifications with 2+ contact points including a mailer (same as in the Control Group) plus an email, text, and/or robo-call depending on what contact information was on file with the water utility. While most customers have a phone number on file, a quarter have an email address on file, and

[10] Leak size appears to be random in the data, however one may argue that it is correlated with factors such as housing age, income, etc.

[11] The water utility is moving towards the blast notification approach for all customers in the future. The original motivation for the randomized field experiment was to help the water utility understand the benefits of investing in the technological system that allows them to send blast notifications.

a quarter have a mobile phone number that can be used for texts on file, not all contact points on file were sent notifications due to internal technical reasons that have since been addressed. As a result, roughly half of contact points on file were randomly used.

- 51% of customers in the Treatment Group were sent 2+ contact methods (i.e. a mailer plus an email, text, and/or robo-call). The remaining 49% of customers were just sent a mailer notification like those in the Control Group. Table A.1 in Appendix A.2 details the exact contact methods sent to each experimental group.
- Customers in the Treatment Group receive their first mode of notification anywhere from 3 to 13 days into their leak depending on the contact points used.

Messaging

The information provided in both groups is identical. It includes the dates of the continuous consumption review period, the size of the potential leak in gallons, the property address, and suggested steps to take to investigate the issue including websites to visit and customer service contact information. There is no information on the cost associated with the leak included in any of the messaging. The only thing that varies across groups is the intensity of the notification delivery moving from just a mailer to a blast. Figure 1.1 shows an example of the letter notification.

The messaging for both groups advises customers to take first-steps towards addressing their potential leak by:

1. Logging into their online account portal (MyAccount)
2. Visiting the leak tips webpage
3. Calling or emailing SF Water customer service

The water utility tracked if customers did any of the suggested first-steps. Both of the URLs included in the messaging had links that allowed tracking of whether the websites were visited.[12] The water utility also made note of the time and content of household calls and emails to customer service in regards to potential leaks.

Outcome Variables

I observe a set of outcome variables that fall into two categories: Resource-Use and Follow-Through. The Resource-Use outcomes measure if the treated customers are more likely to do any one of the first-steps suggested in the messaging (outlined above). The Follow-Through outcome measures if treated customers ended up fixing their leak or behavior any faster.

[12]The email and text messaging included customized links for each household, so it was possible to track exactly what household visited the link and at what time. The mailer notifications included only group-specific links allowing me to track whether the household using the links from these notifications was from the Control Group or the Treatment Group, but does not allow tracking of the exact household. The robo-calls provided general links only and were not trackable.

Follow-up Survey

Once the experimental period was over, a follow-up survey was sent to all participants in the experiment. The survey included questions about the source of the nonstop water use, the role of the notification in discovering the problem, and the time, resources, and money associated with addressing it. The survey also included questions regarding household specific characteristics such as income, education, age, and gender. This data is used as qualitative evidence to help better understand the mechanisms behind the findings from the subsequent analysis. A copy of the follow-up survey is included in Appendix A.3 for reference, as well as a detailed description of how the surveys were administered.

1.4 Data

For this analysis I compile a data set including household-level billing history for all SF Water single-family residential customers (October 2015 - February 2019), leak incidence and notification data, and follow-up survey results. I use the 2015 US Census American Community Survey data at the zip code level for additional demographic information.

I first compare the study sample of customers experiencing leaks between September 2017 and February 2019 to all single-family residential customers serviced by SF Water to understand how representative the sample is of the overall population. Next, I compare the observable characteristics across treatment groups for the natural experiment to assure that the groups are comparable for the analysis. I then do the same comparison across treatment groups in the randomized field experiment.

1.4.1 Comparing the Study Sample to the Population

Table 1.1 presents a comparison of demographics for customers in the study sample (i.e. single-family residential customers that experience leaks during the study period) to the demographics for all single-family residential customers. It also includes demographics for the single-family residential customers that experienced leaks and voluntarily responded to the follow-up survey. The first four rows represent household-level data provided by the water utility. The remaining variables are zip code level data from the 2015 US Census American Community Survey. The average monthly consumption, shown in hundred cubic feet (ccf), and the monthly bills are based on 2016 calendar year billing data, which predates the observed leaks, so these variables are baseline comparisons and do not include consumption associated with the leaks during the study period.[13]

As seen in Table 1.1, and confirmed by a statistical comparison, customers with leaks tend to consume more water overall, prior to the study period, and accordingly have higher monthly bills as compared to the average single-family residential customer. This is partially explained by the fact that customers with leaks tend to have more occupants, however they also have higher consumption per occupant. The fact that the households that end up

[13]The bill includes both water costs and sewage costs. Sewage consumption is assumed to be 90% of water use.

leaking have more occupants suggests that there may be more fixtures in the home that could spring a leak, or there is more fixture-use in general that could cause leaks. The fact that the households that end up leaking consume more per occupant may suggest that these customers live in homes with less efficient, older appliances, and have likely leaked in the past, which would be captured within the baseline consumption.

Table 1.1: Summary Statistics for All Single-Family Residential Customers, Single-Family Residential Customers with Leaks, and Single-Family Residential Customers with Leaks that Volunarily Responded to the Survey

	All SFR (Population)	SFR w/ Leaks (Study Sample)	SFR w/ Leaks (Survey)
Number of Customers	110,875	5,792	572
Monthly Consumption (ccf)	5.10	6.83	6.74
	(3.59)	(5.19)	(5.08)
Monthly Bill ($)	91.42	121.55	119.25
	(61.81)	(91.46)	(89.13)
Number of Occupants	3.35	3.60	3.34
	(2.23)	(2.54)	(2.03)
Monthly Bill Per Occup ($)	32.38	40.18	41.18
	(24.74)	(33.30)	(31.30)
Median Income ($)	90,704.25	92,867.72	95,472.85
	(22,781.87)	(24,266.48)	(23,921.40)
Median Age	40.10	39.85	40.06
	(3.24)	(3.42)	(3.45)
Proportion Male	0.50	0.50	0.50
	(0.03)	(0.03)	(0.03)
Proportion Children	0.25	0.24	0.23
	(0.07)	(0.08)	(0.07)

Notes: The table presents the mean for each variable by group with the standard deviation in parenthases. Monthly consumption, monthly bill, number of occupants, and bill per occupant are all on the household-level, based on pre-leak 2016 billing data. The remaining demographics are at the zip code level, based on 2015 Census American Community Survey. Median income is in 2017 dollars.

1.4.2 Balance on Observables for the Natural Experiment

Table 1.2 presents the average of observable characteristics for leak incidences where the customer randomly receives a bill first and for those who randomly receive a notification first. A star appears in the right hand column for instances when customers statistically differ in a characteristic at the 95% significance level. There are more leak incidences than customers that experience leaks because 17% of customers experience leaks multiple times during the year and a half study period.

The first part of the table shows the billing data and demographics discussed earlier, as well as the leak size (shown in cubic feet per hour) and the proportion of customers on automatic bill payment (i.e. autopay).[14] As expected from the natural randomization (discussed in the Research Design section), customers in the two groups are statistically balanced on almost all of the observable characteristics.

The bottom portion of the table shows averages by group for the number of days it takes to receive the first form of contact (i.e. the bill or the notification), for the number of days it takes to receive the second form of contact (i.e. the notification or the bill), and for the proportion of customers sent 2+ notification methods (i.e. mailer + email/text/robo-call). These variables are expected to statistically vary across the two treatment groups by design.

Table 1.2: Summary Statistics by First Contact Type

	Bill First	Notification First	T-Stat	Sig (95%)
Number of Leak Incidences	1,463	5,375	–	–
Leak Size (cf/hr)	3.36	3.60	-1.27	
Monthly Consumption (ccf)	7.05	7.26	-1.32	
Monthly Bill ($)	124.95	129.06	-1.47	
Number of Occupants	3.66	3.68	-0.3	
Monthly Bill Per Occup ($)	40.42	41.50	-1.08	
Proportion Autopay	0.15	0.14	0.81	
Median Income ($)	91,609.41	93,007.31	-1.99	
Median Age	39.68	39.86	-1.77	
Proportion Male	0.50	0.50	-2.03	*
Proportion Children	0.24	0.24	0.28	
Days to First Contact	4.99	6.55	-16.1	*
Days to Second Contact	8.51	22.71	-51.64	*
Proportion 2+ Notifications	0.38	0.63	-17.34	*

Notes: The table presents the mean for each variable by group. (*) signifies a statistical difference in means across groups at the 95% level. There are more leak incidences than customers with leaks because roughly 17% of customers leak multiple times within the study period. Leak size, monthly consumption, monthly bill, number of occupants, bill per occupant, and autopay status are all on the household-level, with all but leak size and autopay status based on pre-leak 2016 billing data. The Leak size and autopay status are measured during the leak. The remaining demographics are at the zip code level, based on 2015 Census American Community Survey. Median income is in 2017 dollars.

1.4.3 Balance on Observables for the Randomized Field Experiment

Table 1.3 presents the average of observable characteristics for those who are randomly assigned to the Control Group versus those assigned to the Treatment Group. The top

[14]Customers on autopay are customers that voluntarily signed up for recurring automatic monthly payments. These customers receive a monthly bill statement, but do not need to take any actions to pay their bill as their account is set-up to automatically charge their pre-specified payment method.

portion of the table compares the same variables as those in the top portion of Table 1.2. As expected from the imposed randomization of the field experiment, customers in the two groups are statistically balanced on almost all of the observable characteristics.

The bottom portion of the table shows experimentation statistics on how many different notification methods were used and how many days until the first notification reaches the customer. It is seen that the average number of notification methods for the Treatment Group is just under 2. This is because the water utility did not have contact information available for all four methods of communication for all customers. These two variables are expected to statistically vary across the two treatment groups by design.

Table 1.3: Summary Statistics by Experimental Group

	Control	Treatment	T-Stat	Sig (95%)
Number of Leak Incidences	960	940	–	–
Leak Size (cf/hr)	3.31	3.47	-0.51	
Monthly Consumption (ccf)	7.31	7.60	-1.05	
Monthly Bill ($)	130.56	134.99	-0.92	
Number of Occupants	3.69	3.70	-0.04	
Monthly Bill Per Occup ($)	42.27	42.49	-0.14	
Proportion Autopay	0.12	0.12	-0.52	
Median Income ($)	92,020.84	92,426.78	-0.36	
Median Age	39.90	39.56	2.15	*
Proportion Male	0.50	0.50	1.56	
Proportion Children	0.24	0.24	0.21	
Number of Notification Methods	1.00	1.96	-26.98	*
Days to First Notification	10.11	7.11	22.34	*

Notes: The table is based on data for the subset of customers that were part of the randomized field experiment. The table presents the mean for each variable by experimental group. (*) signifies a statistical difference in means across groups at the 95% level. Leak size, monthly consumption, monthly bill, number of occupants, bill per occupant, and autopay status are all on the household-level, with all but leak size and autopay status based on pre-leak 2016 billing data. The Leak size and autopay status are measured during the leak. The remaining demographics are at the zip code level, based on 2015 Census American Community Survey. Median income is in 2017 dollars.

1.4.4 Follow-up Survey

A random subset of the study sample were administered surveys. 572 customers responded to the survey, yielding a 28% response rate. Survey response rate was balanced across experimental groups: 51% were those in the Control Group, 49% were in the Treatment Group. A detailed breakdown of response rates is found in Appendix A.3.

Table 1.1 shows that survey respondents have similar demographics to the total study sample of customers with leaks, statistically differing only in the number of occupants, in household income, and in the proportion of children. The fact that the survey respondents are more affluent with less children is consistent with the fact that they have more time to answer a survey. Since the survey respondents are similar to the total study sample on all the other observables, the respondents are at least somewhat representative, and their survey answers can provide useful insight. A summary of survey results are in Tables A.2 and A.3 of Appendix A.3.

1.5 Empirical Analysis and Results

In this section I first test whether informational notifications are successful in influencing customers to fix leaks faster. I next compare the effectiveness of providing information via a notification to providing a financial incentive via a bill increase, and investigate what factors drive the result. I then assess the effect of financial incentives on customers that receive both a notification and a bill increase before fixing their leak, and how the responses vary by autopay status and income level. Finally, I examine what delivery method of information is the most effective.

1.5.1 Do people respond to informational notifications?

To test if informational notifications influence behavior I regress the days leaking on the days until a notification is received. Since the amount of days it takes for a notification to reach a customer is exogenously determined (as described in the Research Design section), if a customer that receives a notification one day earlier leaks for less time overall, then an informational notification is salient and effective.

I estimate the informational treatment effect with the following equation where Y_{itz} is the days leaking and I_i is the amount of days until a customer receives the first informational notification. I control for the amount of days until a customer receives a financial incentive through a bill spike (F_i), and for whether the notification is electronic (E_i). In some specifications I also include day-of-week, day-of-month, and month-of-year fixed effects (λ_t), zip code fixed effects (μ_z), and household covariates (X_i). Household covariates include the leak size, the bill spike size (i.e. percent increase over the previous month's bill), the previous month's bill size, the relative standard deviation (RSD) of month-to-month bill (i.e. standard deviation of monthly bill for the last 6 months divided by the previous month's bill), whether the household is on autopay, the number of occupants, and whether the household has had multiple leaks within the study period.

$$Y_{itz} = \beta I_i + \eta_1 F_i + \eta_2 E_i + \lambda_t + \mu_z + X_i' \gamma + \epsilon_{itz} \tag{1.1}$$

The first row of Table 1.4 provides the coefficient results for β along with the standard errors that are clustered at the zip code level to control for spatial correlation. Column (1) shows the results for a specification without any fixed effects or household controls. Column (2) includes temporal fixed effects, Column (3) also includes spatial fixed effects, and Column (4) includes all fixed effects and household-level covariates.

The results show that a customer that receives an informational notification one day earlier leaks for roughly one day less. This outcome suggests that informational notifications are noticed by customers and are effective in influencing behavior. This result is significant at the 1% level and robust to specification. Further, survey response data supports this finding. 87% of respondents remember receiving a notification. Of those that recall the notification, 77% of respondents were not aware of a leak before receiving the notification, 9% were aware but had not fixed it yet, and the remainder had already fixed the problem. The survey not only shows that notifications are noticed and are helpful in discovering leaks, but also that notifications are wanted. All but one survey respondent requested future notifications.

Table 1.4: Effect of Informational Notifications on Days Leaking

	Dependent variable:			
	Days Leaking			
	(1)	(2)	(3)	(4)
Days to Notification	0.904***	0.950***	0.951***	0.952***
	(0.198)	(0.224)	(0.225)	(0.243)
Days to Bill	−0.034***	−0.035***	−0.038***	−0.032***
	(0.010)	(0.010)	(0.010)	(0.011)
Electronic Notification	0.490	0.105	0.228	0.715
	(1.062)	(1.216)	(1.241)	(1.407)
Time Fixed Effects	No	Yes	Yes	Yes
Zip Fixed Effects	No	No	Yes	Yes
Household Controls	No	No	No	Yes
Observations	6,812	6,812	6,812	6,264
R^2	0.017	0.028	0.036	0.087
Adjusted R^2	0.016	0.021	0.025	0.075

Notes: The dependent variable is days leaking. Time fixed effects include month-of-year, day-of-week, and day-of-month fixed effects. Household controls include bill spike size, leak size, previous monthly bill size, relative standard deviation of the last 6 month bills, whether on autopay, number of occupants, and whether the household has had multiple leaks within the study period. Standard errors are clustered at the zip code level, presented in parenthases. *p<0.1; **p<0.05; ***p<0.01

1.5.2 How does information compare to financial incentives?

The previous results show that informational treatments are effective in influencing customers to fix leaks faster. Next I show how financial incentives (i.e. bill spikes) compare to the informational notifications in encouraging customers. For this analysis, I first compare how customers respond on average to receiving a bill spike versus receiving an informational notification the same amount of days into a leak. I estimate the following equation where D_i is a dummy turned to one if customer i received a bill spike first, and zero if received a notification first. I control for the amount of days to the first contact type of a bill or notification (C_i^1), and for the days to the second contact (C_i^2). The rest of the controls in the equation are the same as those in the previous analysis.

$$Y_{itz} = \beta D_i + \psi_1 C_i^1 + \psi_2 C_i^2 + \psi_3 E_i + \lambda_t + \mu_z + X_i^{'} \gamma + \epsilon_{itz} \tag{1.2}$$

Table 1.5, column (4) shows that customers that receive an unusually high bill versus an informational notification on the same amount of days into a leak end up leaking for over 2 days longer (i.e. a notification is more effective than a bill increase on average). To put this in perspective, the lowest bill spike that a customer experiences is roughly \$15 (10%) increase over the customer's previous monthly bill, while the median bill spike is a \$55 (40%) increase. While the financial incentives received by customers that get the bill first are relatively substantial, the informational notification is still more effective at capturing the customer's attention and influencing behavior change on average. This suggests that informational notifications are more salient than a bill, either due to the fact that bill increases are not always noticed, or that bill increases are noticed but do not provide a clear signal of what they indicate (i.e. a leak that needs to be fixed).

Is there a bill increase amount that is as effective as information?

While informational notifications are more effective than bill increases *on average*, there may be a bill increase amount that is large enough for customers to notice and respond to as they would to a notification. I next ask whether there is a threshold size at which a bill spike becomes as effective as an informational treatment. To analyze this question I repeat the same estimation as above, however this time I redefine D_i with a factor variable that gets a zero for customers that receive a notification first, and a bin value for those who receive a bill spike first. The bin value given to D_i depends on the bill increase size experienced by the customer, where bill increases are calculated as the percent increase on the bill relative to the previous month's bill. The bin values range as follows: 0-20%, 20-40%, 40-60%, 60-80%, 80-100%, 100-200%, 200-300%, 300-400%. The larger bins are wider as there are not as many customers that fall in those bill spike ranges. This approach allows me to find the bin threshold at which consumers start responding to the bill spike at the same rate as the notification.

Figure 1.2 demonstrates the resulting coefficients from the regression. Each coefficient is graphed at the midpoint of the corresponding bin. I connect the coefficients to emphasize the trend. The grey area represents the 90% confidence intervals. The coefficients for each bin show the difference in response time to the leak for customers with bill spikes in that bin range, relative to the response time for those that receive an informational notification the same amount of days into a leak. In other words, coefficients with confidence intervals above the zero y-axis indicate slower response times (i.e. more days leaking) under a bill increase than a notification; these bill increases are not as effective as notifications. Coefficients with confidence intervals that encompass the zero y-intercept indicate that the response to the bill increase is statistically no different than to a notification; these bill spikes have the same effectiveness as a notification.

Figure 1.2 shows that customers below a 50% bill increase take more days to respond to the leak than those who receive a notification (i.e. notifications are more effective below this

Table 1.5: Effect of Receiving Bill Increase vs. Informational Notification on Days Leaking

	Dependent variable:						
	Days Leaking						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Bill First	2.01*	1.90*	1.75	2.44**	1.94**	2.50**	2.71**
	(1.05)	(1.13)	(1.12)	(1.04)	(0.87)	(1.13)	(1.25)
Bill First*Bill RSD					1.77*		
					(0.98)		
Bill First*Autopay						−0.40	
						(1.36)	
Bill First*Big Income							−1.14
							(1.35)
Days to First Contact	0.62***	0.63***	0.63***	0.66***	0.66***	0.66***	0.66***
	(0.12)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)
Days to Second Contact	−0.03***	−0.04***	−0.04***	−0.03**	−0.03**	−0.03**	−0.03**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Electronic Notification	−1.50**	−2.10***	−2.03***	−1.28*	−1.29*	−1.28*	−1.29*
	(0.63)	(0.66)	(0.69)	(0.72)	(0.71)	(0.72)	(0.71)
Time Fixed Effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Zip Fixed Effects	No	No	Yes	Yes	Yes	Yes	Yes
Household Controls	No	No	No	Yes	Yes	Yes	Yes
Observations	6,812	6,812	6,812	6,264	6,264	6,264	6,262
R^2	0.01	0.03	0.03	0.09	0.09	0.09	0.09
Adjusted R^2	0.01	0.02	0.02	0.07	0.07	0.07	0.07

Notes: The dependent variable is days leaking. The independent variable of interest is Bill First, which is a dummy variable equal to 1 if the customer receives a bill increase first, and zero if the customer receives a notification first. Time fixed effects include month-of-year, day-of-week, and day-of-month fixed effects. Household controls include bill spike size, leak size, previous monthly bill size, relative standard deviation of the last 6 month bills, whether on autopay, number of occupants, and whether the household has had multiple leaks within the study period. Bill RSD is the relative standard deviation of the bill calculated as the standard deviation of the consumption amount for the past (pre-leak) 6 months divided by the consumption amount of the most recent pre-leak month. Big income is a dummy variable equal to 1 if income is above median of 100K dollars. Autopay is dummy variable equal to 1 if customer is on autopay. Standard errors are clustered at the zip code level, presented in parenthases. *p<0.1; **p<0.05; ***p<0.01

threshold). Meanwhile customers with a bill increase 50% and above statistically respond to leaks at the same rate as those getting a notification. Importantly, the median bill spike is a 40% increase. Thus, an informational notification is more effective than a bill increase for roughly half of customers (those receiving a bill increase of 40% or less), while both are similarly effective for the other half of customers.

