Essays in Water Conservation and Water Quality Programs

Essays in Water Conservation and Water Quality Programs

A dissertation presented

by

Jonathan Early Baker

to

The Department of Public Policy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Public Policy

Harvard University

Cambridge, Massachusetts

May 2017

Dissertation Advisor:
Robert N. Stavins

Author:
Jonathan Early Baker

Essays in Water Conservation and Water Quality Programs

Abstract

As growing populations continue to drive demand for water, managers of this fundamental resource face the dual challenge of providing both sufficient and clean supplies. In this dissertation, I undertake two analyses exploring a policy aimed at maintaining sufficiency of supply through conservation, and a third analysis evaluating a regulatory approach to promote water quality.

Because lawns comprise a large share of residential water demand, water utilities across the western United States offer subsidies to replace lawns with less water intensive landscape. In my first analysis, I estimate the water savings and property value effects of one such subsidy, the Southern Nevada Water Authority's "Cash-for-Grass" rebate program. Using event studies and panel fixed-effects models, I find that the average conversion reduces baseline water consumption by 21 percent and increases property values by about 1 percent, however I find little evidence of property value spillovers to neighboring properties. I also show that the smallest savings coincide with years in which many conversions took place, suggesting a possible trade-off between program participation and maintaining the effectiveness of individual conversions. I also find that participants with high pre-conversion water demand save more water than participants with lower pre-conversion water demand. Conservation subsidies present attractive alternatives to price-based approaches to water conservation. However, I find that a relatively modest 6 percent price increase may have achieved equivalent savings. Finally, combining my water savings and housing price impact results, I show that the program generates net benefits of $2.00 per square foot of desert landscape converted. My results expand our knowledge of water conservation rebates and,

more generally, contribute to our understanding of the long-term dynamics of conservation rebate program savings as well as how heterogeneous participant characteristics affect conservation rebate program performance.

In my second analysis, which is joint work with Sheila Olmstead, we study the effect that an information disclosure policy has on national water quality violations. Since the 1980's, information disclosure policies have grown in popularity as a means by which to promote policy outcomes where direct regulation is a challenge. In this spirit, the 1996 amendments to the Safe Drinking Water Act require water utilities to disclose drinking water violations to their customers in annual water quality reports. We explore the impact of these reports on health-based drinking water quality violations in a differences-in-differences framework using a nationally comprehensive data set of water quality violations and water systems. Our results suggest that reports published in local newspapers or mailed directly to customers may have reduced violations, but we uncover less evidence that posting reports online had any impact. We also show that reductions in violations remain stable over time, and that the effect of the reports appears to be stronger for those water systems serving higher-income counties. Finally, we provide evidence that implies reports induce reductions in microbial contaminants without increasing disinfection byproducts. Our analysis is among the first to explore the long-term impacts of information disclosure policies, and builds on the small but growing literature exploring heterogeneity in the responsiveness to these policies.

I return to the Las Vegas Cash-for-Grass program in my third analysis. Following the empirical framework and analysis of Bollinger and Gillingham (2012), I estimate the presence of peer effects in the Cash-for-Grass program. Like Bollinger and Gillingham, I find positive peer effects that grow with time when defining the peer network by a zip code, but unlike these authors, I show that only the cumulative conversions rather than the area of conversions drive these effects. Overall, however, my results largely validate the method of Bollinger and Gillingham. Going beyond their results, I also estimate peer effects within zip codes interacted with deciles of assessed home values. At the zip code-decile, I find peer effects to be even stronger than at the zip code, suggesting that a possible driver of the peer effect works through a desire to maintain competitiveness with other homes in an individual's housing market. I also show that my estimated peer effects are at least an order of magnitude larger than the impact of several targeted marketing campaigns administered by the Southern Nevada Water Authority. To my knowledge, mine is the first analysis to compare a peer effect with the impact of advertising efforts.

Overall, my dissertation contributes to environmental economists' understanding of water conservation and water quality policies and aims to improve the management of one of civilizations' most critical resources.

Contents

List of Tables

List of Figures

Acknowledgments

I am grateful to many individuals who supported me throughout my doctoral studies. First, I am grateful for my dissertation committee members, Rob Stavins, Alberto Abadie, and Sheila Olmstead. I have greatly appreciated and benefited much from my conversations with each of them, and I am most grateful for their guidance and input throughout the research, writing, and job search process. Thanks also to Joe Aldy, who provided excellent feedback on chapter 1. I am also grateful to Chris Avery and Marty Weitzman from whom I learned much about teaching economics. Nicole Tateosian makes this program work, and has been a much needed go-to for many phd-induced quandaries...and free candy. Bill Clark provided helpful guidance early in my doctoral studies. Ryan Sheely was also instrumental in guiding me through the early years, introduced me to the fascinating field of institutions, and proved to be a much needed source of encouragement during times of existential angst. And finally, I am very grateful to the late John Briscoe. It was a privilege to serve as one of his teaching fellows, he provided much practical wisdom regarding water policy, and was instrumental in two of the highlights of my doctoral studies; a river management research initiative that took me and others down the Mississippi, and a second set of travels through Nevada that eventually lead to chapter 1 of this dissertation.

I also received much helpful input from Todd Gerarden, Nick Hagerty, and Evan Herrnstadt. I cannot thank these guys enough. My analyses have also benefited from conversions with Patrick Behrer, Rich Sweeney, and Sam Stolper. Rich, Sam, and Liz Walker additionally provided much needed and appreciated advice for navigating the phd.

Many others provided helpful discussions regarding individual analyses. In writing chapter 1, I benefited from discussions with Kirill Borusyak, Rema Hanna (who also proved incredibly helpful throughout the job market process), and Peter Mayer. I am also grateful for comments from seminar participants at Harvard, conference attendees at the AERE 2016 summer conference in Breckenridge, CO, and conference attendees at the AAEA 2016 annual meeting in Boston. Regarding chapter 2, I am grateful to the EPA for providing data, and in particular to EPA staff member Kevin Roland for many helpful discussions and clarifications regarding these data. Thanks also to Simo Goshev of Harvard's Institute for Quantitative Social Science for his advice regarding aspects of our empirical strategy. This analysis has also benefited from input by attendees of the Harvard Environmental Economics Lunch and Harvard Public Policy API-902 seminar participants. Matt Baum, Lori Bennear, and Richard Zeckhauser also provided helpful input. Regarding chapters 1 and 3, thanks to Jeff Blossom of Harvard's Center for Geographic Analysis for his assistance with GIS. I am also grateful to Erik Martinet and the Clark County Assessor's office for their provision of and time spent helping me understand their data. Finally, I am deeply indebted to the Southern Nevada Water Authority staff, in particular Kent Sovocool and Mitchell Morgan, for their invaluable assistance in collecting, providing, and helping me understand their data and the Cash-for-Grass rebate program. I could not have done this work without their generous support and cannot thank them and their colleagues enough. I also gratefully acknowledge financial support from the Vicki Norberg-Bohm Fellowship, which funded my Nevada travels, the Joseph Crump Fellowship, and the Graduate School's dissertation fellowship.

Finally, my doctoral studies would have been unimaginably more challenging without the anchoring relationships of my parents, Mark and Debbie, my sister Rebekah Melton, and many, many friends. At Harvard, in addition to those already mentioned, many thanks to Megan Bailey, Nate Fleming, and Kim Smet. Thanks also to Dave Chorney, Wayne Levy, Jared Markowitz, Ian Nurse, and all my friends on the BAA. I have also been blessed with many grounding relationships made through Parkstreet Church and more recently Church of the Cross. Particular thanks to my COTC neighborhood group, Will Abbott, Matt Carey, Dan Cogswell, Chris and Kara Dodds, Alex and Amy Frenzel, Sarah Haig, Andy Huss, Adam Kurihara, Seth Van Liew, and Andrew Noh. And a final special thanks to John and Jessi ZuHone, and my roommate, Colby Steiner.

Chapter 1

Subsidies for Succulents: Evaluating the Las Vegas Cash-for-Grass Rebate Program

1.1 Introduction

Water is not an abundant resource. In the western United States, municipal water demand threatens to outstrip supplies as new water customers increase service populations and droughts decrease reservoir levels. In the past, water utilities relied on large, often federally funded, water supply augmentation projects. Today, however, these projects are costly and unpopular. And while agriculture holds large amounts of historical water rights, legal institutions, social perceptions and equity concerns appear to inhibit free trade between agricultural and urban users (Libecap, 2007; Howe *et al.*, 1990). Faced with essentially fixed supplies and growing demand, water utilities have responded with various incentive-based and command-and-control demand-side-management strategies. Common strategies include drought awareness campaigns, lawn irrigation and other water use restrictions, and subsidies for water saving capital investments (Price *et al.*, 2014).

I estimate the savings and impacts on property values of the Southern Nevada Water

Authority's Water Smart Landscapes program, or Cash-for-Grass program. The program aims to reduce water demand by offering Las Vegas water customers a subsidy for replacing lawns with desert landscape. Of all the water demand management strategies, those that target lawn irrigation may achieve the greatest reductions in demand. Residential lawn irrigation accounts for over half of all residential water consumption in southwestern cities (Sovocool *et al.*, 2006), and in Las Vegas, nearly 40 percent of total water deliveries goes towards residential outdoor uses.[1] Especially in light of the growing popularity of turf replacement subsidies, comprehensive evaluations of programs targeting outdoor water use are relevant and timely.

I combine over 25 years of single-family monthly water consumption data for the Las Vegas Valley Water District (LVVWD) with single-family rebate participant information since the beginning of the program in 1996 to estimate water savings associated with Cash-for-Grass subsidized conversion to desert landscape. I use event study and panel fixed-effects models to visualize and quantify average water savings. I further test robustness of my results to three additional control samples that endeavor to account for any remaining concerns over selection bias not accounted for in my main specification. First, I estimate a model without non-participants. Second, I build a control sample from non-participants who applied for the rebate, but never completed the conversion process. Third, I develop a matched sample of non-participants, matching on pre-conversion water consumption and lot size.

I find an average conversion reduces monthly water consumption by 21 percent, or 5,000 gallons per month, with this result robust to my three alternative control samples. Savings remain stable over time, suggesting that program participants do not increase other water intensive activities after conversion, but simply reduce outdoor irrigation. I also find that water savings fall over time as rising rebate levels increase the number of annual rebate recipients. The inverse relationship between savings and incentive to participate highlights

[1]The water authority estimates that 59 percent of total water deliveries go to residential customers (SNWA, 2009). Mansur and Olmstead (2012) find that outdoor water use accounts for two thirds of total water use in arid environments. The product of these two values approximately equals 40 percent.

a possible trade-off between expanding the reach of the subsidy and maintaining the effectiveness of conversions. In addition to exploring heterogeneity in savings across time, I also explore heterogeneity in savings across participant type. I find that participants with high pre-conversion water demand save more water (relative to lot size) than participants with low pre-conversion water demand.[2] Finally, using a partial equilibrium framework, I estimate that a 6 percent price increase experienced by the entire service population in place of the rebate program would have achieved the same aggregate savings as did conversions subsidized by the program. [3]

Most evaluations of demand-side-management strategies focus on electricity programs.[4] Furthermore, these studies tend to estimate short-term impacts and there exist few analyses exploring the effect of participant heterogeneity on program outcomes (Allcott and Greenstone, 2012).[5] My results expand our knowledge of water conservation rebates and, more generally, contribute to our understanding of the long-term dynamics of conservation rebate program savings as well as how heterogeneous participant characteristics affect conservation rebate program performance.

I also estimate the effect of conversion to desert landscape on property values in the LVVWD. The impact of conversion on property values will incorporate all private impacts associated with conversion such as water bill savings, reduced lawn maintenance costs, negative energy spillovers from increased urban heat island effects (Klaiber *et al.*, 2015), and any non-monetary impacts such as aesthetic appeal. Because energy spillovers and aesthetic

---

[2]Deoreo *et al.* (2000) come to a similar conclusion.

[3]After completing my analysis, I became aware that another dissertation explores water savings from the Cash-for-Grass program (Brelsford, 2014). While Brelsford uses different methods and census tract consumption data, similar to my results Brelsford finds substantial savings due to the rebate program. But in contrast to my results, Brelsford finds that savings erode over time. Brelsford acknowledges, however, that "a difference in differences approach on household level data would provide a more rigorous estimate of the long term water savings generated through the [Cash-for-Grass] program." I use this approach in my analysis. And while I have not found a paper to reference, I am aware that ongoing work by Brelsford and Josh Abbott of Arizona State University also use this differences-in-differences approach to estimate Cash-for-Grass induced water savings.

[4]For example, see Alberini and Towe (2015), Davis *et al.* (2014), and Arimura *et al.* (2011).

[5]Reflecting the broader demand-side-management program evaluation literature, most studies of which I am aware that give attention to heterogeneity derive their conclusions from energy conservation programs (Allcott, 2011; Allcott *et al.*, 2015).

appeal may impact neighboring homes, I also estimate the spillover effect of conversion to desert landscape by modeling the impact of conversions on adjacent property values. While past studies have explored peer effects in capital investment subsidies (Bollinger and Gillingham, 2012) and researchers have recognized the importance of policy induced externalities in other contexts (Miguel and Kremer, 2004), to my knowledge my analysis is the first empirical investigation to explore the spillover effects of capital investment subsidies on neighboring properties.

I use Cash-for-Grass program enrollment information combined with Clark County Assessor data on property sales to estimate the value of desert landscape within a hedonic property framework. I employ a differences-in-differences strategy, controlling for unobserved characteristics with a rich set of spatial and temporal fixed-effects. I find that conversion increases the value of a home by about 1 percent, or $3,700, and has little impact on neighboring homes. I additionally find that the existence of a conversion, rather than the area converted, drives the increase in property values. I estimate that the present discounted value of annual water bill and lawn maintenance savings approximately equals the increase in home values, suggesting that the hedonic estimate of the direct effect of conversion reflects little more than the monetary savings associated with conversion.

Considering water savings, administrative costs, rebate outlays, and out-of-pocket conversion costs to the rebate recipient, I find that the Cash-for-Grass program costs $4.84/kgal-saved. I estimate the cost of water supply as the sum of the annual water bill for an average single-family customer and the opportunity cost of scarce water, which I base on Nevada agriculture to urban water sales. Comparing the two cost estimates, I find that the program saves water for less than the cost of supply. I also calculate net benefits by subtracting the sum of administrative and conversion costs from the sum of the direct effect of conversion and the value of scarce water. The program generates net benefits equal to $2.00/ft^2 of desert landscape converted.

I organize the remainder of the paper as follows. Section 1.2 provides background on the Cash-for-Grass program. I present my analysis of water savings in section 1.3. Section

1.4 describes my estimates of the direct and spillover effect of conversion to desert landscape on property values. Section 1.5 estimates program costs per gallon saved and program net benefits. Section 1.6 concludes.

1.2 The Cash-for-Grass rebate program

The Southern Nevada Water Authority's (SNWA) Cash-for-Grass rebate program is a voluntary, incentive based demand-side-management program that provides a cash rebate for Las Vegas area water customers that replace their lawns with desert landscape.[6] of The program began as a pilot study in 1996 and was rolled out to all customers beginning in 1998. Though the program has undergone several administrative regime changes,[7] throughout most of the program's history participants have received a one-time check from the water authority determined by the size of the conversion. Currently, program participants receive \$2.00 per square foot of lawn replaced[8] and can receive a maximum rebate of \$300,000.[9] The program also stipulates customers convert a minimum area and requires that conversions

[6]Current and historical program details derive from conversations with and information sent by SNWA staff members.

[7]At the beginning of the program, single-family participants received a \$5 water bill credit for every 1000 gallon reduction relative to baseline average water use. Halfway through the year 2000, however, the water authority began issuing rebates (still in the form of a water bill credit) based on the size of the conversion. Cash-for-Grass program participants continued to receive rebates in the form of a water bill credit until March, 2003, when the water authority began sending one-time checks.

[8]https://www.snwa.com/rebates/wsl.html. The \$2.00 per square foot rebate is valid for conversions of 5,000 square feet or less. Beyond 5,000 square feet, the rebate falls to \$1.00 per square foot.

[9]There has always been a maximum allowable rebate, but this limit does not appear to affect participant decision making, at least since the Fall of 2001 when the limit was set at \$25,000. Since this date, no participant has even approached 75 percent of the limit. Early on in the program, however, limits may have been binding. Initially, single-family participants could receive no more than \$400, and during this stage of the program, 38 percent of participants received a rebate of \$400. Program requirements soon changed such that single-family participants received \$0.40 per square foot of lawn converted to desert landscape for the first 2,500 square feet of lawn replaced. In other words, participants could earn up to \$1,000, even if they converted more than 2,500 square feet. Under this regime, 11 percent of participants converted more than 2,500 square feet. For both the \$400 and \$1,000 limits, however, the distributions of area converted appear to be reasonably continuous, suggesting that these early limits on the allowable rebate had little effect on participants' decision of how large a conversion to undertake.

remain in place in perpetuity.[10] Very few customers renege on their agreement to maintain their conversion.[11]

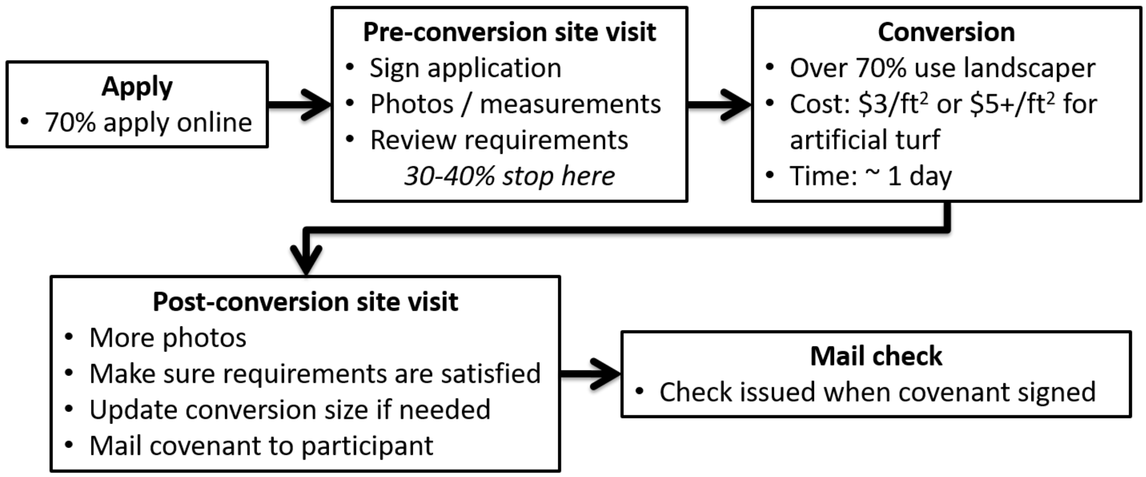

Figure 1.1: *Illustration of the process required for a single-family customer to become enrolled in the Cash-for-Grass program.*

To become a program participant, single-family applicants undergo a multi-step process illustrated in Figure 1.1 that begins with submitting an application and culminates with receiving their check. In between, program applicants review and verify requirements with water authority staff during pre- and post-conversion site visits as well as undertake the conversion itself. Most applicants hire a professional landscaper. In 2014, for example, only about a quarter of applicants performed the conversion themselves.[12] For those that employ a landscaper, conversions cost around $3 per square foot, though the cost can be higher for those installing artificial turf (which does qualify as "desert landscape" under the terms of the Cash-for-Grass program).[13] It takes the average applicant a little over 5

[10] The water authority relaxed the minimum conversion requirement in 2004. I explore the impact of this change in Appendix A.2. Prior to June, 2009, customers agreed to keep the conversion in place for 10 years. I explore the impact of this change in Appendix A.3.

[11] I have "back conversion" data since 2004. These data show that the average number of back conversions occur at an annual rate of 4.5 back conversions per year, or less than one tenth of one percent of all conversions (pers. comm. K. Sovocool, February 2015).

[12] pers. comm. K. Sovocool, February 2015.

[13] Cost estimates derive from conversations the author had with several Las Vegas area landscape professionals

months from the date of application to the receipt of the check.[14] In the empirical models that follow, I attempt to account for the transient behavior likely present between when applicants first signal interest in the program (application) to when applicants complete the process (enrollment date).

1.3 Water savings

1.3.1 Data and summary statistics

I derive estimates of water savings using single-family Cash-for-Grass program participant information and monthly single-family residential water consumption data for the Las Vegas Valley Water District (LVVWD), the largest water utility in the Las Vegas region. The Southern Nevada Water Authority provided both data sets. Program participation data include all participants from the inception of the pilot study in 1996 through June 12, 2014. These data include the parcel identifier, size, and rebate value of the conversion, the participant type (e.g. single-family), for most conversions the date the program participant applied for the rebate, and for all conversions the date the program participant became enrolled in the program.

I focus on single-family participants that undertake one conversion. Single-family participants comprise nearly 90 percent of all conversions, and nearly 90 percent of single-family participants perform one conversion.[15] 80 percent of the observations of single-family participants undertaking one conversion include both the application and enrollment dates. Among these observations, the average period between application and enrollment date spans about 5 months (150.6 days). Since I use time relative to application date in my event studies, I proxy for missing application dates by subtracting the average 5-month time

in March, 2016.

[14]pers. comm. M. Morgan, October 27, 2015.

[15]About 10 percent of single-family participants undertake two or more conversions. Golf courses, however, are the most likely participant category to undergo multiple conversions. Of the 33 golf courses that have participated in the Cash-for-Grass program, 27 have undertaken multiple conversions.

period from the enrollment date for observations missing an application date.

Monthly single-family consumption runs from January, 1988 through April, 2014. I restrict consumption data to those service meters that supply a single parcel. Meters that serve multiple parcels or parcels served by multiple meters are primarily associated with larger properties.[16] I reassign negative water consumption values (less than 0.01 percent of observations) to 'zero' upon recommendation from water authority staff. Negative water consumption values can occur due to billing adjustments or corrections for over-estimated meter readings, which sometimes arise if a utility staff member cannot read a meter and must instead estimate that month's consumption.[17] Because I observe water consumption at the service meter level and match service meter identifiers with parcel identifiers, my unit of analysis is a parcel, not an individual. Service meter identifiers do not change when customers move and I do not observe changes in the name of the individual connected to a service meter.

Table 1.1: *Summary of water consumption panel.*

	All parcels	Participating parcels	Non-participating parcels
Monthly water use obs. (N)	64,135,652	6,580,788	57,554,864
Number of parcels	309,608	26,488	283,120
Mean water use (kgal/mo)	15.7	23.8†	15.1

† Derived from pre-enrollment water consumption observations.

Merging the water consumption with program enrollment information yields a panel of over 64 million monthly water use observations from 309,608 parcels. Of the 309,608 parcels, 26,488 parcels participate in the Cash-for-Grass program (about 9 percent). In addition, prior to program enrollment, participating parcels demand more water on average than non-participating parcels (23.8 kgal/month vs. 15.1 kgal/month). Table 1.1 summarizes

[16]pers. comm. M. Morgan, June 30, 2014.

[17]pers. comm. M. Morgan, October 28, 2015.

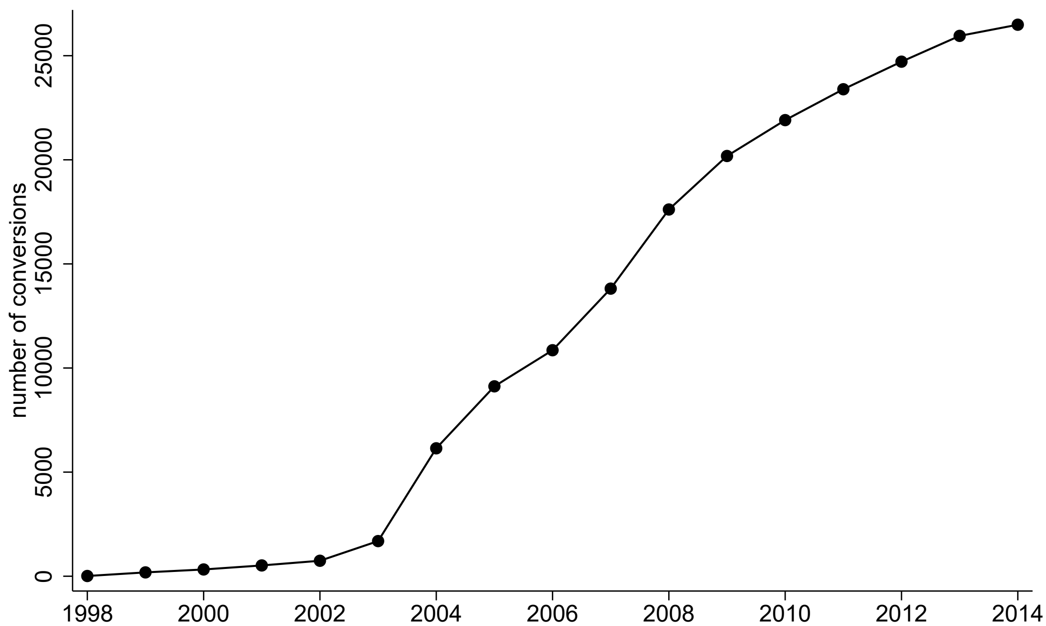

Figure 1.2: *Cumulative number of conversions over time performed under the Cash-for-Grass program (single-family parcels that converted only once).*

these results. Figure 1.2 illustrates the cumulative number of conversions over time. Few conversions took place in the late 1990s and early 2000s, however the number of annual participants increased sharply after 2003. A second jump in participation occurred between 2006 and 2008, but after 2008, program participation has steadily declined. Figure 1.3 illustrates the annual number of conversions and associated major changes in the subsidy rate and average water bill. In February of 2003, the rebate was increased from \$0.40 to \$1.00 per square-foot converted, and in September of the same year, the average water bill increased over 25 percent. These changes were followed by an increase in the number of conversions from 225 in 2002 to 945 in 2003 to 4,456 in 2004. The subsidy rate increased again in December, 2006 to \$2.00 per square-foot. This was followed by an approximate 8 percent increase in the average water bill in February, 2007. The number of conversions subsequently increased from 1,735 in 2006 to 2,958 in 2007. In January, 2008, the subsidy decreased to \$1.50 per square-foot and effectively remained at this level through June, 2014. Water rates, however, continued to increase. In May 2008, January 2010, and January 2011,

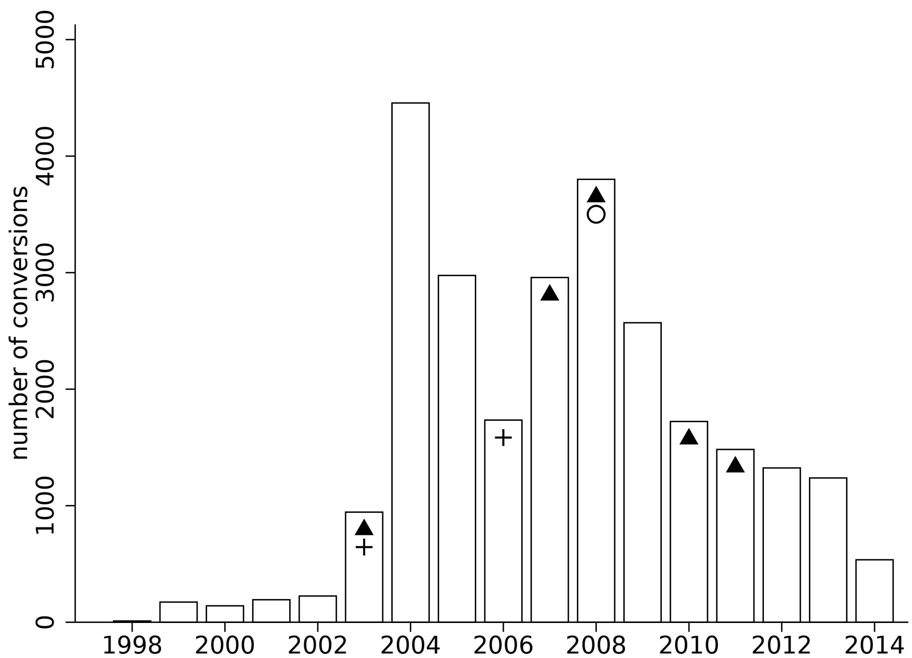

Figure 1.3: *Number of conversions by year performed under the Cash-for-Grass program (single-family parcels that converted only once). 'Plus' signs indicate increases in the subsidy level; the hollow circle indicates a decrease in the subsidy level; solid triangles indicate dates of major increases to a customer's average water bill. 2014 participation data only run through June 12.*

the average water bill increased by approximately 17.5 percent, 6 percent and 5.5 percent respectively. Despite these increases in price, participation steadily declined since 2008.

An average participant converts 1,348 square feet of lawn to desert landscape (approximately 0.03 acres), but average conversion area has fluctuated since program inception. As illustrated in Figure 1.4, in 2003 the average converted area peaked at 1,703 square feet, but declined to just over 1,000 square feet by 2014. Falling converted area could be because recent participants have less grass area to convert. The average property size in the LVVWD has shrunk in the past 30 years, and in 2004, communities began restricting new homes from planting grass in front yards.[18] Both factors would contribute to newer homes having smaller yards.[19]

[18] pers. comm. SNWA staff, March 14, 2016 and May 8, 2017.

[19] Additionally, since 2004, the water authority has allowed participants to convert less than 400 square feet provided the conversion covers an entire front or back yard. Relaxing the minimum conversion size requirement

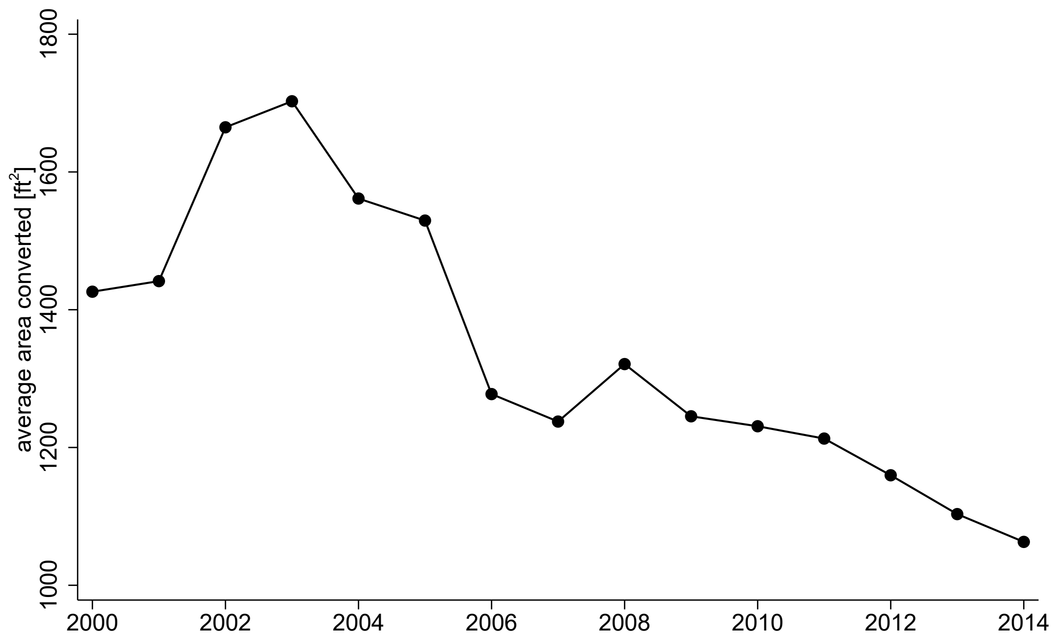

Figure 1.4: *Average size of conversions in each year of the Cash-for-Grass program. No averages are shown before 2000 since only one conversion has a recorded area in 1998 and none of the conversions undertaken in 1999 record the area converted (averages derived from single-family parcels that converted only once).*

Average water use among participating parcels prior to program enrollment runs strikingly parallel to average water use among non-participating parcels, especially before the program begins in 1998. Figure 1.5 summarizes annual water consumption since 1988 for participating parcels prior to Cash-for-Grass program enrollment, and non-participating parcels. Sample sizes for each group change over time due to new home construction, homes being removed from the LVVWD service area,[20] or enrollment of participating parcels into the rebate program. In the empirical models described below, I rely on a differences-in-differences design. In my context, the validity of a differences-in-differences design relies critically on the assumption that average water use among participating parcels would have paralleled average water use among non-participating parcels in the absence of the rebate program. The parallel trends in water use between both groups prior to the inception of the

may also contribute to falling average conversion size.

[20]In 1998, a large number of meters were transferred from the LVVWD to the Henderson water utility (pers. comm. M. Morgan, May 17, 2016).

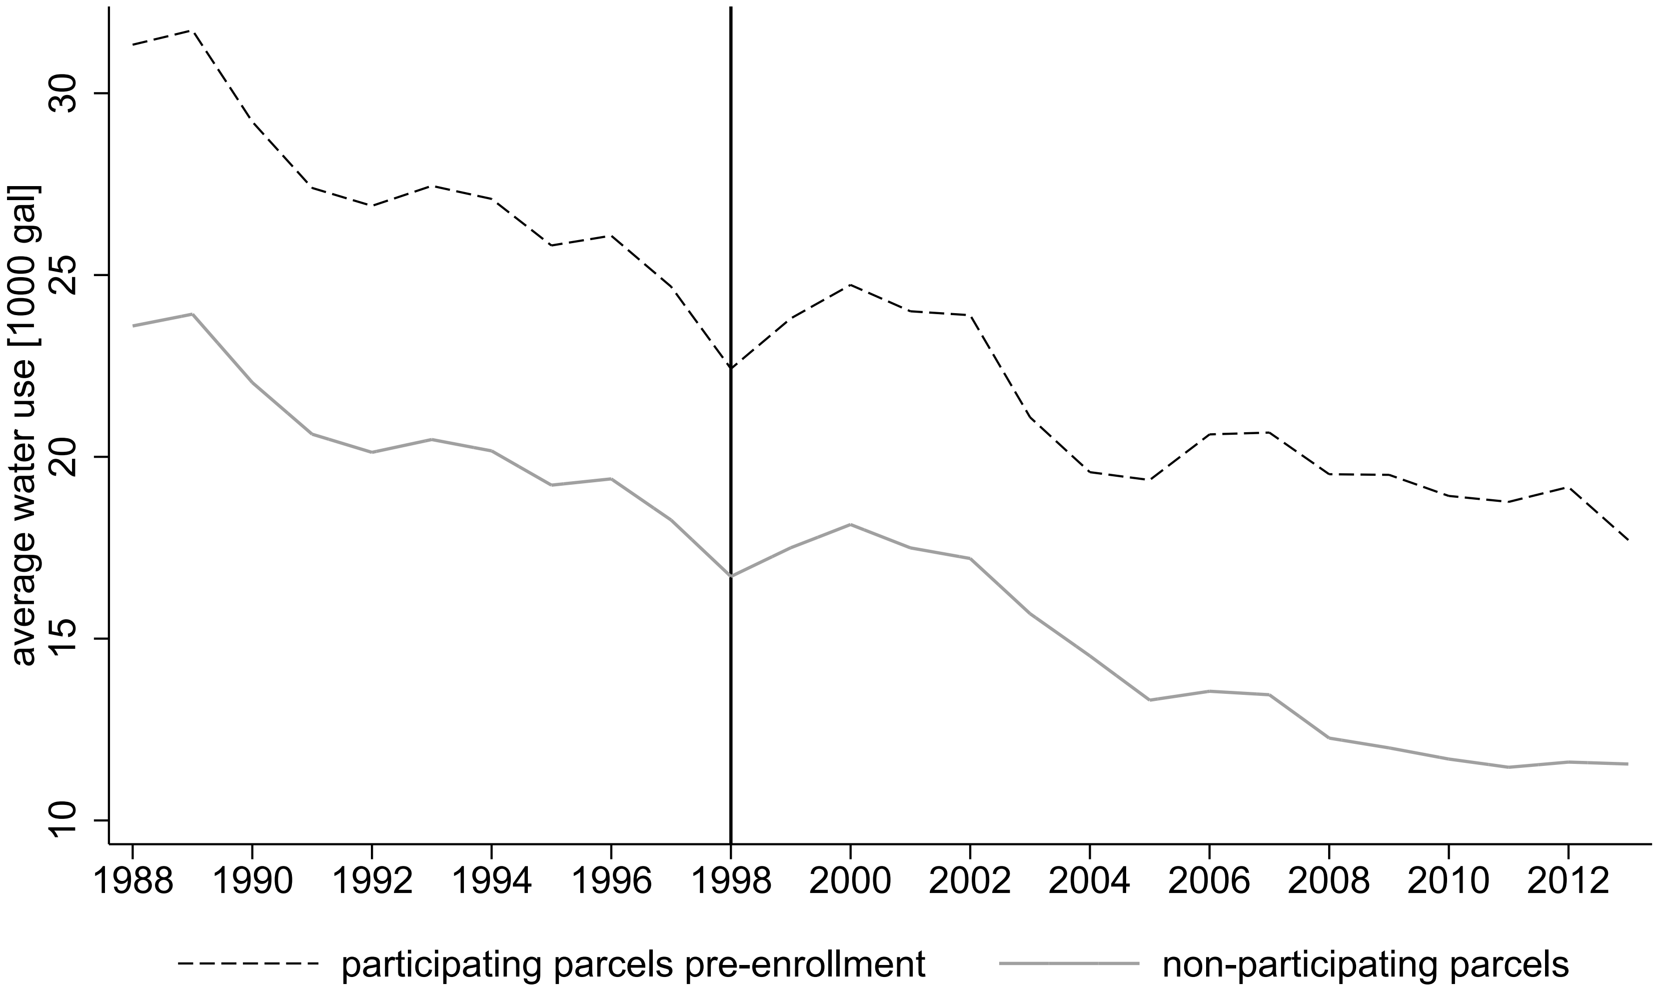

Figure 1.5: *Average annual water use from 1988 through 2013 (the last full year in my panel) for participating parcels prior to Cash-for-Grass program enrollment (dashed black), and non-participating parcels (solid gray). Participating parcels drawn from program participants classified as single-family households who undertook one conversion. Sample sizes change over time as new homes add to the service population, or as pre-enrollment participating parcels enroll (occasionally homes exit the service population as well).*

program illustrated in Figure 1.5 provides at least a necessary (if not sufficient) condition for credibly assuming parallel trends.

1.3.2 Event study

I first illustrate water savings using an event study. Specifically, I estimate the following model:

$$Q_{it} = \sum_{j=-60}^{60} \kappa_j \mathbb{1}\left[\tau_{it} = j\right]_{it} + \mu_i + \delta_t + \epsilon_{it} \tag{1.1}$$

where Q_{it} describes water use in 1000 gallons for parcel i in month of sample t, μ_i are parcel fixed-effects, δ_t are month of sample fixed-effects and ϵ_{it} is the error. I define event time τ_{it} in relation to application date (thus τ is undefined for non-participants). For example, $\tau = -12$ for a parcel observed 12-months before the month of application.[21] κ_j describes average water use across all participants j months relative to the application date (net of parcel and seasonal fixed-effects). To avoid collinearity, I omit $\kappa_{j=0}$. κ_j therefore represents average water use relative to the application month. Month of sample fixed-effects soak up average seasonal fluctuations in water use and parcel fixed-effects control for average water consumption differences across parcels. I select a five-year window around the application month, dropping all participant observations outside the five-year window. Finally, I cluster standard errors at the parcel level.[22]

Figure 1.6 plots resulting point estimates and 95 percent confidence intervals of κ_j from Eq. (1.1), and clearly illustrates a reduction in water use resulting from conversion to desert landscape. Apart from seasonality not fully captured by the time fixed-effects, there does not appear to be any noticeable trends in water consumption prior to program application. An absence of pre-trends lends credibility to fully attributing the drop in water use illustrated in Figure 1.6 to conversion. And while the water savings achieved by conversion looks to be

[21] I could also define event time in relation to enrollment date. However, since the date of application indicates the first time I observe participants signaling interest in the program, defining the event relative to application seems most sensible. I present event study results defining τ in relation to enrollment date in Appendix A.2.