Figure 1.2: Effect of a Bill Increase on Days Leaking
(Relative to Notification)

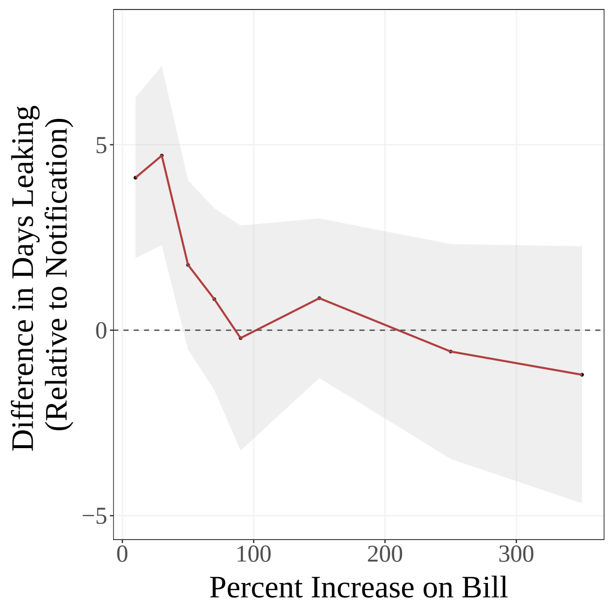

Notes: The points on the plot represent the coefficient estimates for each bill increase bin, plotted at the midpoint of the bin. Coefficients represent the change in days leaking for the associated bill increase relative to days leaking given an informational notification. The less days leaking, the more effective the bill increase. Coefficients above zero represent bill increases that are not as effective as notifications. Coefficients at zero represent bill increases that are as effective as notifications. The points are connected to emphasize the trend. The grey area represents the 90% confidence interval.

Why don't customers respond to lower, but still substantial bill increases?

The previous results indicate that on average a notification is more effective than a bill increase. Moreover, it takes a substantial bill increase of 50% or more for customers to be influenced by a bill spike as they would be by an informational notification. I next explore why customers do not respond to lower, but still substantial bill increases. A couple potential reasons come to mind. First, the customers' month-to-month consumption may have relatively large variation, making it difficult for customers to spot anomalies on their monthly bill, even if they are paying attention. Second, customers that are on autopay may not be paying close attention to the monthly bill, and are unlikely to notice an anomaly unless it is very glaring.[15] Third, customers may have high income levels such that even a substantial bill increase is not considered significant enough to warrant action.

I test how each of these three potential factors affect the effectiveness of a bill increase, relative to an informational notification, by introducing an interaction term into Equation 1.2. To test if consumers with large month-to-month bill variance respond differently to receiving a bill increase I interact the customer's bill relative standard deviation with D_i. Table 1.5, column (5) shows a positive and statistically significant coefficient on the interaction term (second row). This indicates that given the same bill increase, customers that have larger month-to-month variance take longer to fix a leak. Or, in other words, those with less variance are more affected by the bill increase to fix their leak faster.

Figure 1.3 helps solidify the findings visually. The graphic on the left demonstrates customers with low bill variance (i.e. relative standard deviation levels below the median amount of 20%), while those on the right demonstrate customers with high bill variance (i.e. relative standard deviation levels above the median). It is clear from the graphic that customers with low variance respond to lower bill spikes at the same rate as they respond to an informational notification, and even respond more to larger bill spikes than to informational notifications. Meanwhile, customers with high variance respond to bill spikes less than they do to informational notifications unless the bill spike is over a 50% increase. These results suggest that while the financial incentive may be large enough to encourage action if it is known about, perhaps customers are simply not noticing that something is unusual with their bill because they attribute the increase to natural bill variation.

I run similar analyses to test the effect of autopay and income on the response to a bill spike. It's worth noting, neither autopay nor income are correlated with a customer's bill-to-bill variance. Columns (6) and (7) of Table 1.5 show that neither factor statistically affects how customers respond to the bill spikes. Taken together, these findings imply that that the reason customers do not respond to lower, but substantial bill increases has little to do with autopay inattention or with customers having such high incomes that the increase is not considered significant. Instead, the lack of response is mainly due to customers simply not recognizing the bill increase to be an abnormal deviation from their average bill. This finding is consistent with the fact that the average relative standard deviation of the bill is 30%, while customers on average respond to bill increases of 50% or more.

[15]Customers not on autopay receive a monthly bill statement and need to take action to either pay their bill by mail, in person, or online. Customers on autopay do receive a monthly bill statement, but do not need to take any actions to pay their bill as their account is set-up to automatically charge their pre-specified payment method. These customers' bills are paid whether or not the customer looks at the monthly bill statement. It follows that customers with autopay may not pay as much attention to their bill amount as those who actively pay their bills.

Figure 1.3: Effect of a Bill Increase on Days Leaking by Bill Variance Size (Relative to Notification)

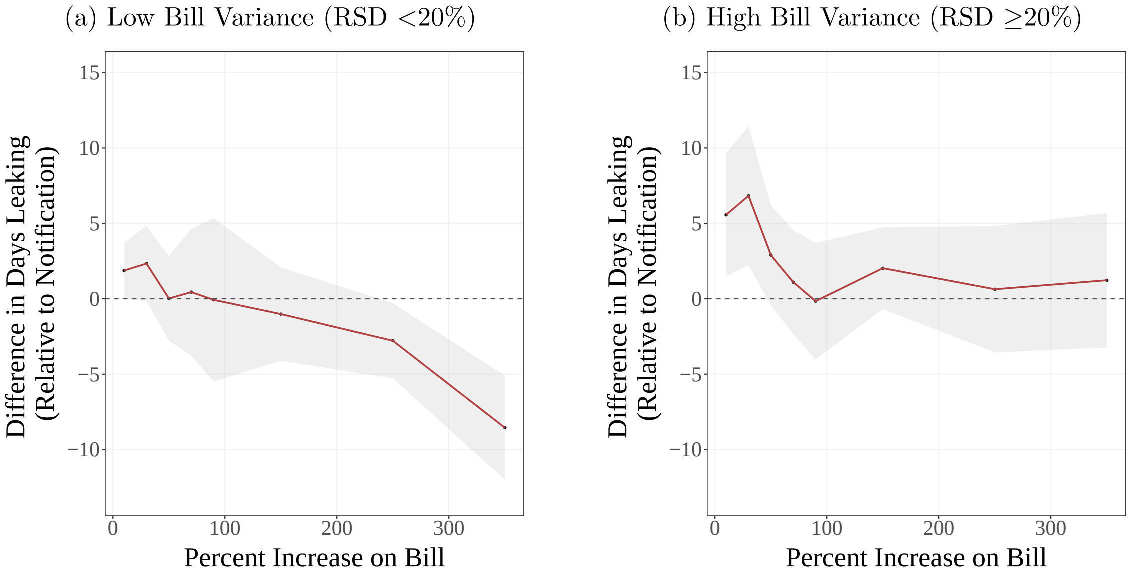

Notes: See notes for Figure 1.2. RSD is the relative standard deviation of the month-to-month bill calculated as the standard deviation of the consumption amount for the past (pre-leak) 6 months divided by the consumption amount of the most recent pre-leak month. The median RSD is 20%.

1.5.3 How do financial incentives affect customers that may already be aware of a leak?

Two-thirds of customers end up fixing their leaks in response to just a physical clue or an informational notification (without an explicit financial incentive). The remaining customers end up receiving a financial incentive through a bill increase. The majority of these customers receive an informational notification shortly before or after they receive a bill increase. Since the analysis above demonstrated that notifications are very effective in capturing attention, these customers are likely aware that they have a leak from the notification. The subsequent analysis is focused on how financial incentives motivate customers to fix leaks when they are likely already aware of the issue, rather than how financial incentives capture attention about the issue in the first place.

To assess how different levels of financial incentives affect days leaking I repeat the same estimation as in Equation 1.2 with a couple adjustments. This time I run the analysis on a subset of customers that receive both a bill increase and an informational notification before they fix their leak. I redefine D_i as a factor variable that takes on the bill increase bin value that customer i experiences.

Figure 1.4 demonstrates the resulting coefficients from the regression and the associated 90% confidence intervals. Similar to the previous figures, the coefficients for each bin show the difference in response time to the leak for customers with bill spikes in that bin range. Unlike in the previous figures, the response time here is relative to the response time for those that receive a bill increase in the 0-20% range. The results show that the greater the financial incentive, the greater the response (i.e. the less days leaking). Customers that receive a financial incentive respond roughly one day faster for every additional 25% increase to the bill.

Figure 1.4: Effect of Financial Incentive on Days Leaking
(Relative to 0-20% Bill Increase)

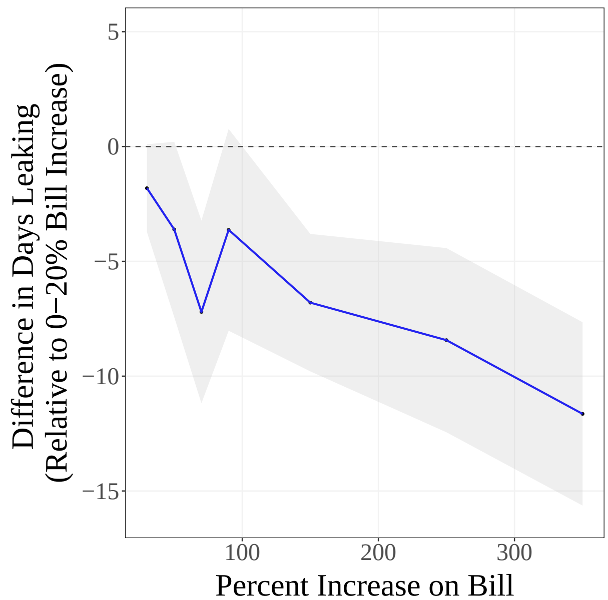

Notes: The points on the plot represent the coefficient estimates for each bill increase bin, plotted at the midpoint of the bin. Coefficients represent the change in days leaking for the associated bill increase relative to days leaking given a bill increase of 0-20%. The less days leaking, the more effective the bill increase. Coefficients at zero represent bill increases that are no more effective than a 0-20% bill increase. Coefficients below zero represent bill increases that are more effective than a 0-20% bill increase. The points are connected to emphasize the trend. The grey area represents the 90% confidence interval.

Do customers on autopay respond to financial incentives?

The literature indicates that customers on autopay are inattentive to their monthly bills (Sexton, 2015). To test the effect of financial incentives on autopay customers in the setting of fixing leaks, I rerun a similar estimation as above, this time interacting D_i with a dummy for whether a customer is on autopay. The results in Figure 1.5 reveal that customers on autopay do respond to financial incentives, but are less sensitive to bill increases than the customers who are not on autopay, particularly at low bill increases. It is unclear if this result is a product of the fact that customers on autopay are less attentive to the their bill, or if customers on autopay are a generally less financially sensitive demographic, which is why they may have signed up for autopay in the first place. Regardless, it is worth knowing that this subset of customers does in fact respond to financial incentives, but may require a higher bill increase to elicit a response.

Figure 1.5: Effect of Financial Incentive on Days Leaking by Autopay Status
(Relative to 0-20% Bill Increase)

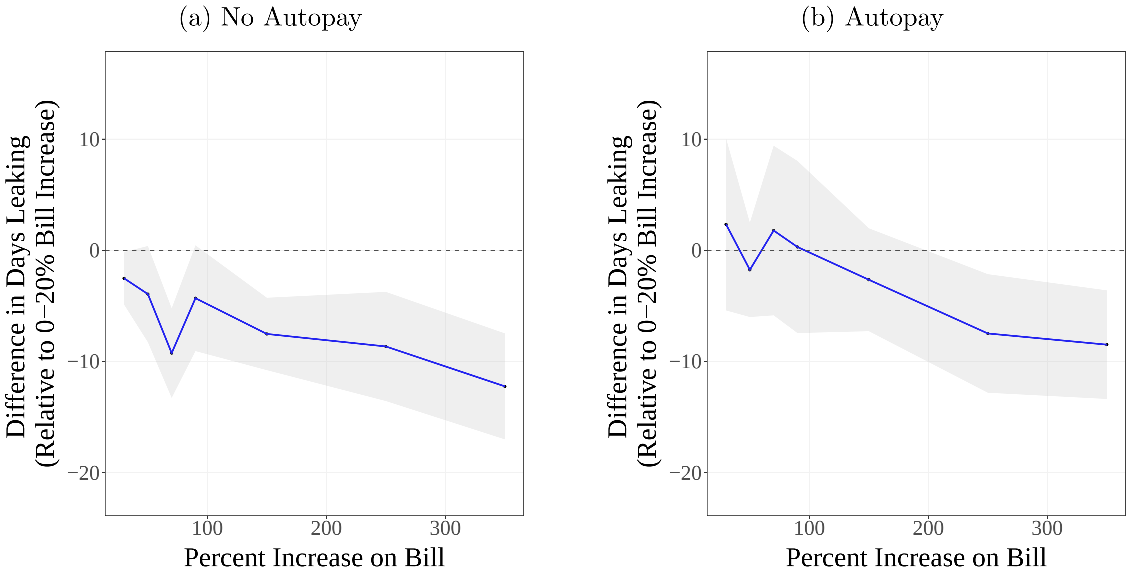

Notes: See notes for Figure 1.4. Roughly 15% of customers are on autopay.

Does the effect of financial incentives vary by income level?

To understand how the effect of financial incentives vary by income level, I repeat the above analysis, this time interacting D_i with a dummy variable indicating whether the customer has an annual household income above the median of \$100K. Figure 1.6 outlines the results. As expected, income level does play a role in the effectiveness of financial incentives. Customers with lower incomes are influenced by bill increases of all sizes, meanwhile customers with higher incomes are statistically influenced by bill increases above 150%. Customers of lower income levels may not only be more price sensitive, but also have a different approach to addressing the leak. The survey results show that the less wealthy are 40% more likely to try fixing the leak themselves, 20% less likely to call a plumber, and report investing more hours to fix the leak overall.

Figure 1.6: Effect of Financial Incentive on Days Leaking by Income Level
(Relative to 0-20% Bill Increase)

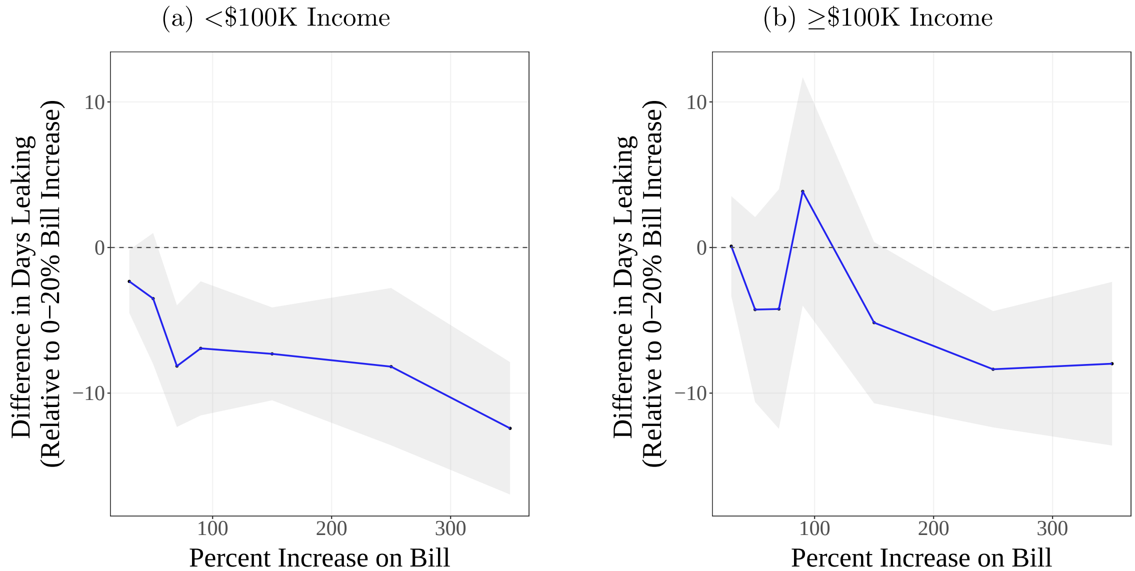

Notes: See notes for Figure 1.4. $100K is the median household income in San Francisco.

1.5.4 What is the most effective way to deliver information?

The analyses above have shown that informational notifications and financial incentives are effective in influencing behavior. I next test whether sending a blast of notifications is more effective than just sending a mailer notification.

Table 1.6 shows the raw Resource-Use and Follow-Through results for each experimental group. The first two rows present the average use of resources described in the notification messaging as first-steps to fixing a leak: contacting customer service and visiting the online account.[16] Customers in the Treatment Group are roughly twice as likely to use resources as those in the Control Group. The last rows of Table 1.6 show the Follow-Through results. Customers in the Control Group take over 20 days to fix a leak while customers in the Treatment Group take just under 17 days. Almost all customers in both groups eventually fix their leaks.

[16]The table does not include proportion of customers that visited the leak tips webpage because the tracking data was inconclusive.

Table 1.6: Raw Results by Experimental Group

	Control	Treatment	T-Stat	Sig (95%)
Proportion Contacted Customer Service	0.12	0.20	-4.96	*
Proportion Visited MyAccount Page	0.11	0.21	-2.28	*
Average of Total Days Leaking	20.41	16.83	2.28	*
Proportion Fixed	0.98	0.98	-1.21	

Notes: The table presents the mean for each variable by experimental group. (*) signifies a statistical difference in means across groups at the 95% level. The proportion visiting MyAccount page variable represents the customers that visit the landing page, but is not necessarly reflective of the proportion of customers that logged in. The proportion fixed variable is based on 6 months after the leak. Only a few customers in each group remain with unfixed leaks one year after the leak.

Since the Control Group and Treatment Group are randomly assigned and are balanced on observable characteristics, as shown in Table 1.3, just looking at the raw results tells most of the story. To rule out any additional variations in the two groups that might explain the changes in the total days leaking I estimate the following equation:

$$Y_{itz} = \beta T_i + \lambda_t + \mu_z + X_i^{'}\gamma + \epsilon_{itz} \tag{1.3}$$

T_i is a dummy that takes on zero for customers in the Control Group and one for customers in the Treatment Group. All other variables in the equation are the same fixed effects and household covariates as described in the previous estimations.

Table 1.7 shows the results for the estimated coefficient β. As before, column (1) shows a simple difference in means, columns (2)-(4) add on fixed effects and household controls. Column (4) is the preferred specification. The outcomes show customers that receive the treatment blast of notifications leak three days less than those who receive just a mailer notification. These results are statistically and economically significant, and robust to specification. Put in perspective, leaking three days less saves a customer roughly $25 on average. Column (5) reruns the preferred specification on the subset of customers that answered the survey. The results for the survey respondents are consistent with those of the entire sample, providing further evidence that the survey respondents are representative of the sample.

It is worth noting that Table 1.7 demonstrates the Intent-to-Treat (ITT) outcomes. These results show the average treatment effect across all customers in the Treatment Group, and do not account for the fact that only half of customers in the Treatment Group were sent multiple contact points (i.e. only half of customers in the Treatment Group are actually treated). Still, these results are important to estimate in order to understand the average treatment effect of sending a blast given real world constraints; all customer contact information is not always readily available, particularly for institutions that have been around for a long time.

It is, however, instructive to also estimate the Treatment-on-the-Treated (TOT) outcome, which shows that customers in the Treatment Group that actually receive multiple contact points leak over six days less than customers in the Control Group. This is equivalent to savings of $50 on average per customer with multiple contact points. This result gives the upper bound of the savings associated with sending a blast of notifications versus a mailer because this result is for a particular subset of customers that voluntarily had electronic contact points on file. This subset of people are probably more receptive to this mode of communication, so it is likely that the remaining customers would not respond as strongly to an electronic notification blast.

Table 1.7: Effect of Blast Notification Treatment on Days Leaking (Intent-to-Treat)

	Dependent variable:				
	Days Leaking				
	(1)	(2)	(3)	(4)	(5)
Treatment	-3.976^{***}	-3.539^{**}	-3.303^{**}	-3.254^{**}	-3.502^{*}
	(1.441)	(1.385)	(1.419)	(1.511)	(2.086)
Time Fixed Effects	No	Yes	Yes	Yes	Yes
Zip Fixed Effects	No	No	Yes	Yes	Yes
Household-Level Controls	No	No	No	Yes	Yes
Data Subset	All	All	All	All	Survey
Observations	1,889	1,889	1,889	1,732	469
R^2	0.003	0.036	0.053	0.105	0.197
Adjusted R^2	0.002	0.014	0.020	0.067	0.059

Notes: The dependent variable is days leaking. Time fixed effects include month-of-year, day-of-week, and day-of-month fixed effects. Household covariates include bill spike size, leak size, previous monthly bill size, relative standard deviation of the last 6 month bills, whether on autopay, number of occupants, and whether the household has had multiple leaks within the study period. Standard errors are clustered at the zip code level, presented in parenthases. *p<0.1; **p<0.05; ***p<0.01

Are some contact methods better than others?

I next analyze the effect of different types if contact methods (i.e. effect of robo-call, email, text). The experimental design is set up to test the causal effect of receiving two or more contact points, relative to just a mailer, on the amount of days leaking. The design does not explicitly test for the effect of the individual contact methods on days leaking since customers are not randomly assigned contact methods. However, I am able to estimate the average treatment effect of each contact method on customers that provide the contact method on file. In other words, I can compare outcomes for customers in the Treatment Group who have email addresses on file, and thus receive an email notification, versus outcomes for customers in the Control Group that also have email addresses on file, but are not sent an email notification. Such an analysis shows how the type of person who provides an email address would respond to an email notification. It does not necessarily extend to how customers who don't have email addresses on file would respond.

Still, this approach can provide some insight on the effect of robo-calls, emails, and texts in general. Specifically, the results of this analysis provide an upper bound to how all customers would respond to these electronic contact methods as those who currently have these contact points on file are likely to be more technologically savvy and respond the most. Figure 1.7 presents raw data on average days leaking by experimental group for each grouping of contact information available on file. Since only half of contact information on file for customers in the Treatment Group was actually used to send notifications (as described in Section 1.3.2), it is expected that if all contact points on file were sent notifications, then the decrease in average days leaking for the Treatment Group, relative to the Control Group, would be roughly double of what is depicted.

Figure 1.7: Plot of Days Leaking by Contact Information Available

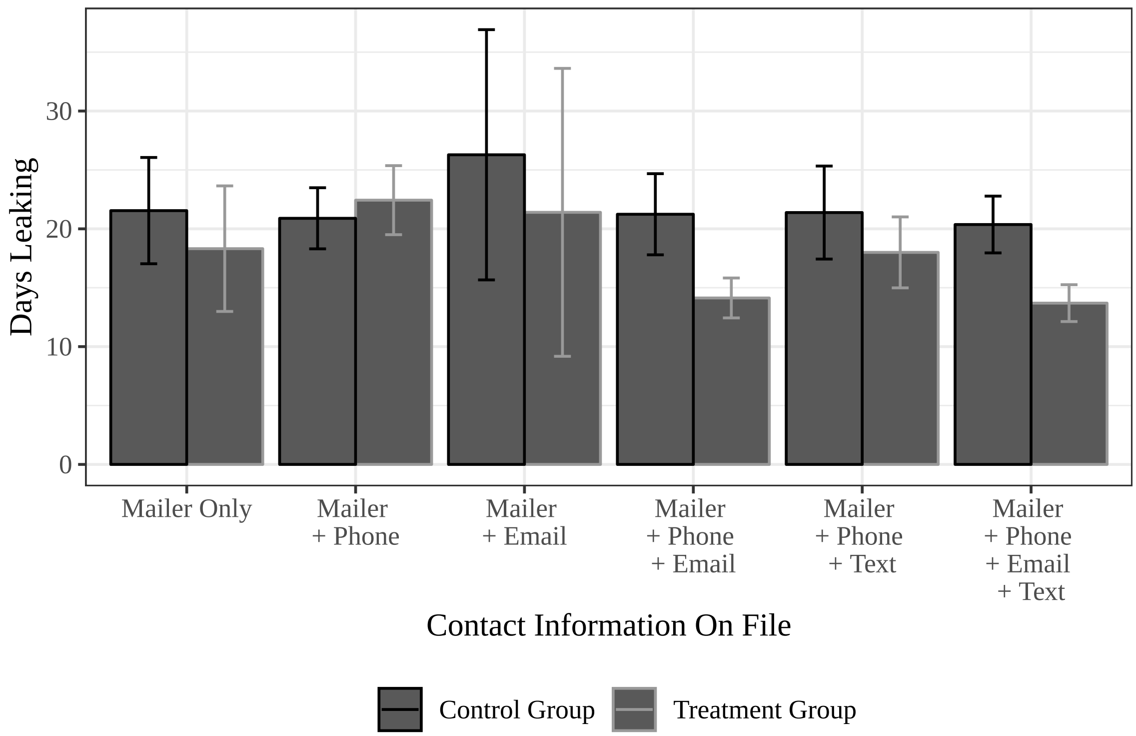

Notes: The graph represents average days leaking by experimental group given contact points on file. The error bars are the standard error of the mean.

To statistically estimate the effect of different contact types I re-estimate Equation 1.3 by interacting the treatment dummy T_i with a dummy of whether a customer has a particular contact point on file. I control for the other contact information the household has on file. This approach yields the causal effect of receiving the contact method on days leaking for the type of customers that would provide contact information for the particular method. As mentioned before, this analysis does not provide causal evidence on how all customers would respond to the type of contact method because the contact methods are not randomly assigned.

Table 1.8 shows the results of the analysis for the different contact methods. Column (1) shows that the effect of a robo-call on customers that provided a phone number to the water utility is not statistically significant, meaning robo-calls are not an effective notification mode for these customers. Column (2) shows that the effect of sending email notifications to customers that provided email contact information is statistically and economically significant at all confidence levels. Column (3) shows that the effect of sending text messages to customers that provide mobile phone numbers is statistically significant at the 10% level. Again, if all contact points on file were used, then the resulting coefficients would likely be double of what is depicted in the table.

Since these results show that robo-calls are not effective, even for customers that provided the contact information, this implies that robo-calls are unlikely to be effective on the remaining customers. Meanwhile, even though the email notifications are very effective on those who provided email addresses, this analysis cannot conclude that email notifications would be effective on all customers. That being said, as email use becomes even more ubiquitous, it is likely that the effect of email notifications on all customers will approach the effect found in the above analysis. A similar trajectory is likely for text message notifications.

Survey results support the theory that robo-calls are not very effective. Only 30% of survey respondents that received a robo-call remember receiving a robo-call. By contrast, 100% of survey respondents that received an email recall the email notification, 55% of respondents that received a text message recall a text notification, and and 66% of respondents that received a mailed letter recall receiving a mailed notification.

While this analysis cannot definitively conclude which contact method is most effective received on the same day, one thing is certain, the speed at which a notification is sent is critical in situations when behavior is time sensitive. Thus, electronic notifications that influence customers (i.e. emails and texts) have this advantage over mailers. In addition, electronic notifications may facilitate resource-use. As seen in Table 1.6, the rate of resource-use by customers is almost double in the Treatment Group. This is likely due to the email and/or text notifications as it is easier for customers to simply click on the link or phone number provided in an email or text message, rather than manually enter the information from a mailer.

When surveyed about preferences for future notifications, where customers could pick as many options as they liked, the majority of respondents requested emails, just under half requested a mailed letter, a third asked for text messages, and a quarter asked for a phone call. The preferences stated are in line with the results found in the analysis.

Table 1.8: Effect of Contact Information Type on Days Leaking

	Dependent variable:		
	Days Leaking		
	(1)	(2)	(3)
Treatment	−1.391	1.160	−0.963
	(6.241)	(2.427)	(2.483)
Phone	1.565	0.718	0.600
	(5.102)	(3.613)	(3.642)
Email	−0.695	3.473	−0.780
	(2.177)	(2.839)	(2.194)
Text	−0.005	−0.078	2.690
	(1.735)	(1.756)	(2.958)
Treatment*Phone	−2.167		
	(7.086)		
Treatment*Email		−8.463***	
		(2.972)	
Treatment*Text			−5.382*
			(3.211)
Time Fixed Effects	Yes	Yes	Yes
Zip Fixed Effects	Yes	Yes	Yes
Household-Level Controls	Yes	Yes	Yes
Observations	1,709	1,709	1,709
R^2	0.106	0.109	0.107
Adjusted R^2	0.065	0.068	0.066

Notes: The dependent variable is days leaking. Time fixed effects include month-of-year, day-of-week, and day-of-month fixed effects. Household covariates include bill spike size, leak size, previous monthly bill size, relative standard deviation of the last 6 month bills, whether on autopay, number of occupants, and whether the household has had multiple leaks within the study period. Standard errors are clustered at the zip code level, presented in parenthases. *p<0.1; **p<0.05; ***p<0.01

1.6 Implied Savings from the Leak Notification Program

The use of informational notifications to encourage customers to fix leaks faster yields significant water and monetary savings. Since all customers eventually receive an informational notification, it is not possible to do a direct comparison of customers that receive a notification to those that do not. However, it is possible to approximate the impact of the informational notification using the results presented in this paper. A conservative, back-of-the-envelope calculation, detailed in Appendix A.4, shows that average savings to customers from an informational mailer comes to roughly \$115 per leak. In contrast, the cost to the water utility of printing and sending mailer notifications is only \$2.74 per leak.[17] The analysis in Section 1.5.4 showed that introducing a blast of notifications, rather than just a mailer, saves customers an additional \$25 per leak on average. As a comparison, the water utility incurs a maintenance and communication cost of \$3.22 per leak (in addition to the aforementioned mailing costs) for the system that sends electronic notifications. Table 1.9 summarizes all the costs and savings.

Taken together, the customer savings associated with providing an electronic informational notification blast, relative to just sending a bill, comes to \$140 per leak. This figure is expected to grow as the water utility collects more contact information per customer, allowing the blast to reach more people. The total cost to the water utility of sending informational notifications is \$6 per leak. There are also costs to the customer associated with time and monetary investments required to fix the leak. It is worth noting, though, that unlike with most other conservation efforts, customers do not receive utility from the leaky water, therefore curbing leaks does not incur additional welfare costs from not consuming the water. In fact, there may be additional benefits of fixing leaks faster by avoiding potential housing damage costs.

1.6.1 Aggregate Savings

The aggregate savings of providing information to customers are economically significant. Absent any informational notifications, water attributed to leaks across all single-family residential customers would add up to 82 thousand ccf (i.e. 188 acre-feet or 61 million gallons) per year, representing 1.2% of all water demanded by single-family residential accounts.[18] On aggregate, this costs customers over \$1.64 million annually. It is important to note that SF Water's leak detection system currently detects only relatively large leaks, >7.5 gal/hr, and since over 50% of leaks are at or just barely above the minimum detected amount, it is likely there are just as many leaks right below the cutoff that are unaccounted for. As a result, when factoring in unaccounted for leaks, leaks in total represent even more of total single-family residential water demand and cost customers substantially more on aggregate.

Sending informational notifications via a mailer saves customers roughly \$590 thousand

[17] This is the average cost of all mailers sent to a customer for a leak. The cost of printing and sending one mailer is \$1.41 per mailer.

[18] The total water demand is based on 2016 billing data. The calculation assumes 100 new leaks per week.

Table 1.9: Summary of Savings and Costs

		Total With No Notifs	Change With Mailer	Additional Change With Blast
Per Leak	Water Loss (ccf)	16	-6	-1
	Water Loss (gal)	11,968	-4,488	-748
	Water Loss Cost ($)	320	-115	-25
	Program Cost ($)		2.74	3.22
Aggregate Annual	Water Loss (ccf)	82,100	-29,500	-6,700
	Water Loss (gal)	61,396,800	-22,083,200	-4,992,500
	Water Loss Cost ($)	1,639,200	-589,600	-128,500
	Program Cost ($)		14,000	16,500
	Program Cost ($/ccf)		0.47	2.47
	% of SFR Demand	1.2%	-0.4%	-0.1%

Notes: The cost of water is based on fiscal year 2017-2018 water and wastewater rates. Program costs are based on realized costs in fiscal year 2017-2018 and are subject to change. One ccf is equal to 748 gallons.

annually with an aggregate printing and mailing cost to the water utility of $14 thousand per year. Sending informational notifications via a blast saves customers an additional $129 thousand annually with an aggregate maintenance and communication cost to the water utility of $16.5 thousand per year.[19] Taken together, sending informational notifications via a blast, relative to just sending a bill, saves customers a total of $0.72 million and 36 thousand ccf (i.e. 83 acre-feet or 27 million gallons) of water annually, while costing the water utility $0.03 million annually. These findings show that by implementing a blast notification system, SF Water is able to save almost half of the water associated with single-family residential in-home leaks. The water savings is equivalent to 0.5% of total single-family water demand, which is enough water to fill 41 Olympic sized swimming pools or cover the annual water demand for almost 600 single-family households.

These annual aggregate savings are expected to grow over time for three reasons. First, as mentioned before, the average savings are expected to increase as the water utility collects more contact information and uses all the information they have on file. Second, as SF Water begins to detect smaller leaks, the water utility will be able to identify additional customers that will benefit from an informational notification. As described above, this is likely to be at least 50% more customers, which would translate into less than a 50% increase in aggregate savings since these customers have smaller leaks. Third, since single-family residential water demand represents only 20% of all water demand serviced by SF Water, savings from the notification program will grow as the water utility expands the notification system to their multi-family customers that experience leaks.[20] While multi-family customers represent even more of overall water consumption at 30% of SF Water demand, it is not clear if response

[19]There is a fixed development cost of $166 thousand associated with setting up the system to send the electronic notifications.

[20]20% of demand is single-family, 30% multi-family, 40% non-residential, and 10% water loss. Source: SFPUC 2015 Urban Water Management Plan, page 4-3.

rates to the notification system will be as high. This is because multi-family customers have one meter for several households, making it harder to detect the source of the leak, as well as potentially diffusing the financial incentive to address the issue.

1.6.2 Cost Effectiveness

The leak notification program is certainly cost effective. Introducing a mailer notification costs the water utility \$0.47 per ccf saved, relative to no notification. The blast notification costs the water utility \$2.47 per ccf saved, relative to the mailer. The cost per ccf for the blast system will decrease as the water utility reaches more customers with leaks while paying a similar annual maintenance and communication system cost. In aggregate, the blast notification system, relative to no notification, costs the water utility \$0.84 per ccf saved (i.e. \$367/acre-foot). To put this figure in perspective, the weighted average cost of some of the other SF Water conservation programs is \$1.60 per ccf saved (i.e. \$697/acre-foot).[21] This weighted cost is almost double that of the leak notification program. The conservation programs that have saved the most water are toilet and clothes washer rebates that cost \$1.83/ccf and \$1.23/ccf, respectively. For reference, the production cost of water for SF Water is \$0.20 per ccf[22], and the cost of water to customers is \$19.97 per ccf.[23]

1.7 Conclusion

This paper is one of the first in the literature that directly compares informational notifications and financial incentives while yielding results with strong external validity. Relying on both a natural experiment and a randomized field experiment, this study provides empirical evidence that information can be a powerful and low-cost tool to influence behavior, and could be even more effective than financial incentives when dealing with inattentive customers.

The findings in this paper are critical for any entity with interest in influencing the behavior of recurring customers – from a utility looking to spur conservation, to a technology company hoping to encourage new product adoption. The key take-away is that a bill change will likely go unnoticed unless it is at least as large as one standard deviation of the bill-to-bill variation. For entities trying to capture customer attention, particularly those with high bill-variance customers, this finding suggests that an informational notification may be preferable. This is especially the case if large price changes are not cost-effective or face political and regulatory barriers. When possible, sending a notification that explicitly states

[21]This statistic includes audits & reports, reuse incentives, and showerhead, toilet, and washer rebate programs. This statistic does not include conservation programs that are focused on disseminating conservation informational materials or outreach efforts. Source: SFPUC 2015 Retail Water Conservation Plan, page 37.

[22]This figure includes power and chemical costs at the water treatment plant, and power used for pumping in the City distribution system. A known missing marginal cost is the water transfer cost from moving water through Hetch Hetchy Regional System. Source: SFPUC AWWA Water Loss Audit spreadsheet for FY 17-18 (publicly available through DWR Water Use Efficiency portal).

[23]Based on SF Water water and wastewater rates FY 17-18.

a financial incentive is optimal as it could both capture the customers' attention and provide financial incentive.

Conversely, if a profit maximizing entity is trying to raise prices without customers noticing, it should keep price increases just under one standard deviation of bill-to-bill variation. Further, since customers with higher incomes and those on autopay are shown to be less sensitive to price changes, these attributes should be taken into consideration. While many entities, such as ride-sharing companies and online retailers, have already embraced the approach of quietly raising prices when possible, and sending notifications to advertise price cuts when needed, having a deeper understanding of how customers are responding to the treatments is valuable for optimal price changes and customer targeting.

This research has environmental and economic significance as well. Providing information saves single-family customers in San Francisco 27 million gallons of water and $0.7 million annually. These aggregate savings are expected to at least double in the next few years as the water utility is able to send informational notifications to more customers experiencing leaks. Further, the use of informational notifications to encourage leak fixing is among the most cost-effective approaches to water conservation used by SF Water, is a sustainable source of water conservation, and is welcome from the customer standpoint. Considering SF Water is just one of hundreds of water utilities in California, the potential environmental and monetary impacts of providing information to even a fraction of customers state-wide would be of high consequence in this drought prone region.

Chapter 2

Flooding Your Conscience – The Effect of Public Appeals on Residential Water Conservation

2.1 Introduction

Traditionally, economists look to market-based incentives to induce behavioral change. While in many situations a market-based solution may be optimal, it is not always allowed due to prohibitive legislation, or it may be simply politically unpalatable (Olmstead and Stavins, 2009). Given these real-world restrictions, other incentives for behavioral change have come to the forefront of research. One of the most popular concepts of non-market based incentives is the idea of Libertarian Paternalism (i.e. Nudge Theory) where positive reinforcement and indirect suggestions are used to alter people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives (Thaler and Sunstein, 2003).

Closely related to this concept is the idea of "moral suasion" (i.e. appealing to morality in order to influence or change behavior). In the context of public goods, where it is easy to free-ride, moral suasion can play an important role. A large number of laboratory experiments have shown that when it comes to contributions to public goods, the self-interested hypothesis is rejected (Ledyard, 1995). This suggests that there is room for morality to be used as a lever to incentivize agents to change their behavior when it is in their interest to free-ride. Some explanations for why morality has an impact on agents' actions are altruism (Becker, 1974), the warm-glow effect (Andreoni, 1989), conditional cooperation (Fischbacher et al., 2001), and reciprocity (Sugden, 1984).

This paper assesses the importance of moral suasion in the setting of residential water consumption during a drought. Droughts are a reoccurring feature of California's climate. The most recent drought of 2011-2017 proved to be particularly severe with the three-year period between fall 2011 and fall 2014 as the driest since recordkeeping began in 1895. Making matters worse, year 2014 was the warmest year on record (PPIC Water Policy Center, 2015). To prepare for water shortages associated with the drought, the state had been taking actions to induce water conservation including regulations, rationing, and public appeals (i.e. moral suasion). Market-based incentives are traditionally not used to incentivize water conservation (Olmstead and Stavins, 2009). In the case of California, there have been regulations against conservation pricing, and only just recently in 2015 has there been a call on local water agencies to adjust their rate structures to implement conservation pricing.[1]

Several years into the 2011-2017 drought, Governor Edmund G. Brown Jr. made several public appeal announcements pleading for Californians to do their part in conserving water during this critical time. This paper estimates the short-term impacts of these public appeals on single-family residential (SFR) water consumption. As a case study, I analyze the water consumption of the SFR customers serviced by San Francisco Water Power Sewer (hereinafter "SF Water"). Using high frequency hourly consumption data at the household-level for the years 2012-2015, I conduct an event study relying on time series variation. The analysis looks at three distinct public appeals in the years 2014 and 2015, and assesses each of the event's effect on the water consumption for the customers in the data.

Controlling for weather, prices, holidays, day-of-week trends, seasonal trends, and household specific time invariant characteristics, I find that there was a statistically and economically significant decrease in water consumption following the two announcements that were successful in producing high drought awareness. Specifically, I find that the Governor's January 2014 declaration of a drought State-of-Emergency was associated with an average decrease in daily water consumption of 1.9% within two weeks of the announcement. I also find an average decrease in daily water consumption of 3.8% within the first week (4.6% within two weeks) following the Governor's April 2015 Executive Order of mandated 25% reductions.