[22] McCrary (2007) suggests clustering at the parcel level to account for bias caused by a changing sample size throughout the event window.

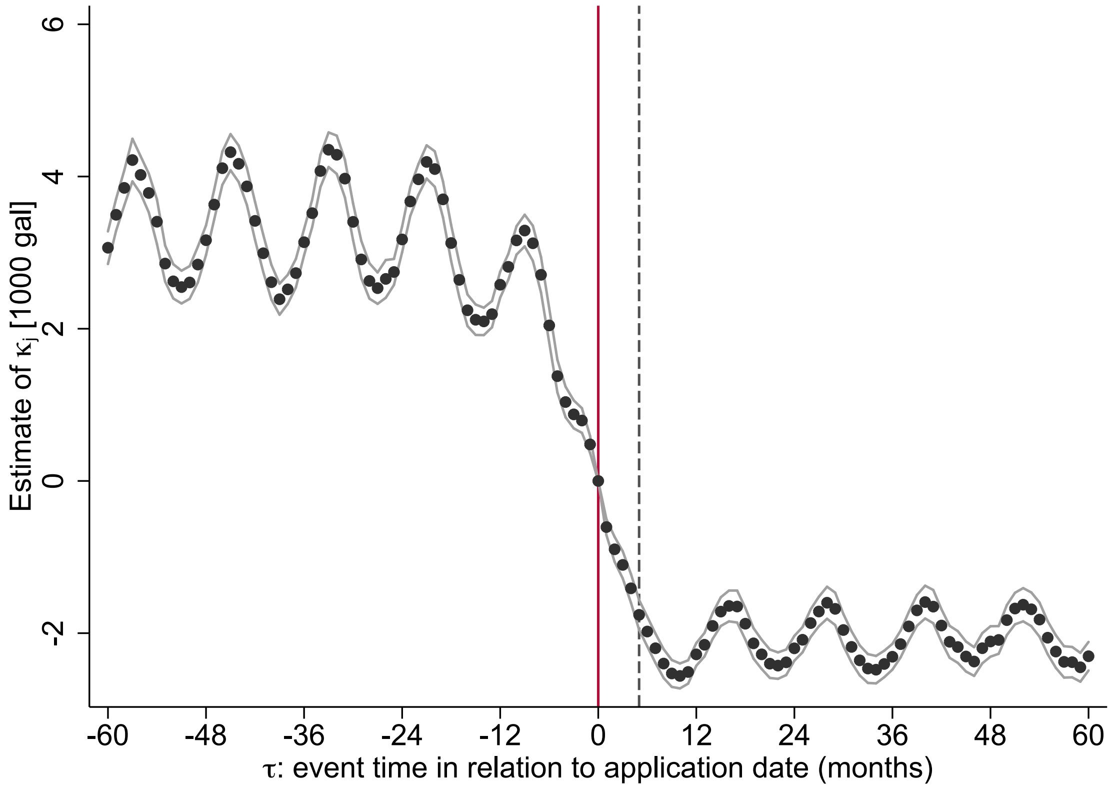

Figure 1.6: *Event study, illustrating water savings from the Cash-for-Grass program. Point estimates and 95 percent confidence intervals of κ_j's are derived from estimating Eq. (1.1). Standard errors are clustered at the parcel level. The omitted category is $\kappa_0 = 0$. Observations are limited to single-family participating parcels that converted only once and all non-participants. Participating parcel observations are further restricted to a five-year window around the month of application; that is $-60 \leq \kappa \leq 60$. The vertical solid line indicates the month of application. The vertical dashed line indicates the average month of enrollment, five months after the application date.*

largely maintained, the central tendency of the coefficient estimates in months following application exhibits a mild increase, suggesting a small erosion in savings in the months following application and subsequent conversion. Overall, Figure 1.6 implies conversion to desert landscape saves about 5,000 gallons per month.

The event study exhibits transient behavior around the month of application. Conversion takes place sometime between the application date (solid vertical line in Figure 1.6) and the enrollment date, which occurs on average five months later (indicated by the dashed vertical line in Figure 1.6). Between application and enrollment, a steady decline in water use can be explained by two factors; conversion, and the possibility that applicants reduce or stop irrigating their lawns following application. But much of the decline I attribute to conversion takes place before the date of application, suggesting some anticipatory behavior on the part of applicants. It could be that applicants stop watering their lawns even prior to application.[23] It could also be that the conversions for which I imputed an application date took much longer than the average five months. In my models that follow, I attempt to eliminate this transient behavior from my savings estimates, both in my model specification and in robustness checks to my main results.

1.3.3 Empirical approach

To quantify average savings, I estimate a panel fixed-effects model, a generalization of the canonical two-period differences-in-differences design.

$$Q_{it} = \alpha[\text{pre-period}]_{it} + \beta[\text{post-enroll}]_{it} + \mu_{im} + \delta_{tc} + \epsilon_{it} \tag{1.2}$$

In Eq. (1.2) Q_{it} is again monthly water consumption (in kgal) for parcel i in month-of-sample, t. Post-enroll is an indicator for months following program enrollment. Estimates of β therefore describe the change in water use due to conversion to desert landscape. Since the event study illustrated some transient behavior prior to enrollment, I further include a pre-period indicator that describes months between the application date and enrollment date. μ_{im} and δ_{tc} are parcel by month-of-calendar year and month-of-sample by cohort fixed-effects, and ϵ_{it} is the error. Parcel by month-of-calendar year fixed-effects control for average seasonal differences across parcels. Month-of-sample by cohort fixed-effects attempt to control for possible compositional differences in parcel characteristics that would invalidate the differences-in-differences design. Newer homes in Las Vegas tend to be built on smaller

[23] But un-watered grass dies quickly in Las Vegas, and to be eligible for the rebate, residents must show that they have been maintaining a lawn. However, water authority staff conducting pre-conversion site visits have to make judgment calls regarding this requirement, and some variation regarding what constitutes a maintained lawn may have allowed for approval of some applicants that ceased watering lawns well before the date of application. Though this is purely speculative, a water authority staff member did explain to me that in earlier years, the standards for what constituted a maintained lawn were not held to as strictly as they are now.

lots and may have more water efficient appliances. Furthermore, since 2003, new home construction regulations preclude lawns in the front yard. Such newer home characteristics may induce differential trends in water use compared to older homes. Including month-of-sample by cohort fixed-effects attempts to control for any such differential trends. Since I match service meters to parcels, new home construction corresponds with the first year in which I observe a parcel in my data. I define five cohorts based on the year a parcel first appears in my panel: 1988-1989 define cohort 1 and comprise 45 percent of observations, 1990-1994 define cohort 2 and comprise 15 percent of observations, 1995-1999 define cohort 3 and comprise 16 percent of observations, 2000-2004 define cohort 4 and comprise 15 percent of observations, and 2004-2014 define cohort 5 and comprise the remaining 9 percent of observations.[24]

The event study results suggest that initial water savings may erode over time. To quantify any erosion in water savings I additionally interact my post-enrollment indicator with a monthly linear time trend describing the number of months past enrollment. I further include a quadratic term to explore the rate at which any erosion in savings takes place.

$$\begin{aligned} Q_{it} &= \alpha[\text{pre-period}]_{it} + \beta_1[\text{post-enroll}]_{it} \\ &\quad + \beta_2[\text{post-enroll}]_{it} T_{it} + \beta_3[\text{post-enroll}]_{it} T_{it}^2 + \mu_{im} + \delta_{tc} + \epsilon_{it} \end{aligned} \tag{1.3}$$

1.3.4 Building alternative control samples

Participants voluntarily join the rebate program and may therefore possess systematically distinct characteristics from non-participants that would bias my water savings estimates. While I believe Figure 1.5 illustrates strong evidence of parallel trends and my month-of-sample by cohort fixed-effects further address concerns regarding compositional differences over time, one may still be concerned that there remains underlying differences between participating and non-participating parcels that would result in biased estimates. To address

[24] The year when a parcel enters the panel will approximate the time of construction for most of the parcels in my sample in cohorts 2 and above. Exceptions include homes previously on groundwater that switch to city water. Many homes in cohort 1, however, will have been built prior to 1988. Cohort 1 homes should therefore be interpreted as homes built in or before 1989.

any remaining concerns, I construct three additional control groups intended to better reflect underlying characteristics of the sample of participating parcels.[25]

First, I drop non-participants. By dropping non-participants, I avoid selection problems since my sample now includes only those who participate in the program. The control sample becomes those participants yet to convert.[26] Second, I build a control sample from a group of non-participating parcels that applied for the rebate, but did not become enrolled. I refer to these non-participating parcels as do-not-finishers, or DNF's. Among all non-participants, these DNF parcels are arguably the most similar to participants since they were not only aware of the rebate, but applied for the rebate as well. DNF's make an imperfect control, however. Among DNFs, over 70 percent do not finish because their application expired or they dropped out of the program and the reasons for dropping out or not following through may be correlated with water use.[27]

Third, I match participating parcels with non-participating parcels on lot size and July water consumption 2, 3, 4, and 5 years prior to enrollment into the Cash-for-Grass rebate program.[28] I use the Mahalanobis nearest-neighbor distance metric and match with replacement. I begin with a balanced panel to ensure that parcels have consumption data 5

[25] In other words, in constructing additional control samples, I attempt to avoid any remaining selection bias not accounted for by my main differences-in-differences specification. To address selection bias, researchers often pursue a matching strategy or seek an appropriate instrument. Matching essentially balances the treatment and control group along observable dimensions, and therefore assumes that unobservable characteristics of the treatment group equal unobservable characteristics of the control group (on average), or that any unobservable characteristics of the treatment group do not affect the selection process and outcome variable. Instrumental variable strategies require instruments uncorrelated with any unobservables (i.e. exogenous instruments). In essence, both approaches select a control group that would have been affected by the policy in the same way as the treated group. Addressing selection, therefore, becomes an exercise in building a valid control.

[26] Dropping non-participants is not without its problems. As shown and discussed by the ongoing work of Borusyak and Jaravel (2016), differences-in-differences estimates derived from samples without a control can underweight long-term impacts, which in my context could lead to over-estimating savings.

[27] A further 16 percent of DNF's ineligible. Among those that are ineligible, the most common reason is a lack of turf. To be approved, applicants must demonstrate that they have been maintaining a lawn. If the grass is dead, or non-existent, the water authority staff may reject the applicant during the pre-conversion site visit. I do not consider ineligible DNF's as part of my control sample.

[28] For example, I match parcels that convert in 2010 with non-participating parcels on lot size and water consumption in July of 2005, 2006, 2007 and 2008. For parcels that convert in 1998, I match on July consumption in 1994, 1995 and 1996. Additionally including July 1993 consumption encountered collinearity issues. Davis *et al.* (2014) also matches on pre-conversion (electricity) consumption in their study of an appliance rebate program in Mexico City.

years prior to conversion, and extract matches by running the Stata `teffects` command on a dummy outcome variable. I do not match on the year immediately prior to conversion since the event study suggests that water consumption may start to fall as much as a year before application and subsequent enrollment. Matching on 4 years of pre-enrollment consumption attempts to capture the downward trend in water use exhibited by Figure 1.5, and I match on July water consumption since parcels with similar peak consumption tend to have corresponding water use patterns throughout the other months of the year. Also, choosing only one month each year keeps the number of matching variables to a minimum. I build my matched control sample by pooling all non-participating parcel matches, keeping track of parcels matched more than once and weighting such parcels by the appropriate frequency in my regressions.

Figure 1.7 and Figure 1.8 compares annual water use of participating and non-participating parcels for the DNF and matched control samples, respectively. Both the DNF and matched control sample exhibit more similar average water consumption patterns than the full sample of non-participating parcels, and generally exhibit parallel trends prior to the beginning of the program in 1998.

1.3.5 Results

Water savings: main results Table 1.2 shows results from estimating my main model, Eq. (1.2).[29] In all models, I report parcel clustered standard errors in parentheses and suppress the coefficient estimate on the pre-period indicator. The pre-period indicator does little more than control for transient water behavior and is unimportant for understanding savings.

Focusing on the first five columns, the negative point estimates of the post-enroll indicator imply that conversion to desert landscape saves water. Column 1 shows results from estimating Eq. (1.2) with the full panel described in section 1.3.1. In column 2 and 3, I drop observations of program participants within 12 and 24 months of the month of

[29] I implement these and all following panel fixed-effect models, as well as my hedonic models discussed in section 1.4, using `reghdfe` (Correia, 2016).

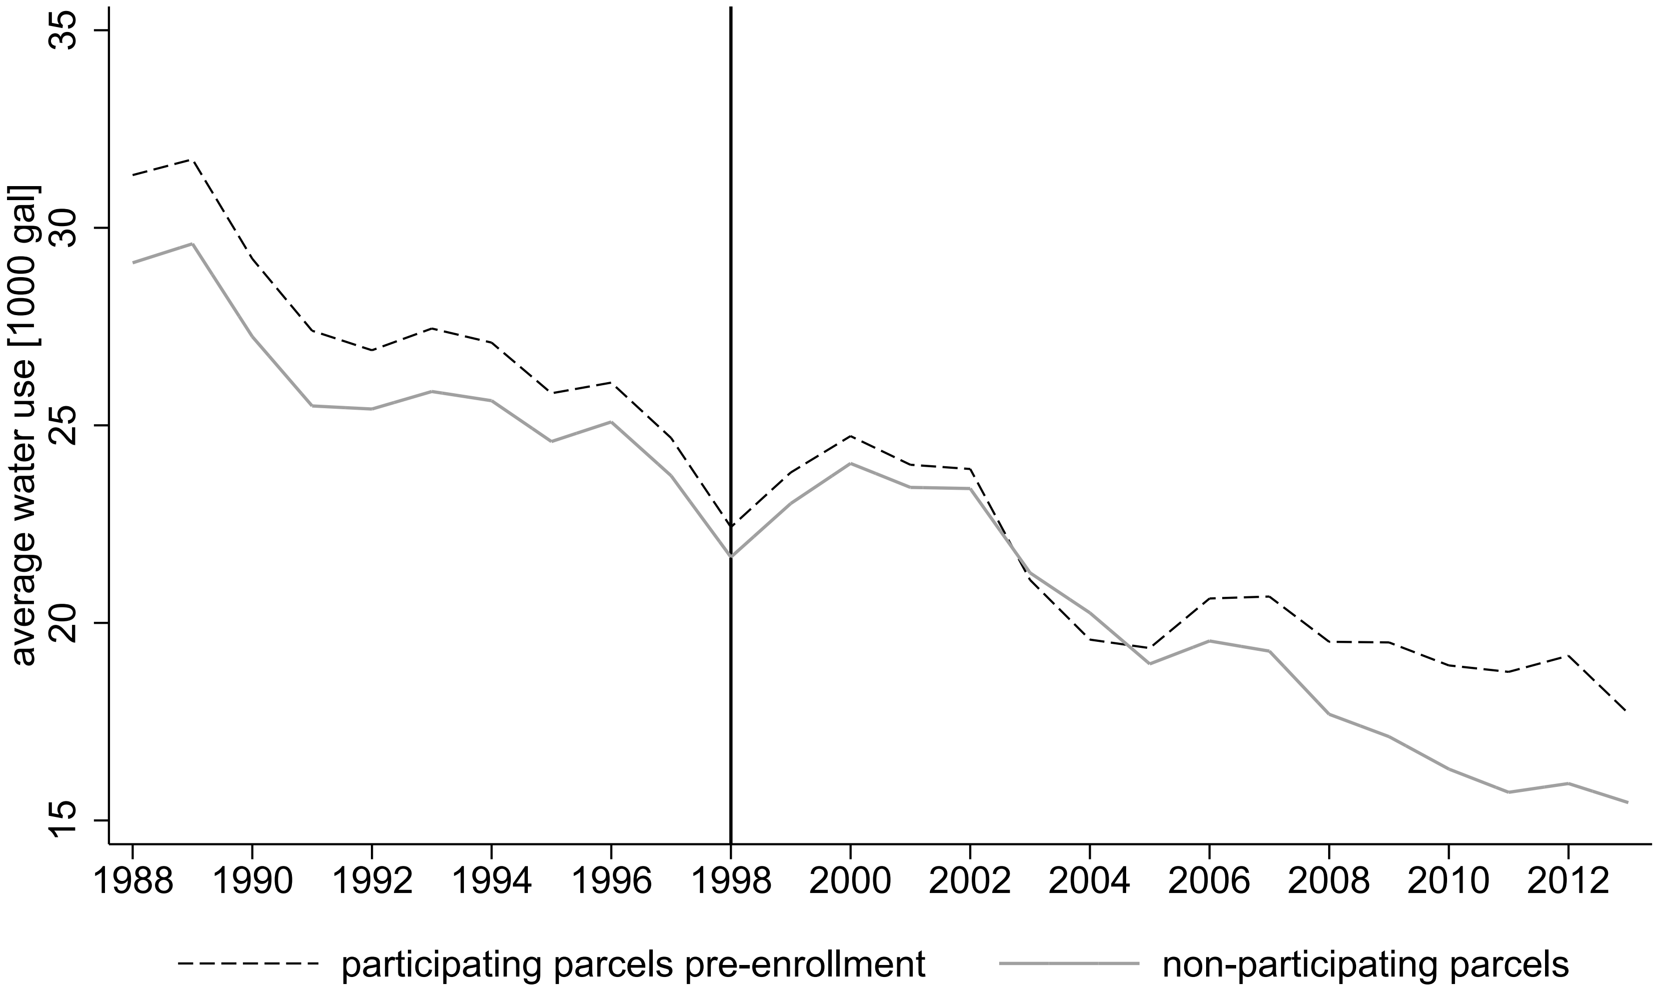

Figure 1.7: *Average annual water use from 1988 through 2013 (the last full year in my panel) for participating parcels prior to Cash-for-Grass program enrollment (dashed black), and non-participating parcels that applied for the program, but never become enrolled (solid gray), i.e. the DNF control sample. Participating parcels drawn from program participants classified as single-family households who undertook one conversion. Sample sizes change over time as new homes add to the service population, or as pre-enrollment participating parcels enroll (occasionally homes exit the service population as well).*

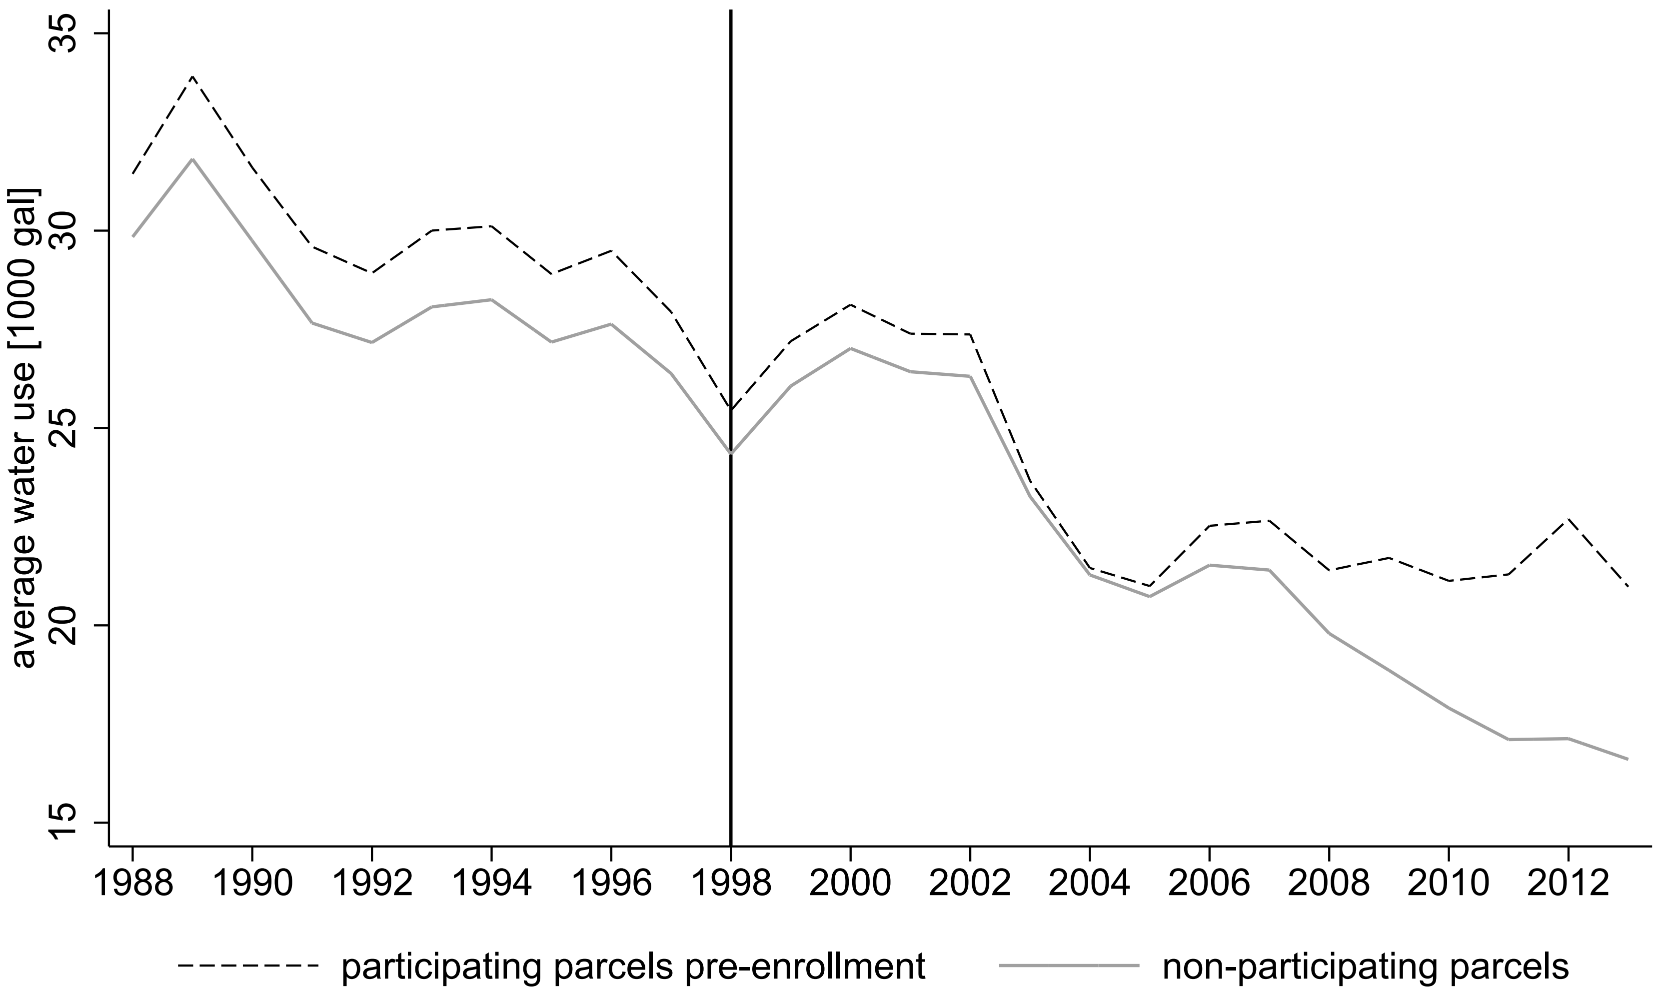

Figure 1.8: *Average annual water use from 1988 through 2013 (the last full year in my panel) for participating parcels prior to Cash-for-Grass program enrollment (dashed black), and non-participating parcels that match with participating parcels on lot size and pre-enrollment July water use (solid gray), i.e. the matched control sample. Participating parcels drawn from program participants classified as single-family households who undertook one conversion. Sample sizes change over time as new homes add to the service population, or as pre-enrollment participating parcels enroll (occasionally homes exit the service population as well).*

Table 1.2: *Water savings from converting to desert landscape (results from main specification).*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
post-enroll	-4.92	-4.97	-5.01	-4.90	-5.19	-5.16	-5.17
	(0.06)***	(0.07)***	(0.07)***	(0.12)***	(0.14)***	(0.06)***	(0.06)***
post-enroll $\times T_{it}$						0.005	0.006
						(0.001)***	(0.002)***
post-enroll $\times T_{it}^2$							-5.4e-06
							(1.7e-05)
Sample	full	1yr window	2yr window	balanced	no zero	full	full
Fixed-effects							
μ_{im}	yes	yes	yes	yes	yes	yes	yes
δ_t	-	-	-	yes	-	-	-
δ_{tc}	yes	yes	yes	-	yes	yes	yes
adj. R^2	0.30	0.30	0.30	0.53	0.71	0.30	0.30
Parcels	309,201	309,198	309,193	81,476	47,771	309,201	309,201
Observations	64,120,344	63,472,767	62,867,660	25,746,416	13,775,144	64,120,344	64,120,344

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Parcel clustered standard errors (reported in parentheses).

μ_{im} parcel by month-of-calendar year fixed-effects.

δ_t, δ_{tc} month-of-sample and month-of-sample by cohort fixed-effects.

Cohorts based on parcel's first year in sample: 88-89, 90-94, 95-99, 00-04, 05-14.

Xyr window: drop participant observations within X years around enrollment.

balanced: includes parcels with non-missing water consumption in each month from Jan. 1988 to Apr. 2014.

no zero: includes parcels with all positive consumption only.

enrollment. These two "donut" specifications attempt to capture the difference between steady state behavior before and after conversion to desert landscape by avoiding any transient behavior around the time of conversion not already accounted for by the pre-period indicator. The consistency between the estimates in columns 1, 2, and 3 indicate that there remains little biasing transient behavior around the time of conversion after controlling for the pre-period between application and enrollment dates. In column 4, I estimate Eq. (1.2) using a balanced sample. Column 4 aims to test for the robustness of water savings to differences across users not captured by the month-of-sample by cohort fixed-effects. The similarity between the estimates in column 1 and 4 indicate that any differences across cohorts do not affect my results, or that my time-cohort fixed-effects adequately control for any such differences. In Column 5, I limit observations to parcels that only have positive values of consumption throughout the panel. Excluding parcels experiencing zero water consumption attempts to control for properties under foreclosure that may have been vacant for some time. Though less precise, the estimate of savings in column 5 compare favorably to the savings estimate in column 1.

In columns 6 and 7 of Table 1.2 I test the stability of savings for a given conversion by estimating Eq. (1.3) with the full sample. Results from these two models demonstrate that the average conversion experiences about a 5 gallon per month erosion in water savings, but that this erosion rate decreases over time (row 3, col. 7). The erosion rate is statistically significant, and possibly makes sense as larger, mature plants will require more water than younger plants. However, at one tenth of a percent of total savings, the erosion rate hardly seems practically relevant. I conclude that conversions to desert landscape maintain their savings over the long term.

Overall, water savings estimates in Table 1.2 display remarkable consistency across specifications, and demonstrate an average conversion to desert landscape saves about 5,000 gal/month, confirming the results from the event study.[30] 5,000 gal/month represents

[30] I also run a set of falsification tests to examine the validity of the parallel trends identifying assumption behind Eq. (1.2). See footnote 148.

an approximate 21 percent decrease in baseline water use, and for a customer in 2013, corresponds to about $150 in annual water bill savings, or a 30 percent reduction. Appendix A.1 provides details of these calculations. Appendix A.2 includes further water savings results. In particular, I explore the effects of various fixed-effects specifications, robustness to parcels that exit before the end of the sample, and the effect of two program policy changes.

Table 1.3: *Water savings from converting to desert landscape (results from alternative control samples).*

	(1)	(2)	(3)	(4)
post-enroll	-4.92	-5.22	-4.15	-4.11
	(0.06)***	(0.08)***	(0.09)***	(0.14)***
Sample	full	participants	DNF	match
Fixed-effects				
μ_{im}	yes	yes	yes	yes
δ_t	-	-	-	yes
δ_{tc}	yes	yes	yes	-
adj. R^2	0.30	0.64	0.33	0.70
Parcels	309,201	26,414	32,368	16,774
Observations	64,120,344	6,579,892	8,022,520	5,568,552

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$
Parcel clustered standard errors (reported in parentheses).
μ_{im} parcel by month-of-calendar year fixed-effects.
δ_t, δ_{tc} month-of-sample and month-of-sample by cohort fixed-effects.
Cohorts based on parcel's first year in sample: 88-89, 90-94, 95-99, 00-04, 05-14.
participants: includes only rebate program participants.
DNF: uses a control sample of non-participants that apply, but do not finish.
match: uses a control sample of matched non-participating parcels.

Water savings: additional control samples Table 1.3 shows results from estimating Eq. (1.2) with each of the three alternative control samples discussed above. All estimates are significant at the 1 percent level, and I cluster standard errors at the parcel level. For purposes of comparison, column 1 replicates the main specification result with the full sample. Column 2 includes only participating parcels. Column 3 shows estimates based on a control sample constructed from DNFs. This sample of DNF's includes only DNF's that do not later become rebate program participants. Finally, column 4 presents results derived from my matched sample. The estimates derived from the DNF and matched control samples fall about 15 percent lower than the estimate from the main specification. However, each alternative control sample still demonstrates a clear reduction in water use due to conversion to desert landscape. I therefore conclude that my estimates of savings remain generally consistent across different control samples.

Heterogeneous effects across time Since the program has been in place for over fifteen years, one might expect savings to have fluctuated over time. Landscape professionals may have become more skilled at installing water saving landscapes (a learning-by-doing argument), and variation in the subsidy level may induce participants with heterogeneous characteristics that differentially affect program outcomes. To explore the impact of time, I estimate the following model:

$$Q_{it} = \alpha[\text{pre-period}]_{it} + \sum_k \beta_k[\text{post-enroll}]_{it} \times \mathbb{1}\,[k] + \mu_{im} + \delta_{tc} + \epsilon_{it} \tag{1.4}$$

where the index k represents program enrollment years. The coefficient estimate on the interaction of the post-enrollment indicator and the enrollment year describes savings achieved by conversions taking place for that enrollment year. Absolute savings, however, depend upon the conversion size, and as Figure 1.4 demonstrates, the average conversion area has changed over time. In order to make appropriate comparisons across years, therefore, I normalize the estimates of savings by the average converted area in each year. Figure 1.9 shows the results of this exercise, illustrating savings achieved in each year of the program normalized by the corresponding annual average conversion area, scaled up to a per annum basis in order to compare my estimates with water authority estimates.[31] The

[31] Normalized savings equals $f(\hat{\beta}_k, A_k) = c\frac{\hat{\beta}_k}{A_k}$, where $\hat{\beta}_k$ represents the estimate of savings in 1000 gal/month per average conversion in year k derived from Eq. (1.4), A_k represents average converted area in year k, and $c = -12,000$, which converts a negative change in water use in kgal/month to a positive savings in gal/year. Because I have the universe of conversion records within the LVVWD, I consider A_k a fixed parameter and calculate standard errors as: $\text{Var}\left[\frac{c}{A_k}\hat{\beta}_k\right] = \frac{c^2}{A_k^2}\text{Var}\left[\hat{\beta}_k\right] = \frac{c^2}{A_k^2}\hat{\sigma}^2_{\beta_k} \implies \hat{se}_k = \sqrt{\frac{c^2}{A_k^2}\hat{\sigma}^2_{\beta_k}} = \frac{c}{A_k}\hat{\sigma}_{\beta_k}$. If instead, I consider A_k a random variable I would apply the Delta method to estimate standard errors for $f(\hat{\beta}_k, \hat{A}_k) = c\frac{\hat{\beta}_k}{\hat{A}_k}$. In particular, if $\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, V)$, then $\sqrt{n}(f(\hat{\theta}) - f(\theta_0)) \xrightarrow{d} \mathcal{N}(0, AVA')$. Dropping hats, k subscripts, and c for clarity, in my context, $f(\cdot) = \frac{\beta}{A}$, $\theta = (\beta \quad A)'$, $V = \begin{pmatrix} \sigma^2_\beta & \sigma_{\beta A} \\ \sigma_{\beta A} & \sigma^2_A \end{pmatrix}$, and $A = \begin{pmatrix} \frac{\delta f}{\delta \beta} & \frac{\delta f}{\delta A} \end{pmatrix}$.

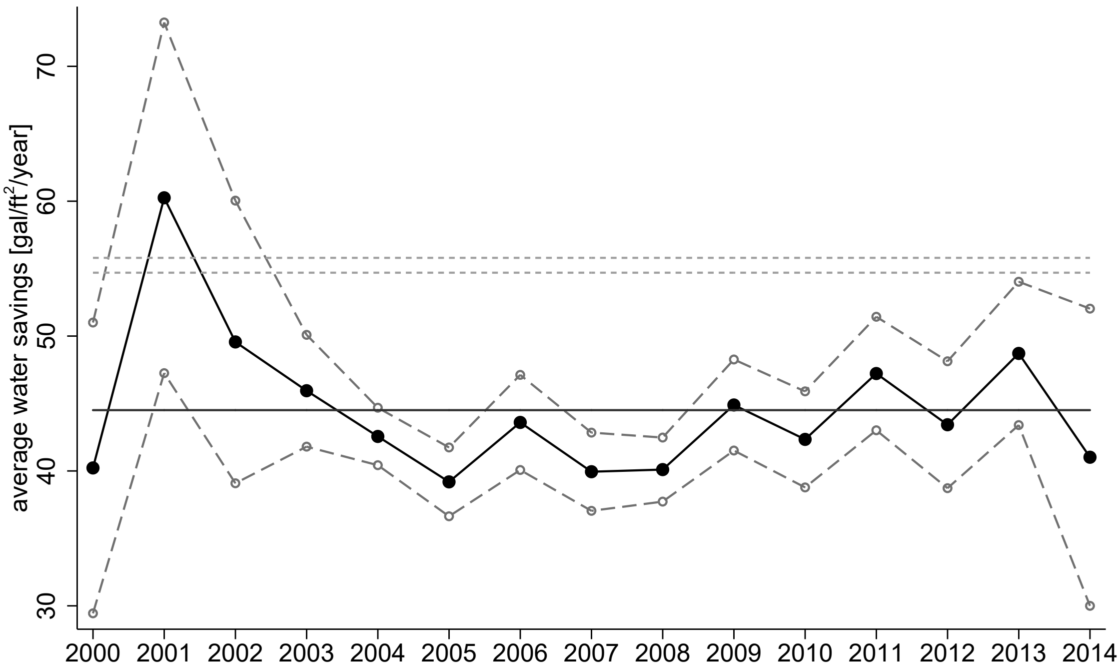

Figure 1.9: *Average water savings achieved by participants in each year of the program (derived from Eq. (1.4)) and normalized by the corresponding average annual conversion area (see Figure 1.4). 95 percent confidence intervals are derived considering average converted area a fixed parameter. Horizontal gray dotted lines show the water authority's two point estimates of normalized savings: 55.8 and 54.7 gal/ft^2/year. The horizontal solid line represents my estimate of 5,000 gal/month normalized by the overall average converted area, 1,348 ft^2, and then scaled up to gal/ft^2/year.*

figure demonstrates a distinct 'U'-shaped pattern, achieving an initial peak of 60 gal/ft^2/year in 2001, early in the program. Savings then fall quickly to a low of 39 gal/ft^2/year in 2005. Since 2008, however, normalized savings have a exhibited steady upward trend.[32]

The 'U'-shaped pattern observed in Figure 1.9 does not appear consistent with a learning-by-doing story among landscape professionals because savings achieved a peak value early

---

$AVA' = \begin{pmatrix} \frac{1}{A} & \frac{-\beta}{A^2} \end{pmatrix} \begin{pmatrix} \sigma_\beta^2 & \sigma_{\beta A} \\ \sigma_{\beta A} & \sigma_A^2 \end{pmatrix} \begin{pmatrix} \frac{1}{A} \\ \frac{-\beta}{A^2} \end{pmatrix}$, implying that $AVA' = \frac{1}{A}\left(\frac{\sigma_\beta^2}{A} - \frac{\beta\sigma_{\beta A}}{A^2}\right) - \frac{\beta}{A^2}\left(\frac{\sigma_{\beta A}}{A} - \frac{\beta\sigma_A^2}{A^2}\right) = \frac{1}{A^2}\sigma_\beta^2 + \frac{\beta^2}{A^4}\sigma_A^2 - 2\frac{\beta}{A^3}\sigma_{\beta A}$ and $\hat{se}^{alt} = \sqrt{\frac{1}{A^2}\sigma_\beta^2 + \frac{\beta^2}{A^4}\sigma_A^2 - 2\frac{\beta}{A^3}\sigma_{\beta A}}$. Note that $\hat{se}^{alt}$ reduces to the expression for $\hat{se}_k$ for fixed A_k, since then $\sigma_A = \sigma_{\beta A} = 0$. How $\hat{se}_k$ compares to $\hat{se}^{alt}$ depends upon the relative size of $\sigma_{\beta A}$ and σ_A^2. For example, $\sigma_{\beta A} \gg \sigma_A^2$ implies that my simple estimate of the standard error, considering A_k fixed, may be larger than $\hat{se}^{alt}$. A similar argument holds for Figure 1.10.

[32]Though not shown, I observe similar behavior for the participant only, DNF, and matched control samples. Appendix A.2 provides further discussion.

in the program, implying that landscape professionals already knew how to design and implement efficient conversions.[33] Rather, the 'U'-shaped pattern seems most consistent with heterogeneous participants effecting program outcomes over time. Middle period participants received the largest subsidies. In order to participate, therefore, early and late period participants may have derived some (or more) non-monetary benefits not experienced by middle period participants. Such non-monetary benefits could be correlated with greater awareness regarding water scarcity or a stronger desire to conserve, both of which could explain more efficient conversions. This discussion also highlights a trade-off between increasing participation, and increasing program efficiency. Participation was high in the mid to late 2000's (Figure 1.3), but this period also corresponds to the lowest per unit savings, implying the lowest "bang for the buck".

Figure 1.9 also presents a comparison between my estimates and two estimates derived by the water authority (Sovocool *et al.*, 2006). Beginning in 1995, the water authority conducted a pilot study that recruited participants and measured irrigation specific application rates. Sovocool *et al.* conclude that conversion to desert landscape saves 55.8 gal/ft^2/year. In a follow-up analysis, Sovocool *et al.* re-estimate savings, finding that desert landscape saves 54.7 gal/ft^2/year. The re-estimate draws from participants converting in 2003 and not recruited for the pilot program. I illustrate these two estimates in Figure 1.9 with horizontal dashed lines. The horizontal solid line represents my estimate of 5,000 gal/month normalized by the overall average conversion area (1,348 ft^2) and scaled to gal/ft^2/year. While the Sovocool *et al.* estimates lie well above my overall average estimate, their estimates fall within the confidence intervals for my annual estimates in 2001 and 2002. Therefore, water authority estimates may not be overstated, at least for early program participants.[34]

[33]Or highly skilled do-it-yourselfers performed the early conversions, and landscape professionals did not expand their services to include conversions until the early to middle 2000's, after which they improved their ability to install efficient desert landscape.