The paper proceeds as follows: Section 2.2 reviews the literature. Section 2.3 details the events of interest and Section 2.4 describes the data. Section 2.5 explains the empirical method. Sections 2.6 and 2.7 present and discuss the results, respectively. Section 2.8 concludes.

2.2 Literature

In practice, public appeals (i.e. moral suasion) are used all the time[2], yet there is little evidence of how effective real-world public appeals really are (Reiss and White, 2008; Dal Bo and Dal Bo, 2014). Reiss and White (2008) analyze the effects of a 6-month public appeals media campaign in San Diego during the California energy crisis in the early 2000s. They find a steady decrease of energy consumption of 7% over a 6-month period, suggesting "well-orchestrated mass public appeals can be an effective means to avert rationing when the price mechanism is unable (or, in this instance, not permitted) to equilibrate the market" (Reiss and White, 2008). Ito et al. (2018) analyzes energy demand in Japan and finds that

[1] A call for conservation pricing came as part of the April 2015 Executive Order B-29-15. Details found at: `https://www.gov.ca.gov/news.php?id=18913`

[2] Examples: President Jimmy Carter urged Americans to reduce oil use with Oval Office broadcasts, California issues "Spare the Air" declarations on smoggy days that encourage people to refrain from driving, and local utilities broadcast appeals for consumers to shut off air conditioners in hot weather when electricity supplies are tight (Reiss and White, 2008).

while moral suasion was not as effective as economic incentives, it did have an economic and statistically significant impact on energy consumption in the short-run.

Allcott (2011) looks at the famous Opower experiments to understand the effects of moral suasion on energy consumption through the mechanism of social comparison, finding that non-price interventions can substantially and cost effectively change consumer behavior, decreasing energy consumption by roughly 2% on average. Ferraro and Price (2013) do a similar analysis in the water sector assessing the impacts of moral suasion on water consumption through different types of messaging, finding that social comparison messaging had a greater influence on behavior than simple pro-social messaging or technical information alone.

This paper contibutes to the literature of the effects of moral suasion by analyzing the short-term impacts of a public appeal. This analysis differs from that of Ito et al. (2018), Allcott (2011), and Ferraro and Price (2013) in that it is assessing the affect of moral suasion that reaches customers indirectly through media rather than through direct targeted messaging that comes straight to the customer. While this paper is most similar to that of Reiss and White (2008), it differs in that this analysis looks at short-term effects of an isolated public appeals announcement rather than the long-term effects of a 6-month public appeals campaign.

2.3 Events of Interest

This analysis looks at three distinct public appeal announcements made by Governor Brown in 2014 and 2015. The following subsections describe the contents of these appeals and SF Water's responding actions. To understand how these events may have impacted drought awareness, I examine historical Google search volume intensity as found on Google Trends.[3]

2.3.1 Event 1: Drought State-of-Emergency Declaration

On Friday, January 17, 2014, Governor Brown declared a drought State-of-Emergency.[4] Snowpacks had fallen to roughly 20% of normal, and it had been projected that 2014 would be the driest year on record. In addition, California's largest water reservoirs had very low water levels for this time of year, California's major river systems had significantly reduced surface water flows, and groundwater levels throughout the state were notably depressed. As part of the announcement, the Governor called on Californians to voluntarily reduce their water use by 20%.

While the State-of-Emergency announcement was not the first action taken by the Governor to address the impending water shortage[5], it was the first largely publicized announce-

[3]Google Trends is a Google database that keeps track of historical trends of Google search volume for specific terms.

[4]The State-of-Emergency Declaration details can be found here: `https://www.gov.ca.gov/news.php?id=18368`

[5]This declaration had followed a series of actions that the administration had taken to ensure that California was prepared for record dry conditions including: issuing and Executive Order in May 2013 to direct state water officials to expedite the review and processing of voluntary transfers of water and water rights; and forming a Drought Task Force in December 2013 to review expected water allocations, California's

ment that appealed to the everyday consumer. Governor Brown:

> *"We can't make it rain, but we can be much better prepared for the terrible consequences that California's drought now threatens, including dramatically less water for our farms and communities and increased fires in both urban and rural areas. I've declared this emergency and I'm calling all Californians to conserve water in every way possible."*

On the day of the annoucement, SF Water released a statement on Governor Brown's declaration. Two weeks later, SF Water had their first big announcement to the press and public in regards to conservation actions, calling for 10% voluntary system-wide water reduction.[6]

Governor Brown's State-of-Emergency announcement was followed by a huge spike in Google searches for the word "drought" and phrase "water conservation" by people located in the San Francisco Bay-Area (see Figure 2.1). According to Google Trends, the word "drought" was searched 10 times more during the week of the State-of-Emergency announcement than it had been in any week prior, even though the drought had already been happening for over two years at this point. An extended historical look at Google Trends shows that the week of this announcement was the greatest search of the word "drought" in the San Francisco Bay-Area on record.

The search trends for both "drought" and "water conservation" reveal similar patterns of sharp increases in the first two weeks after the announcement, followed by steady decline in the latter two weeks. Unlike the search trend for "drought", the "water conservation" search trend peaks at the second week rather than the first week. Perhaps the initial news of the announcement sparked interest in the topic of the drought, but it took a little extra time for people to digest what the drought implied and turn it into the search for information on water conservation. This suggests there may be a delay in the action of water conservation after the initial announcement.

preparedness for water scarcity and whether conditions merit a drought declaration.

[6]Timeline of SF Water response to the Governor announcements: `http://www.sfwater.org/modules/showdocument.aspx?documentid=7228`

Figure 2.1: Google Search Interest Over Time in the San Francisco Bay-Area

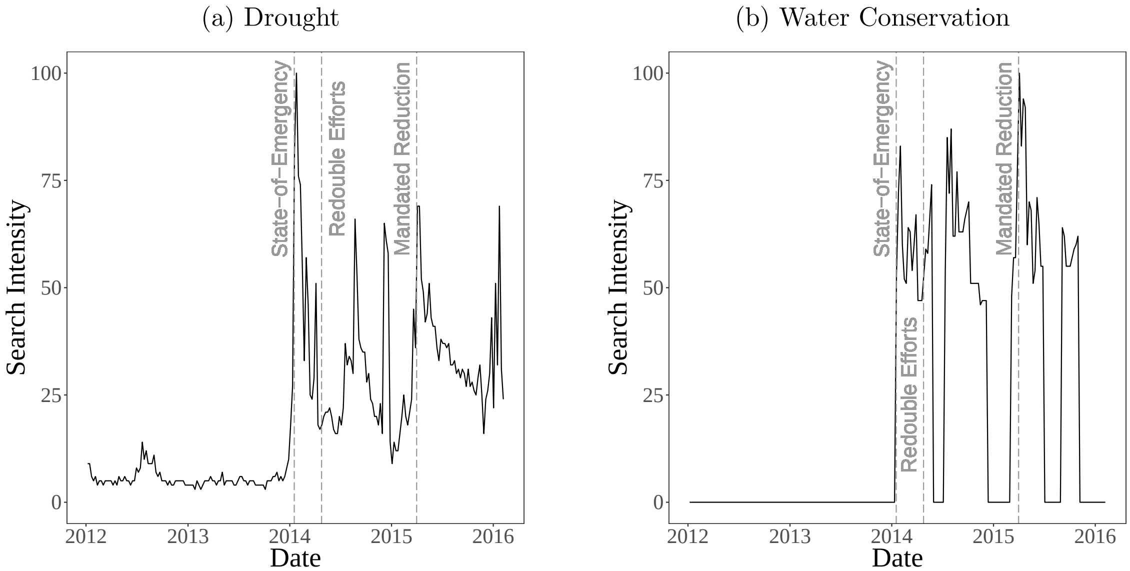

Notes: These graphics are taken from Google Trends. Search intensity is defined as the number of searches on that day divided by maximum searches in any one day on record, multiplied by 100. Searches are reported weekly. Graphic reports data on end of week date. The two graphics are not directly comparable to each other since the y-axis is in percentage terms. The magnitude of searches of "water conservation" is much smaller than for "drought". When search levels are low, Google may report the search levels as zero, explaining the dramatic drops to zero in the "water conservation" graphic. The grey dotted lines represent the dates of the public appeals.

2.3.2 Event 2: Executive Order to Redouble State Drought Efforts

On Friday, April 25, 2014, Governor Brown issued an executive order of continued State-of-Emergency. The purpose of this order was to strengthen the state's ability to manage water and habitat effectively in drought conditions, and called on all Californians to redouble their efforts to conserve water.[7] While much of the order is focused on water management improvements, the order also calls on Californians and California businesses to take specific actions to avoid wasting water, including limiting lawn watering and car washing.[8] Governor Brown:

> *"The driest months are still to come in California and extreme drought conditions will get worse. This order cuts red tape to help get water to farmers more quickly, ensure communities have safe drinking water, protect vulnerable species and prepare for an extreme fire season. I call on every city, every community,*

[7]The Redoubling Executive Order details can be found here: `https://www.gov.ca.gov/news.php?id=18496`

[8]The order also recommends that schools, parks and golf courses limit the use of potable water for irrigation; and asks that hotels and restaurants give customers options to conserve water by only serving water upon request and other measures. The order also prevents homeowner associations from fining residents that limit their lawn watering and take other conservation measures.

every Californian to conserve water in every way possible."

This event proved to spark very little interest in terms of Google searches. Since there was no increase in search volume for the word "drought" in the San Francisco Bay-Area following the announcement, it is reasonable to assume that not many people were aware of the appeal. A discussion with SF Water revealed that this announcement was not considered extraordinary from the perspective of the utility and was likely even less noticed by the public. SF Water did not send out any press releases regarding this announcement in particular.[9] There is speculation that this event may have been overshadowed by a media dialogue switch that had happened around that time to the politically charged discussion of water right curtailment.[10] This evidence suggests that one should not expect a change in water consumption in response to this event.

2.3.3 Event 3: Executive Order of 25% Statewide Mandatory Water Reductions

On Wednesday, April 1, 2015, Governor Brown issued an executive order of the first ever statewide mandatory water reductions.[11] This order directed the State Water Resources Control Board to implement mandatory water reductions in cities and towns across California to reduce water usage by an average of 25 percent, relative to 2013 consumption levels. The State Water Board assigned ranging mandatory rationing levels across the utilities in California such that utilities with high gallons per capita daily (GPCD) consumption (e.g. City of Beverly Hills with 235.9 GPCD) would have to conserve a higher percentage of 36%, in comparison to a low consuming utility (e.g. SF Water with 45.4 GPCD), which would have to conserve 8% in its retail service area.[12] Utilities with higher GPCD are required to conserve at higher levels since they have more room for low cost conservation. Once the reduction levels were assigned by the Board (roughly three weeks after the initial announcement), it was up to the utilities themselves to instill regulations to achieve their specific goals. Governor Brown:

"Today we are standing on dry grass where there should be five feet of snow. This historic drought demands unprecedented action. Therefore, I'm issuing an

[9]Timeline of SF Water response to the Governor announcements: `http://www.sfwater.org/modules/showdocument.aspx?documentid=7228`

[10]Water rights curtailment news: http://www.sacbee.com/news/local/article2597560.html

[11]The order also included a series of actions to help save water, increase enforcement to prevent wasteful water use, streamline the state's drought response and invest in new technologies that will make California more drought resilient. The Executive Order details can be found here: `https://www.gov.ca.gov/news.php?id=18913`

[12]The 2013 baseline consumption was calculated off of consumption during 9 months of the year (Jan, Feb, June - Dec). For specific calculations and rationing distributions visit: `http://www.waterboards.ca.gov/waterrights/water_issues/programs/drought/docs/emergency_regulations/supplier_tiers_20150428.pdf`

executive order mandating substantial water reductions across our state. As Californians, we must pull together and save water in every way possible."

Four weeks after the Governors announcement, SF Water imposed additional water use restrictions consistent with the State Water Board's emergency regulations.[13]

The search intensity for "drought" and "water conservation" spiked again, with "water conservation" hitting an all time high in the week following this announcement. Unlike with the State-of-Emergency announcement, there was no delay in search intensity for "water conservation". One explanation for the more immediate spike in "water conservation" searches is that this announcement was not just a declaration of the state-of-affairs, but a clear actionable Executive Order. This suggests that consumers may have acted more quickly in response to this event than they did in response to the State-of-Emergency event.

2.4 Data and Summary Statistics

SF Water has 180,000 accounts throughout its retail service area across all sectors (single-family residential, multi-family residential, commercial, instituational, and irrigation). Roughly 100,000 of these accounts are single-family residential. While the research question of interest is applicable to all sectors, this analysis is focused on the single-family residential sector. I chose to focus on residential consumers as water is purely a consumption good in this sector, while water may also be considered an input in the other sectors, complicating the analysis. I focus specifically on the single-family accounts within residential as these accounts represent the overwhelming majority of the residential accounts, and are metered on the household level. Multi-family accounts may have many households per meter, making it more difficult to estimate precise changes in consumption on the household level.

This analysis relies on 2,500 randomly sampled accounts for the SFR sector (roughly 2.5% of all SFR accounts). The data contains hourly water consumption reads for each account for the years 2011-2015, as well as number of occupants, zip code of the account, final consumption level billed per billing cycle, and dollar value billed. Consumption is measured in cubic feet (cf), where one cubic foot is equal to 7.48 gallons.

SF Water began installing smart meters in 2011 and finished deployment in 2013. As a result, only a small subset of the accounts had automated meters that transmitted hourly data in 2011. Given this limitation, this analysis only uses data starting in 2012 when a substantial percentage of the accounts in the sample had automated meters. For the purposes of this analysis I aggregate household consumption to the daily level.

Table 2.1 shows summary statistics by year for the main variable of interest (daily water consumption in cubic feet), as well as the average price of water per cubic foot, the marginal price of water per cubic foot, the average number of occupants per household, average maximum daily temperature in Fahrenheit, average daily precipitation in inches, and number of accounts.

[13]Timeline of SF Water response to the Governor announcements: `http://www.sfwater.org/modules/showdocument.aspx?documentid=7228`

Table 2.1: Average Covariates by Year

	2012	2013	2014	2015
Average Daily Consumption (cf)	18.46	18.52	17.46	16.52
Average Price of Water ($/cf)	0.13	0.14	0.15	0.17
Marginal Price of Water ($/cf)	0.15	0.16	0.17	0.18
Number of Occupants	3.35	3.35	3.35	3.35
Maximum Daily Temperature (°F)	64.95	66.01	69.68	67.89
Daily Precipitation (in)	0.06	0.01	0.06	0.02
Number of Accounts	2,105	2,398	2,392	2,395

Notes: One cubic foot (cf) is equal to 7.5 gallons. Single-family residential rate schedule consists of a fixed charge and a two tier volumetric block rate with one price for the first 4 hundred cubic feet of water, and a higher price for any additional units consumed per month. Average price is calculated as household specific total dollars billed divided by total cubic feet consumed. Marginal price is taken as the second tier rate since the average household consumes over 5 hundred cubic feet per month. Average and marginal price of water includes water and associated sewage costs. For current prices and further explanation visit: `http://www.sfwater.org/index.aspx?page=168`. The number of occupants is from the water district billing data. The temperature and precipitation data is from Weather Underground (weather station: KSFO): `https://www.wunderground.com/us/ca/san-francisco`.

Average daily consumption decreased in the latter two years, with a decrease of over 10% from 2012 to 2015. The average and marginal price of water has steadily increased over this time frame with increases ranging from 8.8-10.6% annually and 6.4-8.2% annually, respectively.[14]. The average number of occupants remains stable at 3.35 over all years. The maximum daily temperature peaks in 2014, and the driest year is in 2013.

Tables 2.2, 2.3, and 2.4 show the summary statistics for the 14 days before and after each event. Prices do not change from the before to after period of any of the events. Temperature remains constant for Event 1 across both periods. Temperature increases by less than 4 degrees for Event 2, and decreases by 4.5 degrees for Event 3. While these are statistically significant changes in temperature, they are unlikely to be economically significant changes. Precipitation remains constant for Event 1 across the before and after period, and increases for both Events 2 and 3, but is still relatively low. The number of accounts remains generally constant before and after for all three events.

[14]Rate increases are implemented annually on July 1st. The displayed rates in the table are averages of the rates in the first half of the year and the new rate implemented for the second half of the year.

Table 2.2: Average Covariates for Event 1

	14 Days Before Event	14 Days After Event
Average Daily Consumption (cf)	18.15	17.91
Average Price of Water ($/cf)	0.15	0.15
Marginal Price of Water ($/cf)	0.16	0.16
Maximum Daily Temperature (°F)	64.23	64.58
Daily Precipitation (in)	0.00	0.00
Number of Accounts	2,361	2,364

Notes: See notes for Table 2.1.

Table 2.3: Average Covariates for Event 2

	14 Days Before Event	14 Days After Event
Average Daily Consumption (cf)	17.67	17.82
Average Price of Water ($/cf)	0.15	0.15
Marginal Price of Water ($/cf)	0.16	0.16
Maximum Daily Temperature (°F)	67.24	71.09
Daily Precipitation (in)	0.00	0.03
Number of Accounts	2,364	2,364

Notes: See notes for Table 2.1.

Table 2.4: Average Covariates for Event 3

	14 Days Before Event	14 Days After Event
Average Daily Consumption (cf)	17.15	16.44
Average Price of Water ($/cf)	0.16	0.16
Marginal Price of Water ($/cf)	0.18	0.18
Maximum Daily Temperature (°F)	69.43	64.92
Daily Precipitation (in)	0.00	0.05
Number of Accounts	2,356	2,361

Notes: See notes for Table 2.1.

Figure 2.2 demonstrates the mean daily consumption averaged over all the households in the sample. This figure shows the general trend of consumption over time. The vertical dashed lines represent the 3 events central to the analysis.

Figure 2.2: Daily Household Consumption Trend

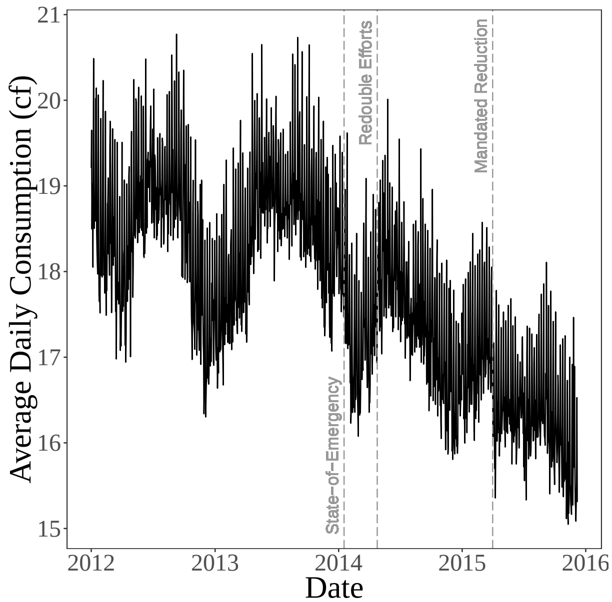

Notes: The graphic shows household daily consumption averaged across all the households in the study sample. The grey dotted lines represent the dates of the public appeals.

There is a clear trend of average consumption declining in the latter two years as compared to the former two years. The general decline in consumption cannot be explained by price increases alone. Since elasticity of demand for SF Water is -0.182 (Buck et al., 2016), and there is a decrease of about 11% in consumption from 2013 to 2015, prices would have to increase by almost 60% to explain the trend. Prices only increase by roughly 20%, explaining about 1/3 of the declining trend.

Keeping in mind the general declining trend, there is a visually detectable drop at the State-of-Emergency event (first line) and even a more clear drop at the Mandatory Reduction event (third line). There does not seem to be any obvious break in average consumption around the Redouble Efforts event (second line). In the following sections I analyze these drops to understand if there is in fact a significant decrease in consumption in the weeks following the event that cannot be explained by any other factors other than the event itself.

2.5 Empirical Method

To estimate the effect of a public appeal on water consumption, I conduct an event study using time-series variation. Specifically, I am interested in understanding the average change in water consumption during the time period following the appeal as compared to the average water consumption in the time period preceding the appeal.

Not all households may have been aware of the appeal right away. Some households may have seen the Governor's announcement live, while others may have learned of the news from a co-worker weeks later. Thus, for each event, I define an event window length following the announcement that is assumed to represent the time span for which most households are likely to have been informed of the public appeal. A discussion of the event window size comes later in this section.

Given a defined event window size (e.g. 2 weeks following the announcement), I am interested in comparing consumption in this event window to the consumption in a window of the same length directly preceding the announcement. The combination of these two windows, before and after the announcement, are from here on out referred to as the *analysis time frame*.

I define the estimating equation as follows:

$$ln(q_{it}) = \beta_0 + \beta_1 event_t + \delta_1 pre_t + \delta_2 post_t + \theta_1 temp_t + \theta_2 precip_t + \lambda price_t + \gamma^{'}\mathbf{X_{it}} + \epsilon_{\mathbf{it}}$$

The dependent variable, $ln(q_{it})$, is the natural log of water consumption by household i in day t. The explanatory variable of interest is the event indicator $event_t$, equal to 1 if the time interval t falls in the defined event window following the announcement. Since the dependent variable is in logs, the coefficient of interest, β_1, can be interpreted as a percent change.

The specification also includes two dummy variables, pre_t and $post_t$, indicating the time interval t that happens before the *analysis time frame*, and another indicating the time interval t that happens after the *analysis time frame*. These dummies allow me to isolate the trends in the *analysis time frame* from the rest of the data. By isolating the *analyis time frame* I am able to interpret the coefficient of interest as the change in consumption in the second half of the *analyis time frame* (i.e. days following the appeal) as compared to the consumption in the first half of the *analyis time frame* (i.e. days preceding the appeal). I keep the remaining *pre* and *post* data in the analysis for more accurate estimates of the controls and fixed effects (described below). I run a robustness check in the Appendix using a restricted dataset that only contains data during the analysis time frame and the same days-of-year in the other years of the data (i.e., for Event 1 the restricted data would only contain the month of January for the years 2012, 2013, 2014, 2015).

The estimating equation also contains several controls and fixed effects. The $temp_t$ and $precip_t$ variables contain maximum daily temperature and total daily precipitation, respectively. The $price_t$ variable is average price over time. The preferred specification controls for average price. I include robustness checks in the Appendix where I control for marginal price in one specification and for both average and marginal price in another specification.

The matrix $\mathbf{X_{it}}$ contains all the fixed effects controlled for in the preferred specification. I include household fixed effects to control for time invarying household characteristics; day-of-year fixed effects to control for seasonality; and the interaction of the two to capture household specific seasonal trends. I also include day-of-week fixed effects to control for consumption variation over the days of the week; holiday fixed effects (one for each holiday) to control for holiday specific water consumption patterns; and the interaction of day-of-week fixed effects and holiday fixed effects to capture variation in day-of-week consumption that coincides with a holiday. Finally, I include year fixed effects to account for year to year consumption trends. Standard errors are clustered by household to account for household level correlation of consumption over time.

I conduct several sensitivity checks and present results for a variety of estimating specifications in the Appendix. As an additional check of the results, I run the preferred specification on a set of placebo events at every one-week interval of the data. Specifically, I re-run the preferred regression specification 206 times assuming an event has happened at a different week for each run. This process allows me to recover the distribution of general consumption fluctuations and sheds light on how the unexplained variation at the events of interest compares to typical unexplained variation in consumption at other points in time. This test instills confidence in the notion that the unexplained variation at the public appeal events is in fact exceptional, and not simply due to chance from typical consumption fluctuations.

2.5.1 Unexplained Variation

Figure 2.3 is a graph of the residuals from the preferred specification. This figure gives a sense of how much of the variation is explained by the controls in the estimating equation. For the purposes of this graphic, I exclude the $event_t$, the pre_t, and $post_t$ indicators from the regression in order to show the unexplained variation at the events. I also run the regression with the dependent variable of consumption as levels instead of logs for easy comparison to the raw consumption trend figure displayed in the Data section. Figure 2.3 shows average residuals across all households for each day.

Figure 2.3: Unexplained Variation in Daily Household Consumption Trend

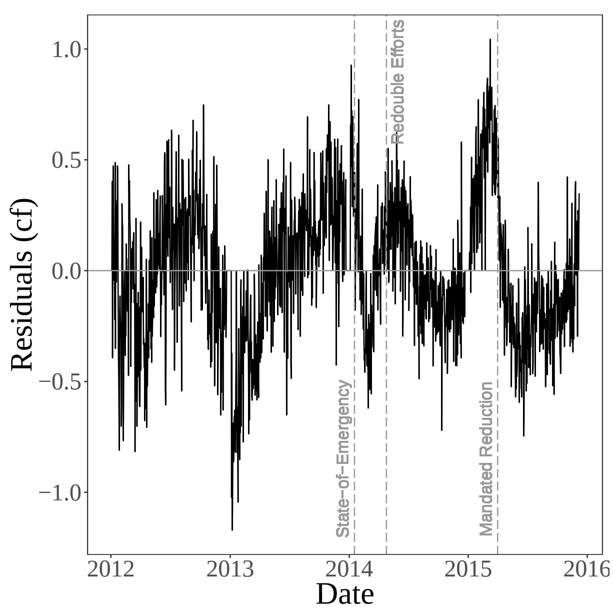

Notes: The graphic shows the average residuals across all households in the study sample for each day. The grey dotted lines represent the dates of the public appeals.

From this figure, it is clear that there is particularly high unexplained variation at the State-of-Emergency event and at the Mandated Reduction event. This suggests that there must be some exogenous factor, not explained by the model, driving consumption to decrease so rapidly at these points in time. According to this graphic, there is nothing extraordinary about water consumption during the Redoubling Efforts event.

I use the placebo analysis described earlier to understand how the unexplained variation at these events compares to the unexplained variation at other times in the span of the data.

2.5.2 Event Window Size

The choice for event window lengths is guided by the Google Trends described in the Events section. Implicitly, the public appeal announcements lead to increased drought awareness, which then affects water consumption. I take the Google Trends search intensity levels as evidence of increased drought awareness generated by the public appeal announcements. I assume that while Google search intensity is high, households are still learning about the public appeal, which is the time span of interest.

For both the State-of-Emergency event and the Mandatory Reduction event the Google Trends on "drought" and "water conservation" showed very high search intensity for the first two weeks, tapering off during the third and fourth week. Given this manifestation I analyze event window sizes that range from 3 days to 30 days following the announcement as these are the most likely window lengths for which customers are still learning of the event. My preferred event window size is two weeks, as this is the length of time that proves to have the most awareness intensity.

2.5.3 Identification

Unfortunately, the nature of a public appeals event as described in this paper is such that all households are exposed to the event (i.e. treatment), and therefore there is no "control" group in the traditional sense. Thus, the counterfactual used for the consumption of a treated household day-of-year is the average of household consumption for the same day-of-year across all years in the data (controlling for weather, price, and all other fixed effects). For example, the counterfactual consumption of the treated day January 17th, 2014 is the conditional average of consumption across all January 17ths over all the years in the data. The identifying assumption is that absent the appeal, departures in household consumption from the conditional household day-of-year average are identically distributed before and after what would be the event day. Thus, any systematic departure from the conditional household day-of-year average is driven by the announcement.

Of course, one must be wary of threats to identification validity. Such threats include other external factors that may have caused a sudden drop in consumption (i.e. departure from the conditional average). An example could be that a large portion of the sampled households had a family member move out the days following the announcement. While household composition changes are possible in a short period of time (e.g. children come home for spring break), I control for these systematic fluctuations that happen every year with day-of-year fixed effects and holiday fixed effects. Otherwise, it is hard to imagine that suddenly a large fraction of the sampled households had changes in composition within the two week window of the event.

2.6 Results

For each event, I show the estimated percent change in water consumption for varying sizes of event windows (ranging from 3 days to 30 days). The estimated percent change is the average change in daily consumption across the window size after the event as compared to the average daily consumption in the same window length before the event.

Given the preferred window size of two weeks, I run robustness checks for varying specifications to understand the sensitivity of the result. These tables are in the Appendix. I find that the results are not sensitive to specification and the preferred specification is on the conservative side. Lastly, I compare the results to other unexplained consumption fluctuations in the data using the aforementioned placebo test.

2.6.1 Event 1: Drought State-of-Emergency Declaration

The Governor's declaration of a drought State-of-Emergency is associated with a statistically significant decrease in consumption in the SF Water retail service area. Table 2.5 shows the average percent decrease in daily consumption for each window size. The decrease in water consumption becomes statistically significant at the 14 day window size with an average decrease of 1.9%. Looking at the window size to 30 days, the average decrease doubles at 3.8%.

Roughly 5% of the placebo events show a statistically significant decrease of the same size or more for the two week window. In other words, according to the placebo test, I am 95% confident that the decrease in consumption at the two week window is due to the public appeal and not typical fluctuations in consumption. This result matches the statistical significance found by the regression analysis as seen in Table 2.5.

The delayed decrease in consumption until the second week following the announcement is consistent with my predictions associated with the Google Trends analysis (described in the Events section). While search intensity for the word "drought" peaked during the first week, the search intensity for "water consumption" peaked in the second week. This suggests that while people may have had heightened awareness of the drought during the first week, it may have taken them a little longer to translate the information into actionable news.

Table 2.5: State-of-Emergency Declaration (Event 1)

	Varying Event Windows				
	Log(Consumption)				
	(1)	(2)	(3)	(4)	(5)
Event Indicator	-0.006	-0.012	-0.019**	-0.028***	-0.038***
	(0.013)	(0.009)	(0.008)	(0.008)	(0.008)
Window Size	3 days	7 days	14 days	21 days	30 days
R^2	0.693	0.693	0.693	0.693	0.693
Adjusted R^2	0.539	0.539	0.539	0.539	0.539
Residual Std. Error	0.613	0.613	0.613	0.613	0.613

Notes: The dependent variable is the natural log of daily consumption at the household level. The resulting coefficient can be interpreted as a percent change in daily water consumption. All regressions include weather controls, price controls, and the following fixed effects: household, day-of-week, day-of-year, year, holiday, household by day-of-year, and holiday by day-of-week. Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01

2.6.2 Event 2: Executive Order to Redouble State Drought Efforts

According to the results in Table 2.6, there is no statistically significant change in consumption at any window size following Governor Brown's public appeal of redoubling the state's drought actions. This is consistent with the lack of noticeable change in Google Trend searches during this event. It is also consistent with the anecdotal evidence suggesting the moderate nature of this event and the potential drought media focus shift from water conservation efforts to the politically charged water rights curtailment discussion. In addition, since this announcement came just three months after the State-of-Emergency announcement, this public appeal may have felt like old news by this point.

Table 2.6: Redouble State Drought Actions (Event 2)

	Varying Event Windows				
	Log(Consumption)				
	(1)	(2)	(3)	(4)	(5)
Event Indicator	-0.023*	-0.007	-0.001	0.009	0.009
	(0.012)	(0.010)	(0.008)	(0.007)	(0.007)
Window Size	3 days	7 days	14 days	21 days	30 days
R^2	0.693	0.693	0.693	0.693	0.693
Adjusted R^2	0.539	0.539	0.539	0.539	0.539
Residual Std. Error	0.613	0.613	0.613	0.613	0.613

Notes: The dependent variable is the natural log of daily consumption at the household level. The resulting coefficient can be interpreted as a percent change in daily water consumption. All regressions include weather controls, price controls, and the following fixed effects: household, day-of-week, day-of-year, year, holiday, household by day-of-year, and holiday by day-of-week. Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01

2.6.3 Event 3: Executive Order of 25% Statewide Mandatory Water Reductions

The Mandatory Water Reduction announcement was followed by a statistically significant decrease in water consumption at all window sizes above 3 days. Table 2.7 shows a statistically significant decrease in daily water consumption of 3.8% at the 7 day window size, a decrease of 4.6% at the 14 day window, and 6.2% at the 30 day window. Less than 1% of placebo events showed a statistically significant decrease of the same size for the two week window. In other words, the placebo test shows a 99% confidence that the decrease in consumption at the two week window is due to the public appeal and not typical fluctuations in consumption. This result matches the statistical significance found in the regression result as seen in Table 2.7.

It is not surprising that the change in consumption following this announcement comes earlier, and is greater, than the change following the State-of-Emergency annoucement. Google Trends shows the search intensity for “water conservation” reaches an all time high during the week following this announcement. A hypothesis is that since this announcement was directly focused at water conservation, there was less of a delay in translating the news into conservation. It is also important to remember that this announcement came a little over a year following the State-of-Emergency declaration, so people were likely primed this time around to take more immediate action.

Table 2.7: 25 Percent Mandated Reduction (Event 3)

	Varying Event Windows				
	Log(Consumption)				
	(1)	(2)	(3)	(4)	(5)
Event Indicator	-0.012	-0.038***	-0.046***	-0.054***	-0.062***
	(0.014)	(0.011)	(0.009)	(0.008)	(0.008)
Window Size	3 days	7 days	14 days	21 days	30 days
R^2	0.693	0.693	0.693	0.693	0.693
Adjusted R^2	0.539	0.539	0.539	0.539	0.539
Residual Std. Error	0.613	0.613	0.613	0.613	0.613

Notes: The dependent variable is the natural log of daily consumption at the household level. The resulting coefficient can be interpreted as a percent change in daily water consumption. All regressions include weather controls, price controls, and the following fixed effects: household, day-of-week, day-of-year, year, holiday, household by day-of-year, and holiday by day-of-week. Standard errors are clustered at the household level. *p<0.1; **p<0.05; ***p<0.01

2.7 Discussion

One must be cautious in interpreting these results and be mindful of limitations.

2.7.1 Limitations

First, one must keep in mind the identification assumptions discussed in the Empirical Methods section. Given these assumptions, the decrease in consumption following an announcement can be interpreted as the average treatment effect of the event.

Second, the effect of any individual public appeal will likely be affected by its time placement relative to other public appeals. One cannot directly compare effects of the State-of-Emergency appeal to the Mandatory Reduction appeal because the effect of the latter is influenced by the existence of the former. It is not clear if the existence of a previous appeal primes the audience and makes the following appeal more effective, or if the existence of the previous appeal makes the following appeal feel like old news and weakens its effect. I hypothesize that the third event benefits from the existence of the first event by preparing people to think about the drought a year before. Meanwhile, the effect of the second event may have been dampened by the fact that it came just over three months after the first event.

Third, the effect a public appeal depends on the level of publicity associated with the announcement. The second event is a prime example of an appeal that likely had modest media coverage, cultivating low levels of drought awareness, and thus did not result in a change in consumption.

Fourth, one must be careful in extrapolating results as there may be limited external validity. San Francisco water consumption inherently differs from that of other cities due to the fact that there is very little residential outdoor irrigation, and San Francisco residents are a particularly environmentally conscious population in general.[15]

It is not clear how these two idiosyncratic characteristics of San Francisco bias the results. On the one hand, the effect of the public appeal may be greater in an environmentally conscious city, biasing the estimates up. On the other hand, with a city that has little outdoor irrigation and already falls in the top three most water conservative cities in California[16], there is much less room for easy conservation (e.g. stop watering lawns or filling swimming pool). The idea of "demand hardening" kicks in, which is when it becomes more and more difficult to conserve once superfluous consumption has been eliminated and one is left with mostly vital consumption needs. For this reason, it is possible that the estimates are attenuated given the baseline level of San Francisco consumption and the limitations on further conservation.

2.7.2 Interpretations

The first State-of-Emergency Declaration event is mainly an appeal of moral suasion because the Governor's central message is a call for voluntary water consumption reductions. Thus one can interpret the decrease in consumption from this event as an effect of moral suasion. However, the second and the third event are announcements that are a mix of moral suasion and implied regulation change. While the mandated reductions from the third event do not become active until four weeks after the announcement[17] (and are mandates on the utility as a whole, not on households directly[18]), I cannot rule out that the decrease in consumption following this event may be a combination of people reacting to moral suasion as well as early compliance with expected regulations and/or price changes.

I argue that since SF Water is very transparent and explicit about regulation and price changes with their customers in advance (e.g. price changes are clearly posted several years in advance online[19]), SF Water customers are less likely to be worried about receiving an unexpected fine without forewarning. Since there is an implicit cost to the consumer associated with decreasing consumption, and if the customer is not worried about monetary repercussions of consuming the usual amount of water, then the resulting decrease in consumption is likely to be in response to moral suasion. It is worth noting, SF Water never

[15]San Francisco ranks as the greenest of 27 large cities in North America in a survey released by Siemens Corp.: http://usatoday30.usatoday.com/news/nation/environment/2011-06-29-green-cities-environment-recycling-San-Francisco_n.htm

[16]As calculated by the State Water Board, based on 2013 consumption: http://www.waterboards.ca.gov/waterrights/water_issues/programs/drought/docs/emergency_regulations/supplier_tiers_20150428.pdf

[17]Timeline of SF Water restrictions in response to the Governor announcements: http://www.sfwater.org/modules/showdocument.aspx?documentid=7228

[18]The mandated reductions are placed on the utility as a whole and the utility is responsible for instilling regulations across their sectors to comply on average with the mandated reduction level.

[19]Future water rates: http://www.sfwater.org/index.aspx?page=168

imposed mandatory reductions on its single-family customers during this drought.[20]

In the event that customers are still concerned about unexpected charges associated with price increases or regulation changes, I am still able to attribute a portion of the decrease in water consumption to moral suasion. As described in the Implications section below, customers would have to expect a price increase of 25% following Event 3 to respond with such a drastic decrease in consumption in the two week window. As seen in the Data section, the typical annual price increase is less than 10%. Thus, even if the customer believed there to be an unwarned price hike of 20%, double the usual annual amount, the price increase would only decrease consumption by 3.6% in the two week window of Event 3. Since there was a total of 4.6% decrease during this time, I would still be able to attribute the remaining 1% decrease in consumption to moral suasion.

2.7.3 Put in Perspective

Despite the limitations and interpretations described above, these results are still evidence of a statistically and economically significant short-term response to a public appeals announcement. Put in perspective, a 1.9% consumption decrease over 2 weeks, as there was from the State-of-Emergency event, amounts to an average savings of 4.9 cubic feet per household (or 37 gallons). These average savings are roughly equivalent to a decrease of one toilet flush per day per household over the two weeks[21], or forgoing a one-time 10 minute shower per household in the two week span[22]. A decrease of 4.6%, as in the Mandatory Reduction event, amounts to an average savings of 11.3 cubic feet per household over the two weeks. It is important to keep in mind that it is likely that not all households take action to conserve, meaning the effects on those who do must be much greater (Allcott, 2011).

In aggregate, given there are roughly 100,000 SFR households in the SF Water retail service area, the decrease in consumption following the State-of-Emergency event amounts to a total of almost half a million cubic feet (over 3.7 million gallons or 11 acre-feet) over the two week span. This is enough water to fill over 5.5 Olympic size swimming pools[23], or supply the water demand of almost 1000 households for a whole month. The decrease in consumption following the Mandatory Reduction event amounts to 1.1 million cubic feet (8.4 million gallons or 26 acre-feet), over twice the water savings following the State-of-Emergency event.

[20] The only mandatory reduction imposed was on irrigation account holders and very few single-family properties that have dedicated irrigation accounts.

[21] One flush varies from 1.28 gallons to 7 gallons. New federal plumbing standards specify a 1.6 gallons per flush minimum. More details at: `http://www.conserveh2o.org/toilet-water-use`

[22] Assuming a 3.5 gallons per minute shower head. More details at: `http://www.allianceforwaterefficiency.org/Residential_Shower_Introduction.aspx`

[23] There are 660,430 gallons of water in an Olympic size swimming pool. Details at: `http://www.patagoniaalliance.org/wp-content/uploads/2014/08/How-much-water-does-an-Olympic-sized-swimming-pool-hold.pdf`

2.7.4 Implications

Given a price elasticity of demand of -0.182 (Buck et al., 2016), SF Water would have had to raise prices by 10% to induce the same percent decrease in water consumption that is seen under the State-of-Emergency event for the two week window. Prices would need to be raised by 25% to match the decrease in consumption under the Mandatory Reduction event for the two week window. An important distinction between the price mechanism and moral suasion mechanism is that a price increase would likely induce a permanent demand reduction while moral suasion could lead to reductions that attenuate over time (Ito et al., 2018). In addition, the effect of moral suasion on consumer surplus is not as clear as it would be under a price mechanism. Under moral suasion one must consider the additional losses from a guilt factor, or alternatively, benefits from a warm-glow effect (Ito et al., 2018).

2.8 Conclusion

This paper quantifies the short-term effects of Governor Brown's public appeals on water conservation. I find statistically and economically significant decreases in water consumption of 1.9% to 4.6% in the two weeks following a well publicized public appeal announcement. Given certain assumptions about identification, I argue that this is evidence that moral suasion in the form of a public appeal can be an effective tool to induce conservation. It is important for policy-makers to understanding the effectiveness of various mechanisms that can influence conservation as it allows them to weigh the cost and benefits of using non-pecuniary actions versus market-based approaches. This can be particularly important in a climate where there are legislative and political challenges of employing market-based incentives for behavioral change.