[34]Arguing for the importance of including confidence bounds on estimates derived from statistical methods, Auffhammer *et al.* (2008) make a similar finding in the context of energy demand-side-management programs. That is, utility based estimates may not be misstated once researchers calculate bounds on their analyses (Auffhammer *et al.*, 2008).

These results stand in contrast to a general finding in the demand-side-management program evaluation literature that utility based estimates often overstate savings (Joskow and Marron, 1992; Loughran and Kulick, 2004; Allcott and Greenstone, 2012). But my results also demonstrate the importance of continued program assessment; as Figure 1.9 illustrates, savings may fluctuate throughout the life of the program.

Heterogeneous effects across pre-treatment consumption I also estimate a form of Eq. (1.4) for savings achieved by participants within pre-enrollment consumption deciles. I derive pre-enrollment consumption deciles from a 12-month average of water use beginning 24 months prior to the month of enrollment. I normalize water use, Q_{it}, by lot size, and define pre-enrollment consumption based on this normalized water use. Since I define lot size in 1000 ft^2, my outcome variable becomes monthly water use in gal/ft^2 of lot size. Compared to high demand-small lot consumers, similarly high demand-large lot consumers use water more efficiently, and therefore may not achieve the savings realized by their high demand-small lot counterparts from an equally sized conversion to desert landscape. Normalizing water use by lot size distinguishes the high demand-small lot consumers from the high-demand large lot consumers and avoids potentially downward biasing estimates of savings achieved by higher pre-enrollment consumption deciles.

Figure 1.10 illustrates that high-demand consumers, relative to lot size, achieve the greatest savings.[35] I find this result robust to a series of pre-enrollment consumption decile definitions and to the different control samples discussed above.[36] To the extent that consumers in the high pre-enrollment consumption deciles over-water relative to their

[35] Investigating the Cash-for-Grass pilot study conducted in the late 1990's, Deoreo *et al.* (2000) make a similar finding, and in the context of the OPOWER experiment, Allcott (2011) also uncovers a positive relationship between savings and pre-treatment consumption.

[36] I additionally derive pre-enrollment consumption deciles based on a 24, 36, and 48-month average and find a similar relationship between normalized water use and pre-enrollment deciles. I also run models that consider non-normalized water use, deriving pre-enrollment consumption deciles from a 12, 24, 36, and 48-month average. Similar to my procedure for deriving Figure 1.9, I then normalize my resulting savings estimates by the average conversion area within each pre-enrollment consumption decile. I again find a positive relationship between savings and pre-enrollment consumption decile. Finally, I perform the above analyses with my participant only, DNF, and matched samples. For each sample and normalization method, savings continue to be positively related to pre-enrollment consumption decile. I provide further discussion in Appendix A.2.

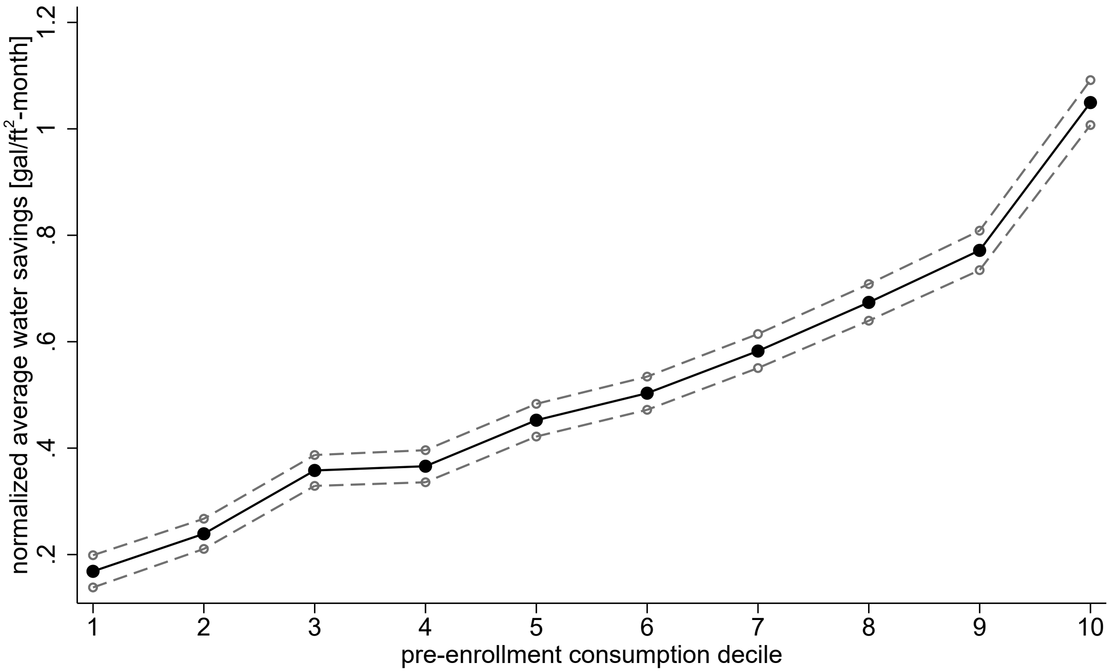

Figure 1.10: *Average water savings per square foot of lot size achieved within each pre-enrollment water consumption decile. Point estimates and 95 percent confidence intervals are derived from a model based on Eq. (1.4) with the dependent variable, Q_{it}, normalized by 1000 ft^2 of lot size. Pre-enrollment consumption deciles are defined based on a 12-month average of water use for participating parcels (normalized by lot size) beginning 24 months prior to the month of enrollment.*

low decile counterparts, high decile users will save more water than low decile users with the same conversion to desert landscape. The positive relationship between savings and pre-enrollment consumption decile therefore suggests that consumers in the higher deciles use water inefficiently relative to their low decile counterparts. This further implies that the highest savings from a program like Cash-for-Grass arise from the least efficient water using customers.[37]

1.3.6 Rebates versus prices

If instead of implementing the Cash-for-Grass program the water authority had raised prices, what price increase would have induced the same aggregate savings? To answer this

[37] I am indebted to Peter Mayer for a clarifying discussion regarding the explanation for the relationship between savings and pre-enrollment consumption decile.

question, I begin with a partial equilibrium (i.e. supply-demand) framework, shown in Eq. (1.5).

$$\%\Delta P = \frac{\%\Delta Q}{\epsilon} \approx \frac{\Delta Q}{\epsilon Q} \tag{1.5}$$

Since the first conversions in my panel begin in 1998 I take the baseline quantity of water, Q, to be the average water use among all single-family residential customers in 1997, or 19,047 gallons/month. In 1997, 12 percent of the service population were future Cash-for-Grass program participants. But since all customers experience a price increase, to reduce consumption by the same aggregate amount as eventually achieved by the share of the service population that would participate, $\Delta Q = 5,000 \times 0.12 = 600$, the product of the share of the service population that became program participants and the savings these program participants achieved.[38] In their respective meta-analyses of residential water demand price elasticities, Espey *et al.* (1997) find a mean elasticity of -0.51 and Dalhuisen *et al.* (2003) find a mean elasticity of -0.41. Long-term elasticities are generally higher, and since this analysis asks how consumers in 1997 would have responded over a nearly 20-year period, I take the elasticity to be -0.5.[39] Based these values, Eq. (1.5) predicts that a 6 percent price increase would have achieved equivalent aggregate savings.

This hypothetical percentage price increase is relatively modest. Since 1958, customers of the LVVWD have experienced actual average price increases between 5.6 percent in 2011 and 26.8 percent in 2003.[40] Since all customers experience and respond to a price increase, large aggregate savings require comparatively small individual cutbacks. This fact drives the modest hypothetical price increase estimated by Eq. (1.5). As the share of the service population that participates in the rebate program increases (and continues to achieve the same average savings), so does the price increase required to induce the same aggregate

[38]Consider a service population of N and define aggregate monthly savings as $N\sigma$, where σ is the per-parcel monthly savings. The aggregate savings of 12 percent of this service population saving 5,000 gallons per month is $0.12 \times N \times 5,000$. Equating the two implies: $N\sigma = 0.12 \times N \times 5,000 \implies \sigma = 0.12 \times 5000 = 600$.

[39]It would be preferable to derive an estimate of elasticity using Las Vegas specific data. While my data include historical water rates and consumption, I do not have individual income data. For this reason, I select elasticity estimates from the literature.

[40]pers. comm. SNWA staff, July 2014.

savings. If the entire service population participates in the program, Eq. (1.5) predicts that a 53 percent price increase is needed to realize the savings achieved by the rebate.

The price increase analysis assumes homogeneous customers. But lower income households with little to no landscape may not be able to achieve much if any savings regardless of the price increase. Eliminating these households from the analysis above would increase ΔQ since it would decrease the effective service population saving water due to the price hike. Demand elasticity would also decrease, since wealthier customers tend to have lower water demand elasticities than poorer customers (Mansur and Olmstead, 2012). Both factors would drive up the estimate of the hypothetical price increase.

Price increases may also cause regressive outcomes. If both wealthy customers and poorer customers experience the same price increase, the ratio of wealthy customer demand elasticity to poorer customer demand elasticity roughly approximates the ratio of the change in Marshallian surplus of poorer to wealthier customers.[41] Taking at face value the point estimates of wealthy and poor residential water demand elasticities derived by Mansur and Olmstead implies that the reduction in surplus for poorer customers could be as much as 67 percent greater than the reduction in surplus for wealthier customers. So despite the fact that modest price increases could achieve large aggregate savings, heterogeneity in customer characteristics may increase the estimated price hike and create large differential, and arguably inequitable, welfare effects across customer types.

[41] In a standard partial equilibrium framework that assumes linear demand, the reduction in Marshallian surplus from reducing water consumption given a change in price is approximated by $1/2 \times (Q/P)\Delta P^2 e^{-1}$, where ΔP is the change in the price of water, P and Q are initial prices and quantities of water consumed, respectively, and e is the demand elasticity. If wealthy and poorer customers experience similar unit prices of water, i.e. $P_w/Q_w \approx P_p/Q_p$, where w refers to wealthy and p refers to poor, then the ratio of the two changes in Marshallian surplus is given by e_w/e_p. Under block pricing, it may be that $P_w/Q_w > P_p/Q_p$ since wealthier customers tend to use more water than poorer customers, and this increased usage is priced at a higher marginal rate. If $P_w/Q_w > P_p/Q_p$, the ratio of Marshallian surplus is $(e_w/e_p) \times \frac{Q_p/P_p}{Q_w/P_w} > e_w/e_p$. The estimate of e_w/e_p could therefore be thought of as a lower bound on the differential welfare effect.

1.4 Value of desert landscape

1.4.1 Motivation

Desert landscape can affect participants' utility outside water bill savings. Those choosing to convert may find desert landscapes aesthetically pleasing (Walls *et al.*, 2015) or value signaling a commitment to environmental stewardship (Mustafa *et al.*, 2010). Desert landscapes also tend to require less maintenance than lawns.[42] However, replacing grass with drought tolerant flora may lead to increased energy costs. Because of the higher evapotranspiration rates[43] of lawns versus desert landscapes, conversion to desert landscape could increase local air temperatures (Bonan, 2000), leading converting properties to demand more air-conditioning. Importantly, neighbors of desert landscaped properties may also experience these negative energy spillovers and derive aesthetic utility from neighboring desert landscapes. Therefore, simply calculating the value of saved water will not necessarily capture the full value of desert landscape to participants and will fail to reflect any externalities that conversions impose upon neighbors.

Formalized by Rosen (1974), the hedonic property method provides a theoretically consistent method for estimating the private benefits of converting to desert landscape. Modeling property values as a function of the property's individual characteristics, Rosen showed that the effect of an individual characteristic on property values represents the benefit a consumer receives from the characteristic. Since Rosen, hedonics has become a widely used strategy for valuing non-market goods, especially environmental characteristics of properties (Davis, 2004; Greenstone and Gallagher, 2008; Muehlenbachs *et al.*, 2015). Hedonic estimates, however, will fail to capture benefits not communicated through housing prices, such as benefits from reduced water-transport costs[44] and the ability to reallocate

[42] Cash-for-Grass program staff noted that reduced lawn maintenance appears to be a primary driver for individuals who apply for the rebate (pers. comm. K. Sovocool, February 2015).

[43] A process of simultaneous evaporation and plant transpiration.

[44] Water requires substantial energy to deliver, and reducing water use decreases greenhouse gas emissions and other pollutants through lower energy production. Reduced energy consumption also implies lower energy bills for the utility, freeing up funds for alternative uses.

scarce water to alternative or future uses. My hedonic estimates, therefore, may understate the true effect of conversion.

In the following analysis, I estimate the private, direct and spillover effects generated by conversions to desert landscape subsidized by the Cash-for-Grass program.[45] To estimate the direct effect of conversion, I make use of the variation in conversion status across properties. To estimate the spillover effect, I make use of the variation in a property's adjacency to a conversion. In other words, I characterize properties by their conversion status and whether they lie adjacent to properties that convert. I focus on single-family participants as they account for over 90 percent of all conversions, have received the largest share of rebate monies (41 percent) and are responsible for the largest share of area converted (36 percent).[46] To maintain consistency with the water savings analysis, I limit my analysis to the LVVWD service area.

1.4.2 Data

I construct a panel of residential sales occurring within the LVVWD from historical Las Vegas area sales data provided by the Clark County Assessor's Office. I include sale price and date, home age, parcel, home, garage, and pool square footage, and finally the number bedrooms and bathrooms (full and half bathrooms). I convert all sale prices to $2014 using the CPI housing index, and drop observations outside of the first and ninety-ninth percentiles of the sale price distribution. I restrict my observations from January 1, 1996 to June 12, 2014, the latest program enrollment date.[47] I keep arms-length, single-family

[45]I am not the first to estimate the impact of desert landscape on property values. Both Baker (2004) and Rollins (2008) find that desert landscape increases Las Vegas home values. In particular, using a hedonic framework with neighborhood characteristics defined at the zip code level, Rollins finds that desert landscape increases home values by about 7 percent. Using a larger data set which includes more recent home sales as well as a highly spatial and temporally refined set of fixed-effects (quarter of sample by census block fixed-effects), I find smaller effects on home values due to desert landscape.

[46]Multi-family participants have received 29 percent of total rebate monies and have converted 31.5 percent of the total converted area; golf courses received 20 percent of total rebate monies and converted 21.5 percent of the total converted area, and commercial and industrial participants have received 10 percent of total rebate monies and converted the remaining 11 percent of total converted area.

[47]The rebate pilot program begins in 1996, and the first enrollment occurred on May 6, 1996.

transactions, and further drop all parcels with a negative age, a parcel area of zero, or a home size of zero. Finally, I associate parcels with 2010 U.S. Census block areas using GIS software. On average, approximately 29 properties fall within each block.

Including properties that undergo unobserved structural changes may bias my analysis. Since the assessor data do not record changes in property characteristics, I develop three criteria for assessing whether a property may have undergone a structural change. The assessor data include the square footage of any additions made to a parcel, which often occurs if the property owner converts a garage to living space.[48] I drop parcels with positive addition area. The assessor data also provide the year of home construction, as well as the effective year of home construction. For most homes, these years are equivalent. If the construction year does not equal the effective construction year, I conclude the parcel likely underwent an addition, and drop such parcels. I drop remaining parcels that have a detached garage if the year built or effective year built of the detached garage does not equal the home construction year.

I merge the resulting panel of sales with the enrollment panel described in section 1.3 and a "neighbors" dataset that I construct using GIS software. I consider all single-family participating parcels, some of which undertook more than one conversion.[49] Illustrated in Figure 1.11, I define neighbors as parcels that lie directly adjacent to any single-family participating parcel. Importantly, neighbors may themselves be participating parcels. The neighbors dataset contains enrollment dates and converted areas of conversions *adjacent* to the parcel.[50]

Six variables characterize conversion to desert landscape. P_{it} takes one if parcel i has converted by the sale date, t. N_{it} takes one if parcel i lies adjacent to a conversion by the sale date, t. DP_i and DN_i describe parcel i's status as an eventual program participant (i.e. the

[48]pers. comm. E. Martinet, May 2016.

[49]A small number of enrollments that the water authority considered non-single-family match with the single-family assessor data. I drop these observations. In general, though, the water authority's classification of single-family agrees closely with that of the assessor's office.

[50]Some neighboring enrollments occurred on the same date. For these cases, I include the total rebate and converted area for that particular enrollment date.

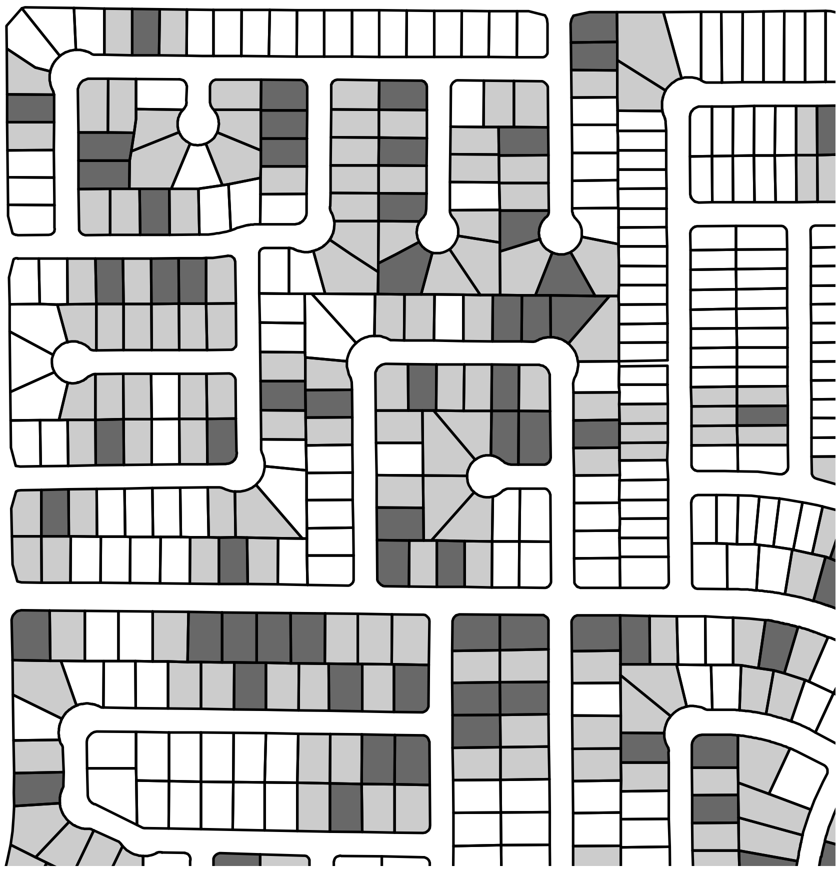

Figure 1.11: *Illustration of neighboring parcels. Dark shaded parcels eventually participate in the Cash-for-Grass program. Lightly shaded parcels neighbor participating parcels. Empty parcels are neither participating parcels nor neighbors, and thus make up my control sample. Note that neighboring participating parcels are neighbors of the other, participating parcel.*

parcel converts at some point in the panel) or eventual neighbor of a program participant, respectively.[51] Pa_{it} and Na_{it} describe the total converted area and total adjacent converted area for parcel i at sale date t.[52] If a parcel neighbors multiple participating parcels prior to the sale date, the total converted area from all adjacent participating parcels define Na_{it}.

[51]In the language of differences-in-differences, P_{it} and N_{it} describe an interaction between the treatment indicator DP_i or DN_i, and a post-treatment period indicator.

[52]For example, consider a parcel that sells three times and undergoes two conversions of 500 square feet between the first and second sale, and another conversion of 500 square feet prior to the third sale, and lies adjacent to a property that converts 500 square feet between the second and third sale. In this illustrative example, $DP_i = DN_i = 1$ for each sale observation. In addition, for the first sale observation, $P_{it} = N_{it} = Pa_{it} = Na_{it} = 0$. For the second sale observation, $P_{it} = 1$ and $Pa_{it} = 1000$, and $N_{it} = Na_{it} = 0$. Finally, for the third sale observation, $P_{it} = N_{it} = 1$, $Pa_{it} = 1500$, and $Na_{it} = 500$.

Table 1.4: *Summary statistics for hedonic panel: sale years 1996 - 2014*

(a) All sales		*Participant*			*Neighbor*		
	Full	$P = 0$	$P = 1$	p-val	$N = 0$	$N = 1$	p-val
sale price	309,077	309,823	268,839	0.00	311,356	274,165	0.00
age	6.3	6.1	17.3	0.00	5.6	16.6	0.00
lot sqft	5,755	5,725	7,361	0.00	5,662	7,181	0.00
home sqft	2,042	2,042	2,017	0.04	2,042	2,035	0.32
bedrooms	3.39	3.39	3.43	0.00	3.39	3.45	0.00
full bath	2.26	2.26	2.26	0.65	2.26	2.27	0.20
half bath	0.53	0.53	0.36	0.00	0.54	0.37	0.00
pool	0.19	0.19	0.33	0.00	0.19	0.32	0.00
garage sqft	465	464	505	0.00	463	498	0.00
Observations	199,037	195,410	3,627		186,840	12,197	

(b) Repeat sales		*Participant*			*Neighbor*		
	Full	$P = 0$	$P = 1$	p-val	$N = 0$	$N = 1$	p-val
sale price	295,867	297,054	239,803	0.00	299,783	241,908	0.00
age	5.9	5.7	15.1	0.00	5.3	14.5	0.00
lot sqft	5,072	5,037	6,743	0.00	4,970	6,484	0.00
home sqft	2,047	2,047	2,055	0.75	2,047	2,057	0.44
bedrooms	3.36	3.36	3.46	0.00	3.35	3.45	0.00
full bath	2.26	2.25	2.30	0.02	2.25	2.27	0.06
half bath	0.63	0.63	0.42	0.00	0.64	0.44	0.00
pool	0.16	0.16	0.33	0.00	0.16	0.29	0.00
garage sqft	452	450	500	0.00	449	487	0.00
Observations	40,755	39,910	845		37,998	2,757	

Summary statistics in panel (a) reflect the sample used in estimating Eq. (1.6). Panel (b) reflect summary statistics from estimating Eq. (1.6), additionally including parcel fixed-effects (and dropping time-invariant property specific controls). Prices adjusted to 2014 dollars and further restricted to the 1^{st} and 99^{th} percentiles of the sale price distribution. I further drop parcels meeting criteria for undertaking additions.

Table 1.4 displays summary statistics for the resulting panel described above, comparing average home characteristics across non-participants and participants ($P = 0$ vs. $P = 1$) and non-neighbors and neighbors ($N = 0$ vs. $N = 1$). Panel (a) shows results from estimating Eq. (1.6). Panel (b) shows results from estimating a model with repeat sales (i.e. including parcel fixed-effects in Eq. (1.6)). The repeat sales model contains fewer observations because not every home observed in the assessor data sells multiple times.

Compared to non-participants or non-neighbors, panel (a) and panel (b) demonstrate that participating homes and homes neighboring participants are lower priced, older, sit on larger lots, have larger garages, and more likely have a pool. In both panels, the average magnitude of the remaining home characteristics excepting half bathrooms compare similarly across participants and non-participants, and across neighbors and non-neighbors. Despite the large difference in means for the age and lot size variables, I show in Appendix A.3 that the distributions of age and lot size overlap for participants and non-participants and neighbors and non-neighbors.

Table 1.5 displays summary statistics after restricting sales to pre-2007.[53] Similar to Table 1.4, Table 1.5 shows that participating and neighboring homes selling before 2007 are older, sit on larger lots, have larger garages, and more likely include a pool, compared to their non-participating or non-neighboring counterparts. But unlike in Table 1.4, participating and neighboring homes sell for higher average prices, suggesting that the housing crash may have disproportionately affected areas concentrated with rebate program participants. The average magnitude of the remaining home characteristics excepting half bathrooms compare similarly across participants and non-participants, and across neighbors and non-neighbors.

Table 1.4 and Table 1.5 suggest that participants in the Cash-for-Grass rebate program and their immediate neighbors reside in older sections of Las Vegas. Older areas of Las Vegas may be correlated with unobserved characteristics that influence the value of desert landscape. Below I propose an empirical strategy with a rich set of spatial and temporal

[53] I exclude summary statistics for a pre-2007 repeat sales model since so few observations exist. Under repeat sales, $N_{P=1} = 15$ and $N_{N=1} = 87$.

Table 1.5: *Summary statistics for hedonic panel: sale years 1996 - 2006.*

All sales		*Participant*			*Neighbor*		
	Full	$P = 0$	$P = 1$	p-val	$N = 0$	$N = 1$	p-val
sale price	344,873	344,449	399,708	0.00	343,415	393,853	0.00
age	5.3	5.2	13.8	0.00	5.0	13.9	0.00
lot sqft	5,959	5,949	7,219	0.00	5,920	7,275	0.00
home sqft	2,024	2,025	1,910	0.00	2,027	1,939	0.00
bedrooms	3.40	3.40	3.37	0.22	3.40	3.41	0.86
full bath	2.26	2.26	2.20	0.00	2.26	2.21	0.00
half bath	0.48	0.48	0.34	0.00	0.49	0.34	0.00
pool	0.21	0.21	0.30	0.00	0.21	0.31	0.00
garage sqft	467	467	491	0.00	466	487	0.00
Observations	124,742	123,785	957		121,136	3,606	

Summary statistics reflect the sample used in estimating Eq. (1.6). Insufficient observations preclude estimating a model with parcel fixed-effects. Prices adjusted to 2014 dollars and further restricted to the 1^{st} and 99^{th} percentiles of the sale price distribution. I further drop parcels meeting criteria for undertaking additions.

fixed-effects to address concerns over unobserved characteristics.

1.4.3 Empirical strategy

I estimate the direct and spillover effect of conversion to desert landscape with the panel fixed-effects model shown in Eq. (1.6). I regress the natural log of sale price in \$2014 for parcel i on sale date t on the indicators characterizing conversion to desert landscape described above, home characteristics, Z_i and census block-by-quarter fixed-effects, b_{iq}. For each property, the vector Z_i includes parcel, home, pool, and garage square footage, home age, and the number of bedrooms, full bathrooms, and half bathrooms.

$$\ln p_{it} = \alpha_1 DP_i + \beta_1 P_{it} + \alpha_2 DN_i + \beta_2 N_{it} + \delta Z_i + b_{iq} + \epsilon_{it} \tag{1.6}$$

The intuition behind Eq. (1.6) is differences-in-differences.[54] β_1 describes the approximate[55] average percentage change in the value of a parcel that converts to desert landscape. β_2 describes the approximate average percentage change in the value of a home that lies directly adjacent to a property that converts to desert landscape. β_2 therefore describes the spillover effect, or externality associated with neighboring conversions.

Census block-by-quarter fixed-effects control for average differences across 2010 U.S. Census block boundaries in each quarter of the sample. I am able to group parcels within 6,819 blocks, with the average block in each sample containing about 29 parcels. By controlling for neighborhood effects at such a refined geographic level, I endeavor to minimize concerns that unobserved neighborhood fixed-effects will bias results.

Larger conversions may have stronger impacts than smaller conversions. To explore this possibility, I interact the variables describing area converted or area adjacent at the time of sale (Pa_{it} and Na_{it}) with P_{it} and N_{it}, as shown in Eq. (1.7). θ_1 and θ_2 describe the percentage change in the value of a home from an additional square foot of desert landscape, or an additional adjacent square foot of desert landscape.

$$\begin{aligned} \ln p_{it} &= \alpha_1 DP_i + \beta_1 P_{it} + \theta_1 (P_{it} \times Pa_{it}) \\ &\quad + \alpha_2 DN_i + \beta_2 N_{it} + \theta_2 (N_{it} \times Na_{it}) + \delta Z_i + b_{iq} + \epsilon_{it} \end{aligned} \tag{1.7}$$

1.4.4 Results

In the results below, I show that Cash-for-Grass subsidized conversion to desert landscape increases property values by about 1 percent with little evidence for any spillovers. In appendix A.3, I explore the potential effect of two policy changes on the value of desert

[54] In their investigations into the impact of crime on property values, Linden and Rockoff (2008) and Pope (2008) employ a differences-in-differences strategy that considers "treatment" properties to be within a 0.1 mile radius of a residence of a sex offender, and "control" properties to be between 0.1 and 0.3 miles of a sex offender. In principle, my exploration of the spillover effect of desert landscape mirrors this strategy, but instead of defining treatment and control based on distance, I define treatment and control based on adjacency. My modeling of the direct and spillover effects of desert landscape is in part inspired by the model of health and education externalities proposed by Miguel and Kremer (2004).

[55] In semi-log models with explanatory indicators variables, Halvorsen and Palmquist (1980) show that the coefficients on these indicators do not directly describe percentage effects. However, for small coefficient estimates, the bias is minimal.

landscape, investigate heterogeneous effects over time, and demonstrate robustness of my results to additional specifications.

Main results Table 1.6 shows estimates of the direct and spillover effects of conversion to desert landscape. In all models, the coefficient estimates on the control variables effect housing prices in expected ways[56] and I cluster standard errors at the block level. Column 1, my preferred specification, shows results from estimating Eq. (1.6) with the full range of sale years, 1996 to 2014. Column 2 runs a similar model, but for sales restricted from 1996 to 2006. Restricting sales to pre-housing crisis dates tests the robustness of my column 1 estimates to any additional mortgage spillovers associated with the housing market crash not absorbed by the quarter-block fixed-effects. In column 3, I further address concerns regarding potential unobserved neighborhood and household characteristics by estimating a model that additionally includes parcel fixed-effects (i.e. a repeat-sales model). Across each specification, Table 1.6 demonstrates similarly positive estimates of the direct effect, suggesting my results are robust to concerns over unobserved neighborhood effects and housing crisis impacts.

Cash-for-Grass program requirements preclude conversions to barren landscapes. Furthermore, property owners in Las Vegas often leave potentially landscaped areas uncultivated, or covered in rock. Preference for desert landscape over barren landscape could partially explain my positive coefficient estimates for the direct effect. In columns 4 and 5, I test whether barren landscapes drive results by re-estimating the models in columns 1 and 2 with a subset of non-participants and non-neighbors that applied for the rebate program, were approved for the rebate, but never became enrolled. Since the water authority only approves applicants that have been maintaining a lawn prior to application, estimating my hedonic model with the control group of non-enrolled applicants creates a more direct comparison between desert landscape and turf landscape. The estimate of the direct effect

[56]The two exceptions involve the negative coefficient on full bathrooms in columns 1, 2, and 4, and the negative coefficient on half bathrooms in column 1. Toilets make up the largest share of indoor water use (Bennear *et al.*, 2013), and the negative coefficient on bathrooms may reflect consumers' recognition of higher water bills associated with an increased number of water-intensive fixtures.

Table 1.6: *Regression results for the effect of conversion to desert landscape on home property values.*

	(1)	(2)	(3)	(4)	(5)
DP (ever converts)	0.0048	0.0040		0.0027	
	(0.0013)***	(0.0013)***		(0.0019)	
Direct effect	0.012	0.018	0.021	0.015	0.014
	(0.0034)***	(0.0040)***	(0.012)*	(0.0044)***	(0.025)
DN (neighbors DP)	0.0016	0.0019		0.0012	
	(0.0010)	(0.0010)*		(0.0022)	
Spillover effect	-0.0018	0.00040	-0.011	-0.0018	-0.0040
	(0.0022)	(0.0026)	(0.0072)	(0.0031)	(0.026)
age (years)	-0.0090	-0.0042		-0.0069	
	(0.00070)***	(0.00055)***		(0.00093)***	
parcel sqft	1.7e-05	1.5e-05		1.4e-05	
	(7.6e-07)***	(6.9e-07)***		(8.3e-07)***	
house sqft	2.5e-04	2.3e-04		2.5e-04	
	(4.0e-06)***	(3.7e-06)***		(4.8e-06)***	
bedrooms	0.0041	0.0054		-0.00072	
	(0.0016)**	(0.0015)***		(0.0020)	
full bath	-0.018	-0.0091		-0.020	
	(0.0028)***	(0.0026)***		(0.0050)***	
half bath	-0.0064	0.0043		0.0037	
	(0.0027)**	(0.0024)*		(0.0036)	
pool sqft	1.1e-04	9.1e-05		1.2e-04	
	(3.2e-06)***	(3.1e-06)***		(5.0e-06)***	
garage sqft	2.0e-04	1.9e-04		1.6e-04	
	(1.3e-05)***	(1.5e-05)***		(1.5e-05)***	
Control group	All	All	All	DNF	DNF
Sale years	1996-2014	1996-2006	1996-2014	1996-2014	1996-2014
Fixed-effects					
quarter-block	yes	yes	yes	yes	yes
parcel	-	-	yes	-	yes
adj. R^2	0.95	0.95	0.93	0.94	0.73
Observations	199,037	124,742	40,755	36,608	2,639

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. 2014 adjusted sale prices trimmed at the 1^{st} and 99^{th} percentiles. Fixed-effects: 2010 U.S. Census Blocks (blocks) by quarter of sample (e.g. 1^{st} quarter of 1997 is quarter 5). Sample excludes parcels undertaking additions (see section 1.4.2). Block clustered standard errors.

in column 4, however, compares quite favorably to the estimates in columns 1 through 3. And while the estimate of the direct effect in column 5 is not statistically significant due to the reduced sample size, the magnitude of the point estimate falls within the range of the estimates in columns 1 through 3. These two results provide evidence that barren landscapes do not drive the positive, direct effect of desert landscape on property values.[57] In absolute terms, my results in column 1 imply that conversion to desert landscape increase the value of a home by $3,700, with 95 percent confidence bounds approximately between $1,600 and $5,700.[58]

Turning to the spillover effect, results in Table 1.6 show statistically insignificant and fairly precisely estimated zero effects across specifications. Klaiber *et al.* (2015) find positive spillovers associated with lawns in Phoenix and that consumers value cooler temperatures. One might therefore expect negative spillovers from converting grass to desert landscape, which may increase ambient temperatures around a property. My finding of no spillovers of desert landscape imply that conversion to desert landscape has no effect on micro-climates in Las Vegas, or that a combination of positive aesthetic or other positive spillover effects and negative micro-climates counterbalance each other. I am currently working towards securing electricity demand data to test for an effect of conversion to desert landscape on electricity consumption.

---

[57]Importantly, I assume that non-enrolled applicants continued to maintain a lawn. But non-enrolled applicants may have converted to desert landscape and not taken the rebate. This seems unlikely to have taken place on a large scale. First, if approved, applicants would have little incentive to decline the rebate unless the terms of the rebate were sufficiently burdensome to them. My discussions with water authority staff suggest that only a very few individuals do not follow through on account of program requirements. Second, if non-enrolled applications were converting on a large scale, one would not expect positive point estimates on the direct effect, since there would be little difference between treated observations and control observations. To develop a more precise sense of the relationship between subsidized conversions and total conversions, the water authority provided me with a summary of aerial footage analysis done in 2006, 2008 and 2010. Between 2006 and 2008, the total converted areas are about 50 percent of the change in turf determined by aerial footage. But between 2008 and 2010, aerial footage detected only about 25 percent of the total converted area subsidized by the Cash-for-Grass program (Brand, J. ASPRS Annual Conference, May 3, 2011 - slide deck). The accuracy of aerial footage may therefore make it challenging to directly assess the assumption I make that only those that enroll in the program undertake a conversion to desert landscape.

[58]Absolute increases represent the product of the average home sale and percentage effect corrected for the fact that in log-linear models, estimates on indicator variables do not have a direct interpretation as percentage effects (Halvorsen and Palmquist, 1980). $\$309,077 \times \left(e^{\beta} - 1\right) \approx \$3,700$.

Effect of an additional square foot of desert landscape I present results of estimating the direct and spillover effects of an additional square foot of desert landscape in Table 1.7. The point estimates for the additional direct and spillover impact of an extra square foot of desert landscape (rows 3 and 6 of Table 1.7, respectively) are statistically indistinguishable from zero at the 5 percent level, implying that the presence, rather than the size, of the conversion primarily drives the effect of conversion to desert landscape. Furthermore, the estimates for the direct and spillover effects (rows 2 and 5 respectively), generally agree with those of Table 1.6, reinforcing the conclusion that the direct effect of desert landscape raises the value of a home by small a percentage.[59] And while not shown, the coefficient estimates on each covariate effect housing prices in expected ways.[60]

Relationship between water savings and capitalization In principle, hedonic estimates will capture all private benefits associated with conversion to desert landscape. Two obvious benefits include water bill savings and reduce lawn maintenance costs. Saving 5,000 gal/month yields $150 in annual savings, and I estimate reduced lawn maintenance to be $79 per year.[61] The present discounted value of an infinite stream of these savings with a 5 percent discount rate equals $4,800. This value falls within the confidence interval of the increase in value from converting to desert landscape derived by the hedonic model above. For the present discounted value of water bill and maintenance savings to equal $3,700, the consumers would need to be applying a 7 percent, which seems reasonable. These results suggest that prices of desert landscaped homes reflect monetary savings, but little to no

[59] I do estimate positive, though very small, spillover effects in column 2, significant at the 10 percent level, and negative area effects in columns 1 and 2, also significant at the 10 percent level. While these results provide some evidence for a positive spillover effect that decreases in the size of the conversion, the weight of the evidence I present seems to point towards the conclusion of zero spillover effects.

[60] The two exceptions involve negative coefficient estimates on full and half bathrooms in columns 1 and 2. As explained above, this could arise from consumers' recognition that more water-intensive fixtures, like toilets, may lead to higher water bills.

[61] A quick internet search for average lawn care costs in the Las Vegas area revealed that mowing and maintenance required $37.52 per visit per quarter acre and a fertilization visit per quarter acre cost $61.57. I normalize these values to costs per square feet, and multiply by the average conversion size, 1,348 ft^2. I further assume 3 fertilization applications (based on a posting about lawn care) and 12 mowing visits. The total annual cost comes out to $78.60.

Table 1.7: *Regression results for the effect (in percentage terms) of an extra square-foot of conversion to desert landscape on home property values conditional on the presence of desert landscape.*

	(1)	(2)	(3)
DP (ever converts)	0.0049	0.0040	
	(0.0013)***	(0.0013)***	
Direct effect	0.016	0.029	0.018
	(0.0053)***	(0.0074)***	(0.015)
Direct×area effect	-4.3e-06	-8.3e-06	5.2e-06
	(4.1e-06)	(5.3e-06)	(1.7e-05)
DN (neighbors DP)	0.0015	0.0018	
	(0.0010)	(0.0010)*	
Spillover effect	0.0030	0.0071	-0.015
	(0.0031)	(0.0039)*	(0.0093)
Spillover×area effect	-3.7e-06	-4.8e-06	5.6e-06
	(1.9e-06)*	(2.6e-06)*	(8.3e-06)
Sale years	1996-2014	1996-2006	1996-2014
Fixed-effects			
quarter-block	yes	yes	yes
parcel	-	-	yes
adj. R^2	0.95	0.95	0.93
Observations	198,519	124,375	40,685

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Standard errors clustered at the block level.

blocks: 2010 United States Census Block boundaries.

quarter: quarter of sample (e.g. 1^{st} quarter of 1997 is quarter 5).

2014 adjusted sale prices trimmed at the 1^{st} and 99^{th} percentiles.

Sample excludes parcels undertaking additions (see section 1.4.2).

other benefits associated with conversions. This finding is consistent with Myers (2016), who shows that home buyers fully capitalize differences in energy costs between homes utilizing different heating sources.

1.5 Estimates of cost per gallons saved and net benefits

Cost per gallons saved Since the beginning of the program, the water authority has rebated single-family residents converting once a total of \$53M (\$2014 dollars). I estimate that the resources required to administer the program equal 22.5 percent of total outlays, or about \$12M.[62] The average conversion saves 5,000 gallons per month, implying that total savings for the 26,488 conversions equal 1.6M kgal/year. Following Bennear *et al.* (2013), I convert total rebate outlays plus administrative costs to an annual equivalent by annuitizing \$65M using a 5 percent discount rate and a 30-year time horizon.[64] Dividing annual equivalent expenditures by total annual savings yields an annual program cost of \$2.65/kgal-saved.[65] If I additionally include out-of-pocket conversion costs to the rebate recipient (about \$54M) and repeat the calculations above, annual program costs increase from \$2.65/kgal-saved to \$4.84/kgal-saved. The water bill for an average customer in 2013 equals \$3.54/kgal. If the bill reasonably reflects the cost of supplying water, from the water authority's perspective, the Cash-for-Grass program saves \$0.89/kgal every year. But from a societal perspective, conserving water with the Cash-for-Grass program costs more than supplying the same amount. Appendix A.4 provides further details.

Two caveats bear discussion. On the one hand, water bills tend not to incorporate the opportunity cost of scarce water supplies (Griffin, 2001), which Griffin suggests can be estimated using water market transactions. Edwards and Libecap (2015) report that the

---

[62]I calculate the percentage based on the ratio of total labor costs and overhead to the budget for rebate outlays during the 2015/2016 fiscal year.[63] Since most of the rebates were funded through one-time connection charges, I ignore financing costs. Appendix A.4 provides further discussion.

[64]Similar to the argument that Bennear *et al.* make for toilets, it is unlikely that savings from desert landscape will last indefinitely. I calculate annualized rebate outlays using the following formula: $\frac{r \times \text{Costs}}{1-(1+r)^{-t}}$, where $r = 0.05$ refers to the discount rate and $t = 30$ equals the number of years of assumed sustained savings.

[65]Bennear *et al.* (2013) find a high-efficiency toilet rebate program costs \$7.33/kgal-saved, and Price *et al.* (2014) find that a range of rebate programs cost \$0.39/kgal-saved (for low flow shower heads) to \$8.33/kgal-saved (for an additional low flow toilet), and that a desert landscape rebate costs \$4.51/kgal-saved, assuming the desert landscape stays in place for 25 years (assuming 25 years of sustained savings, the Cash-for-Grass program costs \$2.89/kgal-saved, or \$5.28/kgal-saved if I include the opportunity cost of scare water). Price *et al.* estimate cost-effectiveness using a levelized-cost method which divides the present value of costs by the present value of savings.

median price of a series of agriculture to urban water sales in Nevada during the 2000s equals \$0.06/gal. Most water sales grant the buyer a perpetual right to a certain amount of water each year. Therefore, I assume \$0.06/gal reflects the present discounted value of an infinite stream of benefits arising from the right to withdraw a gallon of water each year in perpetuity. Assuming a 5 percent discount rate, \$0.06/gal implies an annual measure of about \$3/kgal.[66] Including an estimate of the opportunity cost of scarce water therefore increases the real annual cost of supply from \$3.54/kgal to \$6.54/kgal. The program now appears economical, since it saves water for less than the cost of supply.

On the other hand, the measure of program costs per gallon saved ignores the possibility that some rebate recipients may have converted to desert landscape without the subsidy. Savings and associated out-of-pocket conversion costs from any such free-riders should not be attributed to program savings or program costs. I estimate that if free-riders account for over 39 percent of program participants, the program costs more than the cost of supply. In other contexts, free-riding has been found to account for at least half of all rebate recipients (Bennear *et al.*, 2013; Houde and Aldy, 2014; Boomhower and Davis, 2014). These studies, however, explore free-riding in rebates for relatively small ticket items (Bennear *et al.* investigate toilet rebates and Houde and Aldy and Boomhower and Davis examine energy appliance rebates). In contrast, re-landscaping poses substantial costs and occurs infrequently. It therefore seems reasonable to hypothesize a small share of individuals planning to convert to desert landscape apart from the subsidy. However, estimating free-riding would be a valuable extension to this research.[67]

Net benefits While hedonic theory promises to recover willingness to pay for a non-market good, in practice, identifying welfare effects from hedonic analyses pose empirical challenges. Because I exploit panel variation, my estimates of the effect of desert landscape on property values measure a capitalization rate, not necessarily real benefits (Kuminoff

[66] Present discounted value (PDV) $= \frac{B(1+r)}{r} \implies B = \frac{0.05\times 0.06}{1+0.05} = \$2.84$/kgal/year.

[67] One way forward is to estimate a demand model for desert landscape, and then extrapolate to demand at zero subsidy. Boomhower and Davis (2014) essentially take this approach.

and Pope, 2014; Muehlenbachs *et al.*, 2015). Kuminoff and Pope explain this capitalization rate can only be interpreted as a measure of benefits if the hedonic price schedule remains fixed throughout time.[68] Following the suggestion of Kuminoff and Pope, I test for a stable hedonic price schedule by estimating the direct effect of conversion in each year. Figure A.5 in Appendix A.3 illustrates consistent point estimates of the direct effect across time. I therefore assume that the hedonic price schedule is stable and the estimated direct effect reflects real benefits. I estimate the benefits per square foot of desert landscape by dividing the average increase in home values by the average conversion area.[69] I find benefits equal \$2.67/ft^2, with 95 percent confidence bounds ranging from \$1.18/ft^2 to \$4.16/ft^2.

Because the hedonic estimate of the direct effect of conversion only captures monetary savings from reduced water bills and maintenance expenditures (section 1.4), my measure of benefits does not reflect the value of scarce water. As above, I assume the \$0.06/gal median sale price reported by Edwards and Libecap reflects the value of the right to use a gallon of water every year in perpetuity. If one assumes that savings do not erode over time, then my results in section 1.3 imply that conversions save 44.5 gal/ft^2 every year in perpetuity. The product of the sale price (\$0.06/gal) and the average savings per square foot (44.5 gal/ft^2) therefore yields a rough approximation of the scarcity value of water embedded in converting one square foot to desert landscape, or \$2.66/ft^2.

I estimate total costs to be \$3.33/ft^2. The largest share of program costs arises from conversion costs, \$3.00/ft^2. I again estimate administrative cost to be 22.5 percent of total rebate outlays, then divide by the number of conversions and the average conversion area.[70] Administrative costs come to \$0.33/ft^2. Rebates are simply transfers from the utility to the customer, and therefore do not represent costs from a societal perspective.

---

[68] A second challenge associated with hedonic estimation of welfare effects involves estimating demand parameters for the purpose of deriving aggregate welfare impacts. Epple (1987) and Bartik (1987) provide clarifying discussions.

[69] Since my hedonic estimates include parcels undertaking more than one conversion, I include data from this larger set of participants in my analysis of net benefits. Total rebate outlays equal \$62M, the number of conversions equals 31,049, and the average conversion size equals 1,379 ft^2.

[70] See footnote 69.

Ignoring scarcity, program net benefits equal -\$0.66/ft^2, with a 95 percent confidence interval of -\$2.15/ft^2 to \$0.84/ft^2. Including scarcity, net benefits increase to about \$2.00/ft^2. My estimate of net benefits, however, may not accurately reflect true net benefits. One the one hand, I have not incorporated any positive health and climate externalities arising from reducing the energy needed to treat and distribute water. Including positive externalities would increase net benefits.[71] On the other hand, I have ignored free-riders. Including free-riders would reduce the benefits arising from conversions as well as the associated conversion costs, leading to lower estimates of net benefits. But provided most rebate recipients would not convert to desert landscape without the subsidy, my analysis suggests that the Cash-for-Grass program may enhance welfare.

1.6 Conclusion

I have analyzed the water savings and net benefits generated by the Southern Nevada Water Authority's Cash-for-Grass rebate program. Using event studies and panel fixed-effects models, I estimate an average conversion saves about 5,000 gallons per month. Furthermore, these savings remain stable over time, and are robust to a series of specifications and control samples. The stability of savings suggests that program participants do not substitute to other water intensive activities, but simply cut back on outdoor irrigation. Finally, since savings per square foot are inversely related to incentives to participate, encouraging greater participation over the life of the program may have come at the expense of program cost-effectiveness.

Consumers value conversions to desert landscape. I find that a conversion increases the value of a Las Vegas single-family home by \$3,700 (about 1 percent), and that the present discounted value of estimated water bill and lawn maintenance savings essentially accounts for the entire increase in home values. I also find little evidence that conversions to desert landscape have any net impact on neighboring properties. To further explore

[71] I may also be understating net benefits if reductions in utility revenue are less than benefits from reduced operating costs. Appendix A.4 provides further discussion.

desert landscape externalities, I am securing electricity demand data for the Las Vegas area and plan to test for any impact that conversion to desert landscape has on electricity consumption, both on converting and neighboring homes.

When I incorporate an estimate of the value of scarce water, the Cash-for-Grass program yields positive net benefits. But I also find that a modest 6 percent price increase would have achieved similar savings over the life of the program. Increasing prices raise distributional concerns, but subsidies also pose issues of equity. I show that program participants with higher pre-conversion water demand achieve the most savings, and thus targeting these individuals would increase program cost-effectiveness (Allcott, 2011). However, higher water users also tend to earn larger incomes. Public utility managers deciding between prices and subsidies may therefore face a trade-off between regressive price policies and subsidizing wealthy individuals.

Chapter 2

Impacts of Information Disclosure on Drinking Water Violations[72]

2.1 Introduction

The recent crisis in Flint, Michigan concerning toxic levels of lead in the city's drinking water thrust water quality into the national spotlight. Ever since the passage of the Safe Drinking Water Act (SDWA) in 1974, however, the U.S. Environmental Protection Agency has endeavored to limit the occurrence of Flint-like crises by regulating drinking water quality across the thousands of the country's public water systems (Tiemann, 2014). In 1996, Congress substantially amended the SDWA. One of the new requirements mandated community water systems–a category of public water system serving a non-transient population more than 6 months per year–to submit annual water quality reports to their customers. The move in part reflected a growing trend among regulators to achieve policy goals through information disclosure, especially in contexts where direct, centralized regulatory control pose challenges (Fung *et al.*, 2007). While all community water systems are required to generate reports, the method of disclosure required by the 1996 amendments depends upon the water system service population. Systems serving up to 500 customers

[72] Co-authored with my committee member, Sheila Olmstead.

must make a report available upon request; systems serving between 501 and 9,999 customers must publish a report in a local newspaper; systems serving at least 10,000 customers must mail the report; and systems serving at least 100,000 customers must additionally make the report available on the internet. The amendments first required water systems to submit reports in 1998.

In our analysis, we estimate the change in national, health-based water quality violations due to the annual disclosure of water quality reports required by the 1996 SDWA amendments. We collect data on all water systems and all water quality violations in the United States from the U.S. Environmental Protection Agency (EPA) through a Freedom of Information Act request. We build a panel of community water systems active from 1990 through 2001, matching the number of health-based water quality violations with the violating water system in each year of our panel. We use these data in differences-in-differences models that exploit the timing of the information disclosure policy as well as the discontinuity in the method of disclosure to estimate the impact of the publication, mailing, and online posting disclosure requirements. We find that the publication and mailing requirements reduce water quality violations by around 30 percent relative to pre-disclosure policy levels. We further show that these reductions remain stable over time. We find less evidence that the online posting requirement had any effect on water quality violations.

We also explore heterogeneity in water system response along dimensions of service population income and contaminant types. We match water systems with the county in which the water system operates, and then use median income census data by county to categorize water systems in median income deciles. We find that water systems serving higher income counties respond more strongly to the publishing and mailing requirements compared to systems serving lower income counties. Regarding heterogeneity in response across contaminant types, we focus on violations pertaining to microbial contaminants (such as cryptosporidium or e. coli) and contaminants created as byproducts of the disinfection process (primarily trihalomethanes). Using a seemingly unrelated regression framework, we show that the publication and mailing requirements reduce violations, but that violations from disinfection byproducts do not increase as a result. While this could indicate only a weak link between the disinfection process and disinfection byproducts, it could also indicate that water utilities used disinfection processes that created disinfection byproducts other than those regulated by the EPA.[73]

Our analysis contributes to two main areas in the literature. Our first area of contribution lies with empirical investigations of information disclosure policies. The effects of information disclosure have been studied in a wide range of contexts, including finance, health, and education (Dranove and Jin, 2010), energy (Allcott and Sweeney, 2016), and the environment (Mastromonaco, 2015). But little empirical work exists regarding the long-term impacts of information disclosure policies (Dranove and Jin, 2010). In our analysis, we begin to fill this gap by estimating the impacts of water quality reports several years after the reports were first required. We also offer some insight into the most salient methods of information disclosure by exploring the impacts of the publishing, mailing, and online information disclosure methods. Finally, we contribute to the growing, but still somewhat limited literature investigating heterogeneity in responsiveness to information disclosure.[74]

Our second area of contribution lies within the broad economics literature covering water quality. Much of this literature explores issues regarding ambient water quality (Leggett and Bockstael, 2000; Cho *et al.*, 2011), with some analyses focusing specifically on the Clean Water Act (Chakraborti and McConnell, 2012; Keiser and Shapiro, 2017).[75] A second component of the water quality literature explores issues surrounding drinking water quality. Research points to the substantial benefits from providing drinking water infrastructure (Olmstead, 2010), and investigates the impacts of drinking water contaminants (Muehlenbachs *et al.*, 2015; Wrenn *et al.*, 2016). But at least in developed countries, there seems to be less work

[73]Until 2002, the only disinfection byproducts the EPA regulated were trihalomethanes. Depending upon the disinfectant used, however, bromate and chlorite can also be byproducts of the disinfection process. The EPA began regulating these two chemicals in 2002.

[74]Examples of analyses exploring heterogeneity in responsiveness include Shimshack *et al.* (2007), Delmas *et al.* (2010), Powers *et al.* (2011), and Doshi *et al.* (2013).

[75]The Clean Water Act regulates the surface waters of the United States.

exploring efficacy of drinking water quality regulations.[76] We assess a specific intervention designed to improve drinking water quality at the national scale.

We organize the remainder of the paper as follows. In section 2.2 we discuss our policy context, the amendments to the Safe Drinking Water Act. We describe our water system and water quality violations data in section 2.3 and our differences-in-differences empirical models in section 2.4. We present and discuss results in section 2.5. In section 2.6 we conclude.

2.2 Policy Context: 1996 Safe Drinking Water Act Amendments

Ever since the passage of the Safe Drinking Water Act (SDWA) in 1974, the EPA has regulated the quality of drinking water delivered by water utilities across the United States. Initially, the SDWA regulated 22 contaminants, such as lead, arsenic, coliform bacteria, and mercury. New contaminants have been added to the regulated list as the act has been amended over time, such as radioactive isotopes, copper, and cryptosporidium. As of 2013, the EPA regulates 91 contaminants under the SDWA.[77]

In 1996 Congress undertook a major overhaul of the SDWA. Among other requirements, the 1996 Amendments instituted a rule stipulating that all community water systems[78] submit annual water quality reports to their customers. Sometimes referred to as Consumer Confidence Reports (CCR), these annual water quality report cards are intended to "Improve public health protection by providing educational material to allow consumers to make educated decisions regarding any potential health risks pertaining to the quality, treatment,

---

[76]Innes and Cory (2001) explore the optimal design of public notification in the context of the SDWA. They explore these issues in a theoretical context, rather than in an empirical setting as in our analysis.

[77]A timeline of regulated contaminants can be found at: https://www.epa.gov/sites/production/files/2015-10/documents/dw_regulation_timeline.pdf.

[78]Though discussed in more detail in section 2.3, community water systems comprise one category of public water systems, the other two being transient non-community water systems, and non-transient, non-community water systems. Community water systems serve the vast majority of the American population.

and management of their drinking water supply".[79] In addition to including information about the water system and contaminants, the water quality report must make note of any water quality violations that occurred during the reporting period.[80]

The required method of disseminating the annual water quality report depends upon the service population of the water system, with requirements changing at service populations of 501, 10,000, and 100,000 customers. Water systems serving few customers face the least stringent requirements; below a service population of 501, water systems "may provide a notice stating the CCR is available upon request".[81] In fact, systems of all sizes face this requirement,[82] making the "available upon request" requirement the baseline disclosure requirement to which we make comparisons. Between a service population of 501 and 9,999 customers, water systems "may publish their CCR in a local newspaper".[83] Water systems must mail their report if they serve 10,000 or more customers, and systems "serving 100,000 or more persons must also post its current year's report on a publicly accessible site on the Internet".[84] We refer to these three information disclosure requirements as the publishing, mailing, and online posting disclosure requirements, respectively.[85]

Why might these information disclosure methods reduce water quality violations? Bennear and Olmstead (2008) explore the impact of the mailing requirement on Massachusetts drinking water violations. They postulate that the most likely mechanism through which information disclosure has an effect is a desire by water systems to avoid customer complaints. Many water systems are publicly owned, and elected officials often sit on the

[79]EPA Quick Reference Guide No. 816-F-09-009. http://nepis.epa.gov/Exe/ZyPDF.cgi?Dockey=P100529A.txt

[80]EPA 816-F-09-009

[81]EPA 816-F-09-009

[82]See EPA 816-F-09-009.

[83]EPA 816-F-09-009

[84]EPA 816-F-09-009

[85]To be precise, the default disclosure method involves mailing, with the additional online disclosure required for large systems, and the less stringent disclosure requirements available to smaller systems. See EPA 816-F-09-009.

public utility commissions that govern non-public, investor-owned utilities. Thus whether directly or indirectly, water systems are beholden to their customer base, and demonstrating compliance with EPA regulations in the required annual water quality reports would be one way to maintain customer satisfaction.

2.3 Community water system and water quality violation data

We obtained characteristics of public water systems as well as associated violations from the EPA via a Freedom of Information Act request in July 2014. The water system data covers the universe of U.S. public water systems and provides information such as population served, whether the system is operational, system deactivation date if applicable, system location by state, whether the system uses surface water or groundwater, and the ownership structure (e.g. either public or private). The violation data contains every violation recorded by the EPA since the agency began keeping records.

In our analysis, we focus on health-based water quality violations from community water systems serving all fifty U.S. States.[86] Community water systems comprise a subset of public water systems that serve a non-transient community for more than six months out of the year. As of July 2014, community water systems make up about 34 percent of all active water systems, but serve nearly 300 million persons. In contrast, non-community water systems serve about 19 million persons. So while community water systems make up only a third of all public water systems, they serve the majority of the American population.

Health-based water quality violations occur either because the concentration of a regulated contaminant exceeds a specified threshold or a water system fails to follow an established technique for treating a contaminant. A water system would incur a maximum contaminant level (MCL) violation if a regulated contaminant such as arsenic exceeds the maximum contaminant level set by the EPA. The EPA regulates other contaminants, such as lead and copper, through treatment techniques rather than contaminant thresholds. A water

[86] We therefore exclude water systems managed by Native American tribes and other territories.

system failing to apply the appropriate treatment technique would result in a treatment technique (TT) violation. While not the focus of our analysis, the EPA also collects information on two additional types of violations; monitoring and reporting (MR) violations and "Other" violations. MR violations make up the vast majority of all violations and essentially amount to procedural violations. For example, if a utility fails to take a water quality reading at a specified time or does not monitor the quality of its water source, it incurs an MR violation. "Other" violations also involve procedural violations. Water systems that do not submit an annual water quality report, or fail to notify its customers of any acute health risks discovered in its water would incur violations classified as "Other".

We match water system information with violations data, generating a panel of water quality violations in each system-year from 1990 to 2001 for systems that remained active throughout this time frame. Beyond 2001, changes in water quality standards for various contaminants complicate identifying an effect of the disclosure policy. By keeping water systems that have remained active between 1990 and 2001, we also ensure that we have a balanced panel. Some water utility systems report deactivation dates in the early 1900s. Since the EPA did not exist before 1970, we reason these dates could be a product of a Y2K bug, and drop such utilities.[87]

The full panel includes 46,900 water systems observed over 12 years, for a total of 562,800 observations. Table 2.1 provides summary statistics for water systems in the full panel. The average water system serves 6,104 customers. Most water systems, however, are small. About 55 percent of systems (25,847 systems) serve 500 or fewer individuals. A further 36 percent (16,906 systems) of systems serve between 501 and 9,999 customers, 8 percent (3,732 systems) serve between 10,000 and 99,999 customers, and the final 1 percent (415 systems) serve at least 100,000 customers.[88] Public institutions govern about half of the water systems in our panel, and most systems source their water from groundwater. We also categorize

[87]Specifically, we drop 21 utilities with deactivation dates prior to 1950.

[88]We also describe our systems in terms of the three information disclosure thresholds. 45 percent of systems serve over 500 customers, 9 percent serve at least 10,000 customers, and the final 1 percent serve at least 100,000 customers.

Table 2.1: *Summary of water system characteristics.*

	Mean	SD	Min	Max	Obs
service population (all systems)	6,104	59,471	0	8,271,000	46,900
service population (0-500)	160	132	0	500	25,847
service population (501-9,999)	2,630	2,288	501	9,999	16,906
service population (10k-99,999)	28,285	19,684	10,000	99,750	3,732
service population (≥100k)	318,394	540,420	100,000	8,271,000	415
$T_{pub} = \mathbb{1}(\text{service pop.} > 500)$	0.45	0.50	0	1	46,900
$T_{mail} = \mathbb{1}(\text{service pop.} \geq 10\text{k})$	0.09	0.28	0	1	46,900
$T_{web} = \mathbb{1}(\text{service pop.} \geq 100\text{k})$	0.01	0.09	0	1	46,900
publicly owned systems	0.50	0.50	0	1	46,661
systems using surface water	0.23	0.42	0	1	46,869
median income (decile 1)	24,878	2,054	17,434	27,209	1,801
median income (decile 2)	28,582	703	27,217	29,714	2,230
median income (decile 3)	30,753	579	29,724	31,645	2,753
median income (decile 4)	32,549	466	31,649	33,285	3,159
median income (decile 5)	34,417	547	33,297	35,260	3,528
median income (decile 6)	36,257	559	35,268	37,293	3,809
median income (decile 7)	38,379	652	37,296	39,497	4,791
median income (decile 8)	41,206	975	39,508	42,723	5,236
median income (decile 9)	45,389	1,611	42,742	48,452	6,529
median income (decile 10)	57,246	7,568	48,515	93,383	7,681

median income: 2003 median income of the county where the water system operates.

our panel of water systems based on the median income of the county in which the water system operates. We first categorize United States' counties into deciles based on 2003 median income values (thus county income in Table 2.1 is in 2003 dollars) and then match water systems to these deciles based on the county in which the water system operates.[89]

[89]We are unable to match all water systems with a county, and even though some water systems serve more

As Table 2.1 shows, water systems are concentrated in the higher income deciles.

Table 2.2: *Summary of violations for system-years.*

	Totals	Mean	SD	Min	Max	$\Pr(v_{it} > 0)$	Obs
Health violations	63,922	0.114	0.59	0	29	0.069	562,800
systems serving 0-500	33,933	0.109	0.55	0	21	0.067	310,164
systems serving 501-9,999	23,350	0.115	0.63	0	29	0.067	202,872
systems serving 10k-99,999	6,244	0.139	0.63	0	24	0.087	44,784
systems serving $\geq$100k	395	0.079	0.42	0	10	0.053	4,980
$T_{pub} = \mathbb{1}(\text{service pop.} > 500)$	29,989	0.119	0.63	0	29	0.071	252,636
$T_{mail} = \mathbb{1}(\text{service pop.} \geq 10\text{k})$	6,639	0.133	0.61	0	24	0.084	49,764
MCL violations	47,452	0.084	0.40	0	17	0.058	562,800
TT violations	16,470	0.029	0.42	0	29	0.011	562,800
Microbial violations	53,398	0.095	0.54	0	29	0.058	562,800
DBP violations	642	0.001	0.05	0	8	0.001	562,800
MR violations	485,282	0.862	6.71	0	1,247	0.167	562,800
Other violations	26,393	0.047	0.52	0	59	0.032	562,800
All violations	575,597	1.023	6.80	0	1,247	0.228	562,800

DBP violations are disinfection byproducts, primarily Trihalomethanes.

For each water system-year in our panel, we record the number of health-based violations incurred by each system during that calendar year. Table 2.2 summarizes these health-based violations. Health-based violations take place infrequently, occurring in 6.9 percent of system-years. And though the smaller systems generate the most violations, smaller systems are not appreciably more likely to incur a violation than larger systems. In fact, the class of system most likely to incur a violation are those systems serving between 10,000 and 99,999 customers. MCL violations account for 74 percent (47,452 violations) of health-based violations, with TT violations making up the remaining 26 percent (16,470

than one county, our data associates only one county with a water system. See section 2.4.

violations). Breaking down violations another way, most health-based violations represent violations of microbial standards (53,398 violations). Disinfectants used to remove pathogens from drinking water react with other substances, producing harmful byproducts regulated by the EPA. Violations associated with disinfection byproducts (DBP violations in Table 2.2) are rare in our sample. In fact, health violations in general make up a small share of all violations reported to the EPA. MR violations comprise the largest share with 485,282 recorded violations. A further 26,393 are "Other" violations. Depending upon the specific nature of the violation and response of the water utility, violations can remain "open" for several months or even a year or more before the water system comes back into compliance. If a water system first reports a health-based violation to the EPA in March 1999 and came back into compliance in January 2000, we count this health-based violation as a single violation recorded in 1999.[90]

Figure 2.1 visualizes our panel, illustrating total violations in each year for four categories of community water systems; systems serving between 0 and 500 customers, systems serving between 501 and 9,999 customers, systems serving between 10,000 and 99,999 customers, and systems serving over 100,000 customers. Since there are many more systems serving fewer customers and since regulatory stringency has changed over time, we normalize violations by the number of systems in the four categories described above as well as the number of MCL regulations systems must comply with in each year of our panel. For all categories of water systems, normalized violations increase sharply after 1990. Becoming effective at the end of 1990, both the Surface Water Treatment Rule and the Total Coliform Rule drive this early increase in normalized violations. Beginning in the early to mid-1990s, normalized violations fall for all categories of water systems, with the largest declines experienced by those water systems serving the most customers, and the smallest declines experienced by the systems serving the fewest customers.

We motivate our empirical models below with two additional summaries of our panel. First, in Figure 2.2, we plot average violations as a function of water system size (in logs) for

[90]Considering how the disclosure policy impacted the duration of violations is the subject of future work.

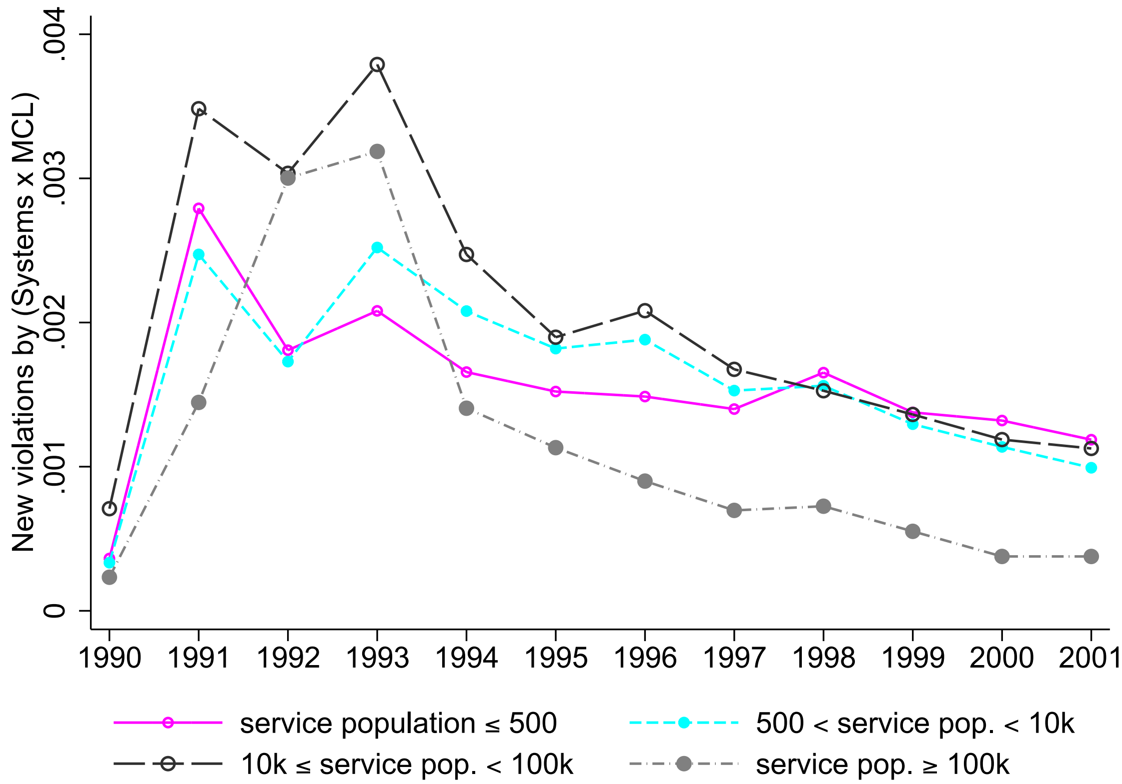

Figure 2.1: *Number of health-based water quality violations in each year of the panel for community water systems serving between 0 and 500 customers, between 501 and 9,999 customers, between 10,000 and 99,999 customers, and over 100,000 customers. We normalize violations by the total number of water systems within each population service category and the annual number of MCL rules (the number of MCL rules grew from 31 in 1990 to 83 in 2001).*

1997 and 1998. We bin water systems within each of the four service population categories, calculating the average service population for each bin as well as the average number of violations for that bin in 1997 and 1998. We then plot average violations as a function of the average water system service population for each bin for each year.[91] Blue vertical lines illustrate the three information disclosure thresholds. Since the information disclosure policy takes effect in 1998, we may expect to observe substantial changes in violations in 1998 relative to 1997 on either side of each disclosure threshold. Probably the most noticeable pattern, however, is the noisy nature of the data.

[91]For ease of visualization, we remove water systems serving under 25 customers and over 1 million customers.

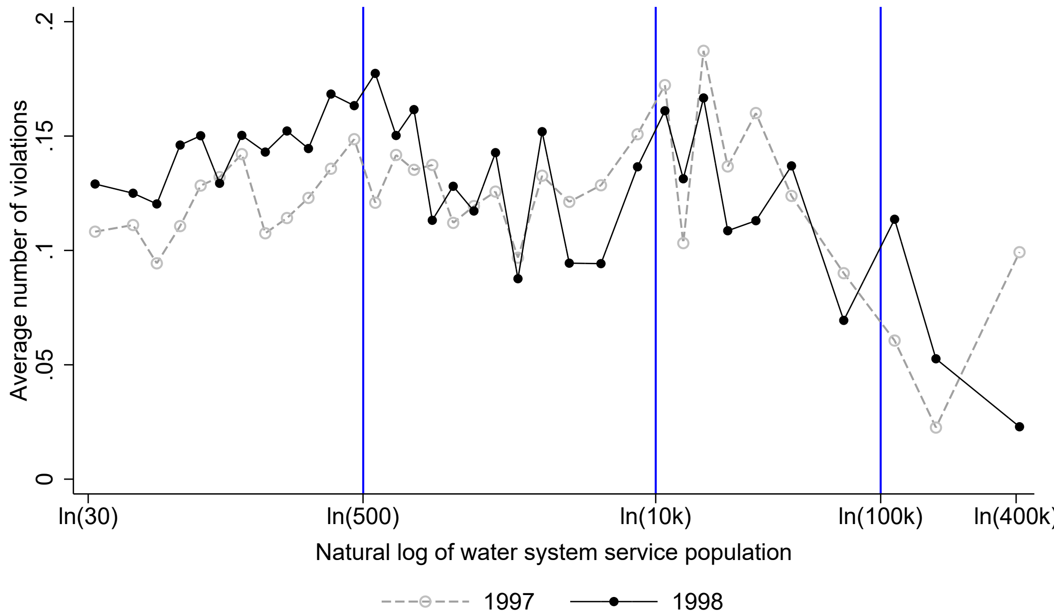

Figure 2.2: *Average number of violations as a function of the natural log of a water system's service population. We first categorize systems into four groups defined by the information disclosure service population thresholds; 501, 10,000 and 100,000 customers. Within each group, we bin water systems, calculating the average service population for each bin as well as the average number of violations. We then plot average violations as a function of water system service population. Blue vertical lines illustrate the three information disclosure thresholds. Before categorizing systems into our four groups, we remove water systems serving under 25 customers and over 1 million customers.*

Second, we explore the raw differences in means on either side of each information disclosure threshold. As such, these calculations represent the most basic form of a differences-in-differences design testing for an effect of the information disclosure policy. Table 2.3 illustrates our results. At the publishing threshold, average violations for those systems serving over 500 customers fall by -0.023 violations after 1998. In contrast, average violations increase slightly for systems serving 500 or fewer customers after 1998, leading to a differences-in-differences estimate of -0.032 violations for systems serving over 500 customers. Raw differences in means also imply reductions in violations at the mailing and online posting thresholds. The empirical tests in section 2.4 examine whether these patterns in the raw data are borne out in quasi-experimental models testing for causal impact.

Table 2.3: *Raw differences in means for each disclosure threshold.*

	Above threshold			*Below threshold*			
	Post-1998	Pre-1998	Δ_1	Post-1998	Pre-1998	Δ_2	$\Delta_1 - \Delta_2$
Publishing	0.103	0.127	-0.023	0.115	0.107	0.008	-0.032
Mailing	0.101	0.149	-0.048	0.110	0.112	-0.002	-0.046
Online Posting	0.042	0.098	-0.056	0.110	0.116	-0.006	-0.050

2.4 Empirical strategy

2.4.1 Differences-in-differences empirical approach

We employ a differences-in-differences model to assess the impact of the publishing, mailing, and online disclosure thresholds on health-based drinking water quality violations. Our empirical strategy therefore exploits the variation in violations before and after the policy for those systems "treated" with a method of information disclosure relative to those systems that remain untreated. Our main model, Eq. (2.1), follows that of Bennear and Olmstead (2008). In particular, we regress the number of health-based violations, v_{it}, from water system i in year t on an interaction between a treatment indicator, T_i, and an indicator $Post_t$ that describes observations in or after 1998, a flexible function of system size, $\sum_{j=0}^{k} \delta_j \left(Post_t \times size_i\right)^j$, where *size* describes the service population in 100,000s, system fixed-effects, u_i, state-by-year fixed-effects, $st_i \times d_t$, and an error term, ϵ_{it}. The treatment indicator T_i equals one if the system serves a population exceeding one of the three disclosure thresholds; 501, 10,000, or 100,000 customers.

$$v_{it} = \theta \left(T_i \times Post_t\right)_{it} + \left(\sum_{j=0}^{k} \delta_j \left(Post_t \times size_i\right)^j\right) + \left(st_i \times d_t\right) + u_i + \epsilon_{it} \qquad (2.1)$$

In Eq. (2.1), θ describes the differences-in-differences estimate of the impact of requiring the respective disclosure method (publishing, mailing, or online posting) on health-based drinking water quality violations, net of the differential effect of system size. To see this mathematically, considering the following derivation, where for the sake of example, we choose a quadratic function of system size:

$$
\begin{aligned}
E[v_{it}|T_i = 1, Post_t = 1] &= \theta + \delta_1 E[size_i|T_i = 1] + \delta_2 E[size_i^2|T_i = 1] + \\
&\qquad E[st_i \times d_t|Post_t = 1] + E[u_i|T_i = 1] \\
E[v_{it}|T_i = 1, Post_t = 0] &= E[st_i \times d_t|Post_t = 0] + E[u_i|T_i = 1] \\
E[v_{it}|T_i = 0, Post_t = 1] &= \delta_1 E[size_i|T_i = 0] + \delta_2 E[size_i^2|T_i = 0] + \\
&\qquad E[st_i \times d_t|Post_t = 1] + E[u_i|T_i = 0] \\
E[v_{it}|T_i = 0, Post_t = 0] &= E[st_i \times d_t|Post_t = 0] + E[u_i|T_i = 0] \\
DD &= \Delta_1 - \Delta_2 \\
\Delta_1 &= E[v_{it}|T_i = 1, Post_t = 1] - E[v_{it}|T_i = 1, Post_t = 0] \\
\Delta_2 &= E[v_{it}|T_i = 0, Post_t = 1] - E[v_{it}|T_i = 0, Post_t = 0]
\end{aligned}
$$

$$
\begin{aligned}
\Delta_1 &= \theta + \delta_1 E[size_i|T_i = 1] + \delta_2 E[size_i^2|T_i = 1] + E[st_i \times d_t|Post_t = 1] + E[u_i|T_i = 1] \\
&\quad - E[st_i \times d_t|Post_t = 0] - E[u_i|T_i = 1] \\
&= \theta + \delta_1 E[size_i|T_i = 1] + \delta_2 E[size_i^2|T_i = 1] + E[st_i \times d_t|Post_t = 1] - E[st_i \times d_t|Post_t = 0] \\
\Delta_2 &= \delta_1 E[size_i|T_i = 0] + \delta_2 E[size_i^2|T_i = 0] + E[st_i \times d_t|Post_t = 1] + E[u_i|T_i = 0] \\
&\quad - E[st_i \times d_t|Post_t = 0] + E[u_i|T_i = 0] \\
&= \delta_1 E[size_i|T_i = 0] + \delta_2 E[size_i^2|T_i = 0] + E[st_i \times d_t|Post_t = 1] - E[st_i \times d_t|Post_t = 0]
\end{aligned}
$$

$$
\begin{aligned}
&DD = \Delta_1 - \Delta_2 \implies \\
&\quad \theta + \delta_1 E[size_i|T_i = 1] + \delta_2 E[size_i^2|T_i = 1] + E[st_i \times d_t|Post_t = 1] - E[st_i \times d_t|Post_t = 0] \\
&\quad - \big(\delta_1 E[size_i|T_i = 0] + \delta_2 E[size_i^2|T_i = 0] + E[st_i \times d_t|Post_t = 1] - E[st_i \times d_t|Post_t = 0]\big) \\
&DD = \theta + \underbrace{\delta_1 \big(E[size_i|T_i = 1] - E[size_i|T_i = 0]\big) + \delta_2 \big(E[size_i^2|T_i = 1] - E[size_i^2|T_i = 0]\big)}_{\text{differential effect of system size}}
\end{aligned}
$$

If we derive the above without the flexible function $f(size)$, θ takes on the standard differences-in-differences interpretation. With the flexible function of system size, the differences-in-differences estimator includes θ, as well as an estimate of the differential effect of system size. The differential effect captures any post-treatment differences across systems due to size, apart from the disclosure requirement. For example, larger systems may be equipped with additional administrative resources enabling compliance with water quality standards, or economies of scale in drinking water treatment may reduce compliance costs and make compliance more likely (Bennear and Olmstead, 2008).[92] By including the flexible function of system size, θ isolates the effect of the disclosure requirement immediately around the threshold, similar to a regression discontinuity design (Bennear and Olmstead, 2008).