Chapter 3

Assessing the Impact of Drought on Agriculture: Ex Post Evidence from California

3.1 Introduction

Droughts are a recurring phenomenon across the western United States, and are expected to increase in frequency and intensity as a result of climate change (Jones, 2015). Because water is a key input into a number of sectors of the economy, it is of interest to understand whether drought has measurable economic consequences, or whether water users are able to manage supply reductions with only small changes in outcomes such as production or employment.

This paper uses impact assessment methods to measure the economic consequences of drought on the agricultural sector. Despite significant population gains in the western states, farming accounts for the largest share of water use in the region. Further, agriculture is the primary source of household income for some socioeconomically vulnerable populations. Farmworkers are overwhelmingly Latinx and predominately recent immigrants or undocumented workers. According to the American Community Survey, farmworkers have the lowest family income of any occupation surveyed, the highest poverty rate of any surveyed population, the lowest level of educational attainment, and have one of the lowest rates of health insurance coverage (USCB, 2012). Changes in economic activity in agriculture thus have the potential to cause significant economic hardship, dislocation and distributional impacts.

While understanding the effect of drought on agricultural employment and economic activity is crucial, the question is relatively novel in the peer-reviewed literature. The studies that do exist are based on ex ante assessment models that make a large number of structural assumptions about agricultural production, capital investment and land allocation (Howitt et al., 2009a; Howitt et al., 2009b; Howitt et al., 2009c; Michael, 2009a; Michael, 2009b; Michael et al., 2010). The contribution of this paper is to empirically estimate the response of employment and harvested acreage to changes in irrigation water deliveries using an *ex post* difference-in-differences strategy commonly employed in the causal inference literature.

We have constructed a detailed dataset which matches both employment and harvested acreage with irrigation water deliveries to counties located in the agricultural heartland of California and the American West – the San Joaquin Valley.

To preview the results, we find a statistically and economically significant impact of surface water imports on agricultural employment and harvested acreage in the study area. We conclude that during the 2009 California drought, direct farm employment and employment by agricultural service providers in the San Joaquin Valley was reduced by about 9,800 jobs as compared to the normal water year of 2005. This estimate is almost *double* that of estimates described in previous papers based on ex ante assessment models (Michael et al., 2010). We also show that job losses resulting from water supply reductions in California agriculture are mediated by reductions in harvested acreage during shortage years. We conclude that the 2009 water supply reductions reduced harvested acreage in the seven-county study area by roughly 6% as compared to 2005.

We further show statistical evidence that employment and area harvested in counties with superior access to groundwater are less sensitive to drought. Moreover, the estimated effects of water deliveries on employment and land cultivated have declined over time, a finding consistent with the continued development of groundwater resources in the San Joaquin Valley. However, relying on groundwater as a backstop water supply may not be sustainable in light of the State's recent efforts to limit groundwater overdraft. These results suggest that absent other interventions, the future effects of drought on economic outcomes in California agriculture could be larger than those observed in the recent past.

3.2 Why Study California?

California is a natural setting in which to measure the economic impact of drought. The state has experienced nine large-scale, multi-year droughts since 1900. The two most recent of these occurred in 2007-2009 and 2011-2017 (Jones, 2015). These droughts, and other less pronounced dry spells, decrease water inflow to the Sacramento-San Joaquin River Delta, the largest delta and estuary in the western United States. Besides being an important habitat for a number of threatened and endangered species, the Delta is the source of water exported by the State Water Project (SWP) and the federal Central Valley Project (CVP) to more than 25 million Californians, including many farmers in the agricultural heartland of California. California is the largest agricultural producing state in the U.S and the state and federal projects are the largest sources of surface water imported into the San Joaquin Valley (Hanak et al., 2015).

Regulatory protection of several fish species, including Chinook salmon and Delta smelt, combined with these drought-induced decreases in inflows have led to significantly reduced water deliveries to farmers in recent years. In 2009, for instance, only 40% of SWP and 10% of CVP contracted water delivery amounts were allocated leading to a severe shortage of irrigation water in the San Joaquin Valley (CDWR, 2011; USBR, 2009). More recently, in 2014 and 2015, the SWP received an allocation of 5% and 20% of its contracted amount, respectively (CDWR, 2011). Meanwhile, the CVP received none (0%) of the contracted amount for both 2014 and 2015 (USBR, 2009).

In one of the most sensitive areas supplied by the CVP, the west side of the San Joaquin Valley, federal water contractors have only received their full water supply contracted amount four times since 1990. They have received 75% or more of their contracted amount only nine times during this period (USBR, 2009). While some farmers in this area have access to groundwater, many do not. Beyond the CVP, there are no other meaningful sources of surface water in the western San Joaquin Valley. A somewhat different picture emerges in Kern County located in the southern part of the study area. Farmers in Kern County are serviced by the SWP, which also exports water from the Delta and is subject to the same fluctuations as CVP, however farmers in Kern also have access to high-quality groundwater reserves and local sources of surface water that they use for groundwater recharge to smooth consumption (CDWR, 2013). Accordingly, as described below, our preferred specification of the impact assessment model allows for the effect of surface water exports on economic outcomes to vary by location.

3.3 Data

The data used in this analysis are comprised of a newly compiled annual, county-level panel covering the years 1980 to 2009. This data period evidences significant variation both in employment and surface water availability. It also includes two of the largest droughts in the recent past (1987-1992 and 2007-2009).

Counties included in the study area (Fresno, Kern, Kings, Merced, San Joaquin, Stanislaus and Tulare) receive surface water exports from the Delta via the CVP and/or SWP. These projects are the most important sources of surface water available to growers in the study area and are highly influenced by drought conditions. Other water supplies available to farmers in the study area include groundwater and local sources of surface water, the largest being the Kern and Tuolumne Rivers. In California, groundwater extraction is only infrequently monitored at the farm level and thus data on observed groundwater extraction do not exist across a broad enough geography and date range to be useful for our purposes. Similarly, surface water supplies from local sources in the study area are privately controlled and comprehensive data on diversion and usage aggregated to the county level are difficult to obtain. To control for the influence of local water sources, the econometric model includes county fixed effects both entered directly and interacted with CVP and SWP supplies imported from the Delta into each county. Further discussion of the influence of local supply sources is contained in the results section. As an additional control to capture factors affecting the agricultural economy over time, we consider six California agricultural counties that do not receive surface water exports from the Delta as a control area: Imperial, Madera, Monterey, Sutter, Yolo and Yuba. Figure 3.1 shows the study area including the seven counties of interest and the six control counties.

The panel data set also includes variables for direct farm employment, total agricultural employment (i.e., sum of direct farm and agricultural service sector employment including employment by farm labor contractors), and harvested acreage. It also includes water deliveries from both the CVP and SWP. Table 3.1 displays average employment, water deliveries and harvested acreage by county from 1980-2009. Further details on the source of the data,

Figure 3.1: Study area

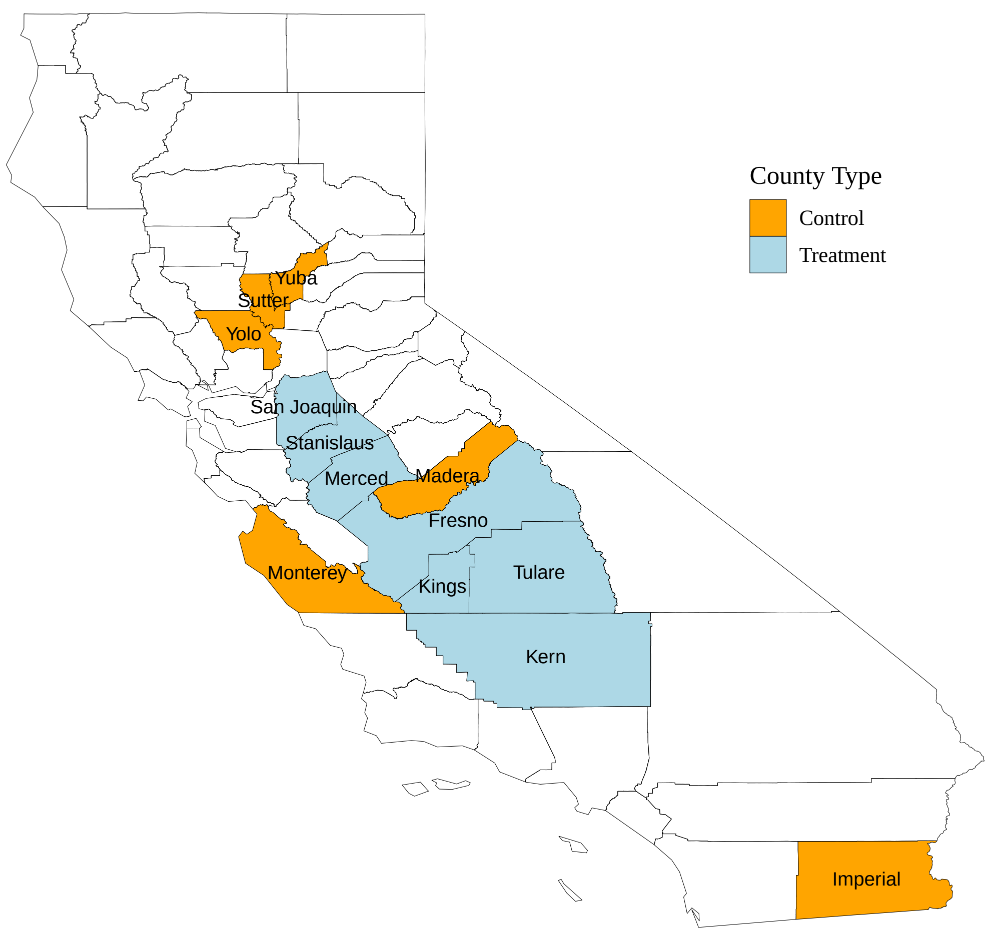

Source: County boundaries from the US Census Bureau's 2016 MAF/TIGER data base.

the construction of the variables, and variable definitions are described below.

Table 3.1: Summary statistics by county

County	Total Employment	Direct Farm Employment	Total Agricultural Employment	Acres Harvested	Surface Water Deliveries
Fresno	363	29	58	844	1029
Kern	279	17	40	597	985
Kings	45	5	8	300	285
Merced	78	10	14	282	144
San Joaquin	228	13	20	261	37
Stanislaus	178	12	18	200	100
Tulare	153	18	37	419	11

Notes: Employment is in thousands of jobs per year. Total agricultural employment includes direct farm employment and agricultural service sector employment. Acres harvested is in thousands of acres per year and account for roughly two-thirds of total harvested acreage in the San Joaquin Valley. Surface water deliveries are measured in acre-feet per year and include Delta exports delivered from the CVP and SWP.

3.3.1 Employment Data

County-level employment data are publicly available, and we obtained them from the Bureau of Economic Analysis (BEA, 2011b). We distinguish between direct farm employment, and total agricultural employment. Direct farm employment includes anyone who works in the direct production of agricultural commodities, including crops and livestock (SIC codes 01 – 02; NAICS code 111 - 112) (BEA, 2011a). Total agricultural employment is the sum of direct farm employment and employment in the agricultural services sector (SIC code 07; NAICS code 113 - 115). The agricultural services sector includes farm labor contractors.

The data we used from 1980 - 2000 are categorized in the Standard Industrial Classification (SIC) system. In the 1990s, a new classification system (North American Industrial Classification System (NAICS)) was introduced, in part to facilitate accounting under the North American Free Trade Agreement. The SIC data series was discontinued in 2000. In that year, the BEA shifted to reporting sectoral employment based on the SIC industry classification to reports based on the NAICS classification. The BEA provides a concordance to match industry descriptions between the two coding systems. As we control for year fixed effects in our preferred specification, if there are year-to-year differences in employment that are due to the new classification, our method implicitly controls for these differences.

3.3.2 Water Deliveries Data

Government water delivery data include both state deliveries from the State Water Project and federal deliveries from the Central Valley Project. The state water delivery data come from the California Department of Water Resources' Bulletin 132 (CDWR, 2011) and the Kern County Water Agency (KCWA, 2011). The federal water deliveries data are from the Bureau of Reclamation (USBR, 2009). The water delivery data is at the water district level and captures deliveries made on the CVP's Delta Mendota Canal, Cross Valley Canal, Mendota Pool and San Luis Unit; this data covers deliveries to the Exchange Contractors and the South of Delta Agricultural Water Service Contractors from the CVP's Delta and West San Joaquin Divisions. We used a Geographic Information System to allocate water deliveries to counties. We first took the intersection of the boundaries of each of the water districts and counties.[1] We then calculated the acreage of the district-county intersection and divided that by the acreage of each of the districts. We multiplied this ratio by the water deliveries in each water district and summed the share of water deliveries in the district-county intersection over counties. Thus, water deliveries are allocated to the county level according to the share of acres of each water district that falls within each county. Annual deliveries are reported in acre-feet.

3.3.3 Acreage Harvested Data

The data set also includes harvested acres by county. These data come from the Agricultural Commissioners' Offices of Fresno, Imperial, Kern, Kings, Madera, Merced, Monterey, San Joaquin, Stanislaus, Sutter, Tulare, Yolo and Yuba counties for the years 1980 through 2009. We consider land allocated to a subset of crops: almonds, avocados, broccoli, cotton, grapes, hay, lemons, lettuce, oranges, pistachios, rice, strawberries, tomatoes and walnuts. We use a subset of crops for this analysis because acreages are more consistently defined for these individual crops than for total harvested acreage. For example, some counties include rangeland in total area statistics in some years, but not in other years. The crops in our analysis account for roughly two-thirds of total harvested acreage in the San Joaquin Valley.

3.4 Empirical Approach

To understand the effects of drought on our three outcome variables of interest, direct farm employment, total agricultural employment, and acres harvested, we use an ex post difference-in-differences strategy. We run four variations of a standard statistical model for each outcome variable to show the robustness of our results. Model (1) is a basic specification estimating the correlation between water deliveries and the outcome variables, not controlling for any other observable or unobservable confounding variables. Model (2) builds on Model (1) including county fixed effects, which control for unobservable confounding variables that vary across counties but not over time such as soil quality, groundwater quality,

[1] Cal-Atlas Geospatial Clearing House, boundaries of "Federal," "State" and "Private" water districts. Available at: http://www.atlas.ca.gov/download.html. Accessed May 26, 2009. Boundaries of Counties obtained from ESRI ArcGIS basemap layers.

climatic conditions, transportation distances and rights to non-project sources of surface water. Model (3) builds on Model (2) by including time fixed effects, controlling for unobservable confounding factors that vary across time but affect all counties simultaneously such as statewide business cycle fluctuations, interest rates, exchange rates and commodity price trends.

Models (1)-(3) are estimated using data from the seven counties in our treatment area receiving surface water deliveries from the CVP and/or SWP. The identifying source of variation for Model (3) is within county time series variation. Model (4) is similar to the third, but is estimated with data from all counties including those that receive deliveries from the Delta and the control counties that do not receive these deliveries. The identifying source of variation for this specification is within county variation relative to the control group within county variation. The model details are described below.

3.4.1 Model Details

Let N_{it} be the agricultural employment in county i during year t. D_{it} are the deliveries from the federal and state water projects to the water districts in county i in year t. Without controlling for any other observable and unobservable confounders, a basic statistical model estimating the correlation between N_{it} and D_{it} is given by Model (1):

$$N_{it} = \alpha + \beta_1 D_{it} + \varepsilon_{it} \tag{3.1}$$

The identifying assumption required in order for Ordinary Least Squares to provide consistent estimates of the parameter of interest (β_1) is that $E[\varepsilon_{it}|D_{it}] = 0$. Any factor not included in this simple model, which is correlated with deliveries *and* employment would violate this assumption (e.g. soil quality and availability of groundwater). One could explicitly control for unobservable confounders varying across counties but not time by including county fixed effects in a regression as given in Model (2) below:

$$N_{it} = \alpha_i + \beta_1 D_{it} + \varepsilon_{it} \tag{3.2}$$

The α_i capture the effects of unobservable confounders which vary across counties and are time invariant. Failing to control for the α_i via this fixed effects strategy will lead to biased coefficient estimates of β_1. Further, there are certainly factors, which vary across time, but affect all counties simultaneously and might be correlated with employment and deliveries (e.g. national business cycle fluctuations). To deal with these (un)observable confounders which do not vary across space, but time, one uses the well-established approach of adding year fixed effects:

$$N_{it} = \alpha_i + \phi_t + \beta_1 D_{it} + \varepsilon_{it} \tag{3.3}$$

where the α_i are again the county specific time invariant confounders and the ϕ_t are the shocks common to all counties. The identifying assumption then becomes $E[\varepsilon_{it}|D_{it}, \alpha_i, \phi_t] = 0$. This assumption would be violated if one failed to include any confounders that are correlated with deliveries over time within a county.

One could estimate this equation on a sample containing just the counties receiving deliveries (as we do in Model (3)) or a sample of counties receiving deliveries and include counties, which do not receive deliveries as a control group (Model (4)). The estimation results are robust to using either sample. In the first sample, the identifying source of variation is within county time series variation. For the larger sample it is within county variation relative to the control group county variation, which identifies the coefficient of interest β_1.

3.5 Estimation Results

3.5.1 Employment

We first estimate the effects of surface water exports from the Delta on direct farm employment using the four models described above on the data from 1980 - 2009. The results are summarized in Table 3.2. The estimated coefficients for all four models are statistically significant where the standard errors are clustered at the county level (Wooldridge, 2002). Model (4) is our preferred specification since it includes both county and year fixed effects, and is estimated on a dataset including a group of control counties. The estimated coefficient for this model is 0.0036 and is statistically significant at the 5% level in a two-sided test. To put this estimated coefficient in perspective, it suggests that reducing water deliveries by 278 acre-feet results in the loss of one direct farm job.

Table 3.2: Effect of changes in surface water deliveries on direct farm employment

	Direct Farm Employment			
	(1)	(2)	(3)	(4)
Deliveries	0.0098**	0.0041**	0.0037*	0.0036**
	(0.0048)	(0.0020)	(0.0019)	(0.0018)
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Number of Counties	7	7	7	13
Observations	210	210	210	390
Adjusted R^2	0.3366	0.8840	0.9334	0.9534

Notes: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable in each regression is direct farm employment at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level, presented in parentheses. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

Next we consider the influence of CVP and SWP surface water exports on *total* agricultural employment by county. Table 3.3 displays the results of this analysis. The models correspond to those discussed above, with the exception that the dependent variable is the sum of direct farm employment and agricultural service sector employment. As before, the estimated coefficient on deliveries for Model (4) is positive, this time with a value of 0.0049. It is significant at the 1% level. To put this estimated coefficient into perspective, it indicates that reducing water deliveries by 205 acre-feet results in the loss of one agricultural job.

Table 3.3: Effect of changes in surface water deliveries on total agricultural employment

	Total Agricultural Employment			
	(1)	(2)	(3)	(4)
Deliveries	0.0246***	0.0047***	0.0046***	0.0049***
	(0.0081)	(0.0005)	(0.0013)	(0.0011)
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Number of Counties	7	7	7	13
Observations	193	193	193	356
Adjusted R^2	0.4282	0.9222	0.9533	0.9608

Notes: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable in each regression is total agricultural employment at the county level, defined as the sum of direct farm employment and employment in the agricultural services industry. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level, presented in parentheses. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

As described above, while surface water exports from the Delta account for the majority of surface water used in the study area, it is necessary to account for the influence of local water supply sources. Rights to the Kern River, for example, are held by a small number of water districts in Kern County who use it for groundwater recharge rather than direct irrigation. Supplies of irrigation water from the other largest local supplies such as the Tuolumne River do not exhibit the same inter-annual variation as the CVP and SWP owing to the seniority of rights held by growers and the presence of significant surface storage capacity in Don Pedro Reservoir and other facilities. In both cases, county-level fixed effects should largely control for the influence of non-project, local water sources. We also note that even if local surface water supplies are not perfectly captured by county fixed effects, they are likely to be positively correlated with surface water exports from the Delta. In this case, the correlation between Delta exports and our outcome variables will be overestimated. However, our estimates of the main relationships of interest, namely the overall impact of drought on employment and harvested acreage in the study region, should be unbiased.

We turn now to the aggregate impact of changes in Delta exports on employment in the study area. Total state and federal surface water exports from the Delta to these counties in 2009 were roughly 1.1 million acre-feet, which was about one-third that of deliveries in 2005, a normal water-supply year. We calculate the implied drought-induced reduction in farm employment using our preferred Model (4) in Table 3.2. Considering just direct farm employment, the reduction in 2009 deliveries causes an estimated loss of about 7,240 jobs, which is equivalent to an *8% decline* in direct farm employment in the seven-county study area.

Reductions in water deliveries in 2009 caused even larger losses in total agricultural employment. Using the estimated coefficient in Model (4) found in the right portion of Table 3.3, we conclude that the 2009 drought caused an estimated loss of about 9,830 jobs. This is equivalent to almost a 6% decline in total agricultural employment in the study area. Our county-level model therefore is consistent with the hypothesis that reductions in water supplies in 2009 caused economically and statistically significant losses in employment in California's agricultural sector.

3.5.2 Harvested Acreage

While the models above do not formally test the mechanism of how changes in water deliveries influence job losses, one would expect that acreage planted to crops would decrease if water supplies fall short due to drought, which would reduce the demand for labor. We therefore test whether deliveries are correlated with total harvested acreage in the counties in our sample.

We find that there is a strong and statistically robust relationship between surface water availability and area harvested in the San Joaquin Valley. The model specifications are the same as those used in the models explaining direct farm employment and total agricultural employment, only that we use area cropped in acres as the dependent variable. The estimated coefficient on deliveries in Model (4) of Table 3.4 is 0.0817, which is significantly different from zero at the 1% level with clustered standard errors. This finding suggests that increasing surface water deliveries in a given year significantly increases the amount of land under cultivation in the relevant counties.

The estimated coefficient in Model (4) suggests that over the historical record from 1980 to 2009, a reduction in water exports of roughly 12 acre-feet causes one additional acre to be fallowed. Recall that the model is estimated based on plantings of a subset of crops accounting for roughly two-thirds of total harvested acreage in the San Joaquin Valley. Accounting for the crops not in the sample, and assuming the same acreage response to changes in water deliveries, it follows that a reduction in water deliveries of around 8 acre-feet would result in an extra acre of total fallowing.

It is also instructive to estimate the amount of fallowing caused by the water delivery reductions of 2009 as compared to 2005. Model (4) indicates that an estimated 247,000 acres were fallowed in our seven-county study area in 2009 as a result of the water delivery reductions. This figure is calculated by multiplying the change in deliveries between these two years by the coefficient on deliveries in Model (4) and then adjusting for the fact that Model (4) is based on a subset of crops accounting for two-thirds of total acreage. This is equivalent to almost a 6% decline in harvested acres in the seven-county study area as

Table 3.4: Effect of changes in surface water deliveries on area harvested

	Acres Harvested			
	(1)	(2)	(3)	(4)
Deliveries	0.4040***	0.0741***	0.0827**	0.0817***
	(0.0959)	(0.0137)	(0.0320)	(0.0220)
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Number of Counties	7	7	7	13
Observations	210	210	210	390
Adjusted R^2	0.6815	0.9674	0.9684	0.9786

Notes: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable for each regression is harvested acreage for a subset of crops at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level, presented in parentheses. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

compared to 2005.

3.5.3 Groundwater Backstop Attenuates the Impacts of Drought

It is evident from our data that periodic droughts in California reduce the amount of surface water available to agriculture. It is also the case that many farmers in the study area have the ability to extract groundwater, both to enhance mean water supplies and stabilize consumption. In this section, we consider the degree to which the effects of drought on agriculture are mitigated by the existence of backstop supplies of groundwater. Within the study area, there is cross-sectional variation in groundwater conditions that will allow us to test this proposition by introducing heterogeneity in the treatment effect via interaction terms into our basic model. We further explore the effect of groundwater by examining the trend of the estimated treatment effect over time.

To test whether the presence of groundwater reserves attenuates the impact of drought on economic activity in agriculture, we rerun our analysis of direct farm employment allowing for the influence of water deliveries on farm jobs to vary between Kern and Kings counties and the rest of the study area. The Kern County Subbasin of the San Joaquin Valley Groundwater Basin is one of the largest and most productive aquifers in the state, and growers in Kern and Kings county have better access to groundwater supplies than the rest of the study region (CDWR, 2013). As described above, Kern River rights are not directly used in agriculture but rather are used to recharge the aquifer. The ground waters of the west side of the San Joaquin Valley, by contrast, are of generally lower quality and can be more difficult to access (CDWR, 2006). In western Fresno County, groundwater in the upper aquifer ranges as high as 2,000 to 3,000 mg/l of TDS (Davis et al., 1959); dissolved solids in shallow groundwater can be as high as 10,000 mg/l in some lower fan areas (Dubrovsky et al., 1993). The confined lower aquifer in the western part of our study area is separated from the upper semi-confined aquifer by an aquitard known as the Corcoran Clay layer. To access the lower aquifer, water users have established wells up to 3,500 feet deep, meaning that these backstop groundwater supplies come at a high cost to users (USGS, 1995). Because of these significant differences in groundwater availability and quality among parts of the study area, it is of interest to measure the differential impact of drought in Kern and Kings Counties compared to the overall study area.

Table 3.5 displays the results of this analysis. The model formulation and variable definitions are exactly as in Table 3.2, with the addition of an interaction term on deliveries that allows deliveries to have a different effect on farm employment in Kern and Kings counties than in the rest of the sample. The coefficient shown on surface water deliveries for Kern and Kings is relative to the other counties. This suggests the overall effect of deliveries on direct farm employment is significantly lower in Kern and Kings. This finding is consistent with the hypothesis that groundwater reserves act as a buffer stock, attenuating the short-run impacts of drought.

Next, we rerun our analyses of direct farm employment allowing for the influence of water deliveries to vary across years. This formulation results in a year-specific treatment effect with coefficients showing the relationship of surface water supplies and direct farm jobs over time. Figure 3.2(a) displays these coefficients graphically with a 95% confidence interval. This figure shows the coefficients tending towards zero as the years increase. This result is also consistent with the fact that groundwater use has been increasing in the San Joaquin

Table 3.5: Effect of changes in surface water deliveries on direct farm employment including interaction term of deliveries with Kern and Kings counties

	Direct Farm Employment			
	(1)	(2)	(3)	(4)
Deliveries	0.0140***	0.0064***	0.0058***	0.0058***
	(0.0022)	(0.0010)	(0.0006)	(0.0006)
Deliveries * Kern/Kings	−0.0011	−0.0046***	−0.0043***	−0.0045***
	(0.0028)	(0.0010)	(0.0008)	(0.0006)
County FEs	No	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes
Control Counties	No	No	No	Yes
Number of Counties	7	7	7	13
Observations	210	210	210	390
Adjusted R^2	0.6626	0.8862	0.9356	0.9547

Notes: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009. The dependent variable for each regression is direct farm employment at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. Standard errors are clustered at the county level, presented in parentheses. Coefficients are significantly different from zero at the 1% (***), 5% (**) or 10% (*) level.

Valley region over time (CDWR, 2015; Hanak et al., 2015). Figure 3.2(b) demonstrates results of the same analysis of year-specific treatment effects but this time showing the correlation of deliveries and acres harvested over time. The results for acres harvested show a similar trend as employment where coefficients tend towards zero as the years increase, implying that access to groundwater likely mitigates effects of water delivery fluctuations on acres harvested.

It is also of interest that the coefficient on the year 1991 in Figure 3.2(a) is particularly large. This year is one of the three driest on record in California, and the driest in our dataset. This year is also the last of a six-year drought (Jones, 2015). This result is consistent with the fact that groundwater acts as a buffer stock since groundwater reserves were particularly low by this year (Cooley et al., 2015) and so there was less availability to counteract the employment effects of surface water reductions.

We close this section by noting that while groundwater development has increased over time in the study area, there has also been a related shift in crop mix towards perennials, particularly almonds and pistachios (Lobell and Field, 2009). Although irrigated acreage in the San Joaquin Valley remained relatively stable from 1980 to 2012, perennial crops grew from 21 percent to 36 percent of total acreage over this period (Hanak et al., 2017). Acres of almonds grown in California increased from just over 300,000 in 1980 to over 1.3 million in 2017 (CDFA, 2018). More than three-quarters of the statewide total is in the study area. For purposes of this paper, it is important to note that perennial crops have a more stable pattern of economic activity than annuals as annual crops can be more easily fallowed in response to drought. This shift toward perennials (which partly explains the increased development of groundwater resources in the San Joaquin Valley) is likely to be another factor that explains the attenuation of the economic impact of drought over time as shown in Figure 3.2.

Figure 3.2: Effect of changes in surface water deliveries on direct farm employment and acres harvested over time

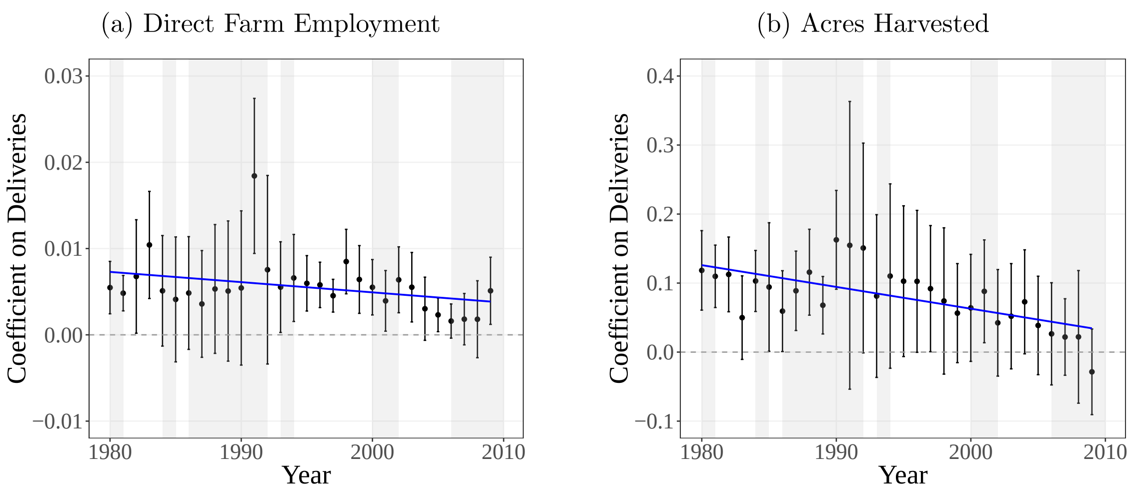

Notes: The dataset used in these regressions is comprised of annual data for the period 1980 – 2009 and includes both treatment and control counties. The dependent variable in these regressions are direct farm employment (left) and acres harvested (right) at the county level. Deliveries are acre-feet of water delivered to the districts within a county by the Central Valley Project or State Water Project. The regressions include county and year fixed effects. Coefficients displayed are for each year and show a 95% confidence interval. The coefficients are interpreted as a change in the dependent variable given a one acre-foot change in water deliveries in the particular year. Standard errors are clustered at the county level. The shaded areas represent dry years.

3.5.4 Sustainability Considerations

This paper has presented statistical evidence suggesting that the effect of drought on California agriculture is statistically significant and substantial, but is decreasing over time, perhaps as a result of the continued development of groundwater resources in the study area. Groundwater is a limited resource, however, which raises the question of whether relying on groundwater supplies is a sustainable strategy for agricultural water users. If it is not, then the effect of drought on economic activity can be expected to rise in the future.

The sustainability of groundwater supplies is certainly under question as groundwater overdraft has been a concern in the San Joaquin Valley for decades. Indeed, groundwater overdraft was a primary motivation for the formation of the Westlands Water District in the western portion of the study area (Westlands is the largest irrigation district in the nation with roughly 450,000 acres under cultivation in a typical year). In recent years, groundwater overdraft has continued in the San Joaquin Valley, especially in dry years. According to the Public Policy Institute of California (2018), the estimated average annual overdraft since the mid-1980s is equal to 13 percent of net water use, and the rate of pumping has only accelerated in recent years during the 2012-16 drought. By 2014, most of the San Joaquin Valley groundwater basins were critically overdrafted, where pumping exceeds replenishment (Hanak et al., 2018). Overdraft has caused groundwater elevations in the study area to decline, resulting in attendant problems such as municipal and irrigation wells running dry and land subsidence (Hanak et al., 2015). The State of California estimates that in portions of the western San Joaquin Valley, groundwater pumping has caused the land surface to decline as much as 8.5 meters since the 1920s (NASA, 2016).

To address groundwater overdraft and promote sustainable management of this critical resource, in 2014 California implemented the Sustainable Groundwater Management Act (SGMA). It mandates the creation of local Groundwater Sustainability Agencies that are responsible for managing groundwater such that overdraft is limited in dry years. They have been given until the year 2020 to create sustainable groundwater management plans, and until 2040 to reach sustainability (Hanak et al., 2015). Our empirical results suggest that as the groundwater backstop supply is limited due to SGMA, the impacts of drought on economic activity in agriculture may increase.

3.6 Conclusions

In this paper we show evidence of an economically and statistically significant effect of irrigation water deliveries from California's state and federal water projects on both county level employment as well as area harvested. We show that for a shortage similar to that experienced in 2009, California's total agricultural employment in these counties would be lowered by roughly 9,830 jobs. For comparison, estimates of direct *and* indirect job losses based on ex ante assessment models were between 4,700 and 4,900, meaning that these ex ante estimates are lower than our ex post econometric estimates by almost a factor of two (Michael et al., 2010). Further, we show that the likely mechanism through which this effect operates is the fallowing of land during drought. For the same reduction in irrigation water deliveries we estimate that 247,000 acres were fallowed in 2009.

Importantly, we find that employment and cropped acreage in areas with high-quality groundwater resources appear to be less sensitive to fluctuations in irrigation water deliveries than areas with less favorable groundwater. Similarly, we find that over time, as reliance on groundwater resources increases, and there is a shift towards perennial crops in the study area, the effect of irrigation water deliveries on employment and acreage is mitigated. This finding, however, should not be interpreted as a sign that drought is of less importance to irrigated agriculture in California; withdrawal of groundwater has both private and social costs, and the State has put in place measures to require sustainable groundwater use, limiting overdraft in dry years. Thus, relying on groundwater backstop water supply in the event of a drought may not be sustainable. Absent other interventions, the effects of drought on economic outcomes in California agriculture could be even larger than those observed in the recent past.

References

Allcott, Hunt (2011). "Social norms and energy conservation". In: *Journal of Public Economics* 95.9-10, pp. 1082–1095.

Andreoni, James (1989). "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence". In: *Journal of Political Economy* 97.6, pp. 1447–58.

BEA (2011a). *Local Area Personal Income Methodology, Appendix: Concordance between BEA industry descriptions and SIC codes.* Tech. rep. U.S. Bureau of Economic Analysis.

— (2011b). *Regional Economic Accounts, Local Area Personal Income, Table CA25-Total employment by industry.* Tech. rep. U.S. Bureau of Economic Analysis.

Becker, Gary S. (1974). "A Theory of Social Interactions". en. In: *The Journal of Political Economy* 82.6, pp. 1063–1093.

Bertrand, Marianne and Adair Morse (2011). "Information Disclosure, Cognitive Biases, and Payday Borrowing". In: *The Journal of Finance* 66.6, pp. 1865–1893.

Borenstein, Severin (2009). "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing". In: *Working Paper*.

Buck, Steven, Maximilian Auffhammer, Stephen Hamilton, and David Sunding (July 2016). "Measuring Welfare Losses from Urban Water Supply Disruptions". In: *Journal of the Association of Environmental and Resource Economists* 3.3, pp. 743–778.

CDFA (2018). *2017 California Almond Acreage Report.* Tech. rep. California Department of Food and Agriculture.

CDWR (2006). *Ground Water Basins in California, Bulletin 118.* Tech. rep. California Department of Water Resources.

— (2011). *State Water Project Analysis Office, Bulletin 132, Appendix B.* Tech. rep. California Department of Water Resources.

— (2013). *California Water Plan, Tulare Lake Hydrologic Region, Volume 2.* Tech. rep. California Department of Water Resources.

— (2015). *California's Groundwater Update 2013: A Compilation of Enhanced Content for California Water Plan Update 2013.* Tech. rep. California Department of Water Resources.

Charness, Gary and Uri Gneezy (2009). “Incentives to Exercise”. In: *Econometrica* 77.3, pp. 909–931.

Chetty, Raj, Emmanuel Saez, and László Sándor (Aug. 2014). “What Policies Increase Prosocial Behavior? An Experiment with Referees at the *Journal of Public Economics*”. en. In: *Journal of Economic Perspectives* 28.3, pp. 169–188.

Cooley, Heather, Kristina Donnelly, Rapichan Phurisamban, and Madhyama Subramanian (2015). *Impacts of California’s Ongoing Drought: Agriculture*. Tech. rep. Pacific Institute, p. 30.

Dal Bo, Ernesto and Pedro Dal Bo (2014). “”Do the Right Thing:” The Effects of Moral Suasion on Cooperation”. In: *Journal of Public Economics* 117, pp. 28–38.

Davis, G.H., J.H. Green, F.H. Olmsted, and D.W. Brown (1959). *Groundwater conditions and storage capacity in the San Joaquin Valley, California*. USGS Numbered Series 1469. U.S. Govt. Print. Off.,

Davis, Lucas W. and Gilbert E. Metcalf (Sept. 2016). “Does Better Information Lead to Better Choices? Evidence from Energy-Efficiency Labels”. In: *Journal of the Association of Environmental and Resource Economists* 3.3, pp. 589–625.

Dubrovsky, Neil M., Steven J. Deverel, and D.W. Brown (1993). *Regional Ground-Water Quality*. New York: Van Nostrand Reinhold.

Ferraro, Paul J. and Michael K. Price (2011). *Using Non-Pecuniary Strategies to Influence Behavior: Evidence from a Large Scale Field Experiment*. NBER Working Paper 17189. National Bureau of Economic Research.

— (2013). “Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment”. en. In: *The Review of Economics and Statistics* 95.1, pp. 64–73.

Finkelstein, Amy (2009). “E-ztax: Tax Salience and Tax Rates”. In: *The Quarterly Journal of Economics* 124.3, pp. 969–1010.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr (2001). “Are people conditionally cooperative? Evidence from a public goods experiment”. In: *Economics Letters* 71.3, pp. 397–404.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg (2006). “Costly information acquisition: Experimental analysis of a boundedly rational model”. In: *The American Economic Review* 96.4, pp. 1043–1068.

Gosnell, Greer K, John A List, and Robert Metcalfe (June 2016). *A New Approach to an Age-Old Problem: Solving Externalities by Incenting Workers Directly*. Working Paper 22316. National Bureau of Economic Research.

Hanak, Ellen et al. (2015). *What If California’s Drought Continues?* Tech. rep. Public Policy Institute of California.

Hanak, Ellen et al. (2017). *Water Stress and a Changing San Joaquin Valley.* Tech. rep. Public Policy Institute of California.

Hanak, Ellen, Jelena Jezdimirovic, Sarge Green, and Alvar Escriva-Bou (2018). *Replenishing Groundwater in the San Joaquin Valley.* Tech. rep. Public Policy Institute of California.

Handel, Benjamin, Allyson Barnett, Igal Handel, Jonathan T. Kolstad, and Michael D. Whinston (2014). "Information or Compensation? Understanding the Role of Information Technology in Physician Response to Pay-For-Performance". In: *Working Paper.*

Hastings, Justine S. and Jeffrey M. Weinstein (Nov. 2008). "Information, School Choice, and Academic Achievement: Evidence from Two Experiments". In: *The Quarterly Journal of Economics* 123.4, pp. 1373–1414.

Howitt, Richard, Duncan MacEwan, and Josué Medellín-Azuara (2009b). "Economic Impact of Reductions in Delta Exports on the Central Valley Agriculture". In: *Agricultural and Resource Economics Update.* edited.

— (2009a). "Economic Impact of Reductions in Delta Exports on the Central Valley Agriculture". In: *Agricultural and Resource Economics Update* 12.3.

Howitt, Richard, Josué Medellín, and Duncan MacEwan (2009c). *Measuring the Employment Impact of Water Reductions.* Tech. rep. Department of Agricultural, Resource Economics, and Center for Watershed Sciences. University of California, Davis, California, p. 10.

Hoxby, Caroline and Sarah Turner (2013). "Expanding College Opportunities for High-Achieving, Low-Income Students". In: *Working Paper.*

Ito, Koichiro (Feb. 2014). "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing". In: *American Economic Review* 104.2, pp. 537–563.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka (Feb. 2018). "Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand". In: *American Economic Journal: Economic Policy* 10.1, pp. 240–267.

Jessoe, Katrina and David Rapson (2014). "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use". In: *American Economic Review* 104.4, pp. 1417–38.

Jones, Jeanine (2015). *California's Most Significant Drought: Comparing Historical and Recent Conditions.* Tech. rep. California Department of Water Resources.

KCWA (2011). *SWP Supply and Delivery Summary. Years 1970- 2009.* Tech. rep. Kern County Water Agency.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian V. Wrobel (Feb. 2012). "Comparison Friction: Experimental Evidence from Medicare Drug Plans". In: *The quarterly journal of economics* 127.1, pp. 199–235.