Because systems below information disclosure thresholds may voluntarily undertake more comprehensive disclosure than required, θ represents a lower bound on the effect of the information disclosure. That is, it is entirely possible that some systems serving fewer than 10,000 customers elect to mail their water quality reports. While we believe that a water system would have little incentive to voluntarily mail its water quality report on account of the additional costs incurred, any such voluntary behavior by systems serving less than 10,000 customers would dilute the impact of the mailing requirement picked up by θ in Eq. (2.1) leading to estimates that understate the true impact of mailing. The same argument also applies to the publishing and online posting requirements.

In addition to quantifying the individual effect of each disclosure requirement, we also explore the impact of the information disclosure requirements simultaneously. For the publishing requirement, $T_i = 1$ includes systems that must publish, mail, and post online the water quality report. Therefore, the estimate on the interaction between the publishing treatment indicator and the post-1998 information disclosure requirement indicator will potentially pick up some of the effect of the mailing and online posting requirement.

[92]See also Table 1 in Raucher *et al.* (2011). Here the authors show the costs per household associated with complying with the updated arsenic MCL. The costs range from $407/household for the smallest systems to $1/household for the largest system.

Similarly, for the mailing requirement, $T_i = 1$ includes systems that must mail the water quality report, as well as some systems that must post the report online. Therefore, the interaction between the mailing treatment indicator and the post-1998 information disclosure requirement indicator may contain some of the effect of the online posting requirement. To distinguish the effects of publishing, mailing, and online posting requirements, we estimate a model that includes an interaction between each treatment indicator and the post-1998 indicator, as shown in Eq. (2.2). The subscripts on θ and T now refer to specific publishing (501) mailing (10k) or online posting (100k) requirement thresholds.

$$\begin{aligned} v_{it} = & \ \theta_{501} \left(T_{501} \times Post\right)_{it} + \theta_{10k} \left(T_{10k} \times Post\right)_{it} + \theta_{100k} \left(T_{100k} \times Post\right)_{it} \\ & + \left(\sum_{j=0}^{k} \delta_j \left(Post_t \times size_i\right)^j \right) + u_i + \left(st_i \times d_t\right) + \epsilon_{it} \end{aligned} \tag{2.2}$$

2.4.2 Persistence of information disclosure

We also explore the impact of the information disclosure policy over time. At the outset, it is unclear whether we should expect any effect to erode, persist, or strengthen over time. One the one hand, a service population may become accustomed to poor water quality reports and instead of voice concern, engage in some form of avoidance behavior (Zivin *et al.*, 2011). The poorly performing water system may then learn over time that less than stellar reports have few consequences, thereby reducing or eliminating any effect associated with being required to distribute water quality reports. On the other hand, water quality reports may increase water system operators' attentiveness to violations, leading to increased expertise associated with reducing violations. This learning based response would then lead to an increasing impact of the disclosure policy over time.[93]

To assess the persistence of any impacts for each disclosure threshold, we estimate Eq. (2.3) and Eq. (2.4). In Eq. (2.3), we interact our treatment indicator, T_i, with an annual linear time trend describing the number of years since 1998. The estimate on this

[93]We should distinguish between learning how to address violations from learning about the existence of violations. Water systems have been reporting violations to the EPA for many years prior to 1998, so it seems unlikely that the information disclosure would cause utilities to learn about the existence of violations (Bennear and Olmstead, 2008).

interaction describes whether the incentive to reduce water quality violations due to the information disclosure requirement increases ($\beta < 0$) or decreases ($\beta > 0$) with time. In Eq. (2.4), we interact our treatment with annual dummy variables beginning in 1998. Rather than providing an average annual effect, Eq. (2.4) estimates the impact of the information disclosure policy in each post-disclosure year.

$$v_{it} = \theta\left(T \times Post\right)_{it} + \beta(yrsPost_t \times T_i) + \delta\left(f(\text{size})\right) + (st_i \times d_t) + u_i + \epsilon_{it} \tag{2.3}$$

$$v_{it} = \sum_{t=1998}^{2001} \beta_t(year_t \times T_i) + \delta\left(f(\text{size})\right) + (st_i \times d_t) + u_i + \epsilon_{it} \tag{2.4}$$

2.4.3 Heterogeneity in water system response

Empirical analyses of information disclosure policies have begun to explore response heterogeneity (Dranove and Jin, 2010; Doshi *et al.*, 2013). For example, Delmas *et al.* (2010) show that investor owned utilities shift away from fossil fuels to cleaner fuels when required to disclosure their fuel mixture, and that this effect becomes more pronounced in utilities serving primarily residential customers. Powers *et al.* (2011) finds that large pulp and paper manufactures in India reduce pollution in response to an environmental quality rating program, especially in higher income areas. Building on these studies, we test for whether water systems heterogeneously respond to information disclosure as a function of service population income levels. We specifically choose service population income for two reasons. First, it allows us to test whether the results found by Powers *et al.* generalize to a developed country context. Second, income is a readily available proxy for education and newspaper readership, high levels of which Shimshack *et al.* (2007) suggest drive stronger responses to information disclosure. Based on the empirical finding of Powers *et al.* and the fact that income tends to be correlated with education levels, which appear to drive responsiveness to information, we would expect water systems to reduce violations more in higher income areas as a response to the water quality reporting requirements.

To explore the impact of customer income level we estimate Eq. (2.1) separately for each income decile of the county within which the water system resides.[94] We source income data from the 2003 census data[96] and categorize each county into income deciles based on the median county income. Finally, we match these county-decile data to our water system information, assigning each water system an income decile based on the water system-county match.[97]

We also hypothesize that there may exist heterogeneity in response at the level of individual contaminants. Adamowicz *et al.* (2007) find that individuals have higher willingness to pay to reduce microbial contaminants compared to known carcinogens, and as a consequence, water systems may be more eager to reduce microbial contaminants than other contaminants that increase the risk of cancer or contribute to other long-term deleterious health conditions. This is particularly relevant to the problem of drinking water regulations, because the disinfectants used to reduce the presence of disease-carrying pathogens (such as chlorine and chloramine) are, themselves, the sources of carcinogenic disinfection byproducts, which are also regulated drinking water contaminants. Thus, there is a direct trade-off between disinfecting drinking water supplies so as to reduce the probability that end-users are exposed to bacterial, viral and other pathogens, and the degree of exposure to carcinogens resulting from the disinfection process.

To explore this trade-off, we employ the following seemingly unrelated regression framework shown by Eq. (2.5). All the coefficients are analogous to those of Eq. (2.1), however each equation now refers to either microbial health-based violations, denoted by superscript m, or disinfection byproducts,[98] denoted by superscript dbp. In addition

[94]We take these water system-county matches from EPA's 2013 GPRA pivot tables. For large water systems that serve multiple counties, the GPRA table simply picks one county at random and assigns that county to the water system.[95] While this possibly introduces mismatches between water systems and the appropriate income decile, we believe the error introduced to be small since only a few water systems are sufficiently large to span multiple counties.

[96]Accessed in January, 2017 at https://www.census.gov/did/www/saipe/data/statecounty/data/2003.html. We selected 2003 since these data were the closest to 1998 (the date of the policy change) provided in a usable format.

[97]We are unable to match 5,383 systems (about 11 percent of our sample) with county names.

[98]Trihalomethanes are the most relevant disinfection byproduct in our sample. Other disinfection byproducts such as bromate, chlorite, and haloacetic acids were not regulated until after 2001.

to estimating the impact of information disclosure on microbial contaminants, θ^m, and disinfection byproducts, θ^{dbp}, our model also includes model-specific flexible functions of system size, state-by-year fixed-effects, and water system fixed-effects. To estimate this system of equations, we follow the method suggested by Blackwell (2005).[99]

$$\begin{aligned} v_{it}^{m} &= \theta^{m}\left(T_i \times Post_t\right)_{it} + \delta^{m} f(\text{size}) + (st_i \times d_t)^{m} + u_i^{m} + \epsilon_{it}^{m} \\ v_{it}^{dbp} &= \theta^{dbp}\left(T_i \times Post_t\right)_{it} + \delta^{dbp} f(\text{size}) + (st_i \times d_t)^{dbp} + u_i^{dbp} + \epsilon_{it}^{dbp} \end{aligned} \tag{2.5}$$

Evidence of water systems reducing microbial contaminants at the expense of disinfection byproducts depends upon the sign and relative magnitudes of θ^m and θ^{dbp}. If water systems reduce both types of violations in response to the information disclosure requirement equally, we expect $\theta^m \approx \theta^{dbp} < 0$ (relative to their respective baseline levels of violations). If water systems reduce microbial violations at the expense of disinfection byproducts, and do so differentially in response to disclosure, we would expect $\theta^m < 0 < \theta^{dbp}$. If both coefficients indicate reductions in violations, but the magnitude of microbial violations is larger than that of disinfection byproducts (for example, $\theta^m < \theta^{dbp} < 0$), then there may be simultaneously real reductions in water quality violations along with some trade-off of disinfection byproducts for fewer microbial contaminants.[100]

---

[99] In particular, we construct a matrix of regressors for each model, X^m and X^{dbp}, and then place these matrices along the diagonal of a larger matrix as follows: $\mathbb{X} = \begin{bmatrix} X^m & 0 \\ 0 & X^{dbp} \end{bmatrix}$. In Stata, we then regress $\mathbb{v} = [v_{it}^m \quad v_{it}^{dbp}]'$ on $\mathbb{X}$, using the `reghdfe` command (Correia, 2016) (we also use `reghdfe` to implement our other differences-in-differences models with the exception of the simple pooled model.).

[100] Water systems may also differ in their response across management structure. Konar and Cohen (2001) find that firms with higher environmental performance achieve higher stock market values. To the extent that water quality violations reflect a water system's environmental performance, a publicly managed water system would not enjoy a boost in market value if it incurred few water quality violations, but a privately managed firm might. If the incentive to reduce violations induced by the water quality report and the incentive induced by market value identified by Konar and Cohen are substitutes, publicly managed firms may reduce violations more than privately managed firms on account of the water quality report (this is because privately managed firms would not experience any additional incentive to reduce water quality violations on account of the annual water quality report). On the other hand, if the annual water quality report and market incentive are compliments, privately managed water systems may reduce water quality violations more than publicly managed water systems. Exploring the heterogeneity in response across management structure is the subject of ongoing work.

2.5 Results

2.5.1 Main results

Table 2.4: *Regression results illustrating the impact of the publishing requirement.*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{pub} \times Post$	-0.032	-0.032	-0.032	-0.034	-0.033	-0.032
	(0.004)***	(0.004)***	(0.004)***	(0.004)***	(0.004)***	(0.004)***
$Post$	0.008	0.008				
	(0.003)***	(0.003)***				
T_{pub}	0.020					
	(0.003)***					
$f(size)$					-0.007	-0.011
					(0.001)***	(0.003)***
$f(size)^2$						7.9e-05
						(4.0e-05)**
Fixed-effects						
u_i	-	yes	yes	yes	yes	yes
d_t	-	-	yes	-	-	-
$st_i \times d_t$	-	-	-	yes	yes	yes
adj. R^2	0.00	0.23	0.24	0.25	0.25	0.25
Systems	46,900	46,900	46,900	46,900	46,900	46,900
Observations	562,800	562,800	562,800	562,800	562,800	562,800

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$
Water system clustered standard errors (reported in parentheses).
Fixed-effects: system (u_i), year (d_t), and state-by-year ($st_i \times d_t$).
Flexible function of system size: $f(size)^n = Post_t \times size_i^n$
size refers to the water system service population in 100,000s.

We present our main results of the impact of the publication, mailing, and online posting requirement in Table 2.4, Table 2.5, and Table 2.6, respectively. For each table, in column 1 we show results from a simple pooled model where we include indicators for post-1998 years ($Post$) and systems above the disclosure threshold (T_{pub}, T_{mail}, or T_{web}). In column 2 we show results from a model that includes system fixed-effects (thus T cannot be independently identified). In column 3 we add annual fixed-effects (thus $Post$ cannot be independently identified), and column 4 includes state-by-year fixed effects. In columns 5 and 6 we present results from estimating our main model, Eq. (2.1), including linear and quadratic forms of our flexible function of system size.[101]

Table 2.5: *Regression results illustrating the impact of the mailing requirement.*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{mail} \times Post$	-0.046	-0.046	-0.046	-0.043	-0.040	-0.040
	(0.006)***	(0.006)***	(0.006)***	(0.006)***	(0.006)***	(0.007)***
$Post$	-0.002	-0.002				
	(0.002)	(0.002)				
T_{mail}	0.037					
	(0.006)***					
$f(size)$					-0.005	-0.006
					(0.001)***	(0.003)**
$f(size)^2$						1.4e-05
						(3.7e-05)
Fixed-effects						
u_i	-	yes	yes	yes	yes	yes
d_t	-	-	yes	-	-	-
$st_i \times d_t$	-	-	-	yes	yes	yes
adj. R^2	0.00	0.23	0.24	0.25	0.25	0.25
Systems	46,900	46,900	46,900	46,900	46,900	46,900
Observations	562,800	562,800	562,800	562,800	562,800	562,800

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses).

Fixed-effects: system (u_i), year (d_t), and state-by-year ($st_i \times d_t$).

Flexible function of system size: $f(size)^n = Post_t \times size_i^n$

size refers to the water system service population in 100,000s.

The results illustrated in Table 2.4 demonstrate that the publishing requirement reduces health-based water quality violations by about 0.03 violations per system-year. Given our preferred estimate of 0.032 reductions per system-year, and an average violation count of

[101]While not shown, we include up to a fifth-order form of our flexible function of system size and find results robust to the higher order specifications for the publishing and mailing thresholds. Results associated with the online posting threshold are more sensitive to higher orders of systems size, though robust up through a cubic specification.

Table 2.6: *Regression results illustrating the impact of the online posting requirement.*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{web} \times Post$	-0.050	-0.050	-0.050	-0.056	-0.031	-0.012
	(0.014)***	(0.014)***	(0.014)***	(0.014)***	(0.017)*	(0.021)
$Post$	-0.006	-0.006				
	(0.002)***	(0.002)***				
T_{web}	-0.018					
	(0.012)					
$f(size)$					-0.008	-0.016
					(0.002)***	(0.005)***
$f(size)^2$						1.5e-04
						(6.3e-05)**
Fixed-effects						
u_i	-	yes	yes	yes	yes	yes
d_t	-	-	yes	-	-	-
$st_i \times d_t$	-	-	-	yes	yes	yes
adj. R^2	0.00	0.23	0.24	0.25	0.25	0.25
Systems	46,900	46,900	46,900	46,900	46,900	46,900
Observations	562,800	562,800	562,800	562,800	562,800	562,800

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses).

Fixed-effects: system (u_i), year (d_t), and state-by-year ($st_i \times d_t$).

Flexible function of system size: $f(size)^n = Post_t \times size_i^n$

size refers to the water system service population in 100,000s.

0.119 violations for systems subject to the publishing requirement (see Table 2.2), these reductions correspond to a 27 percent decrease in violations on account of the publishing requirement. The estimate of reductions in violations is robust across model specification, illustrating the stability of these estimates.

Table 2.5 shows that requiring water systems to mail their water quality report reduces violations by around 0.04 violations per system-year, respectively, or about a 30 percent reduction in violations in each system-year.[102] The estimate of reductions in violations is

[102] Like the publishing requirement, we divide our preferred estimate of violation reductions, 0.040, by the average violations experienced by those systems serving at least 10,000 customers, which Table 2.2 reports to be 0.133 violations per system-year.

robust across model specification.

There is less evidence that posting water quality reports online reduces water quality violations. While the first four columns in Table 2.6 suggest that the online requirement reduces water quality violations by about 0.05 violations per system year, the effect dissipates with the inclusion of linear and quadratic forms of the flexible function of system size.

Table 2.7: *Regression results illustrating the impact of considering information disclosure thresholds simultaneously.*

	(1)	(2)	(3)	(4)
$T_{pub} \times Post$			-0.028	-0.028
			(0.004)***	(0.004)***
$T_{mail} \times Post$	-0.041	-0.040	-0.025	-0.024
	(0.006)***	(0.007)***	(0.007)***	(0.007)***
$T_{web} \times Post$	-0.019	-0.003	-0.017	-0.006
	(0.015)	(0.021)	(0.015)	(0.021)
$f(size)$		-0.006		-0.003
		(0.004)		(0.004)
$f(size)^2$		8.0e-06		-2.0e-05
		(4.8e-05)		(4.6e-05)
adj. R^2	0.25	0.25	0.25	0.25
Systems	46,900	46,900	46,900	46,900
Observations	562,800	562,800	562,800	562,800

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$
Water system clustered standard errors (reported in parentheses).
All models include water system and state-by-year fixed effects.
Flexible function of system size: $f(size)^n = Post_t \times size_i^n$
size refers to the water system service population in 100,000s.

In Table 2.7, we consider the effect of including multiple disclosure thresholds simultaneously. In column 1 and column 2, we explore the additional impact of requiring water systems to post their report online. The results in row 3 indicate that requiring water systems that already mail their water quality reports to post the report online does not have a significant additional marginal impact on violations. In column 3 and column 4, we estimate the effect of each disclosure threshold simultaneously. The results confirm the results discussed above; the publishing and mailing requirement appear to reduce water quality violations, but the additional requirement to post water quality reports online has no noticeable impact.

2.5.2 The persistence of the information disclosure

Table 2.8: *Regression results illustrating the persistence of each disclosure method.*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{pub} \times Post$	-0.027	-0.026				
	(0.005)***	(0.005)***				
$T_{pub} \times yrsPost$	-0.004	-0.004				
	(0.002)*	(0.002)*				
$T_{mail} \times Post$			-0.043	-0.040		
			(0.008)***	(0.008)***		
$T_{mail} \times yrsPost$			1.4e-04	1.4e-04		
			(3.6e-03)	(3.6e-03)		
$T_{web} \times Post$					-0.051	-0.007
					(0.019)***	(0.026)
$T_{web} \times yrsPost$					-0.003	-0.003
					(0.006)	(0.006)
$f(size)$		-0.011		-0.006		-0.016
		(0.003)***		(0.003)**		(0.005)***
$f(size)^2$		7.9e-05		1.4e-05		1.5e-04
		(4.0e-05)**		(3.7e-05)		(6.3e-05)**
Threshold	501	501	10k	10k	100k	100k
adj. R^2	0.25	0.25	0.25	0.25	0.25	0.25
Systems	46,900	46,900	46,900	46,900	46,900	46,900
Observations	562,800	562,800	562,800	562,800	562,800	562,800

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses). Water system and state-by-year fixed effects.

Threshold: systems serving $\geq$ 501, 10k, and 100k customers must, respectively, publish in a local venue, mail, or mail and post online the water quality report.

Flexible function of system size: $f(size)^n = Post_t \times size_i^n$

size refers to the water system service population in 100,000s.

We present results of estimating Eq. (2.3) for each disclosure threshold in Table 2.8. In particular, we illustrate estimates for models including system and state-by-year fixed-effects in columns 1, 3, and 5, and then our main specification including the flexible function of system size in columns 2, 4, and 6. Considering first the publishing requirement, the estimate on the interaction between our treatment indicator and the annual linear time trend describing the number of years after 1998 ($yrsPost$) is negative and significant at the 10 percent level. Thus there appears to be a small, but imprecisely estimated increase in the effect of the publication disclosure policy over time. For both the mailing and online posting requirements, the estimate on the interaction between our treatment indicator and the annual linear time trend describing the number of years after 1998 ($yrsPost$) is indistinguishable from zero, implying that the effects of the mailing and online posting requirement have been sustained over time.[103]

Figure 2.3 illustrates results from estimating Eq. (2.4), plotting the estimates and 95 percent confidence intervals of the impact of each disclosure requirement in 1998-2001. Confirming the results from Table 2.8, Figure 2.3 illustrates slightly increasing effects due to the publishing requirement and stable impacts over time for the mailing and online posting requirements.

Taken together, these results imply that the primary impact of the disclosure policy takes place directly after the policy came into effect, with the effect then remaining relatively unchanged through the end of our sample. Johnson (2003) finds that the annual water quality reports have little impact on customers' perceptions of water quality or customers' perceptions of the water system. If a concern over customer complaints drives the reductions we observe in water quality violations, as suggested by Bennear and Olmstead (2008), then the stability of reductions over time may suggest that water systems miscalculate customers' propensity to complain, or that water systems overestimate customers' ability to comprehend the water quality report. More broadly, few empirical analyses study the persistent effects

[103]These results stand in contrast to Fung *et al.* (2007), who postulate that this information disclosure policy will have less of an effect over time on account of few disclosure-associated benefits realized by water systems and the absence of citizen groups to give voice to concerns over water quality.

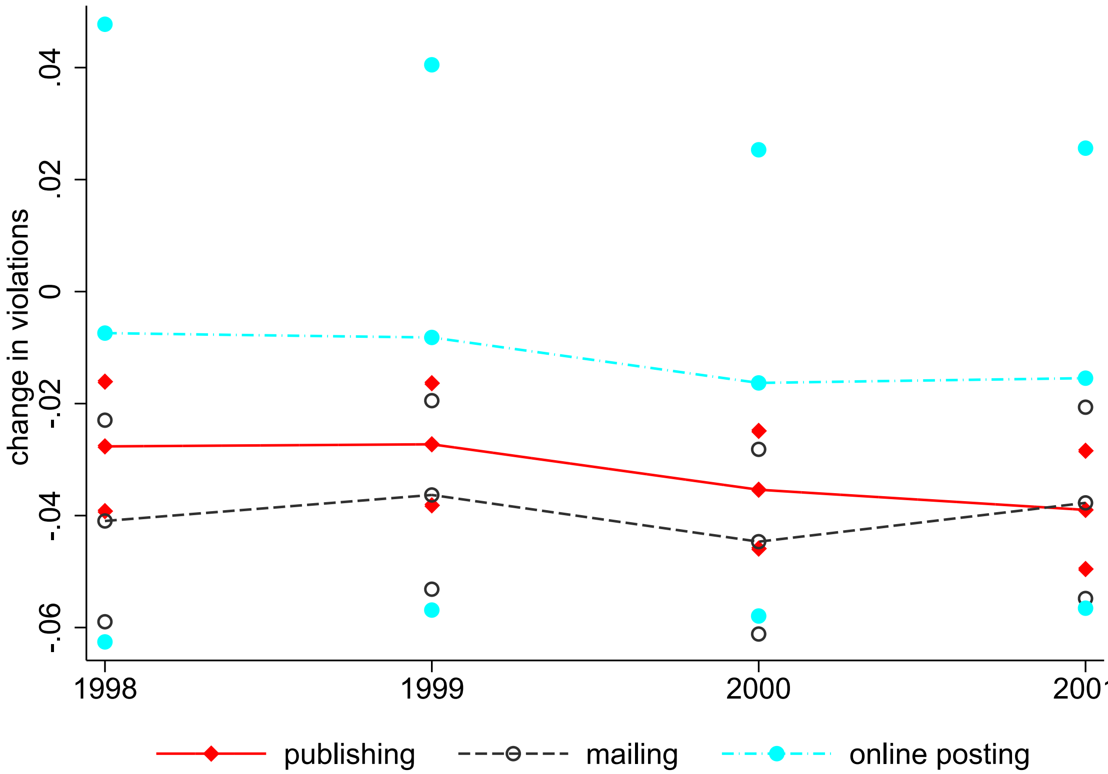

Figure 2.3: *Annual impact of the publishing (diamond), mailing (hollow circle), and online posting (solid circle) disclosure requirements. 95 percent confidence intervals are indicated with the same symbol as the respective point estimate. We derive point estimates and confidence intervals from estimating estimating Eq. (2.4) for each disclosure requirement.*

of information disclosure policies (Dranove and Jin, 2010). By showing that impacts of information disclosure policies can be sustained over time, our results begin to build our understanding of the long-term ramifications of such policies.

2.5.3 Response heterogeneity

We illustrate results of the effect of information disclosure as a function of customer income in Figures 2.4 through 2.6. In particular, we plot point estimates of the main effect ($T \times Post$) and associated 95 percent confidence intervals from estimating Eq. (2.1) for each customer service base income decile. In each plot, we further normalize point estimates by the average number of health-based violations in all system-years for the respective income decile. Normalizing point estimates enables us to compare estimates across deciles. We match water systems with income by first associating water systems with the county served by the water system.[104] We then group counties into deciles based on each county's 2003 Census reported median income.

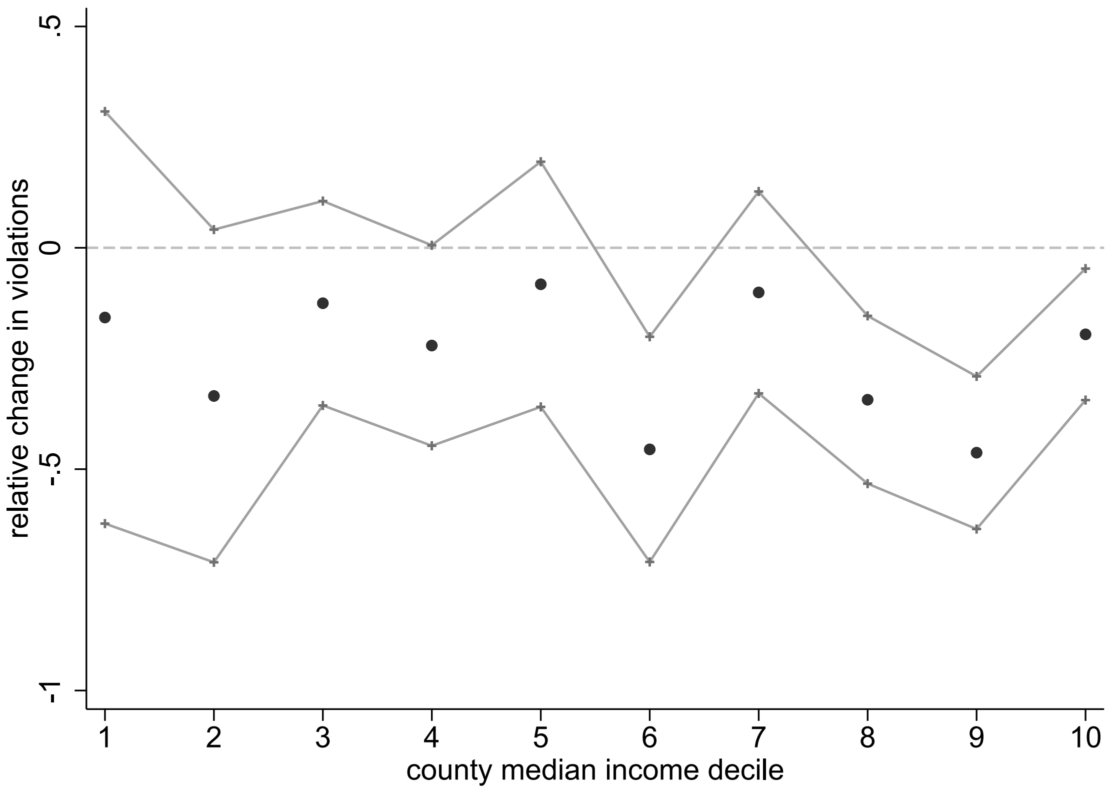

Figure 2.4: *Heterogeneity in response to the publishing requirement by income decile. We illustrate point estimates and 95 percent confidence intervals from estimating Eq. (2.1) for each income decile defined by the median income of the county within which the water system resides. We further normalize point estimates by the average number of health-based violations in all system-years for the respective income decile.*

While there seems to be little heterogeneity in response due to the online posting requirement, there does appear to be a stronger response from the publishing and mailing requirements among water systems serving higher income counties.[105] Regarding Figure

[104]See section 2.4, in particular footnotes 94 and 96.

[105]Our results are therefore generally in agreement with those of Powers *et al.* (2011), who find that an information disclosure policy drove comparatively higher responses from Indian pulp and paper firms located

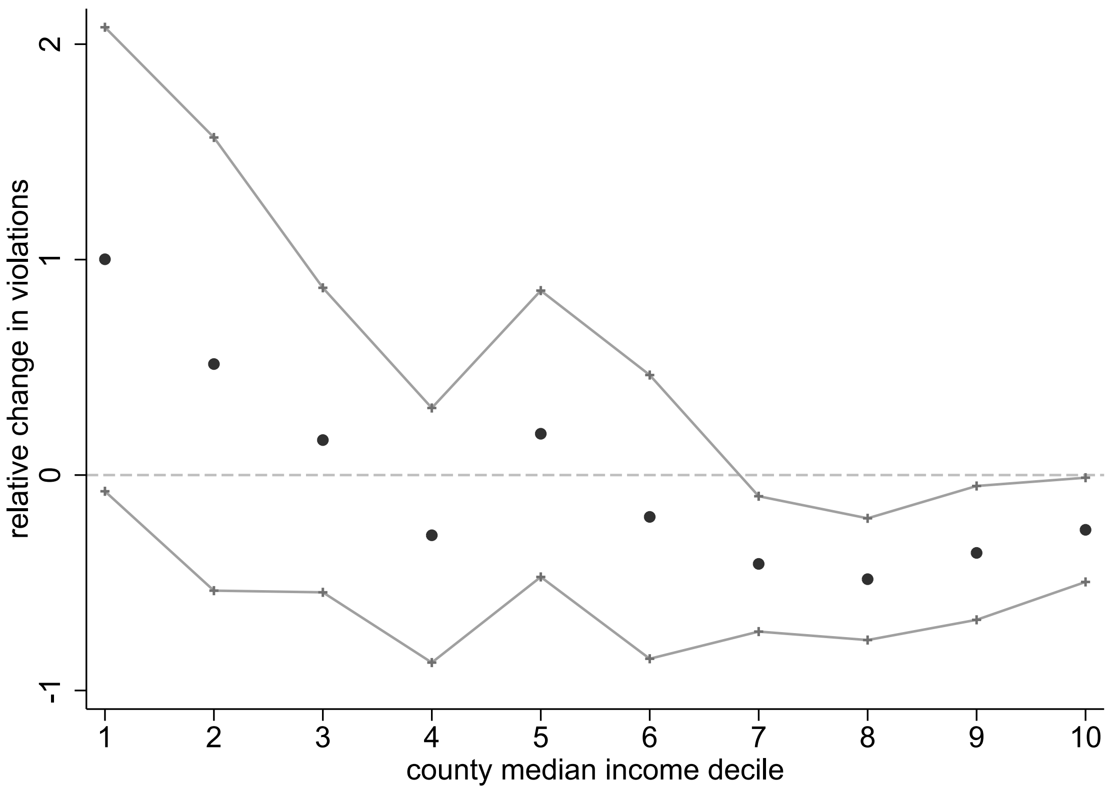

Figure 2.5: *Heterogeneity in response to the mailing requirement by income decile. We illustrate point estimates and 95 percent confidence intervals from estimating Eq. (2.1) for each income decile defined by the median income of the county within which the water system resides. We further normalize point estimates by the average number of health-based violations in all system-years for the respective income decile.*

2.4, which shows the publishing requirement, all estimates associated with the lower income decile counties, deciles 1 through 5, cannot be distinguished from zero. However, the effects in the higher income deciles, especially in deciles 6, 8, 9, and 10, are negative and significant. The effect is even more pronounced for the mailing requirement, shown in Figure 2.5. The highest income deciles (7 through 10) generate the strongest impacts due to the mailing requirement; the impacts in the other deciles are estimated imprecisely.[106] And

in higher-income areas. To the extent that income is correlated with education levels, our results also provide suggestive evidence in support of Shimshack *et al.* (2007)'s findings that education and news readership levels drive responsiveness to information disclosure.

[106]This could in part be driven by the fact that our panel contains fewer observations in the lower income deciles.

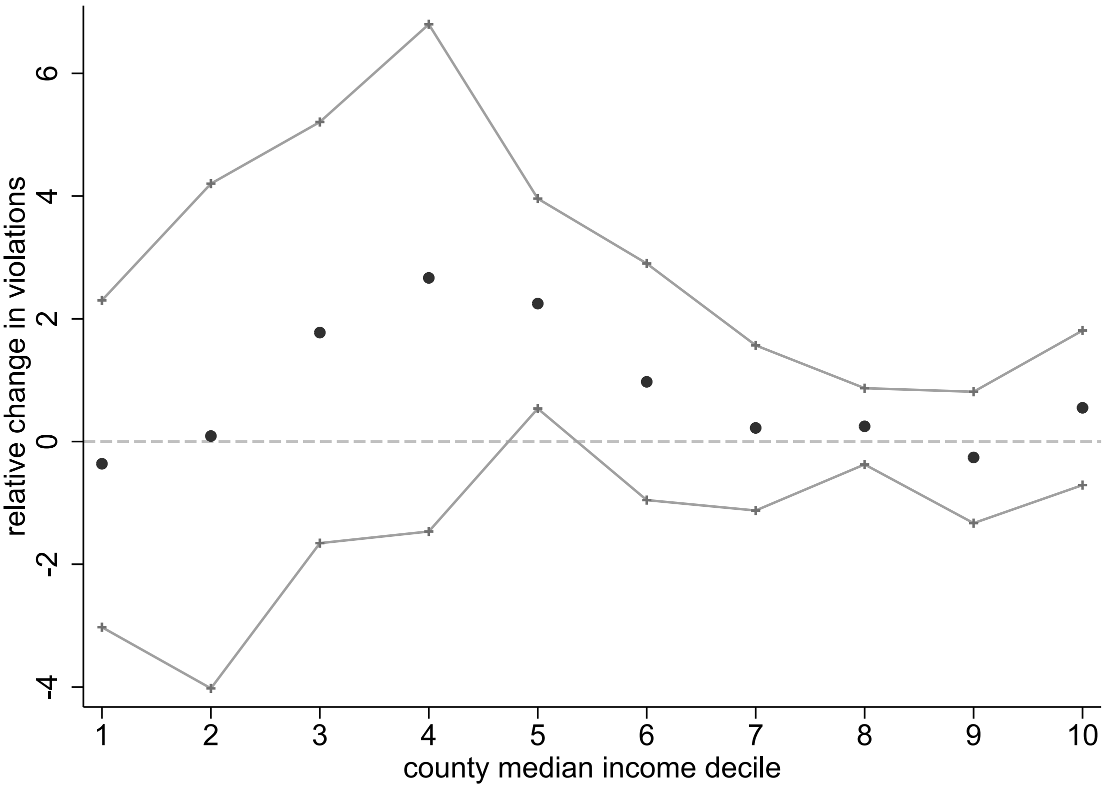

Figure 2.6: *Heterogeneity in response to the online-posting requirement by income decile. We illustrate point estimates and 95 percent confidence intervals from estimating Eq. (2.1) for each income decile defined by the median income of the county within which the water system resides. We further normalize point estimates by the average number of health-based violations in all system-years for the respective income decile.*

finally, Figure 2.6 illustrates that there is no relationship between income and water quality reductions due to the online posting requirement.[107]

In Table 2.9, we illustrate results of running our trade-off analysis, exploring whether water systems decrease microbial violations at the expense of increasing disinfection byproducts. For each threshold, we show regression results from estimating our system of equations, Eq. (2.5), with and without the flexible function of system size. In each specification, we show that microbial violations fall, but that levels of disinfection byproducts remain essentially unchanged. Our results provide at least suggestive evidence that microbial violations and

[107]We also estimate Eq. (2.1) for each decile including only system and state-by-year fixed-effects (i.e. we exclude the flexible function of system size). The resulting figures lead to largely similar conclusions, though the results for the online posting requirement do show precisely estimated negative impacts in deciles 8 and 9.

Table 2.9: *Regression results illustrating how water systems trade off reductions in disinfection byproducts (primarily trihalomethanes) with microbial contaminants.*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{pub} \times Post$ (microbial)	-0.029	-0.028				
	(0.004)***	(0.004)***				
$T_{pub} \times Post$ (DBP)	1.6e-04	1.7e-04				
	(3.6e-04)	(3.5e-04)				
$T_{mail} \times Post$ (microbial)			-0.041	-0.039		
			(0.005)***	(0.006)***		
$T_{mail} \times Post$ (DBP)			1.6e-03	1.9e-03		
			(1.7e-03)	(1.9e-03)		
$T_{web} \times Post$ (microbial)					-0.050	-0.013
					(0.013)***	(0.019)
$T_{web} \times Post$ (DBP)					-8.9e-04	-1.6e-03
					(2.2e-03)	(3.2e-03)
Threshold	501	501	10k	10k	100k	100k
$f(size)$	-	yes	-	yes	-	yes
adj. R^2	0.26	0.26	0.26	0.26	0.26	0.26
Systems	93,800	93,800	93,800	93,800	93,800	93,800
Observations	1,125,600	1,125,600	1,125,600	1,125,600	1,125,600	1,125,600

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses). Water system and state-by-year fixed effects.

Threshold: systems serving $\geq$ 501, 10k, and 100k customers must, respectively, publish in a local venue, mail, or mail and post online the water quality report.

DBP: disinfection byproducts, primarily trihalomethanes.

$f(size)$: 2nd order flexible function of system size: $\sum_{j=1}^{2}(Post_t \times size_i)^j$

size refers to the water system service population in 100,000s.

disinfection byproduct violations are not direct substitutes. That is, reducing microbial violations need not necessarily come at the expense of increasing disinfection byproducts. But it is also possible that water systems employed alternative disinfection processes that created other, at the time unregulated disinfection byproducts, such as chlorite and bromate.[108] While Bennear *et al.* (2009) present evidence of strategic behavior in the context of

[108]These disinfection byproducts were regulated beginning in 2002.

avoiding water quality violations, further research is needed to determine whether strategic avoidance of certain disinfection processes helps to explain our results.

2.5.4 Robustness checks

Because the method of information disclosure depends upon the service population of the water system, systems that gain or lose customers over time may move into different information disclosure regimes. But we only observe the service population in the first quarter of 2014, so it remains possible that we assign water systems to treatments inappropriately. Provided systems do not undergo substantial population growth, however, this potential misassignment of treatment will only affect water systems around our respective service population thresholds. To test the sensitivity to possible treatment misassignment, we run models that exclude systems with service populations (as of 2014) that are within 10 percent, 20 percent, and 30 percent of the mailing threshold service population cutoff. The results of this analysis, illustrated in Table 2.10, show that our estimates are insensitive to dropping systems around the service population threshold. We therefore conclude that any misassignment of treatment does not substantially impact our results.[109]

On the subject of service populations, the EPA defines public water systems as those systems serving at least 25 customers. Many water systems in our sample, however, serve less than 25 customers. To test the sensitivity of our results, we estimate models the exclude all systems that serve fewer than 25 customers (as reported in 2014). While not shown, we find that our results are insensitive to excluding these small systems.

The validity of our identification strategy relies on treated and non-treated water systems exhibiting parallel trends in average annual violations. While impossible to check for the full range of data, we can check for parallel trends in the pre-period. Specifically, we impose a proxy policy in 1994, and run our model (with and without the flexible

[109]We also test for the impact that treatment misassignment may have on our estimates of the publishing and online posting requirements. Though not shown, we find that estimates of the publishing and online posting impacts are insensitive to dropping systems with service populations within 10, 20, and 30 percent of the respective service population cutoffs.

Table 2.10: *Regression results illustrating the effect of dropping systems within 10, 20, and 30 percent of the mailing disclosure threshold.*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$T_{mail} \times Post$	-0.043	-0.045	-0.045	-0.048	-0.040	-0.042	-0.042	-0.045
	(0.006)***	(0.006)***	(0.006)***	(0.007)***	(0.007)***	(0.007)***	(0.007)***	(0.007)***
$f(size)$					-0.006	-0.005	-0.005	-0.004
					(0.003)**	(0.003)*	(0.003)*	(0.003)
$f(size)^2$					1.4e-05	4.8e-06	3.4e-06	-7.0e-06
					(3.7e-05)	(3.6e-05)	(3.7e-05)	(3.6e-05)
Sample	full	10 %	20 %	30 %	full	10 %	20 %	30 %
adj. R^2	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
Systems	46,900	46,244	45,520	44,833	46,900	46,244	45,520	44,833
Observations	562,800	554,928	546,240	537,996	562,800	554,928	546,240	537,996

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses). Water system and state-by-year fixed effects.

Flexible function of system size: $f(size)^n = Post_t \times size_i^n$, where *size* refers to the water system service population in 100,000s.

Sample: systems are excluded that serve within X % of the disclosure threshold; full refers to all systems.

Table 2.11: *Regression results illustrating the effect of a proxy policy applied in 1994 on all systems in the pre-period (i.e. 1990 through 1997).*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{pub} \times Proxy$	0.018	0.021				
	(0.006)***	(0.006)***				
$T_{mail} \times Proxy$			-0.031	-0.024		
			(0.009)***	(0.009)***		
$T_{web} \times Proxy$					-0.061	-0.021
					(0.018)***	(0.025)
$f(size)$		-0.024		-0.013		-0.016
		(0.004)***		(0.004)***		(0.006)***
$f(size)^2$		3.7e-04		2.1e-04		2.7e-04
		(5.5e-05)***		(5.7e-05)***		(7.1e-05)***
Threshold	501	501	10k	10k	100k	100k
adj. R^2	0.26	0.26	0.26	0.26	0.26	0.26
Systems	46,900	46,900	46,900	46,900	46,900	46,900
Observations	375,200	375,200	375,200	375,200	375,200	375,200

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses). Water system and state-by-year fixed effects.

Threshold: systems serving $\geq$ 501, 10k, and 100k customers must, respectively, publish in a local venue, mail, or mail and post online the water quality report.

Flexible function of system size: $f(size)^n = Post_t \times size_i^n$

size refers to the water system service population in 100,000s.

Proxy policy imposed in 1994.

function of system size) on our pre-period sample. Statistically insignificant coefficient estimates on the interaction between the treatment indicator and the proxy policy indicator would provide evidence for parallel trends in the pre-period data. Table 2.11 shows statistically significant coefficient estimates on the treatment-proxy interaction variable for each disclosure requirement.[110]

These results seriously temper the confidence we place in the preceding results. Another way to support our conclusions would be to employ an alternative empirical strategy, and

[110]The one exception is for the model that explores the impact of the online requirement and that includes the flexible function of system size.

towards that end we present results from a linear and quadratic parametric regression discontinuity design in Appendix B.1. Though not without flaws, our RD results demonstrate water quality reductions due to the publishing requirement, but provide less evidence of any impact due to mailing or posting the water quality report online. Thus, while not necessarily conclusive evidence by themselves, taken together the results from the differences-in-differences model and regression discontinuity approach present reinforcing evidence that requiring water systems to publish their water quality report reduces water quality violations.

2.6 Conclusion

The 1996 Amendments to the Safe Drinking Water Act require community water systems to disclose annual water quality reports to their customers. The method of distribution depends upon the water system's service population. Relative to water systems not subject to any direct information disclosure requirements, we have shown that water systems required to publish reports in local newspapers reduce water quality violations by about 27 percent as a direct result of being required to publish the report. In the Appendix B.1, we present evidence from a regression discontinuity design consistent with these results. We also show that the publication requirement is stronger for systems serving higher median income counties. Similar to the publishing requirement, water systems required to mail water quality reports reduce water quality violations by 30 percent due to the mailing requirement. This response appears to be particularly strong for systems serving higher median income counties. However, our regression discontinuity design shows no impact due to mailing. Thus while the mailing requirement appears to be as salient as the publication requirement in differences-in-differences models, we have less confidence in the robustness of our conclusions regarding the mailing requirement. Finally, we show that the requirement to post a water quality report online does not cause any reductions in violations. This is likely because water systems required to post their report online must already mail their reports to customers, and are thus not providing customers with any new information by posting the report online.

Despite there being a direct trade-off between reducing microbial contaminants and increased levels of disinfection byproducts, we have shown that water systems are able to reduce microbial contaminant violations without experiencing increases in disinfection byproduct violations. But among all possible disinfection byproducts, the only ones regulated during our time frame were trihalomethanes. Systems could have used other disinfection methods producing other disinfection byproducts such as haloacetic acids, chlorite, and bromate, which were not regulated until after 2001. We plan to investigate in the future whether non-regulated disinfection byproducts increased coincident with reductions in microbial contaminants.

Finally, we have shown that at least in the three years following the imposition of the information disclosure policy, the reductions in water quality violations appear to be maintained over time. From the overall perspective of information disclosure programs, sustained effects are of great import. It would be desirable, however, to be able to assess the effects in the longer term.

Chapter 3

Does Desert Landscape Encourage More Desert Landscape? Evidence of Peer-Effects in the Las Vegas Cash-for-Grass Rebate Program

3.1 Introduction

In response to growing concerns over water scarcity, water utility managers are increasingly turning to demand-side-management programs to encourage resource conservation. In large part due to the substantial share of water required by outdoor landscaping, many utilities have begun to administer rebate programs that subsidize their customers to replace lawns with less water intensive desert landscape. While the level of the subsidy, potential savings, and conservation preferences are all obvious drivers of customers' decision to participate in such programs, another potentially important driver is the behavior of surrounding individuals. In other words, individuals' preferences for desert landscape may be affected by their peers' decisions to convert (Leibenstein, 1950).

In this analysis, I explore the extent to which peer effects drive participation in the

Southern Nevada Water Authority's Water Smart Landscapes program. Perhaps one of the most well known "Cash-for-Grass" rebate programs, the Water Smart Landscapes program currently offers a $2.00 per square foot subsidy to customers who replace their lawns with desert landscape. I focus my analysis on single-family residents served by the Las Vegas Valley Water District, the largest water district in the Las Vegas valley. The water authority provided me with data concerning program participation from the program's inception in 1996 to about mid-way through 2014, as well as information regarding post card mailings and door hanger advertising efforts administered to selected residents to encourage program participation.

In a closely related context, Bollinger and Gillingham (2012), henceforth referred to as BG, investigate the presence of peer effects in the adoption of solar panels under a California rebate program aimed at promoting solar installations. The authors propose a first-differences empirical model that requires the researcher to observe the timing of the decision to install solar panels as well as the installation. Importantly, their strategy overcomes many of the empirical challenges associated with identifying peer effects (Bollinger and Gillingham, 2012; Narayanan and Nair, 2013; Jackson *et al.*, 2017). As my context is nearly identical to theirs, I adopt BG's empirical framework and several of their analyses in my analysis of the Las Vegas Cash-for-Grass program.

In order to compare my results with those of BG, I define a peer network by zip code. I also define a peer network at a more refined level. In particular, I group single-family homes within the Las Vegas Valley Water District into deciles based on the Clark County Assessors Office's assessed property value (as opposed to sale price). I then define my peer network as a zip code interacted with home value deciles. Within the average zip code or average zip code-decile, conversions to desert landscape take place infrequently. The average number of conversions to desert landscape as a share of eligible converters on any given day is about 20×10^{-6}.

Like BG, I define the peer effect as the change in an unconverted property's application probability due to a new conversion to desert landscape in the same peer network. For a zip code, I find the peer effect to be 2.6×10^{-6}, or about 14 percent of the baseline application probability in a zip code. For a zip code-decile, I find the peer effect to be 12.3×10^{-6}, or about 47 percent of the baseline application probability in a zip code-decile. The larger increase in the conversion probability within a zip code-decile versus a zip code suggests that the peer effect may at least in part operate through homeowners' desire to maintain the competitiveness of their homes to homes in a similar housing market. Similar to what BG show in solar panel installations, I also show that my estimated peer effects increase over time as the number of conversions increases. In contrast to BG's findings, however, I find that the number of conversions rather than the total area of conversions drive my results.[111] As an extension to the work of BG, I show that my estimated peer effect is at least an order of magnitude larger than any of the five marketing campaigns administered by the water authority.

Peer effects encompass a broad range of behavior and mechanisms. Researchers have explored the presence of peer effects in education (Fafchamps and Mo, 2017), entrepreneurial activity (Field *et al.*, 2014), pro-environmental behavior (Narayanan and Nair, 2013) and labor market outcomes (Cornelissen *et al.*, 2017). Regarding mechanisms, some analyses focus specifically on the role of social norms (Arimura *et al.*, 2016; Dolan and Metcalfe, 2015), the impact of social comparison on behavior (Brent *et al.*, 2015), and social pressure (Bursztyn and Jensen, 2016). While my analysis abstracts from the mechanisms driving the peer effect, in the context of pro-environmental behavior my analysis is the first of which I am aware to compare the strength of a peer effect to the impact of alternative interventions designed to encourage the behavior influenced by the peer effect.

I organize the rest of the paper as follows. In section 3.2, I describe the Cash-for-Grass rebate program and summarize program data. In section 3.3, I summarize BG's empirical strategy adopted to my context. I discuss results in section 3.4, and conclude in section 3.5.

[111] To be precise, BG are unable to determine whether size of installation or the number of installations have a larger impact on their estimated peer effect.

3.2 Data

The Southern Nevada Water Authority's (SNWA) Water Smart Landscapes, or "Cash-for-Grass" program offers Las Vegas area water customers a rebate in exchange for replacing grass lawns with desert landscape. The water authority initiated a pilot study in 1996[112] and began to offer the rebate to its entire customer base in 1998. I obtained participation and marketing data from the SNWA from the program's inception through mid-June 2014. I focus on single-family customers within the Las Vegas Valley Water District (LVVWD), the largest water district in the Las Vegas valley.[113]

Adopting BG's empirical strategy requires that I proxy for when customers decide to participate in the program and the date of conversion to desert landscape. Since the rebate application date is the first time I observe a customer express interest in the rebate program, I proxy for the day a customer decides to convert his or her lawn to desert landscape with the application date. After completing a conversion, a customer schedules a post-conversion site visit, during which a water authority staff member verifies that all program requirements have been met. After a successful post-conversion site visit, the customer receives the rebate and the water authority designates the customer as enrolled. Since the post-conversion site visit most closely approximates the date a customer completes a conversion,[114] I proxy for the conversion date with the post-conversion site visit date. For those conversions where I do not observe the post-conversion site visit date, I proxy for the date of conversion with the enrollment date.[115]

The validity of BG's empirical strategy applied to my context hinges upon sufficient time between decisions to apply for the rebate and the conversion date. In Table 3.1, I show that

[112]Sovocool *et al.* (2006) and Deoreo *et al.* (2000) provide analyses of early program results.

[113]Thus I will exclude any "peer" effect that operates on single-family residents via their exposure to other types of properties that convert to desert landscape, for example golf courses or multi-family residences.

[114]pers. comm. water authority conservation staff, February 2016. Staff members explained that while sometimes participants must correct for a minor failure to comply with all program requirements, the post-conversion site visit generally marks the completion of the bulk of most conversions.

[115]I also use the enrollment date for any conversion where the post-conversion site visit is listed after the enrollment date, or for any conversion where the application date is recorded after the post-conversion site visit.

Table 3.1: *Summary statistics for the span of days between application and conversion (as proxied by the post-conversion site visit), between application and enrollment, and between conversion and enrollment.*

	Mean	Stnd. dev.	Min	Max	N
Application to conversion (days)	116.70	84.76	3	607	22,738
Application to enrollment (days)	149.79	90.86	9	800	22,738
Conversion to enrollment (days)	33.09	26.61	0	765	22,738

Values calculated after proxying for conversion date with post-conversion site visit.

the average property undertaking a conversion applies for the rebate a little less than four months (117 days) before completing the conversion. The shortest time between application and conversion is 3 days, however, as Figure 3.1 illustrates, most converting properties for which I observe applications applied for the rebate at least one week if not more before completing the conversion to desert landscape. Finally, I also show in Table 3.1 that it takes an average of one additional month after conversion for properties to become enrolled into the program. As I will show below in my results, using the enrollment date as a proxy for the conversion date severely biases my estimates of the peer effect.

Figure 3.2 illustrates the cumulative number of conversions and cumulative number of applications over time. Though I observe conversions beginning in 1996, I do not observe application dates until July 2003. Because application dates are critical to the empirical strategy, I begin my analysis when I first observe applications. As demonstrated by the figure, the number of conversions increases steadily over time.[116] The last conversion takes place on May 30, 2014. Cumulative applications also track steadily upwards, but the differences in the slope of the cumulative applications and cumulative conversions demonstrate that I do not observe application dates for all conversions. I will thus understate the total number of applications in my analysis.[117] The last observed application occurs on May 3, 2014. In

[116]The marked increase in conversions in 2004 corresponds to an increase in the rebate and water rates.

[117]About 16 percent of post-July 2003 conversions do not have applications linked to them.

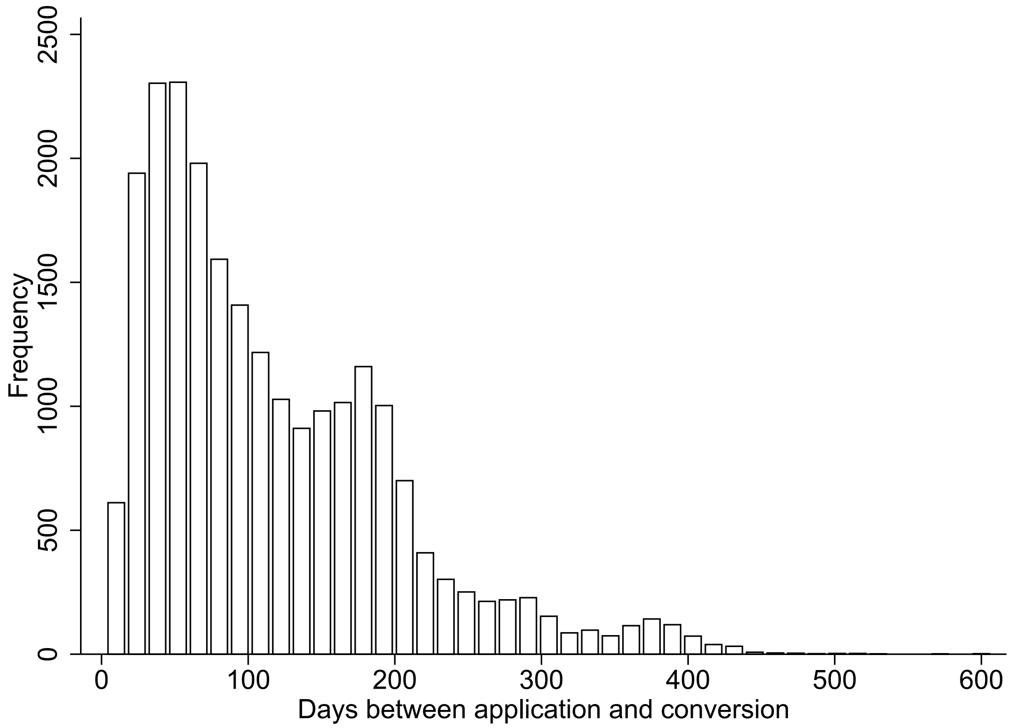

Figure 3.1: *Frequency plot of the distribution of the number of days between the date of application and the date of conversion (as proxied for by the post-conversion site visit).*

total, I observe 22,738 applications out of a total of 28,944 conversions.[118]

Between April 2007 and April 2008, the SNWA ran several marketing campaigns aimed at increasing program participation. Figure 3.2 and Table 3.2 summarize these marketing efforts. In April 2007, the water authority sent out about 15,000 post cards. In August and September of 2007, the water authority distributed around 2,100 door hangers. It sent a second series of post cards in October 2007, and a third series of post cards in April 2008.[119]

[118]I actually observe 22,741 conversions, however I drop these additional three conversions since the year of application takes place before year the assessor indicated the property was built.

[119]The Cash-for-Grass program participation data I received from the SNWA is identified by the Clark County Assessor's Office parcel identifier. The SNWA marketing data is identified by address. Using Clark County Assessor data, which matches parcel id to location address, I endeavored to match SNWA marketing addresses to assessor location addresses, so that I could then link my marketing data addresses with properties served by the LVVWD. While I believe the matching was overall very good, I was unable to match all of the marketing data addresses with assessor addresses. One primary reason is that while the marketing literature was sent to residents within all water districts under SNWA purview, I confine my analysis to the largest water district, the LVVWD. Another reason for imperfectly matching to SNWA marketing addresses is that some marketing

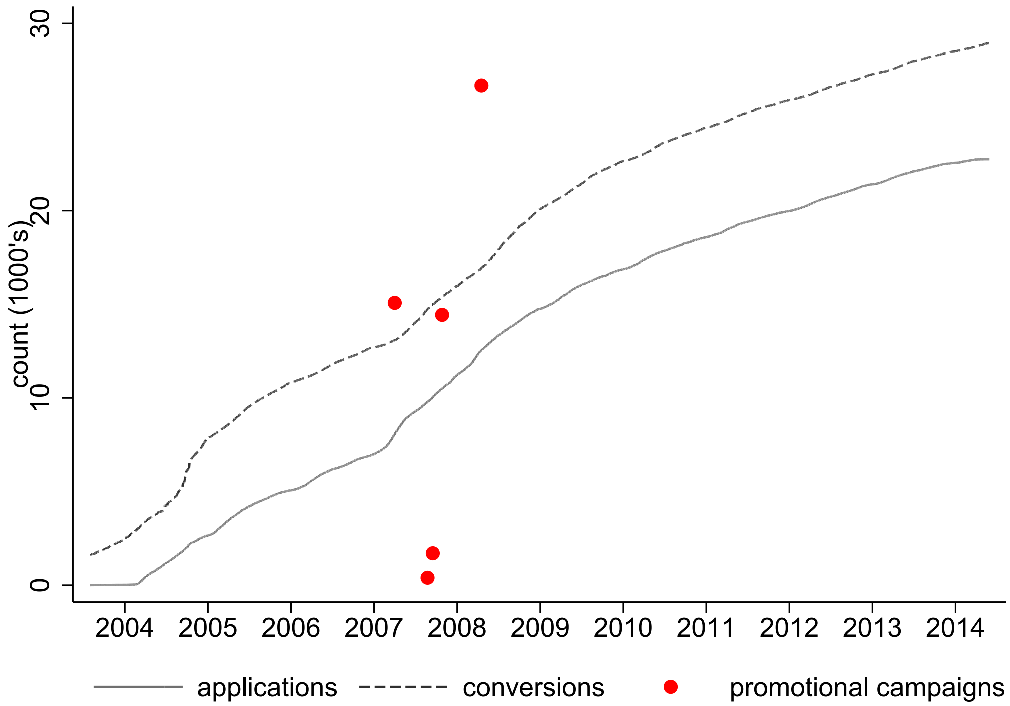

Figure 3.2: *Cumulative applications, cumulative conversions, and total number of promotional campaign materials sent. The first conversion was recorded on 5/29/1996, and cumulative conversions increase nearly linearly until mid-2003. I do not observe applications before 7/31/2003. After this date, most conversions include the associated date of application. I observe five promotional campaigns between April 2007 and April 2008; post cards in April 2007, door hangers in August and September of 2007, and two final postcard campaigns in October 2007 and April 2008. The figure illustrates the number of these materials sent to LVVWD customers.*

Table 3.2: *Ratio of first-floor square footage to property size.*

	Mean	Stnd. dev.	Min	Max	N
All LVVWD properties	0.23	0.068	0.00	0.91	314,307
Receiving April 2007 post card	0.18	0.057	0.02	0.85	15,072
Receiving Aug. 2007 door hanger	0.22	0.068	0.06	0.42	398
Receiving Sep. 2007 door hanger	0.22	0.065	0.05	0.38	1,704
Receiving Oct. 2007 post card	0.18	0.056	0.02	0.85	14,432
Receiving April 2008 post card	0.18	0.060	0.01	0.85	26,668

LVVWD: Las Vegas Valley Water District.

Since the water authority is trying to incentivize lawn removal, one would expect the water authority to send marketing materials to those properties most likely to achieve substantial savings by removing grass. To test this intuition, I first match my list of properties served by the LVVWD to Clark County Assessor property data, and then calculate the ratio of the home footprint (typically just the first-floor square footage) to the property lot size for all LVVWD properties. Small ratios indicate larger potentially landscape-able areas. Table 3.2 shows that on average, the water authority sent post cards to properties with smaller home footprint ratios, suggesting the water authority endeavored to target those homes with the greatest potential for the largest lawns. And while the mean footprint ratio for the two door hanger campaigns approximately equals the average footprint ratio for the entire LVVWD, the maximum footprint ratio is a little less than the maximum footprint ratio I observe among those customers participating in the program. This further suggests the water authority targeted properties with comparatively large lawns.

Using the data described above, I create two panels at the peer network-day following the structure of the panel created by BG. My dates of observation range from July 2003 to June 2014. For each peer network-day, I construct variables describing cumulative conversions, cumulative area converted, cumulative number of marketing materials for each marketing campaign, and finally the ratio of the number of applications to the remaining number of eligible properties that can convert.[120] BG define this last variable as the probability of adoption, or in my case, application probability. Like BG and Narayanan and Nair (2013), I

---

addresses are out of state. There is also the possibility that my matching strategy did not capture all addresses within the LVVWD, either due to address typos, different abbreviations, or other contingencies my address matching strategy does not account for. The numbers illustrated in Table 3.2 should therefore be taken as lower bounds on the number of materials sent out during each marketing campaign.

[120] As properties convert to desert landscape, they are removed from the count of eligible properties. I define eligible properties as the number of properties in a peer network in each year that would have sufficient landscape-able area to consider applying for the rebate program. Using Clark County Assessor data on years in which homes were built, as well as a list of properties served by the LVVWD provided by the SNWA, I develop the number of properties in a peer network in each year, assuming that a property enters my sample on the first of the year. Next, I calculate the ratio of the home footprint to the property lot size (using assessor data). I drop any non-participating properties with footprint ratios greater than the maximum footprint ratio I observe in properties that participated in the rebate program, reasoning that such non-participating properties would not have sufficient landscape-able area to consider applying for the rebate program. Dropping these "infeasible" properties removes less than one half of one percent of properties within the LVVWD.

Table 3.3: *Summary statistics for 2017 Clark County Assessor property values categorized by decile. Considers single-family properties located in the LVVWD service district.*

	Mean	Stnd. dev.	Min	Max	N
1^{st} decile	70,045	12,834	0	89,449	31,453
2^{nd} decile	106,596	9,511	89,451	122,380	31,453
3^{rd} decile	136,079	7,269	122,383	147,820	31,458
4^{th} decile	157,586	5,373	147,823	166,391	31,448
5^{th} decile	175,437	5,241	166,394	184,380	31,455
6^{th} decile	194,079	5,798	184,383	204,157	31,453
7^{th} decile	215,645	7,084	204,160	228,717	31,447
8^{th} decile	246,707	11,246	228,723	267,837	31,453
9^{th} decile	299,449	20,366	267,840	339,797	31,452
10^{th} decile	553,559	495,287	339,803	25,645,306	31,452

define a peer network at the zip code level. But individuals may also tend to interact with those who share similar socioeconomic status, and may be as much or more influenced by the conversion decisions of friends and neighbors who own a home of similar value than the conversion decisions of those whose homes fall well outside the individuals' housing market.[121] Thus I also define the peer network by a zip code interacted with categories of home value. I define home value categories by binning properties within the LVVWD into deciles based on their 2017 assessed value calculated from Clark County Assessor data.[122] Table 3.3 describes the resulting home value deciles for the LVVWD, which range from a mean of $70k for the first decile to about $550k for the tenth decile. Interacting each property's zip code and home value decile defines my zip code-decile peer network.

[121] I am grateful to Patrick Behrer for a helpful discussion on this point.

[122] Following assessor data documentation, I calculate the assessed value by adding total assessed land value to total assessed improvement value, and divide the result by 0.35.

Table 3.4: *Summary of the zip code, and zip code-decile panels.*

	Mean	Stnd. dev.	Min	Max	N
Zip code					
Pr(apply)$\times 10^{-6}$	18.58	94.54	0	7,194	186,901
eligible properties	6,195.87	4,358.25	7	14,544	186,901
cumulative conversions	383.73	450.77	0	2,296	186,901
cumulative area (sq-ft)	548,781.06	679,326.10	0	3,502,223	186,901
post cards (Apr. '07)	212.02	309.79	0	1,239	186,901
door hangers (Aug. '07)	5.29	9.10	0	50	186,901
door hangers (Sep. '07)	22.45	175.73	0	1,543	186,901
post cards (Oct. '07)	186.96	289.89	0	1,191	186,901
post cards (Apr. '08)	320.85	495.31	0	1,949	186,901
Zip code by decile					
Pr(apply)$\times 10^{-6}$	26.06	1,117.93	0	1,000,000	1,445,454
eligible properties	794.89	817.50	1	5,619	1,445,454
cumulative conversions	49.58	66.86	0	538	1,445,454
cumulative area (sq-ft)	70,902.81	105,036.99	0	1,348,976	1,445,454
post cards (Apr. '07)	27.39	55.94	0	484	1,445,454
door hangers (Aug. '07)	0.68	1.88	0	28	1,445,454
door hangers (Sep. '07)	2.90	28.13	0	436	1,445,454
post cards (Oct. '07)	24.15	52.16	0	463	1,445,454
post cards (Apr. '08)	41.45	85.97	0	951	1,445,454

Decile refers to LVVWD 2017 home value decile.

Table 3.4 describes the resulting panels.[123] The average daily application probability is small, about 19×10^{-6} for the zip code panel and 26×10^{-6} for the zip code-decile panel. These daily application probabilities translate to annual application probabilities of 0.68 and 0.95 percentage points, respectively.[124] Thus within any given zip code or zip code-decile, applications take place infrequently. The average number of eligible properties is just over 6,000 in the zip code panel and about 800 for the zip code-decile panel. The average cumulative of conversions in a zip code is just under 400 with an associated area of about 549 thousand square feet. For the zip code-decile panel, the average cumulative of conversions is about 50, with an associated area of just under 71 thousand square feet. And consistent with Figure 3.2 and Table 3.2, zip codes and zip code deciles mostly receive post cards. For example, the water authority sent the average zip code 212 post cards in April 2007, but less than 30 door hangers.

In a departure from BG's analysis, I also construct a panel at the individual property level. This allows me to run models controlling for unobservables at the parcel (i.e. property) level, rather than at the zip code, or zip code-decile level. In particular, I construct a panel of all parcels for which I observe an application and subsequent enrollment into the Cash-for-Grass program.[125] For each parcel, my observations run daily from the first month in which I observe any application, July 2003, to the day the particular parcel applies for the rebate. For each parcel-day, I construct an application indicator that describes whether the parcel applied for the rebate, indicators for whether the parcel had received any marketing material on or before the day of observation, and finally variables describing the cumulative

---

[123] Some properties undertake more than one conversion. However, it seems most reasonable to assume that the impact of any peer effect would be present for the first of any multiple conversions. Thus I consider only those applications associated with an initial conversion to desert landscape.

[124] The probability of not applying each day is of course $\Pr(\text{not apply}) = 1-\Pr(\text{apply})$. The probability of not applying every day for the entire year is the union of all individual probabilities of not applying each day, or $\Pr(\text{not apply})^{365} = (1-\Pr(\text{apply}))^{365}$. Since the probability of applying at some point during the year is the compliment of the probability of never applying during the year, the annual application probability can be calculated by $\left(1-(1-\Pr(\text{apply}))^{365}\right)$. Plugging the values of Pr(apply) for the zip code and zip code-decile reported in Table 3.4 into the preceding formula results in annual application probabilities of 0.68 and 0.95 percentage points, respectively.

[125] As above, I consider only those applications associated with an initial conversion to desert landscape.

number of conversions associated with the parcel's zip code or zip code-decile. Since the resulting panel contains nearly 40 million observations, for computational tractability I draw a random sample of 20 percent of parcels in the full data set.

Table 3.5 describes the resulting panel. The mean value for the application indicator represents the application probability, or about 0.06 percentage points. A converting parcel resides in zip codes with an average of 501 conversions, and within zip code-deciles of about 76 conversions. Consistent with Table 3.4, few converting parcels are exposed to marketing materials. For example, the first round of post cards had arrived on or before 5 percent of parcel-day observations.

Table 3.5: *Summary statistics for the 20 percent sub-sample of the individual level panel.*

	Mean	Stnd. dev.	Min	Max	N
application indicator, Pr(apply)	0.00058	0.024	0	1	7,890,379
zip cumulative conversions	500.9	451.8	0	2,494	7,890,379
zip-dec cumulative conversions	75.5	77.7	0	600	7,890,379
Apr. '07 post card indicator	0.051	0.220	0	1	7,890,379
Aug. '07 door hanger indicator	0.00073	0.027	0	1	7,890,379
Sep. '07 door hanger indicator	0.0017	0.041	0	1	7,890,379
Oct. '07 post card indicator	0.039	0.193	0	1	7,890,379
Apr. '08 post card indicator	0.056	0.230	0	1	7,890,379

Statistics reflect a sub-sample comprising 20 percent of applying parcels.

Lastly, consider Figure 3.3, which illustrates conversions[126] over time within the LVVWD service population. BG present a similar figure, arguing that clustering may indicate the presence of peer effects. While certainly not conclusive, the three panels illustrated in Figure 3.3 provide some visual evidence that areas concentrated with desert landscape conversions become increasingly concentrated with desert landscape conversions over time. In the next

[126]To be precise, I generated Figure 3.3 exclusively using enrollment dates, rather than post-conversion site visit dates.

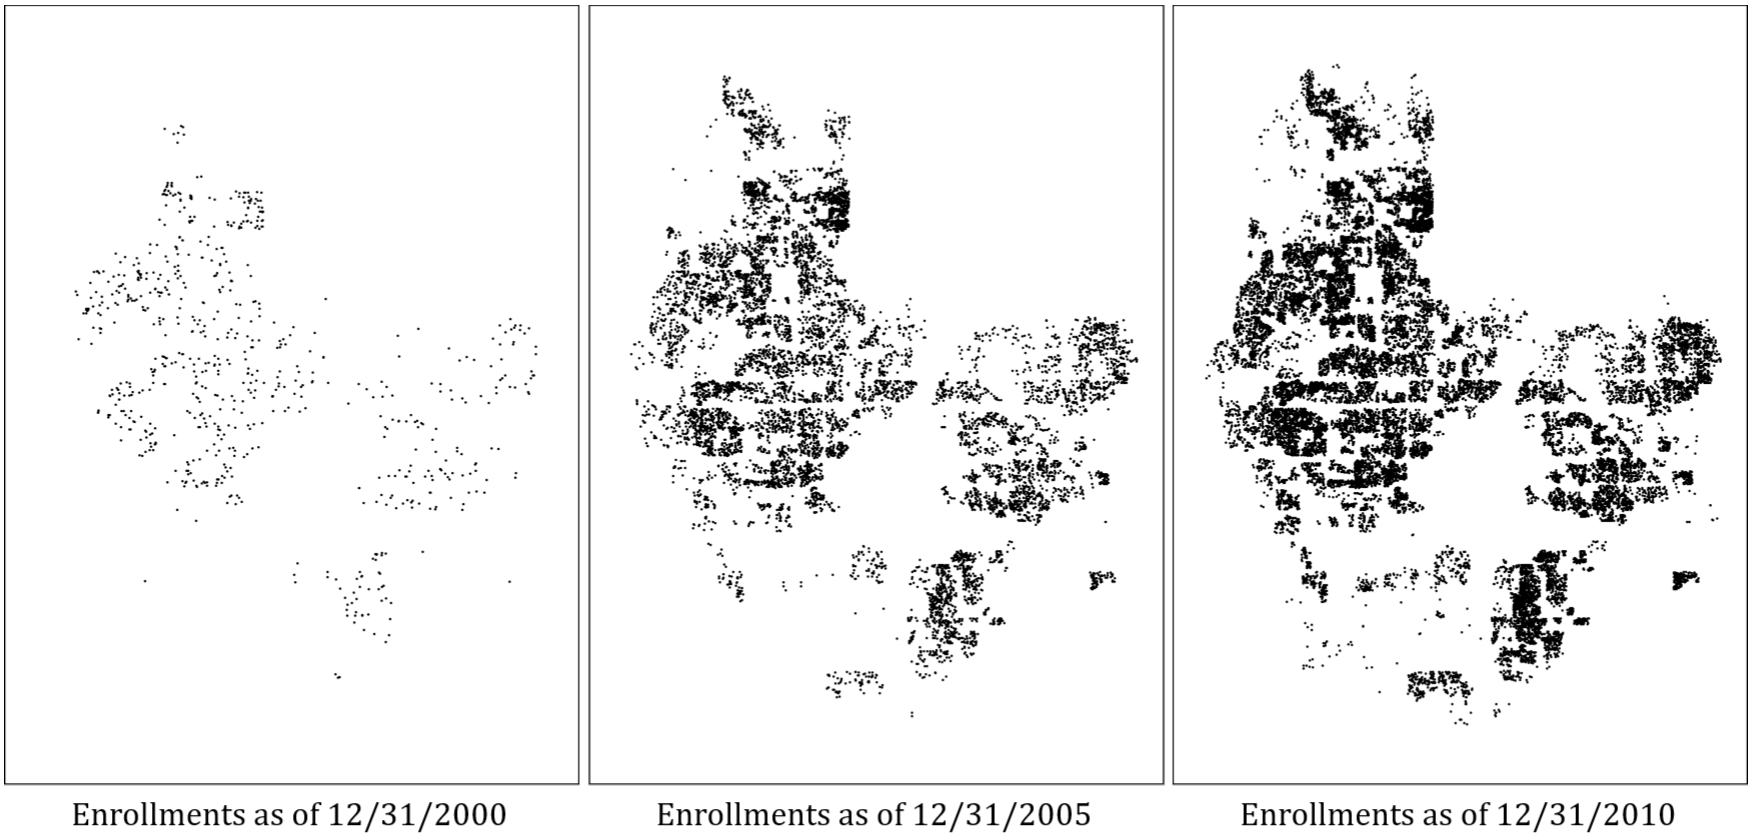

Figure 3.3: *Enrollments within the LVVWD over time. Each dot represents a single enrollment.*

section, I describe an empirical strategy adopted from BG that tests whether peer effects and the marketing campaigns may contribute to the observed pattern of conversions.

3.3 Methods

BG estimate a peer effect by modeling the impact of prior solar panel installations in a zip code, the so-called "installed base", on current decisions to install solar panels within that same zip code. In particular, BG propose a first-differences framework that overcomes many of the empirical challenges that make credibly identifying peer effects a challenge.[127] I adopt their empirical strategy shown below in Eq. (3.1):

$$(Y_{pt} - Y_{pt-1}) = \beta(C_{pt} - C_{pt-1}) + \sum_{j=1}^{5} \mu_j (M_{jpt} - M_{jpt-1}) + (\xi_t - \xi_{t-1}) + (\epsilon_{pt} - \epsilon_{pt-1}) \quad (3.1)$$

where Y_{pt} describes the number of applications[128] for the Cash-for-Grass rebate program divided by the total number of remaining eligible properties in peer network p (either a zip code or zip code-decile as discussed in section 3.2) on day t. C_{pt} describes the cumulative number of conversions to desert landscape in peer network p on day t. The coefficient estimate on β therefore describes the peer effect. ξ_t describes first-differenced year-by-month, day-of-month, and day-of-week fixed effects.[129] BG motivate their first-differences model with a fixed-effects model that includes zip code by quarter fixed-effects. In order for the zip code by quarter indicators to drop out of the first-differences model, BG remove the first observation of each quarter. I also drop the first observation of each quarter, thus implicitly modeling zip code by quarter, or zip code-decile by quarter, fixed-effects in the first-differences specification. Finally, M_{jpt} describes the total number of marketing materials

[127] For an elaboration of these challenges, see BG, Narayanan and Nair (2013), and Jackson *et al.* (2017).

[128] Following BG, I multiply Y_{pt} by 10^6 so that coefficient estimates are more readable.

[129] These fixed-effects will capture unobserved changes over time that are common to all peer networks such as changes in the rebate level. And since I limit my analysis to one water utility service district (the LVVWD), any effects arising from water rate changes will also be absorbed by the time level fixed-effects.

for campaign j in peer network p on day t.[130] Thus I define M_{jpt} analogous to C_{pt}, and the interpretation of the coefficient estimate μ_j is analogous to that of the peer effect, β.[131]

If the order of auto-correlation of ϵ_{pt} is larger than the time between application and conversion, β would be estimated inconsistently (see BG for a proof). The shortest time between application and conversion is three days (Table 3.1), and Figure 3.1 illustrates the vast majority of converting properties apply for the rebate at least one week prior to completing the conversion. Ideally then, the order of auto-correlation would be less than 3. For each peer network, I estimate Eq. (3.1) and then run a Cumby-Huizinga test for auto-correlation using `actest` (Baum and Schaffer, 2015) up through the fifteenth lag, clustering at the peer network level. At the zip code level, the first and seventh lag are statistically significant at the or below the 5 percent level. At the zip code-decile level, only the first lag is significant, and only at the 10 percent level. Since most applications take place at least a week before conversion, I conclude that my data satisfies BG's auto-correlation requirement.[132]

I apply two more of BG's analyses to my context, the first of which explores the impact of conversion area. I define a variable A_{pt}, which describes the cumulative square footage of desert landscape in peer network p on day t, and estimate the following models:

---

[130] $j = 1$ for the first post card campaign, $j = 2$ for the first door hanger campaign, $j = 3$ for the second door hanger campaign, $j = 4$ for the second post card campaign, and $j = 5$ for the third post card campaign.

[131] Alternatively, I could define M_{jpt} with indicators, letting M_{jpt} equal one on the day when peer network p received any of the marketing materials j. However, this would overweight the impact of a peer network receiving only a few materials compared to one that received several hundred. To avoid over-weighting in this way, I model M_{jpt} as the cumulative number of marketing materials in each peer network. But by modeling marketing materials in the same way as the peer effect, I implicitly allow for the same mechanism driving the peer effect to drive the impact of the marketing materials. This may misstate the impact of marketing materials because while everyone can observe a conversion (unless it takes place in a back yard), not everyone can observe which properties receive a post card or door hanger. An alternative modeling approach could be to interact marketing campaign descriptors with cumulative conversions to explore if marketing strengthens the peer effect. Another approach could be to define the marketing materials as some share of eligible properties. The difficulty with modeling impacts of marketing materials in an aggregate model in part motivates the individual level model (presented below) where it becomes possible to model an individual property's receipt of marketing materials, thereby avoiding the problems just discussed.