Landry, Craig, Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp (2005). *Toward an understanding of the economics of charity: Evidence from a field experiment.* Tech. rep. National Bureau of Economic Research.

Ledyard, John (1995). "Public Goods: A Survey of Experimental Research". In: *Handbook of Experimental Economics*. Ed. by J Kagel and A Roth. Princeton University Press.

Lobell, David B. and Christopher B. Field (2009). *California perennial crops in a changing climate*. Tech. rep. California Climate Change Center, pp. 317–333.

Michael, Jeffrey (2009a). *Employment Impacts of Reduced Water Supplies to San Joaquin Valley Agriculture*. Tech. rep. Business Forecasting Center, Eberhardt School of Business, University of the Pacific.

— (2009b). *Unemployment in the San Joaquin Valley in 2009: Fish or Foreclosure?* Tech. rep. Business Forecasting Center, Eberhardt School of Business, University of the Pacific.

Michael, Jeffrey, Richard Howitt, Josué Medellín-Azuara, and Duncan MacEwan (2010). *A Retrospective Estimate of the Economic Impacts of Reduced Water Supplies to the San Joaquin Valley in 2009*. Tech. rep. Business Forecasting Center, Eberhardt School of Business, University of the Pacific, p. 15.

NASA (2016). *San Joaquin Valley is Still Sinking*. Tech. rep. NASA Earth Observatory.

Olmstead, Sheila M. and Robert N. Stavins (2009). "Comparing price and nonprice approaches to urban water conservation". en. In: *Water Resources Research* 45.4.

PPIC Water Policy Center (Apr. 2015). *California's Water*. Tech. rep. Public Policy Institute of California.

Reiss, Peter C. and Matthew W. White (2008). "What Changes Energy Consumption? Prices and Public Pressures". In: *The RAND Journal of Economics* 39.3, pp. 636–663.

Sallee, James (2014). "Rational Inattention and Energy Efficiency". In: *Journal of Law and Economics* 57.3, pp. 781 –820.

Sexton, Steven (May 2015). "Automatic Bill Payment and Salience Effects: Evidence from Electricity Consumption". In: *Review of Economics and Statistics* 97.2, pp. 229–241.

Sugden, Robert (1984). "Reciprocity: The Supply of Public Goods Through Voluntary Contributions". In: *The Economic Journal* 94.376, pp. 772–787.

Thaler, Richard H and Cass R Sunstein (May 2003). "Libertarian Paternalism". en. In: *American Economic Review* 93.2, pp. 175–179.

USBR (2009). *Report of Operations Monthly Delivery Tables. Years 1970-2009*. Tech. rep. US Bureau of Reclamation Mid-Pacific Region Central Valley Operations.

USCB (2012). *American Community Survey 2009-2011 Public Use Microdata*. Tech. rep. U.S. Census Bureau.

USGS (1995). *Ground Water Atlas of the United States, California, Nevada, HA-730-B, 220.* Tech. rep. U.S. Geographic Survey.

Volpp, Kevin G. et al. (Feb. 2009). "A Randomized, Controlled Trial of Financial Incentives for Smoking Cessation". In: *New England Journal of Medicine* 360.7, pp. 699–709.

Wooldridge, Jeffrey M. (2002). *Econometric Analysis of Cross Section and Panel Data.* 1st ed. MIT Press.

Appendix A

Supplemental Materials for Chapter 1

A.1 Notifications

The figure below shows the exterior of the mailed notification.

Notifications were sent in multiple languages:

- The mailers include information in English, Spanish, Chinese, and Tagalog.
- The emails include information in English, Spanish, and Chinese.
- The robo-calls in- clude options to hear the message in English, Spanish, and Chinese.
- Mobile phone texts are in English only.

A.2 Breakdown of Contact Methods Used

Table A.1 shows the breakdown of the type of contact made to each experimental group.

Table A.1: Contact Methods Sent to Each Experimental Group

	Control Group	Treatment Group
Number of Observations	960	940
Mailer	100%	49%
Mailer + Phone	0%	18%
Mailer + Email	0%	3%
Mailer + Phone + Email	0%	7%
Mailer + Phone + Text	0%	11%
Mailer + Email + Text	0%	1%
All 4 Methods	0%	13%

A.3 Follow-up Survey and Results

Surveys were administered one of two ways to a subset of the study sample (i.e. the first 2116 customers that experienced leaks during the study period). Households with email contact info on file received an email with a hyperlink to a survey through Survey Monkey (just under half of customers). All other households received a mailer containing a hard copy of the survey. The hard copy survey was accompanied by a pre-paid return envelope, and also provided a link to the online Survey Monkey form as an option.

- Customers emailed with a survey had a response rate of 32%.
- Customers mailed a survey had a response rate of 24%.
- Overall response rate was 28%.
- Of those who responded 55% filled out the online survey from an emailed link, 8% filled out an online survey from a mailed link, and 38% sent in hard copies.

The follow-up survey questions are presented below, followed by a summary of the responses in Table A.2 and A.3. For reference, the question numbers in the tables correspond to the question numbers in the survey presented.

SFPUC Leak Alert Survey

Thank you for participating in the SFPUC's leak alert survey! Your answers will be kept confidential.

1. Do you recall receiving a notice in the past year from the SFPUC about continuous water use (aka Leak Alert) at your home or property?
   - ☐ Yes
   - ☐ No (if no, please skip to Question #4)

Definition: Common sources of ***continuous water use*** *in homes are leaky toilets, other leaky plumbing fixtures, burst pipes, or faulty irrigation systems.*

2. Which of the following ways did SFPUC contact you about continuous water use? (Please check all that apply)
   - ☐ Email ☐ Text Message
   - ☐ Phone call ☐ Mailed letter
   - ☐ Other (please specify: ______________________________________)

3. Were you aware of the continuous water use before receiving a notice from the SFPUC?
   - ☐ Yes, I had already resolved the situation before receiving the notice
   - ☐ Yes, but I had not resolved the situation yet
   - ☐ No, I was not aware of the continuous usage until I received the notice

4. Do you have a preference of how to be contacted by the SFPUC in the future about continuous water use? (Please check all methods you prefer)
   - ☐ Email ☐ Text Message
   - ☐ Phone call ☐ Mailed letter
   - ☐ No preference
   - ☐ Other (please specify: ______________________________________)

5. Were you able to determine the cause of continuous water use?
   - ☐ Yes, the cause was the following (please check all that apply):
     - ☐ Leaky toilet
     - ☐ Leaky faucet
     - ☐ Leaky showerhead
     - ☐ Leaky irrigation system
     - ☐ Burst Pipe
     - ☐ Irrigation or hose that was left on
     - ☐ Sink, shower, or bath were left running
     - ☐ I prefer not to say
     - ☐ Other (please specify: ________________________________)
   - ☐ No
   - ☐ I don't know

6. What resources were helpful in learning about and trying to identify the cause of continuous water use? (Please check all that apply)
   - ☐ My Account (SFPUC online account)
   - ☐ SFPUC online home leaks tips web page or leak alert guide
   - ☐ SFPUC Customer Service phone line
   - ☐ SFPUC Customer Service email
   - ☐ SFPUC inspection in my home or property
   - ☐ Private plumber
   - ☐ Nothing was helpful
   - ☐ I don't know
   - ☐ Other (please specify: ____________________________________)

7. How did you go about trying to stop the continuous water use? (Please check all that apply)
   - ☐ Called a plumber
   - ☐ Called the landlord
   - ☐ Tried to fix it myself
   - ☐ Replaced an old plumbing fixture or appliance with a new model:
     - ➢ What fixture or appliance did you replace? ______________________________
     - ➢ Is it more water-efficient than the previous one?
       ☐ Yes ☐ No ☐ I don't know
   - ☐ Other (please specify: ____________________________________
     ____________________________________)
   - ☐ I did not try to stop the continuous water usage (if so, please skip to Question #10)
   - ☐ I don't know

8. Were you able to stop the continuous water use at your home or property?
   - ☐ Yes
   - ☐ No
   - ☐ I don't know

9. Did it cost money to stop (or attempt to stop) the continuous water use?
   - ☐ Yes, approximately $___________________
   - ☐ No
   - ☐ I don't know

Please continue on back side >>>

San Francisco Water Power Sewer
Services of the San Francisco Public Utilities Commission

10. Was there any housing damage due to the continuous water use?
    - ☐ Yes, approximately $___________________
    - ☐ No
    - ☐ I don't know

11. Approximately how much total time did you (or others) spend on addressing the continuous water use problem?
    - ☐ Less than 1 hour
    - ☐ 1-3 hours
    - ☐ 3-10 hours
    - ☐ More than 10 hours
    - ☐ I don't know

Questions about your household:

12. Do you live at this property?
    - ☐ Yes
    - ☐ No
    - ☐ I prefer not to say

13. Do you pay the water bill for this property?
    - ☐ Yes
    - ☐ No
    - ☐ I prefer not to say

14. Are you the landlord or the tenant of this property?
    - ☐ Landlord/Owner
    - ☐ Tenant/Renter
    - ☐ Other (please specify: _____________________________________)
    - ☐ I prefer not to say

15. How many people live at this property more than 50% of the time?

    ____

16. What is the relationship of the occupants of this property?
    - ☐ Single family
    - ☐ Housing share among roommates or multiple families
    - ☐ Business operation
    - ☐ Other (please specify: _____________________________________)
    - ☐ I prefer not to say

17. Describe the head decision-maker of the water account that serves this property:
    - ➢ Age:
      - ☐ 18-34 years old
      - ☐ 35-49 years old
      - ☐ 50-64 years old
      - ☐ 65-79 years old
      - ☐ Over 80 years old
      - ☐ I prefer not to say
    - ➢ Gender Identity:
      - ☐ Female
      - ☐ Male
      - ☐ Other: ________________________
      - ☐ I prefer not to say
    - ➢ Education Level:
      - ☐ Some high school
      - ☐ High school / GED
      - ☐ Some college
      - ☐ Bachelor's degree
      - ☐ Graduate degree
      - ☐ I prefer not to say

18. What is the approximate combined total annual income of the occupants living at this property?
    - ☐ Less than $50,000
    - ☐ $50,000 - $99,999
    - ☐ $100,000 - $199,999
    - ☐ $200,000 - $299,999
    - ☐ Over $300,000
    - ☐ Unsure
    - ☐ I prefer not to say

19. Can we reach out to you with further questions in the future?
    - ☐ Yes
      - ➢ Please provide an email address:

        ____
    - ☐ No

Table A.2: Summary of Survey Results

Question	Answers						
1. recall	Yes	No	skipped				
	87%	12%	1%				
2. contact	Email	Text message	Phone call	Mailed letter	Other (please s	skipped	
	22%	8%	10%	66%	2%	15%	
3. aware	Yes, I had alre	Yes, but I had	No, I was not a	skipped			
	13%	9%	62%	16%			
4. pref	Email	Text message	Phone call	Mailed letter	No preference	Other (please s	skipped
	54%	35%	26%	45%	4%	2%	4%
5. cause	Leaky toilet	Leaky faucet	Leaky showerhea	Leaky irrigatio	Burst Pipe	Irrigation or h	Sink, shower, o
	54%	6%	2%	9%	4%	6%	2%
6. resources	My Account (SFP	SFPUC online ho	SFPUC Customer	SFPUC Customer	SFPUC inspectio	Private plumber	Nothing was hel
	16%	10%	9%	7%	10%	24%	8%
7. how	Replaced an old	Called a plumbe	Called the land	Tried to fix it	I did not try t	Other (please s	skipped
	26%	37%	5%	24%	2%	22%	6%
7a. eff	Yes	No	I don't know	skipped			
	16%	3%	5%	76%			
8. stop	Yes	No	I don't know	skipped			
	81%	5%	8%	6%			
9. cost	Yes, approximat	No	I don't know	skipped			
	57%	24%	12%	7%			
10. damage	Yes, approximat	No	I don't know	skipped			
	2%	82%	9%	7%			
11. time	Less than 1 hou	1-3 hours	3-10 hours	More than 10 ho	I don't know	skipped	
	27%	33%	15%	9%	9%	8%	

Table A.3: Summary of Survey Results - Demographics

Question	Answers							
12. live	Yes	No	I prefer not	skipped				
	77%	13%	3%	7%				
13. pay	Yes	No	I prefer not	skipped				
	90%	3%	1%	6%				
14. landlord	Landlord/Owne	Tenant/Renter	Other (please	I prefer not	skipped			
	75%	10%	6%	3%	7%			
16. relations	Single family	Housing share	Business oper	Other (please	I prefer not	skipped		
	68%	10%	1%	6%	5%	9%		
17a. age	18-34 years o	35-49 years o	50-64 years o	65-79 years o	Over 80 years	I prefer not	skipped	
	6%	17%	25%	27%	10%	6%	8%	
17b. gender	Female	Male	Other (please	I prefer not	skipped			
	36%	46%	3%	8%	8%			
17c. educ	Some high sch	High school /	Some college	Bachelor's de	Graduate degr	I prefer not	skipped	
	2%	6%	15%	26%	33%	11%	9%	
18. income	Less than $50	$50,000 - $99	$100,00 - $19	$200,000 - $2	Over $300,000	Unsure	I prefer not	skipped
	11%	9%	15%	5%	9%	6%	35%	9%

A.4 Implied Savings (Back-of-the-Envelope Calculation)

Below details the back-of-the-envelope calculation of the implied average savings to customers from an informational mailer: If customers were not sent an informational notification, they would likely learn about their leak through a signal from an abnormally high bill (aside from a physical leak clue). Based on the previous analysis (Section 1.5.2), customers on average only notice a 50% increase or more on the bill. Since the first bill into a leak has a median increase of 40% only half of customers would notice something was wrong on their upcoming water bill (an average of 16 days from the beginning of the leak) and the other half would notice a month later when they receive the next bill (an average of 46 days from the beginning of the leak). Therefore, customers will notice the leak from the increased bill an average of 31 days into their leak. Meanwhile, customers receiving a mailer notification receive it on average 10 days into the leak, allowing them to learn about the leak 21 days faster than they would if they just relied on signaling from their increased bill. It is important to note that not all customers need a notification to start fixing their leak: 56% of customers in the Control Group fix their leaks before receiving a mailer notification. This means that for 56% of customers, the benefit of receiving a mailer notification is zero. Given that the average leak size for the 44% of customers that do benefit from the mailer is 2.6 cf/hr, leaking for an additional 21 days equates to $261 (i.e. $4.79 per day for 1cf/hr * 2.6 cf/hr * 21 days) on average for these customers. The weighted average across all customers, 44% that get $261 of benefit and 56% of customers that get zero benefit, comes to $115.

Appendix B

Robustness Checks for Chapter 2

Table B.1: State-of-Emergency Declaration (Event 1)

	Varying Specifications (14 Day Event Window)							
	Log(Consumption)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Event Indicator	-0.016***	-0.015***	-0.018**	-0.018**	-0.019**	-0.021***	-0.019**	-0.019**
	(0.006)	(0.005)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Maximum Daily Temperature (F)	0.002***	0.002***	0.002***	0.002***	0.002***	0.001***	0.002***	0.002***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0005)	(0.0001)	(0.0001)
Daily Precipitation (In)	-0.021***	-0.015***	-0.012***	-0.002	-0.010***	-0.066	-0.008**	-0.009**
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.056)	(0.004)	(0.004)
Average Price (Dollars/CF)	-0.058***	-0.011***	-0.014***	-0.014***	-0.012***	-0.029**		-0.012***
	(0.012)	(0.003)	(0.004)	(0.004)	(0.004)	(0.013)		(0.004)
Marginal Price (Dollars/CF)							-0.026***	-0.022***
							(0.008)	(0.008)
Household FE	N	Y	Y	Y	Y	Y	Y	Y
Household X Day-of-Year FE	N	N	Y	Y	Y	Y	Y	Y
Holiday X Day-of-Week FE	N	N	N	Y	Y	Y	Y	Y
Year FE	N	N	N	N	Y	Y	Y	Y
Data Set Used	Full	Full	Full	Full	Full	Restricted	Full	Full
Price Used	Average	Average	Average	Average	Average	Average	Marginal	Both
R^2	0.053	0.547	0.692	0.692	0.693	0.708	0.692	0.693
Adjusted R^2	0.053	0.547	0.537	0.538	0.539	0.547	0.537	0.539
Residual Std. Error	0.878	0.608	0.614	0.613	0.613	0.604	0.614	0.613

Note: *p<0.1; **p<0.05; ***p<0.01
Column (5) represents the preferred specification from the Results section.
Column (6) only uses data restricted to the analysis time frame in each year.

Table B.2: Redoubling State Drought Actions (Event 2)

	Varying Specifications (14 Day Event Window)							
	Log(Consumption)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Event Indicator	0.004	0.004	-0.003	0.003	-0.001	-0.002	-0.001	-0.001
	(0.006)	(0.005)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Maximum Daily Temperature (F)	0.002***	0.002***	0.002***	0.002***	0.002***	0.001***	0.002***	0.002***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0003)	(0.0001)	(0.0001)
Daily Precipitation (In)	-0.022***	-0.016***	-0.016***	-0.006	-0.011***	-0.026	-0.009**	-0.009**
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.028)	(0.004)	(0.004)
Average Price (Dollars/CF)	-0.058***	-0.011***	-0.014***	-0.014***	-0.012***	-0.021**		-0.012***
	(0.012)	(0.003)	(0.004)	(0.004)	(0.004)	(0.009)		(0.004)
Marginal Price (Dollars/CF)							-0.045***	-0.042***
							(0.008)	(0.008)
Household FE	N	Y	Y	Y	Y	Y	Y	Y
Household X Day-of-Year FE	N	N	Y	Y	Y	Y	Y	Y
Holiday X Day-of-Week FE	N	N	N	Y	Y	Y	Y	Y
Year FE	N	N	N	N	Y	Y	Y	Y
Data Set Used	Full	Full	Full	Full	Full	Restricted	Full	Full
Price Used	Average	Average	Average	Average	Average	Average	Marginal	Both
R^2	0.053	0.547	0.692	0.692	0.693	0.694	0.692	0.693
Adjusted R^2	0.053	0.547	0.537	0.538	0.539	0.541	0.537	0.539
Residual Std. Error	0.878	0.608	0.614	0.613	0.613	0.605	0.614	0.613

Note: *p<0.1; **p<0.05; ***p<0.01

Column (5) represents the preferred specification from the Results section.

Column (6) only uses data restricted to the analysis time frame in each year.

Table B.3: 25 Percent Mandatory Reduction (Event 3)

	Varying Specifications (14 Day Event Window)							
	Log(Consumption)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Event Indicator	-0.046***	-0.048***	-0.055***	-0.055***	-0.046***	-0.053***	-0.046***	-0.046***
	(0.008)	(0.006)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Maximum Daily Temperature (F)	0.002***	0.001***	-0.0003	-0.001	0.001***	0.001	0.001***	0.001***
	(0.0003)	(0.0002)	(0.0004)	(0.0004)	(0.0001)	(0.0004)	(0.0001)	(0.0001)
Daily Precipitation (In)	-0.016***	-0.033***	-0.033***	-0.028***	-0.014***	-0.040***	-0.010***	-0.010***
	(0.005)	(0.003)	(0.004)	(0.004)	(0.004)	(0.014)	(0.004)	(0.004)
Average Price (Dollars/CF)	-0.057***	-0.012***	-0.015***	-0.014***	-0.012***	-0.018**		-0.012***
	(0.011)	(0.003)	(0.004)	(0.004)	(0.004)	(0.009)		(0.004)
Marginal Price (Dollars/CF)							-0.048***	-0.042***
							(0.008)	(0.008)
Household FE	N	Y	Y	Y	Y	Y	Y	Y
Household X Day-of-Year FE	N	N	Y	Y	Y	Y	Y	Y
Holiday X Day-of-Week FE	N	N	N	Y	Y	Y	Y	Y
Year FE	N	N	N	N	Y	Y	Y	Y
Data Set Used	Full	Full	Full	Full	Full	Restricted	Full	Full
Price Used	Average	Average	Average	Average	Average	Average	Marginal	Both
R^2	0.053	0.547	0.692	0.692	0.693	0.694	0.692	0.693
Adjusted R^2	0.053	0.546	0.537	0.538	0.539	0.538	0.538	0.539
Residual Std. Error	0.878	0.608	0.614	0.613	0.613	0.615	0.614	0.613

Note: *p<0.1; **p<0.05; ***p<0.01

Column (5) represents the preferred specification from the Results section.

Column (6) only uses data restricted to the analysis time frame in each year.