[132] This discussion follows closely an analogous discussion in BG (see BG section 4.1).

$$\Delta Y_{pt} = \alpha \Delta A_{pt} + \sum_{j=1}^{5} \mu_j \Delta M_{jpt} + \Delta \xi_t + \Delta \epsilon_{pt} \tag{3.2a}$$

$$\Delta Y_{pt} = \alpha \Delta A_{pt} + \beta \Delta C_{pt} + \sum_{j=1}^{5} \mu_j \Delta M_{jpt} + \Delta \xi_t + \Delta \epsilon_{pt} \tag{3.2b}$$

The second analysis I replicate from BG explores the impact of peer effects over time. BG do this in two ways. First, they include a quadratic term in their main estimation framework, arguing that if peer effects are present, the strength of the peer effect should be growing as more solar panels become installed. They also interact annual dummies with their variable describing the number of cumulative solar panel installations to visualize the dynamics of the peer effect. Accordingly, I estimate the following models, where d_y describes annual indicators:

$$\Delta Y_{pt} = \beta \Delta C_{pt} + \beta \Delta C_{pt}^2 + \sum_{i=1}^{5} \mu_i \Delta M_{ipt} + \Delta \xi_t + \Delta \epsilon_{pt} \tag{3.3a}$$

$$\Delta Y_{pt} = \sum_{y=2003}^{2014} \beta_y \Delta (C_{pt} \times d_y) + \sum_{i=1}^{5} \mu_i \Delta M_{ipt} + \Delta \xi_t + \Delta \epsilon_{pt} \tag{3.3b}$$

The preceding models implicitly assume that unobserved drivers of applications are homogeneous at the level of the peer network. Because I observe individual application and conversion decisions, I can model the effect of the cumulative number of conversions on individual decisions to apply for the rebate program, and thereby control for unobserved drivers of applications at the individual level using parcel-level fixed-effects. In particular, I estimate the model in Eq. (3.4):

$$\begin{aligned} apply_{it} - apply_{it-1} = {} & \beta (C_{pt} - C_{pt-1}) + \sum_{j=1}^{5} \mu_j (I^M_{jit} - I^M_{jit-1}) \\ & + (\xi_t - \xi_{t-1}) + (\epsilon_{it} - \epsilon_{it-1}) \end{aligned} \tag{3.4}$$

where $apply_{it}$ indicates parcel i's application status on day t, and I define C_{pt}, and ξ_t as in Eq (3.1). I^M_{jpt} is an indicator describing whether parcel i has received material from marketing campaign j on, or any day prior to, day t. Unlike Eq. (3.1), here I am only concerned with the direct impact of marketing campaign j on a particular parcel. Thus I will not capture any difference in the effect of marketing materials for those parcels located in peer networks that receive many marketing materials compared to those parcels located in peer networks that receive few materials.

I select a first-differences framework, thereby controlling for parcel-level unobservables. Since I only observe a parcel up until the day it applies, I never observe that parcel's conversion date. The order of auto-correlation in my error should therefore always be less than the time period between application to conversion (since in a sense, the time between application and conversion is infinite). Thus BG's first-differences approach, applied here in Eq. (3.4), should produce consistent estimates of β.

3.4 Results and discussion

3.4.1 Main results

Table 3.6 shows my primary results of estimating Eq. (3.1). The model in column 1 defines the peer network by zip codes. The model in column 2 defines the peer network by zip code-deciles. The models in columns 3 and 4 are analogous to 1 and 2, however instead of defining conversions by the post-conversion site visit as described in section 3.2, I define conversions exclusively using the enrollment date. I cluster standard errors at either the zip code (columns 1 and 3) or zip code-decile (columns 2 and 4). Following BG, my results in column 1 imply "that every additional [conversion] increases the probability of a household [application] in the same zip code by" 2.6×10^{-6} (since I define Y in millionths) or a 14 percent change relative to a zip code's average daily application probability of 18.58×10^{-6}. These results compare favorably to those of BG.[133] Column 2 shows that the peer effect within a zip code-decile is an order of magnitude stronger, 12.3×10^{-6}, or 47 percent change relative to a zip code-decile's average application probability.

Two implications follow from my results. First, my estimated peer effects are substantial

[133] They find a peer effect of 1.567×10^{-6}.

Table 3.6: *Effects of cumulative number of conversions and marketing materials on decisions to participate in the Cash-for-Grass rebate program. Estimates derived from Eq. (3.1).*

	(1)	(2)	(3)	(4)
conversions	2.621	12.326	0.392	0.950
	(0.271)***	(1.681)***	(0.129)***	(1.306)
post card (Apr. ’07)	0.158	0.673	0.157	0.669
	(0.051)***	(0.210)***	(0.051)***	(0.209)***
door hanger (Aug. ’07)	4.752	17.389	4.794	17.448
	(3.518)	(10.275)*	(3.502)	(10.252)*
door hanger (Sep. ’07)	0.004	0.008	0.004	0.008
	(0.002)*	(0.018)	(0.002)*	(0.018)
post card (Oct. ’07)	0.056	0.448	0.057	0.447
	(0.026)**	(0.159)***	(0.025)**	(0.158)***
post card (Apr. ’08)	-0.047	-0.357	-0.047	-0.354
	(0.027)*	(0.181)**	(0.027)*	(0.181)*
Conversion date proxy	post-conversion	post-conversion	enrollment	enrollment
adj. R^2	0.01	0.00	0.01	0.00
Cluster level	zip	zip-decile	zip	zip-decile
Number of clusters	47	364	47	364
Observations	184,845	1,429,553	184,845	1,429,637

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Outcome variable, Y_{pt}, defined as Pr(apply)$\times 10^{-6}$. $\bar{Y}_{\text{zip},t} = 18.58 \times 10^{-6}$ $\bar{Y}_{\text{zip-decile},t} = 26.06 \times 10^{-6}$.

Zip code or zip code-decile clustered standard errors reported in parentheses.

Decile refers to deciles of Clark County Assessor assessed property values (2017).

I drop three zip codes with no conversions. I also drop the first observation of each quarter.

All models include (first-differenced) year-by-month, day-of-month, and day-of-week fixed-effects.

compared to baseline application probabilities. Thus not only are my estimated peer effects significant in a statistical sense, they also appear to have a substantial impact on Cash-for-Grass program uptake.[134] Second, the results in column 2 relative to column 1 provide

[134]Though fluctuating somewhat over time, participation (number of conversions in a year) has waned since 2008. More people converting early in the program implies that the change in application probabilities are large when the change in cumulative conversions are large, and small when they are small, leading to a positive dependence of ΔY on ΔC. The coefficient estimates on cumulative conversions could therefore be due to this mechanical correlation and not due to any peer effect. Controlling for time effects helps, but only insofar as time trends are common across peer networks (which while perhaps not the strongest of assumptions, is an assumption nonetheless). In this same spirit, BG note that “if correlated unobservables are time-varying...then

suggestive evidence that at least part of the mechanism behind the peer effect may be operating though people's desire to ensure their home compares favorably to other homes in a similar housing market. But since homes of similar value tend to be located near each other, the larger peer effect in a zip code-decile may simply indicate that peer effects are stronger at more spatially localized levels. BG's street-level analysis supports this explanation.

In columns 3 and 4, I re-estimate Eq. (3.1) using the enrollment date as a proxy for the date of conversion. While columns 3 and 4 show the same general pattern as columns 1 and 2, these estimates fall one to two orders of magnitude below the estimates in the first two columns. As such, columns 3 and 4 demonstrate the bias introduced by using the enrollment date rather than the post-conversion site visit to proxy for the date of conversion. At least for the empirical strategy proposed by BG, my results in column 3 and 4 compared to those of column 1 and 2 highlight the necessity of developing accurate proxies for conversion. Even mis-specifying the conversion by as little as one month (on average) can have substantial impacts on estimates of peer effects.

The second set of results shown by Table 3.6 relate to the impact of the five targeted marketing campaigns.[135] At either the zip code or zip code-decile, I estimate impact of door hangers with little precision, which is unsurprising since so few door hangers were distributed.[136] I do, however, calculate precise estimates for the impact of post cards, with each coefficient estimated at or below the 5 percent level for the first two post card campaigns. Column 1 shows that an additional post card from the first two campaigns increases a property's application probability between 0.06×10^{-6} and 0.16×10^{-6}, respectively, or less than a 1 percent change in the baseline application probability. Compared to the 14 percent change associated with the zip code peer effect, post cards appear to have a comparatively

peer group-time effects would be necessary" and their first-differences model includes zip code by quarter fixed-effects, which I also include. But while zip code by quarter fixed-effects reduce the concern over spurious correlation discussed here, they do not completely eliminate it. My peer effect results should be interpreted with this caveat in mind.

[135]However, due to the challenges associated with modeling the marketing campaigns discussed in footnote 131, the following results should be taken as suggestive only.

[136]The magnitude of the estimates associated with the August 2007 door hanger campaign are on the order of the peer effects shown in row 1, but these door hanger estimates are imprecise.

small, even negligible impact on applications within an average zip code. I find similarly when defining the peer network by a zip code-decile. Column 2 shows that both estimates on the impact of the first two post card campaigns fall two orders of magnitude lower than the estimated peer effect, implying that these post cards had very little impact on the probability of application within an average zip code-decile. Regarding the third post card campaign, my results suggest these may have decreased the probability of application, at least within a zip-code decile. However, at the zip code and zip code-decile level, my estimated impacts due to the third post card campaign imply that the impact of these post cards were at most a 1.4 percent change relative to the baseline application probability. Overall, then, my results imply that despite estimating the impacts of post cards precisely in a statistical sense, post cards appear to have minimal meaningful impact, at least when compared to the strength of the peer effect.[137]

Two broad implications follow from my analysis of the impact of marketing materials. First, my results provide evidence that post cards may have the greatest positive impact. While determining the exact mechanism driving this result is beyond the scope of my analysis, I offer the following speculation. Door hangers can fall off and may never reach the property owner. Post cards, however, have a greater chance at reaching the homeowner and can be quickly and easily internalized if the message is simple and straightforward. Figure 3.4 illustrates an example post card sent to customers in 2015 and 2016, and to the extent that the earlier post cards contained similar design and messaging, it suggests the post cards could be readily digestible. The negative impacts estimated for the third post card campaign, however, suggest that messaging may also deter individuals from conversion.[138]

The second broad implication is that the peer effect dominates the impact of any of the

[137]The negligible effects due to the marketing campaigns could be explained by the fact that the water authority may have sent advertising materials to customers with the least inclination to convert to desert landscape. Thus the small estimated effects due to marketing may be primarily driven by (albeit unobserved) preferences of recipients and as a result, should not necessarily be taken as evidence of the effects of marketing in general.

[138]I do not have precise information regarding the content of these post cards. Thus suggesting that the content of post cards may influence decisions remains speculative at best.

Figure 3.4: *Example of a post card sent by the water authority marketing the Cash-for-Grass rebate program. This particular post card was sent in the Fall of 2015, and the Spring and Summer of 2016. The front matter text reads:* "Still holding on to your lawn? The Southern Nevada Water Authority is now offering a $2 rebate for every square foot of grass you replace with water-smart landscaping. You'll save dollars for yourself and water for our community. Plus, your home will look beautiful while doing your part to conserve." *Post card image provided by the SNWA.*

marketing campaigns. To my knowledge, my results are the first to make such a comparison. But this should not necessarily imply that marketing campaigns achieve little and should be abandoned. If, for example, post cards motivate even a few individuals to convert, these post card induced conversions will then subsequently influence remaining unconverted properties through the comparatively strong peer effect.[139]

Table 3.7: *Effects of the total area of the cumulative number of conversions on decisions to participate in the Cash-for-Grass rebate program. Estimates derived from Eq. (3.2a) and Eq. (3.2b).*

	(1)	(2)	(3)	(4)
total area converted	1.5e-03	-5.3e-04	7.2e-03	3.9e-03
	(1.7e-04)***	(5.1e-04)	(1.3e-03)***	(2.4e-03)
conversions		3.357		6.986
		(0.839)***		(3.416)**
adj. R^2	0.01	0.01	0.00	0.00
Cluster level	zip	zip	zip-decile	zip-decile
Number of zip codes	47	47	364	364
Observations	184,845	184,845	1,429,553	1,429,553

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Outcome variable, Y_{pt}, defined as Pr(apply)$\times 10^{-6}$. $\bar{Y}_{\text{zip},t} = 18.58 \times 10^{-6}$ $\bar{Y}_{\text{zip-decile},t} = 26.06 \times 10^{-6}$.

Zip code or zip code-decile clustered standard errors reported in parentheses.

Decile refers to deciles of Clark County Assessor assessed property values (2017).

I estimate, but do not report, the effect of the marketing campaigns.

Conversion date proxied by the post-conversion site visit.

I drop three zip codes with no conversions. I also drop the first observation of each quarter.

All models include (first-differenced) year-by-month, day-of-month, and day-of-week fixed-effects.

3.4.2 Total converted area's effect on application probability

BG find evidence that larger solar module installations have a stronger peer effect. Is the same true of desert landscape? Table 3.7 shows results from estimating Eq. (3.2a) and Eq. (3.2b). Columns 1 and 2 illustrate results from models defining the peer network by zip code, while columns 3 and 4 illustrate results from models defining the peer network by zip code-decile. Consistent with Table 3.6, defining the peer network by zip code-deciles yields

[139] In my calculations of the variables Y_{pt}, a_{it}, and C_{pt}, I only consider the first conversion associated with any parcel (some properties undertake more than one conversion). First conversions account for 93 percent of all conversions I observe. After converting once, it is difficult to think about how or why any peer effect would be relevant for subsequent conversions. But multiple conversions on other properties may influence an individual just like any other conversion. For all my results, I also run models that redefine C_{pt} to include all conversions, not just first conversions. My results change very little. In addition, I also estimate fixed-effect versions of my main models, and like BG, find my fixed-effect estimates of peer effects to be substantially smaller than my first-differences estimates of fixed effects. While fixed-effect and first-difference models should produce similar results if all the underlying assumptions for model validity are met, BG note their fixed-effect estimates likely suffer from the negative bias that Narayanan and Nair (2013) show to be present in fixed-effects estimates of pre-determined variables (of which BG's solar installations and my cumulative conversions are examples). My fixed-effects estimates likely suffer from this same bias.

larger peer effects than defining the peer network by zip codes. And though columns 1 and 3 suggest that area positively impacts the probability of application, when I additionally include the number of conversions in columns 2 and 4, the precision associated with the impact of area disappears. Furthermore, the estimate on conversions becomes significant at the 5 percent level or lower and appears consistent with the analogous estimates presented in Table 3.6. Opposite to BG's finding, my results suggest that the presence of conversions exclusively drives the peer effect. Conversion size appears not to be a factor when both variables are included in the model.[140]

3.4.3 Impact of an increasing number of cumulative conversions

Table 3.8: *Effects of cumulative number of conversions and the square of cumulative number of conversions, on decisions to participate in the Cash-for-Grass rebate program. Estimates derived from Eq. (3.3a).*

	(1)	(2)
conversions	-0.391	4.247
	(0.247)	(2.157)**
conversions^2	5.6e-03	6.3e-02
	(7.7e-04)***	(9.6e-03)***
adj. R^2	0.01	0.00
Cluster level	zip	zip-decile
Number of clusters	47	364
Observations	184,845	1,429,553

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Outcome variable, Y_{pt}, defined as Pr(apply)$\times 10^{-6}$. $\bar{Y}_{\text{zip},t} = 18.58 \times 10^{-6}$ $\bar{Y}_{\text{zip-decile},t} = 26.06 \times 10^{-6}$.

Zip code or zip code-decile clustered standard errors reported in parentheses.

Decile refers to deciles of Clark County Assessor assessed property values (2017).

I estimate, but do not report, the effect of the marketing campaigns.

Conversion date proxied by the post-conversion site visit.

I drop three zip codes with no conversions. I also drop the first observation of each quarter.

All models include (first-differenced) year-by-month, day-of-month, and day-of-week fixed-effects.

Table 3.8 presents results from estimating Eq. (3.3a). For both definitions of the peer network, the estimate on the square of conversions is positive and significant at the 1 percent

[140] I estimate, but do not show, impacts of the effect of the program advertisements. Effects are similar to those reported in Table 3.6.

level. These results are consistent with the findings of BG, and imply that as the number of conversions increases, each additional conversion has a stronger influence on the application probability.

Figure 3.5 generally reinforces this conclusion. In particular, Figure 3.5 illustrates an increasing annual peer effect. I derive these annual estimates from running Eq. (3.3b), where I interact the cumulative number of conversions with annual indicators.[141] BG show a similarly increasing peer effect in their plot of the annual peer effect.[142]

3.4.4 Modeling peer effects at the individual level

Table 3.9 illustrates results from estimating Eq. (3.4) using the panel described by Table 3.5 where I define cumulative conversions based on zip codes (row 1) and zip code deciles (row 2).[143] I find positive, statistically significant estimates for cumulative conversions at both the zip code and zip code-decile peer network definition. Like my results in Table 3.6, the estimates in Table 3.9 imply that the peer effect is stronger within a zip code-decile compared to a zip code. But for both definitions of a peer network, Table 3.9 implies that the peer effect is about an order of magnitude larger compared to what I estimated using the approach detailed by BG. However, relative to the average application probability, my estimated peer effects represent a 17 percent change when defining a peer network by zip codes and 43 percent change when defining a peer network by zip code-deciles. These

[141] My conversion data for 2014 is incomplete, covering conversions through late May. The significant decline in the peer effect in 2014 could therefore be driven by only observing about 6 months of conversion data.

[142] BG also find evidence that past solar panel installations increase the size of current solar panel installations. They explain this result by suggesting that as more installations take place, the uncertainty surrounding solar installation falls, and customers become more willing to undertake larger installations. I would not expect such a mechanism to be operating in the Cash-for-Grass context. Compared to solar panel installations, conversion to desert landscape involves more tangible resources. Homeowners can see the water they use to irrigate their lawns, whereas energy is not a resource homeowners readily observe. Compared to installing solar panels, the ramifications of converting to desert landscape would therefore be more salient to homeowners. In addition, many if not most homeowners would have re-landscaped their property in the past, or at least heard of others doing so, and thus would have a clearer sense of the necessary steps involved in converting to desert landscape compared to installing solar panels. To test this intuition, I replicate BG's analysis of the effect of solar installations on the size of the installation. While not shown, I find that cumulative conversions have no impact on the area of individual conversions.

[143] Recall this panel is a subset of the full panel, which contains nearly 40 million observations.

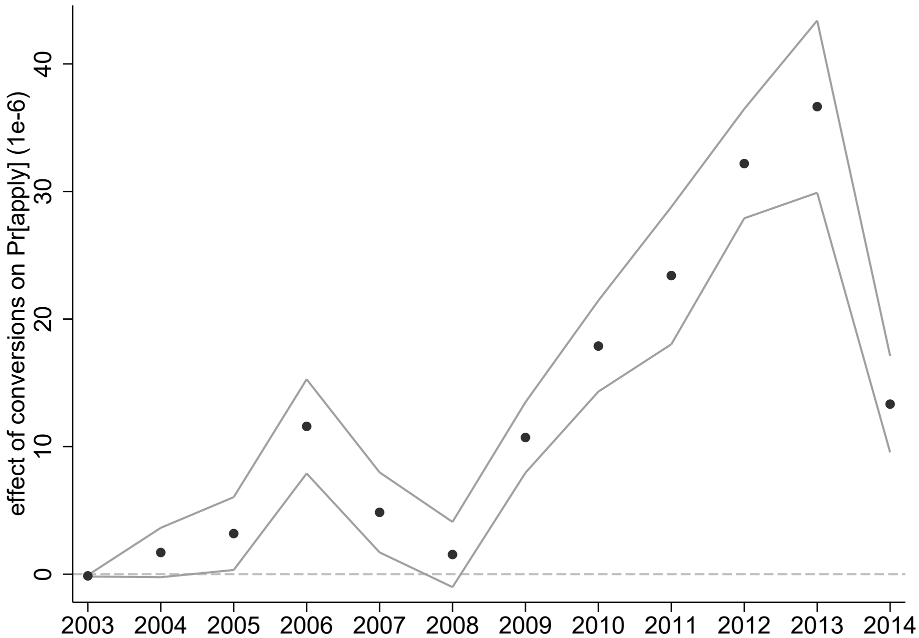

(a) *Peer group: zip code*

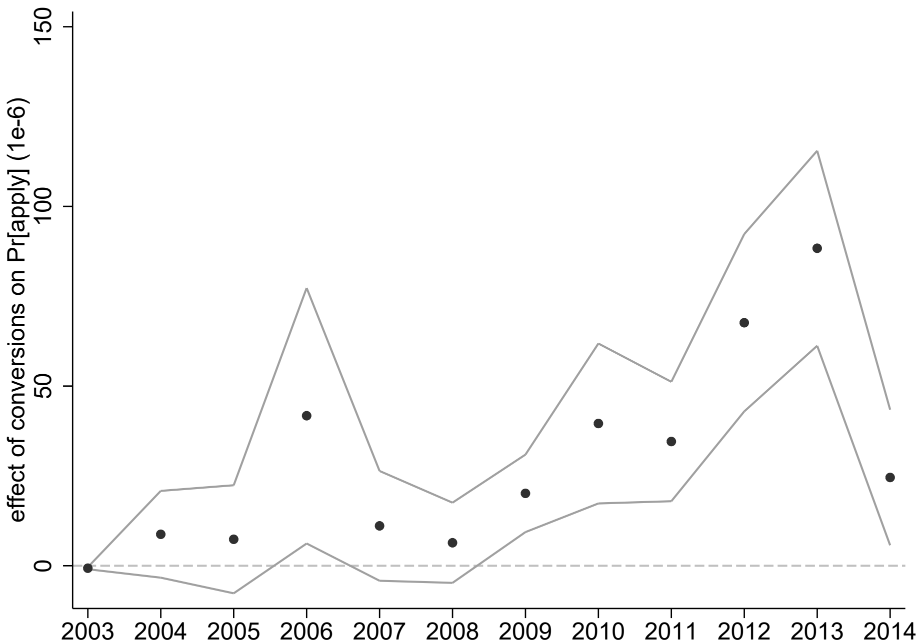

(b) *Peer group: zip code-decile*

Figure 3.5: *Annual peer effect, derived from running the model defined in Eq. (3.3b). In particular, the plot illustrates point estimates and 95 percent confidence intervals of my interaction of cumulative conversions with annual indicators.*

relative effects compare favorably to the relative effects I find in section 3.4.1. I conclude, therefore, that my estimates of peer effects when controlling for unobserved drivers of application at the individual level are broadly consistent with my results derived from BG's approach.

Turning to the impact of the marketing materials, Table 3.9 implies that each set of materials had a negative influence on decisions to apply for the rebate. Furthermore, the

Table 3.9: *Effects of cumulative number of conversions on individual decisions to apply for the Cash-for-Grass rebate program. Estimates derived from Eq. (3.4).*

	(1)	(2)
conversions	9.9e-05	2.5e-04
	(1.2e-05)***	(4.0e-05)***
post card (Apr. '07)	-7.6e-04	-7.8e-04
	(5.4e-05)***	(5.1e-05)***
door hanger (Aug. '07)	-5.0e-04	-5.1e-04
	(5.8e-05)***	(4.7e-05)***
door hanger (Sep. '07)	-1.4e-04	-1.3e-04
	(5.9e-05)**	(5.7e-05)**
post card (Oct. '07)	-3.3e-04	-3.4e-04
	(3.9e-05)***	(4.8e-05)***
post card (Apr. '08)	-7.7e-04	-7.7e-04
	(5.8e-05)***	(5.6e-05)***
adj. R^2	0.00	0.00
Cluster level	zip	zip-decile
Number of clusters	42	323
Observations	7,885,831	7,885,831

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Mean value of dependent variable: $\overline{apply}_{it} = 5.764 \times 10^{-4}$.

Zip code or zip code-decile clustered standard errors reported in parentheses.

Decile refers to deciles of Clark County Assessor assessed property values (2017).

I randomly draw 20 percent of converting parcels for computational tractability.

Conversion date proxied by the post-conversion site visit.

All models include (first-differenced) year-by-month, day-of-month, and day-of-week fixed-effects.

magnitude of these effects are on par with the those of the peer effect (row 1). I would expect my estimates of the impact of marketing to be at least non-negative. Determining the sources of inconsistency between Table 3.9 and my main results in Table 3.6 is the subject of ongoing work.[144]

3.5 Conclusion

In the preceding analysis, I investigate the presence of peer effects in the Southern Nevada Water Authority's Cash-for-Grass rebate program. I find that the peer effect within a zip code is consistent with the peer effect identified by BG in their analysis of solar panel installations. Like BG, I also show that the peer effect grows stronger with time, but unlike BG, I find that the only driver of the peer effect in desert landscape is the number of conversions. The area of cumulative conversions does not drive my results. While this could signal some differences in the specifics of solar panel installations compared to desert landscape conversions, overall my results validate the applicability of BG's proposed method in an alternative setting where peer effects may be important, namely conservation rebate programs.

I also show that the peer effect in a zip code-decile is stronger than the peer effect I find in a zip code. These results provide suggestive evidence that part of the mechanism behind the peer effect may work though an individual's desire to ensure his or her home compares favorably to other homes in a similar housing market. It could also be evidence that peer effects are stronger at more spatially localized scales, since homes of similar value tend to be built near each other.[145]

[144]Since properties do not often apply for the rebate on the day they receive marketing materials, in my first-differences specification estimating peer effects at the individual level, Eq. (3.4), the first-differenced marketing indicator will typically be one when the first-differenced application indicator is zero and vice versa. This could explain the strong negative dependence of the marketing materials on application probability shown in Table 3.9. One possible solution is to define my individual level model at a more aggregate time scale (such as weekly or monthly). Like the marketing campaign results derived from the aggregate level model above, the marketing campaign results derived from the individual model should be taken as suggestive only.

[145]It also bears mentioning that in both the aggregate and individual models, defining a peer group by a zip code or zip code-decile may over-weight the peer effect from some conversions and ignore peer effects from other conversions. Regarding over-weighting of peer effects, a conversion on one end of a sufficiently large peer network may have negligible effect on the probability of conversion for a property located on the other

Finally, I explore the effect that advertising has on Cash-for-Grass program participation. Two of the post card campaigns have the greatest influence on rebate recipients' decision to apply for the rebate, but a third post card campaign may have had the opposite effect. I also show that the estimated peer effect is much stronger than the advertising campaigns. To my knowledge, mine is the first analysis to compare the impact of advertising campaigns to a peer effect.

Conservation rebate programs such as the Southern Nevada Water Authority's Cash-for-Grass program will only grow in popularity as utilities increasingly face imbalances between their resource supply and customer demand. Thus uncovering the drivers of program uptake is not only interesting from an economic perspective, but also contributes to our understanding of the optimal design of subsidy programs.

end of the same peer network. Regarding ignoring peer effects, if a property is located near the border of its peer network, a nearby conversion to desert landscape in another peer network will not, in the models I have presented above, have any impact on the property's application probability. To avoid both problems, a preferred approach may be to define a peer network for each property in my sample, either by some visibility metric, or simply by a geometric boundary such as a circle. By re-estimating the peer effect for increasingly large radii, one could also determine the relationship between the peer effect and distance. By additionally comparing these results to results derived from limiting properties within such peer circles to those of a similar value as the property being analyzed, one could distinguish between the impact of the spatial dimensions of the peer effect and whether the peer effect operates through a competitive housing market channel.

References

Adamowicz, W., Dupont, D., Krupnick, A. and Zhang, J. (2007). *Valuation of Cancer and Microbial Disease Risk Reductions in Municipal Drinking Water: An Analysis of Risk Context Using Multiple Valuation Methods*. Discussion Paper RFF DP 07-39, Resources For the Future, 1616 P St. NW, Washington, DC, 20036.

Alberini, A. and Towe, C. (2015). *Information v. Energy Efficiency Incentives: Evidence from Residential Electricity Consumption in Maryland*. Working Paper 18.2015, Fondazione Eni Enrico Mattei (FEEM), Corso Magenta, 63, 20123 Milano (I).

Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*, **95** (9), 1082–1095.

— and Greenstone, M. (2012). Is There an Energy Efficiency Gap? *Journal of Economic Perspectives*, **26** (1), 3–28.

—, Knittel, C. and Taubinsky, D. (2015). Tagging and Targeting of Energy Efficiency Subsidies. *American Economic Review*, **105** (5), 187–191.

— and Sweeney, R. L. (2016). The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment. *Management Science*.

Arimura, T. H., Katayama, H. and Sakudo, M. (2016). Do Social Norms Matter to Energy-Saving Behavior? Endogenous Social and Correlated Effects. *Journal of the Association of Environmental and Resource Economists*, **3** (3), 525–553.

—, Li, S., Newell, R. G. and Palmer, K. (2011). *Cost-effectiveness of electricity energy efficiency programs*. Working Paper 17556, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Auffhammer, M., Blumstein, C. and Fowlie, M. (2008). Demand-Side Management and Energy Efficiency Revisited. *The Energy Journal*, **29** (3), 91–104.

Baker, A. E. (2004). *Effects of residential xeriscape conversions on property values: A baseline case study*. UNLV Theses/Dissertations/Professional Papers/Capstones Paper 191, University of Nevada Las Vegas.

Bartik, T. J. (1987). The Estimation of Demand Parameters in Hedonic Price Models. *Journal of Political Economy*, **95** (1), 81–88.

Baum, C. F. and Schaffer, M. E. (2015). ACTEST: Stata module to perform Cumby-Huizinga general test for autocorrelation in time series.

Bennear, L. S., Jessoe, K. K. and Olmstead, S. M. (2009). Sampling Out: Regulatory Avoidance and the Total Coliform Rule. *Environmental Science & Technology*, **43** (14), 5176–5182.

—, Lee, J. M. and Taylor, L. O. (2013). Municipal Rebate Programs for Environmental Retrofits: An Evaluation of Additionality and Cost-Effectiveness. *Journal of Policy Analysis and Management*, **32** (2), 350–372.

— and Olmstead, S. M. (2008). The impacts of the "right to know": Information disclosure and the violation of drinking water standards. *Journal of Environmental Economics and Management*, **56** (2), 117–130.

Blackwell, J. L. (2005). Estimation and testing of fixed-effect panel-data systems. *The Stata Journal*, **5** (2), 202–207.

Bollinger, B. and Gillingham, K. (2012). Peer Effects in the Diffusion of Solar Photovoltaic Panels. *Marketing Science*, **31** (6), 900–912.

Bonan, G. B. (2000). The microclimates of a suburban Colorado (USA) landscape and implications for planning and design. *Landscape and Urban Planning*, **49** (3–4), 97–114.

Boomhower, J. and Davis, L. W. (2014). A credible approach for measuring inframarginal participation in energy efficiency programs. *Journal of Public Economics*, **113**, 67–79.

Borusyak, K. and Jaravel, X. (2016). Revisiting Event Study Designs. *Available at SSRN 2826228*.

Brelsford, C. M. (2014). *Whiskey is for Drinking; Water is for Fighting Over: Population Growth, Infrastructure Change, and Conservation Policy as Drivers of Residential Water Demand*. Ph.D. thesis, Arizona State University.

Brent, D. A., Cook, J. H. and Olsen, S. (2015). Social Comparisons, Household Water Use, and Participation in Utility Conservation Programs: Evidence from Three Randomized Trials. *Journal of the Association of Environmental and Resource Economists*, **2** (4), 597–627.

Bursztyn, L. and Jensen, R. (2016). *Social Image and Economic Behavior in the Field: Identifying, Understanding and Shaping Social Pressure*. Working Paper 23013, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Chakraborti, L. and McConnell, K. E. (2012). Does ambient water quality affect the stringency of regulations? Plant-level evidence of the Clean Water Act. *Land Economics*, **88** (3), 518–535.

Cho, S.-H., Roberts, R. K. and Kim, S. G. (2011). Negative externalities on property values resulting from water impairment: The case of the Pigeon River Watershed. *Ecological Economics*, **70** (12), 2390–2399.

Cornelissen, T., Dustmann, C. and Schönberg, U. (2017). Peer Effects in the Workplace. *American Economic Review*, **107** (2), 425–456.

Correia, S. (2016). REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects.

Dalhuisen, J. M., Florax, R. J. G. M., Groot, H. L. F. d. and Nijkamp, P. (2003). Price and Income Elasticities of Residential Water Demand: A Meta-Analysis. *Land Economics*, **79** (2), 292–308.

Davis, L. W. (2004). The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster. *The American Economic Review*, **94** (5), 1693–1704.

—, Fuchs, A. and Gertler, P. (2014). Cash for Coolers: Evaluating a Large-Scale Appliance Replacement Program in Mexico. *American Economic Journal: Economic Policy*, **6** (4), 207–238.

Delmas, M., Montes-Sancho, M. J. and Shimshack, J. P. (2010). Information Disclosure Policies: Evidence from the Electricity Industry. *Economic Inquiry*, **48** (2), 483–498.

Deoreo, W. B., Mayer, P. W. and Rosales, J. (2000). Xeriscape Conversion for Urban Water Conservation. In *Proceedings of the 4th Decennial Symposium*, Phoenix, AZ.

Dolan, P. and Metcalfe, R. D. (2015). *Neighbors, knowledge, and nuggets: two natural field experiments on the role of incentives on energy conservation*. Working Paper 2589269, Becker Friedman Institute for Research in Economics.

Doshi, A. R., Dowell, G. W. S. and Toffel, M. W. (2013). How Firms Respond To Mandatory Information Disclosure. *Strategic Management Journal*, **34** (10), 1209–1231.

Dranove, D. and Jin, G. Z. (2010). Quality Disclosure and Certification: Theory and Practice. *Journal of Economic Literature*, **48** (4), 935–963.

Edwards, E. C. and Libecap, G. D. (2015). Water Institutions and the Law of One Price. In R. Halverson and D. Layton (eds.), *Handbook on the Economics of Natural Resources*, Edward Elgar Publishing, Inc., pp. 442–473.

Epple, D. (1987). Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products. *Journal of political economy*, **95** (1), 59–80.

Espey, M., Espey, J. and Shaw, W. D. (1997). Price elasticity of residential demand for water: A meta-analysis. *Water Resources Research*, **33** (6), 1369–1374.

Fafchamps, M. and Mo, D. (2017). *Peer Effects in Computer Assisted Learning: Evidence from a Randomized Experiment*. Working Paper 23195, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Field, E., Jayachandran, S., Pande, R. and Rigol, N. (2014). Friends at Work: Can Peer Support Stimulate Female Entrepreneurship? *Unpublished manuscript, Northwestern Univ., Evanston, IL.*

Fung, A., Graham, M. and Weil, D. (2007). *Full Disclosure: The Perils and Promise of Transparency*. Cambridge, UK.: Cambridge University Press.

Gelman, A. and Imbens, G. (2014). *Why high-order polynomials should not be used in regression discontinuity designs*. Working Paper 20405, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Greenstone, M. and Gallagher, J. (2008). Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program. *The Quarterly Journal of Economics*, **951-1003**.

Griffin, R. C. (2001). Effective Water Pricing. *Journal of the American Water Resources Association*, **37** (5), 1335–1347.

Halvorsen, R. and Palmquist, R. (1980). The Interpretation of Dummy Variables in Semilogarithmic Equations. *The American Economic Review*, **70** (3), 474–475.

Houde, S. and Aldy, J. E. (2014). *Belt and Suspenders and More: The Incremental Impact of Energy Efficiency Subsidies in the Presence of Existing Policy Instruments*. Working Paper 20541, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Howe, C. W., Lazo, J. K. and Weber, K. R. (1990). The Economic Impacts of Agriculture-to-Urban Water Transfers on the Area of Origin: Case Study of the Arkansas River Valley in Colorado. *American Journal of Agricultural Economics*, **72** (5), 1200–1204.

Innes, R. and Cory, D. (2001). The economics of safe drinking water. *Land Economics*, **77** (1), 94–117.

Jackson, M. O., Rogers, B. W. and Zenou, Y. (2017). The Economic Consequences of Social-Network Structure. *Journal of Economic Literature*, **55** (1), 49–95.

Johnson, B. B. (2003). Do reports on drinking water quality affect customers' concerns? Experiments in report content. *Risk Analysis*, **23** (5), 985–998.

Joskow, P. L. and Marron, D. B. (1992). What Does a Negawatt Really Cost? Evidence from Utility Conservation Programs. *The Energy Journal*, **13** (4).

Keiser, D. A. and Shapiro, J. S. (2017). *Consequences of the Clean Water Act and the demand for water quality*. Working Paper 23070, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Klaiber, H. A., Abbott, J. and Smith, V. K. (2015). *Some Like it (Less) Hot: Extracting Tradeoff Measures for Physically Coupled Amenities*. Working Paper 21051, National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA, 02138.

Konar, S. and Cohen, M. A. (2001). Does the Market Value Environmental Performance? *The Review of Economics and Statistics*, **83** (2), 281–289.

Kuminoff, N. V. and Pope, J. C. (2014). Do "Capitalization Effects" for Public Goods Reveal the Public's Willingness to Pay? *International Economic Review*, **55** (4), 1227–1250.

Leggett, C. G. and Bockstael, N. E. (2000). Evidence of the Effects of Water Quality on Residential Land Prices. *Journal of Environmental Economics and Management*, **39** (2), 121–144.

Leibenstein, H. (1950). Bandwagon, Snob, and Veblen Effects in the Theory of Consumers' Demand. *Quarterly Journal of Economics*, **64** (2), 183–207.

Libecap, G. D. (2007). *Owens Valley Revisited: A Reassessment of the West's First Great Water Transfer*. Stanford, CA: Stanford University Press.

Linden, L. and Rockoff, J. E. (2008). Estimates of the Impact of Crime Risk on Property Values from Megan's Laws. *American Economic Review*, **98** (3), 1103–1127.

Loughran, D. S. and Kulick, J. (2004). Demand-Side Management and Energy Efficiency in the United States. *Energy Journal*, **25** (1), 19–43.

Mansur, E. T. and Olmstead, S. M. (2012). The value of scarce water: Measuring the inefficiency of municipal regulations. *Journal of Urban Economics*, **71** (3), 332–346.

Mastromonaco, R. (2015). Do environmental right-to-know laws affect markets? Capitalization of information in the toxic release inventory. *Journal of Environmental Economics and Management*, **71**, 54–70.

McCrary, J. (2007). The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police. *The American Economic Review*, **97** (1), 318–353.

Miguel, E. and Kremer, M. (2004). Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, **72** (1), 159–217.

Muehlenbachs, L., Spiller, E. and Timmins, C. (2015). The Housing Market Impacts of Shale Gas Development. *American Economic Review*, **105** (12), 3633–3659.

Mustafa, D., Smucker, T. A., Ginn, F., Johns, R. and Connely, S. (2010). Xeriscape people and the cultural politics of turfgrass transformation. *Environment and Planning D: Society and Space*, **28** (4), 600–617.

Myers, E. (2016). *Are Home Buyers Myopic? Evidence From Housing Sales*. E2e Working Paper 024.

Narayanan, S. and Nair, H. S. (2013). Estimating Causal Installed-Base Effects: A Bias-Correction Approach. *Journal of Marketing Research*, **50** (1), 70–94.

Olmstead, S. M. (2010). The Economics of Water Quality. *Review of Environmental Economics and Policy*, **4** (1), 44 – 62.

Pope, J. C. (2008). Fear of crime and housing prices: Household reactions to sex offender registries. *Journal of Urban Economics*, **64** (3), 601–614.

Powers, N., Blackman, A., Lyon, T. P. and Narain, U. (2011). Does Disclosure Reduce Pollution? Evidence from India's Green Rating Project. *Environmental and Resource Economics*, **50** (1), 131–155.

Price, J. I., Chermak, J. M. and Felardo, J. (2014). Low-flow appliances and household water demand: An evaluation of demand-side management policy in Albuquerque, New Mexico. *Journal of Environmental Management*, **133**, 37–44.

Raucher, R. S., Rubin, S. J., Crawford-Brown, D. and Lawson, M. M. (2011). Benefit-Cost Analysis for Drinking Water Standards: Efficiency, Equity, and Affordability Considerations in Small Communities. *Journal of Benefit-Cost Analysis*, **2** (1).

Rollins, C. A. (2008). *Comparing values for a private environmental good, xeriscape: Hedonic price method versus contingent valuation method*. Ph.D. thesis, University of Nevada Las Vegas.

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, **82** (1), 34–55.

Shimshack, J. P., Ward, M. B. and Beatty, T. K. (2007). Mercury advisories: Information, education, and fish consumption. *Journal of Environmental Economics and Management*, **53** (2), 158–179.

SNWA (2009). *Water Resource Plan 09*. Tech. rep., Southern Nevada Water Authority, 100 City Parkway, Suite 700, Las Vegas, Nevada, 89106.

Sovocool, K. A., Morgan, M. and Bennett, D. (2006). An in-depth investigation of Xeriscape as a water conservation measure. *Journal of the American Water Works Association*, **98** (2), 82–93.

Tiemann, M. (2014). *Safe Drinking Water Act (SDWA): A Summary of the Act and Its Major Requirements*. CRS Report RL31243, Congressional Research Service.

Walls, M., Kousky, C. and Chu, Z. (2015). Is What You See What You Get? The Value of Natural Landscape Views. *Land Economics*, **91** (1), 1–19.

Wrenn, D. H., Klaiber, H. A. and Jaenicke, E. C. (2016). Unconventional Shale Gas Development, Risk Perceptions, and Averting Behavior: Evidence from Bottled Water Purchases. *Journal of the Association of Environmental and Resource Economists*, **3** (4), 779–817.

Zivin, J. G., Neidell, M. and Schlenker, W. (2011). Water Quality Violations and Avoidance Behavior: Evidence from Bottled Water Consumption. *The American Economic Review: Papers and Proceedings*, **101** (3), 448–453.

Appendix A

Appendix to Chapter 1

A.1 Volumetric and water bill savings calculation details

I calculate baseline water use from the average water use for all participants converting once prior to enrollment, or 23,818 gallons per month. Dividing my savings estimate of 5,000 gal/month by the baseline value yields the 21 percent reduction reported in section 1.3.

I calculate water bill savings for an average LVVWD customer in 2013 that experiences a constant 5,000 gal/month savings throughout each month of the year. Water charges depend upon the meter size. In 2013, over 99 percent of single-family LVVWD customers in my panel have a 1 inch, 3/4 inch or 5/8 inch meter. Of these customers, 4 percent have a 1 inch meter, 47 percent have a 3/4 inch meter, and the remaining 49 percent have a 5/8 inch meter. Using a bill calculator provided by the water authority, I estimate the annual water bill for each meter size (1″, 3/4″ and 5/8″) for average monthly water use in 2013. I also calculate the water bill for each meter size using average water use net of the 5,000 gal/month savings estimate. Finally, I calculate a weighted average water bill, with weights defined by the share of customers associated with each meter size.

The average customer in 2013 pays a $501.03 water bill. Saving 5,000 gal/month reduces the water bill by $150.44 (about 30 percent) to $350.59. The present discounted value of an infinite stream of these savings equals $3,159 assuming a 5 percent discount rate.

I likely overstate water bill savings. My water use data demonstrate a high degree of variability in demand over the calendar year, and to the extent that most of the increase in summer water use arises from outdoor landscape irrigation, applying average savings evenly throughout the year will likely understate water savings in the summer and overstate water savings in the winter.[146] Since LVVWD customers experience block pricing, I likely underestimate summer bill savings more than I overestimate winter bill savings, producing a net underestimate in the annual water bill. The water bill savings should be interpreted with this caveat in mind, however I do not believe this bias to grossly distort the predicted water bill savings.

A.2 Additional water savings results

A.2.1 Fixed-effects, early exits, and program changes

In this section I present savings estimates that explore the impact of various fixed-effects, parcels that exit before the end of the sample, and two program policy changes.

In July, 2000, the water authority began rebating participants based on the size of their conversion; prior to this the water authority determined rebates based on how much water the participant saved relative to participant specific past average monthly water use. In March, 2004, the water authority relaxed a 400 square foot minimum conversion requirement, allowing conversions less than 400 square feet provided the conversion comprised an entire front or back yard.[147] To study the impact of these program changes, I estimate Eq. (A.1) and Eq. (A.2), where I interact my post-enrollment indicator with indicators for enrollment dates corresponding to dates on or after the change in determining the rebate amount, 'sqft',

[146] Sovocool *et al.* (2006) find that much of the savings due to conversion to desert landscape comes from savings in the summer.

[147] This restriction appears to have constrained participants. Histograms of converted areas for single-family participants (not reported here) show a sharp cutoff at 400 square feet for conversions taking place prior to March 2004. After March 2004, a similar histogram shows a more continuous distribution around 400 square feet.

and relaxing the 400 ft^2 minimum, 'min conv'.

$$Q_{it} = \alpha[\text{pre-period}]_{it} + \beta_1[\text{post-enroll}]_{it} + \beta_2[\text{post-enroll}]_{it} \times [\text{sqft}]_i + \mu_{im} + \delta_{tc} + \epsilon_{it} \tag{A.1}$$

$$Q_{it} = \alpha[\text{pre-period}]_{it} + \beta_1[\text{post-enroll}]_{it} + \beta_2[\text{post-enroll}]_{it} \times [\text{min conv}]_i + \mu_{im} + \delta_{tc} + \epsilon_{it} \tag{A.2}$$

Table A.1 displays results. Compared to the results in column 1, including either month-of-sample by cohort fixed-effects (column 2) or parcel by month-of-calendar year fixed-effects (column 3) reduces the estimate of savings compared to the specification that includes only parcel and month-of-sample fixed-effects (column 1). These results suggest that variation in seasonal fluctuations as well as time-varying characteristics among different aged houses affect results. I therefore control for both in my main specification (column 4).[148] In column 5 I eliminate from my main specification any parcels that drop out before the end of the sample (April, 2014). Dropping early exits has little impact on savings. In column 6 and 7 I explore the impact of the two program administrative changes discussed above. The negative point estimate associated with the effect of the subsidy remuneration method (column 6, row 2) suggests that rebating customers based on the size of their conversion may have increased savings, however I cannot statistically distinguish the effect from zero. The rebate remuneration method therefore likely had little effect.[149] Allowing conversions under 400 square feet, however, appears to have decreased savings by about 1,500 gallons per month per average conversion (column 7, row 3). As more participants convert areas less than 400 square feet, one expects average savings per conversion to fall, since the average conversion area falls. In terms of savings per square foot of converted area, 6,300 gallons per month corresponds to about 46 gal/ft^2/year (the mean conversion area

[148] In addition, results from a falsification test show that a proxy policy defined as a proxy enrollment 5 years prior to actual enrollment has no effect only for the fixed-effect specifications represented in columns 2 and 4. This further validates the importance of controlling for cohort effects.

[149] I am unable to determine the impact on savings per square foot of converted area since the water authority did not record converted area for conversions prior to 2000.

Table A.1: *Additional water savings results: effect of fixed-effects, early exits, and administrative changes.*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
post-enroll	-5.16	-5.00	-5.07	-4.92	-4.93	-4.91	-6.32
	(0.06)***	(0.06)***	(0.06)***	(0.06)***	(0.06)***	(0.61)***	(0.23)***
post-enroll×sqft						-0.02	
						(0.61)	
post-enroll×min conv							1.54
							(0.24)***
Sample	full	full	full	full	no early exits	full	full
Fixed-effects							
μ_i	yes	yes	-	-	-	-	-
μ_{im}	-	-	yes	yes	yes	yes	yes
δ_t	yes	-	yes	-	-	-	-
δ_{tc}	-	yes	-	yes	yes	yes	yes
adj. R^2	0.25	0.26	0.30	0.30	0.30	0.30	0.30
Parcels	309,608	309,608	309,201	309,201	305,068	309,201	309,201
Observations	64,135,652	64,135,652	64,120,344	64,120,344	63,356,472	64,120,344	64,120,344

$^* p < 0.10$, $^{**} p < 0.05$, $^{***} p < 0.01$

Parcel clustered standard errors (reported in parentheses).

μ_i, μ_{im} parcel and parcel by month-of-calendar year fixed-effects.

δ_t, δ_{tc} month-of-sample and month-of-sample by cohort fixed-effects.

Cohorts based on parcel's first year in sample: 88-89, 90-94, 95-99, 00-04, 05-14.

Early exit parcels drop out of sample prior to sample's end.

sqft: indicator for enrollment after June 2000.

min conv: indicator for enrollment after Feb. 2004.

for conversions prior to March, 2004 equals 1,653 square feet). After relaxing the minimum conversion requirement, savings become 4,800 gallons per month ($-6.320 + 1.54 = -4.78$) which corresponds to 43 gal/ft^2/year (the mean conversion area throughout the life of the program equals 1,348 ft^2), however, this value falls within the 95 percent confidence interval of the pre-March, 2004 estimate. Relaxing the minimum conversion requirement does not appear to have affected savings efficiency.

A.2.2 Additional event studies

Figure A.1 illustrates results of an event study derived from Eq. (1.1), but defining τ with respect to the enrollment date rather than the application date. The solid vertical line indicates the date of enrollment, while the dashed vertical line indicates the average month of application, 5 months prior to the date of enrollment. These results exhibit less oscillations prior to the enrollment date compared to Figure 1.6, however overall Figure 1.6 and Figure A.1 lead to the same conclusion; there appears to be an absence of pre-trends prior to the application-enrollment period, savings remain relatively stable after enrollment, and there exists some transient behavior at least one year prior to the enrollment date.[150]

I also implement an alternative event study using my matched sample. In particular, I calculate the difference in water use in each month-of-sample between individual participating parcels and their respective matched non-participating pairs. Then for each participating parcel, I associate the month-of-sample with event time. Finally, I calculate the average of the pairwise differences in water use in each event month.[151] Figure A.2 illustrates results of this procedure, defining event time with respect to application (panel a) and enrollment (panel b). Both figures illustrate an absence of pre-trends prior to the application-enrollment period, and some transient behavior at least one year prior to the enrollment date. In

[150]Though not shown, event studies derived from the three alternative control samples (participants only, DNF, and matched control samples) display largely similar patterns as in Figure 1.6 and Figure A.1, however they show greater instability prior to conversion and exhibit strikingly more pronounced erosion in savings after enrollment. Also, the event studies derived from the DNF and matched control samples illustrate savings closer to 4,000 gal/month, consistent with the estimates shown in Table 1.3.

[151]In particular, I regress my 'difference' variable on a constant, and cluster standard errors at the parcel level.

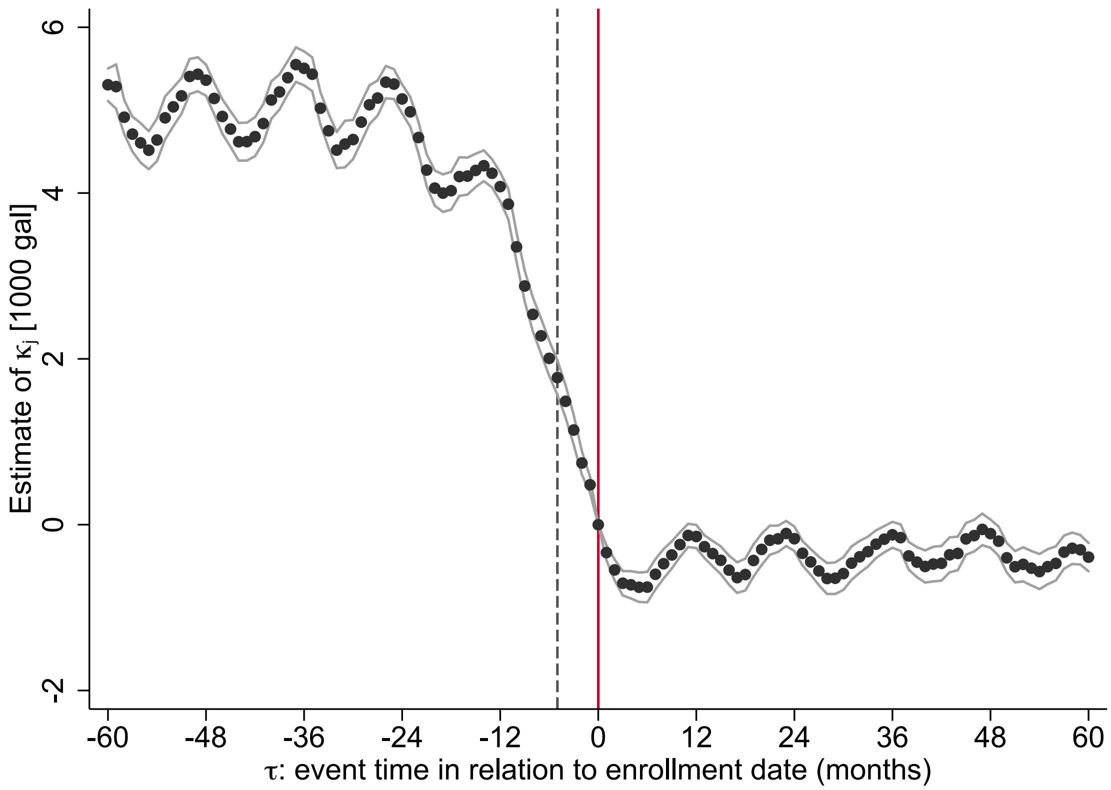

Figure A.1: *Event study, illustrating water savings from the Cash-for-Grass program. Point estimates and 95 percent confidence intervals of κ_j's are derived from estimating Eq. (1.1), with event time defined with respect to the month of enrollment. Standard errors are clustered at the parcel level. The omitted category is $\kappa_0 = 0$. Observations are limited to single-family participating parcels that converted only once and all non-participants. Participating parcel observations are further restricted to a five-year window around the month of enrollment; that is $-60 \leq \kappa \leq 60$. The vertical solid line indicates the month of enrollment. The vertical dashed line indicates the average month of application, five months prior to the enrollment date.*

contrast to Figure 1.6 and Figure A.1, however, Figure A.2 illustrates a more pronounced erosion in savings after conversion.

A.2.3 Additional discussion regarding impact of time and pre-enrollment consumption characteristics

Because estimating Eq. (1.4) essentially compares participants in a given year with participants in all years, one may be concerned that higher water demand early in the program biases estimates of later year annual program savings. To test robustness of my

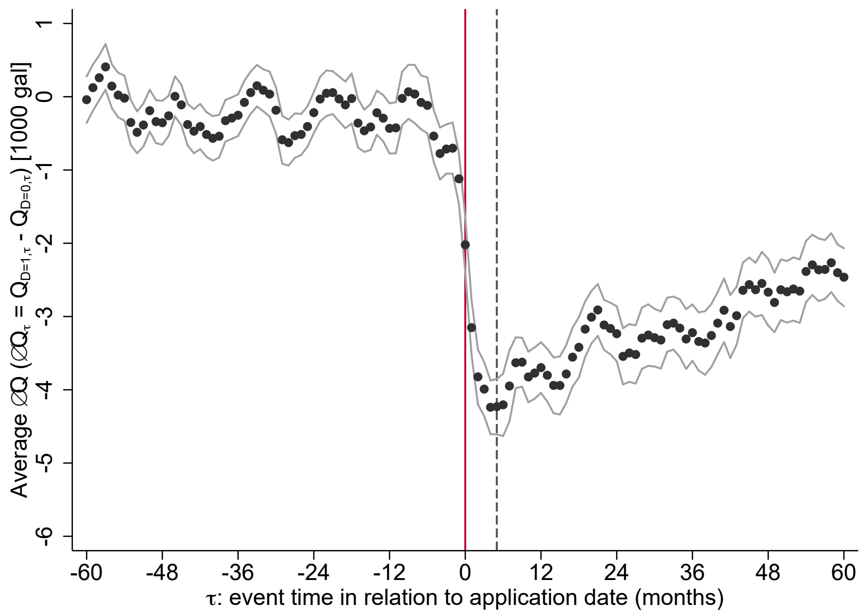

(a) *Event study w.r.t. application date*

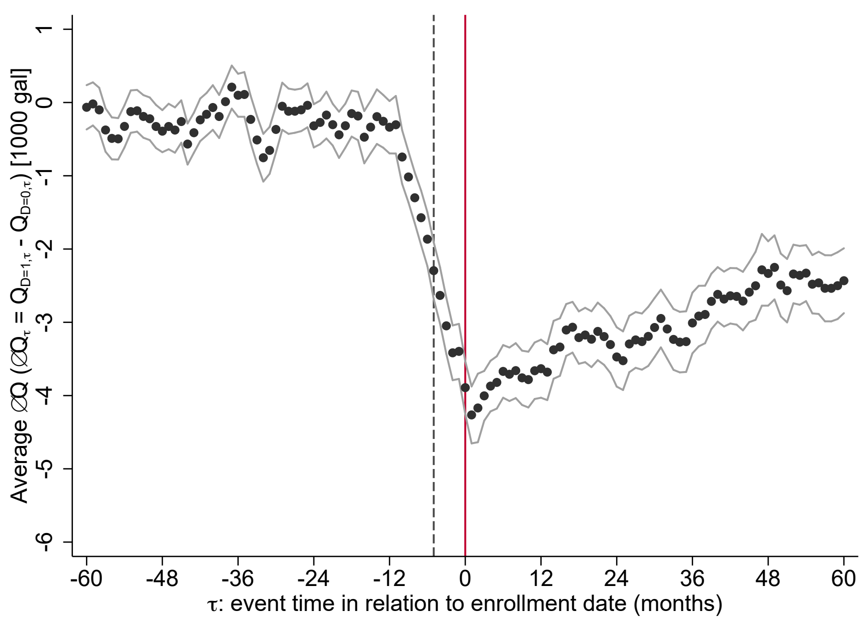

(b) *Event study w.r.t. enrollment date*

Figure A.2: *Event study based on the average difference in water use between treated and matched pairs, defining event time with respect to application date (a) and enrollment date (b). The vertical solid line indicates the month of application (a) or enrollment (b). The vertical dashed line indicates the average month of enrollment (a) or application (b).*

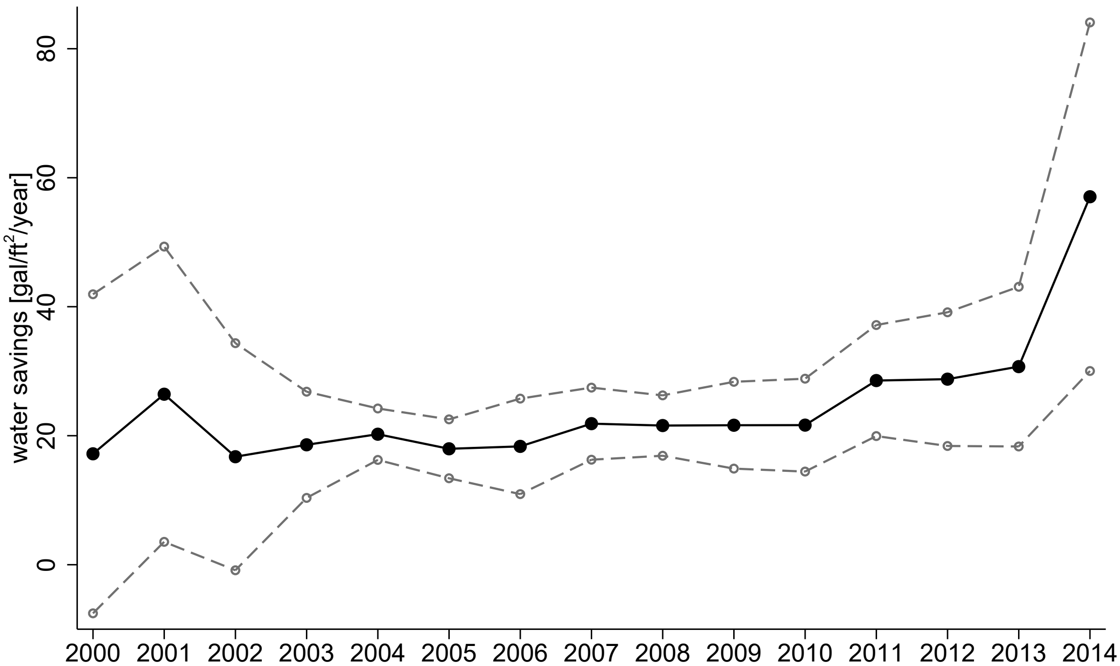

Figure A.3: *Average water savings achieved in each year of the program derived from my main specification model, Eq. (1.2), run separately for each enrollment year. In each year, my sample is comprised from pooling all participating parcels in that year and their respective matched non-participating parcels. Results are further normalized by the corresponding average annual conversion area. I derive 95 percent confidence intervals considering average converted area a fixed parameter.*

results to this concern, I estimate Eq. (1.2) separately for each enrollment year with a pooled sample of parcels enrolling in that year and their matched non-participating parcel pairs (with any duplicate non-participating parcels weighted accordingly by frequency). As in my main results, Figure A.3 shows that savings across years exhibit a 'U'-shaped pattern over time, though less pronounced.[152]

One may be similarly concerned that in deriving savings as a function of pre-enrollment consumption decile (Figure 1.10), comparing high (or low) decile converters with everyone else may bias results. Ideally, one would want to compare high decile converters with comparable non-participants. In Figure A.4, I implement such a procedure. I first derive pre-enrollment consumption deciles (based on a 12-month pre-enrollment average of water

[152]Of course, because of the wide confidence intervals, I cannot rule out the possibility of initially low savings, producing instead of a 'U'-shaped pattern a steady increase in savings across years.

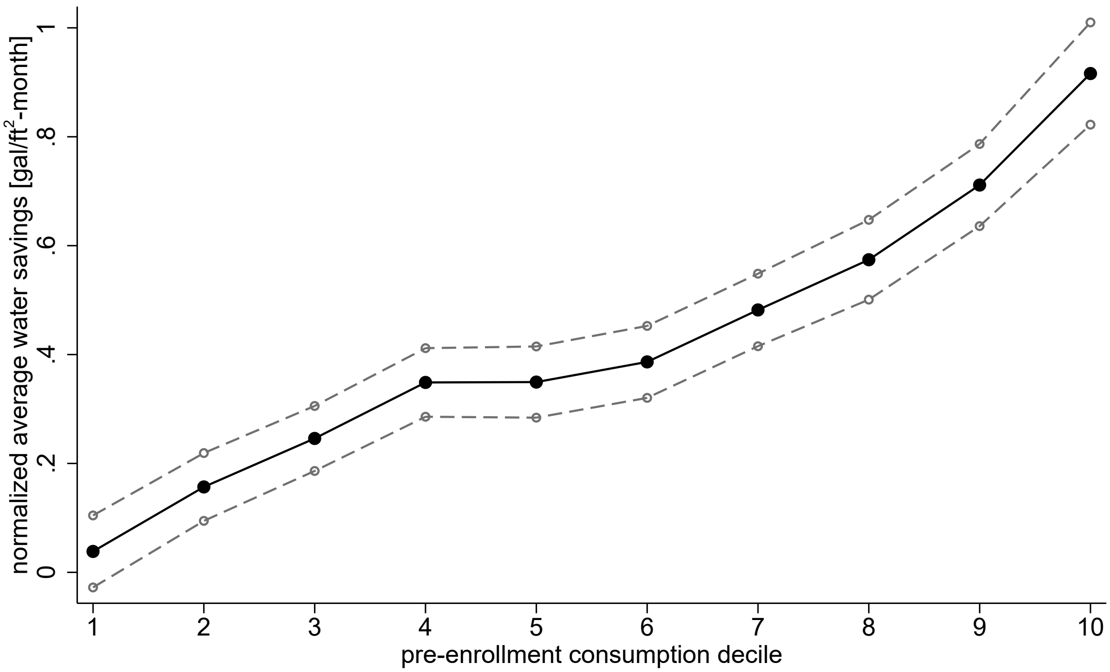

Figure A.4: *Average water savings per square foot of lot size achieved within each pre-enrollment water consumption decile. I derive point estimates and 95 percent confidence intervals by pooling the participating parcels with their matched non-participating parcels, and then estimating Eq. (1.2) separately for each decile. I normalize the dependent variable, Q_{it}, by 1000 ft^2 of lot size. Pre-enrollment consumption deciles are defined based on a 12-month average of water use per lot size beginning 24 months prior to the month of enrollment.*

use normalized by lot size), and then for each decile, pool the participating parcels with their matched non-participating parcels. With this pooled sample, I run Eq. (1.2) separately for each decile. Like my main results, Figure A.4 illustrates a positive relationship between savings and pre-enrollment consumption decile.[153]

[153]This result is robust to whether I define pre-enrollment consumption deciles with a 12, 24, 36, or 48-month average, and whether I estimate Eq. (1.2) with water use normalized by lot size, or simply water use, and then post-estimation normalizing by average conversion area for each given pre-enrollment consumption decile.

A.3 Additional hedonic results and robustness checks

A.3.1 Effect of two policy changes

In 2004, Las Vegas communities restricted new home construction from planting a front lawn.[154] This policy may increase the prevalence of desert landscape after 2004, and thereby effect how the market values properties that participated in the Cash-for-Grass program and/or the neighbors of these properties. I estimate the effect of the 2004 prohibition on front yards with Eq. (A.3), including an indicator for homes built in or after 2004 and interacting this indicator with P_{it} and N_{it}.

$$\begin{aligned}\ln p_{it} &= \psi[\text{post-2004}]_i + \alpha_1 DP_i + \beta_1 P_{it} + \beta_{1p04} P_{it} \times [\text{post-2004}]_i \\ &\quad + \alpha_2 DN_i + \beta_2 N_{it} + \beta_{2p04} N_{it} \times [\text{post-2004}]_i + \delta Z_i + b_{iq} + \epsilon_{it}\end{aligned} \tag{A.3}$$

Since June, 2009, rebate recipients must agree to maintain their conversions in perpetuity. Prior to this date, customers agreed to maintain the conversion for five, then ten years, with the agreement voided upon transfer of ownership.[155] By removing the option value that prospective buyers previously had regarding a pre-existing conversion, the June, 2009 policy may reduce the value of conversions. To explore the effect of this policy, I estimate Eq. (A.4), interacting indicators for enrollments and neighboring enrollments from June, 2009 onward with P_{it} and N_{it}, respectively.

$$\begin{aligned}\ln p_{it} &= \alpha_1 DP_i + \beta_1 P_{it} + \beta_{1jun09} P_{it} \times [\text{post-June '09}]_i \\ &\quad + \alpha_2 DN_i + \beta_2 N_{it} + \beta_{2jun09} N_{it} \times [\text{nbr post-June '09}]_i + \delta Z_i + b_{iq} + \epsilon_{it}\end{aligned} \tag{A.4}$$

Table A.2 shows that neither policy had any additional impact on the direct or spillover effect of conversion to desert landscape. In columns 1 and 2 I present results from estimating

[154] pers. comm. SNWA staff, March 14, 2016 and May 8, 2017.

[155] The November 2008 application states: "The converted area must remain in compliance with all program conditions for a period of ten years. This requirement is void upon transfer of ownership. You agree to return the incentive payment if this requirement is violated." The February 2012 application states: "Rebate is subject to owner's grant of a conservation easement that restricts certain uses of the conversion project areas in perpetuity." The specific language in the applications since June 2009 has changed slightly. The June 2009 and September 2010 applications refer to the agreement as a "restrictive covenant" rather than a conservation easement.

Table A.2: *Regression results for the direct and spillover effect of conversion to desert landscape on home property values and the additional effects of policy changes.*

	(1)	(2)	(3)	(4)
post-2004	-0.046			
	(0.0093)***			
DP (ever converts)	0.0043		0.0048	
	(0.0013)***		(0.0013)***	
Direct effect	0.011	0.017	0.010	0.014
	(0.0035)***	(0.013)	(0.0035)***	(0.014)
Direct×post-2004	-0.0040	0.034		
	(0.015)	(0.028)		
Direct×post-June '09			0.013	0.030
			(0.0095)	(0.024)
DN (neighbors DP)	0.0011		0.0016	
	(0.0010)		(0.0010)	
Spillover effect	-0.0021	-0.016	-0.0016	-0.0067
	(0.0023)	(0.0084)*	(0.0024)	(0.0087)
Spillover×post-2004	-0.010	0.024		
	(0.0095)	(0.020)		
Spillover×nbr post-June '09			-0.0016	-0.014
			(0.0055)	(0.013)
Sale years	1996-2014	1996-2014	1996-2014	1996-2014
Fixed-effects				
quarter-block	yes	yes	yes	yes
parcel	-	yes	-	yes
adj. R^2	0.95	0.93	0.95	0.93
Observations	199,037	40,755	199,037	40,755

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Standard errors clustered at the block level.

blocks: 2010 United States Census Block boundaries.

quarter: quarter of sample (e.g. 1^{st} quarter of 1997 is quarter 5).

2014 adjusted sale prices trimmed at the 1^{st} and 99^{th} percentiles.

Sample excludes parcels undertaking additions (see section 1.4.2).

Eq. (A.3). Both columns include the full sample, and column 2 additionally includes parcel fixed-effects. I cluster standard errors at the Census block level. The statistically insignificant estimates on the interaction of the direct and spillover effect with the post-2004 indicator demonstrate that the front lawn restriction does not affect the direct or spillover value of desert landscape, though the restriction does appear to reduce the overall value of a home by about 5 percent (row 1). In columns 3 and 4 I present results from estimating Eq. (A.4). Both columns include the full sample, and column 4 additionally includes parcel fixed-effects. I again cluster standard errors at the Census block level. The statistically insignificant estimates on the interaction of the direct and spillover effect with their respective indicators for enrollments (or neighboring enrollments) after June 2009 demonstrate that the requirement to keep the conversion in place in perpetuity does not affect the direct or spillover value of desert landscape.[156] And while not the primary coefficient of interest in this exercise, the estimates of the direct and spillover effect demonstrate consistency with those presented in Table 1.6. The one exception is the negative and statistically significant estimate of the spillover effect in column 2. This result raises the possibility that negative spillovers from desert landscape do exist, however the weight of the evidence presented in my analysis points to no spillover effects.

A.3.2 Heterogeneous effects across time

The Cash-for-Grass program has been in place for nearly 20 years, making it reasonable to expect that characteristics of residents changed in ways that impact the value of desert landscape. To estimate the impact of changing consumer preferences, I interact year of sale indicators (yos_j) with P_{it} and N_{it}, as shown in Eq. (A.5). The vector of β_{1j}'s and β_{2j}'s

[156]While not shown, the coefficient estimates on each covariate effect housing prices in expected ways. The two exceptions involve negative coefficient estimates on full bathrooms in columns 1 and 3, significant at the 1 percent level, and the negative coefficient on half bathrooms in column 1 and 3, significant at the 5 percent level. Toilets make up the largest share of indoor water use (Bennear *et al.*, 2013), and the negative coefficient on bathrooms may reflect consumers' recognition of higher water bills associated with an increased number of water-intensive fixtures.

describe the *annual* direct or spillover effect of conversion to desert landscape, respectively.

$$\begin{aligned} \ln p_{it} = \alpha_1 DP_i + \sum_{j=1996}^{2014} \beta_{1j}(yos_j \times P_{it}) \\ + \alpha_2 DN_i + \sum_{j=1996}^{2014} \beta_{2j}(yos_j \times N_{it}) + \delta Z_i + b_{iq} + \epsilon_{it} \end{aligned} \tag{A.5}$$

Figure A.5 illustrates the direct and spillover effect of Cash-for-Grass subsidized conversions to desert landscape in each year. Though imprecise, the estimates show little noticeable pattern over time, generally fluctuating around the point estimate (solid horizontal line) for the overall average direct or spillover effect shown in Table 1.6. Furthermore, the point estimate of the overall average effect generally falls within the confidence intervals of the annual estimates. These results suggest that the effect of conversion to desert landscape under the Cash-for-Grass rebate program has remained stable over time.[157]

A.3.3 Overlap of covariate distributions

Table 1.4 shows very different mean age and lot size between participating or neighboring parcels prior to sale ($P = 1$ and $N = 1$) and parcels not participating or neighboring a conversion prior to sale ($P = 0$ and $N = 0$). The imbalance raises a concern regarding distributional overlap. In the following figures, I illustrate the distributional overlap for age, lot size, and the outcome variable, price. I conclude that the distributions overlap for a substantial portion of each variables' domain.

Age Figure A.6 demonstrates that new homes comprise the majority of non-participating or non-neighboring parcels. The distributions of participating and non-participating parcels, and neighboring and non-neighboring parcels, however, overlap throughout the majority of the distributions' domain. While not shown, I observe a similar distributional overlap for the repeat sales model and the model that restricts sales to pre-2007.

[157] Though not shown, figures derived from estimating Eq. (A.5) with parcel fixed-effects show a similar pattern.

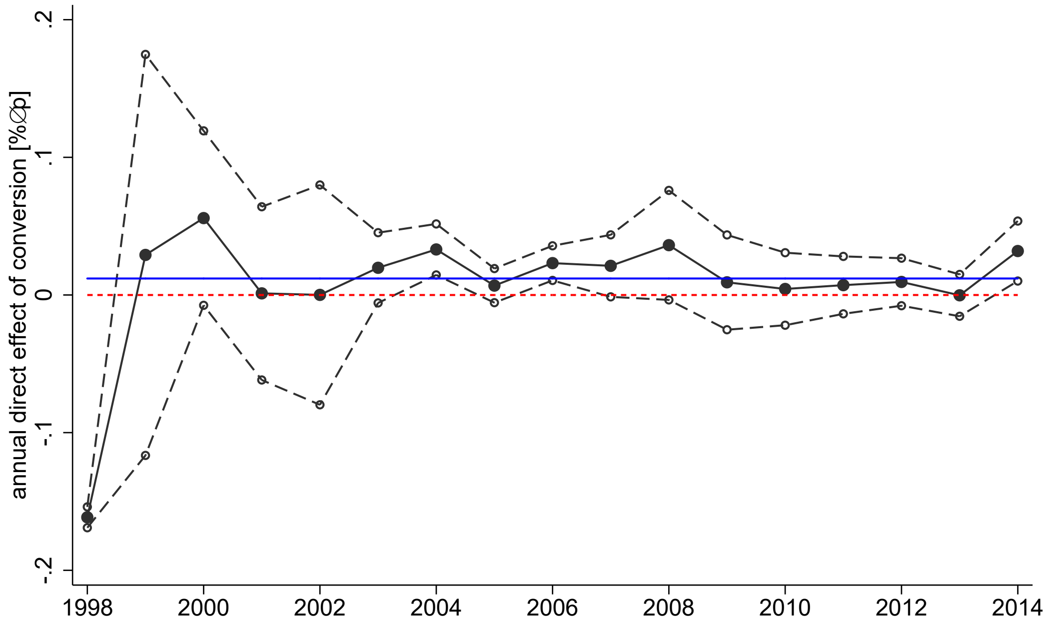

(a) *Annual direct effect*

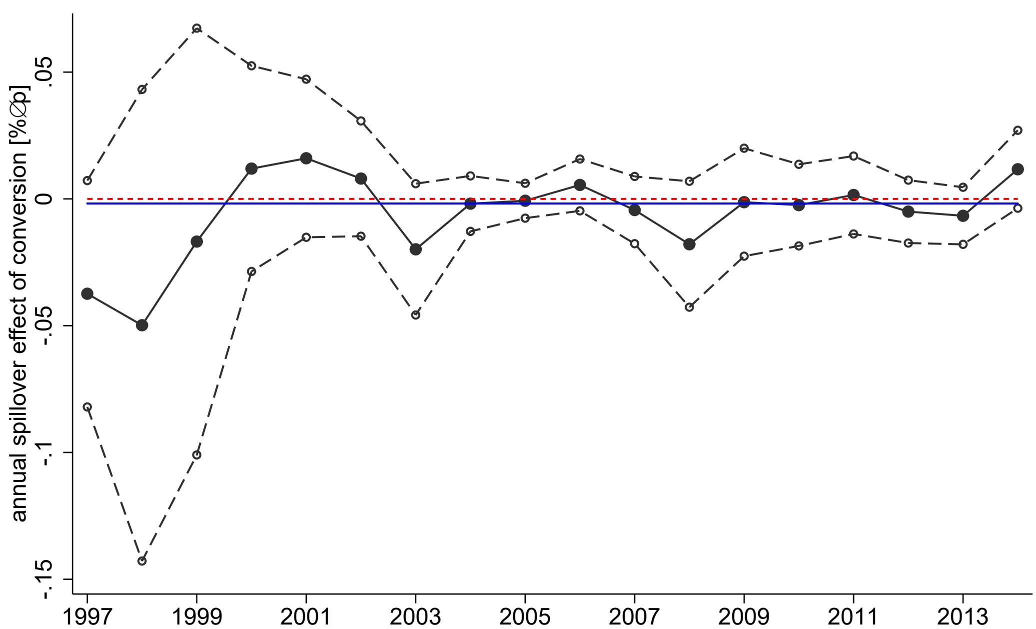

(b) *Annual spillover effect*

Figure A.5: *Annual effect of conversion to desert landscape. Point estimates and 95 % confidence intervals for the direct effect (panel a) and spillover effect (panel b) of Cash-for-Grass subsidized conversion to desert landscape in each year. Results derived from estimating Eq. (A.5). The solid blue line represents the point estimate from Table 1.6, and the dotted red line highlights zero on the y-axis (i.e. no effect). The coefficients for 1996 and 1997 in panel (a) and for 1996 in panel (b) are dropped due to collinearity.*

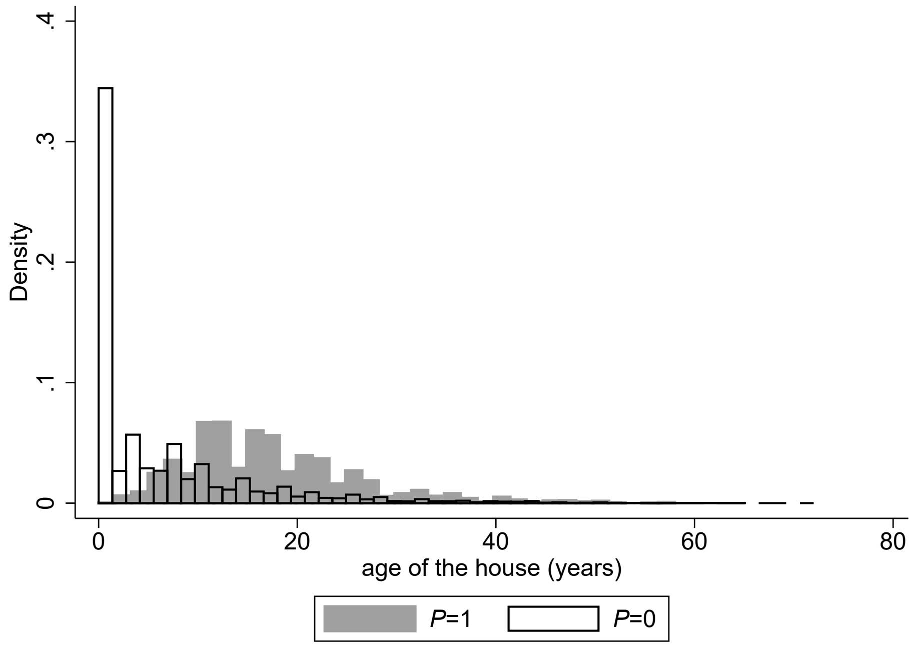

(a) *Participants vs. non-participants*

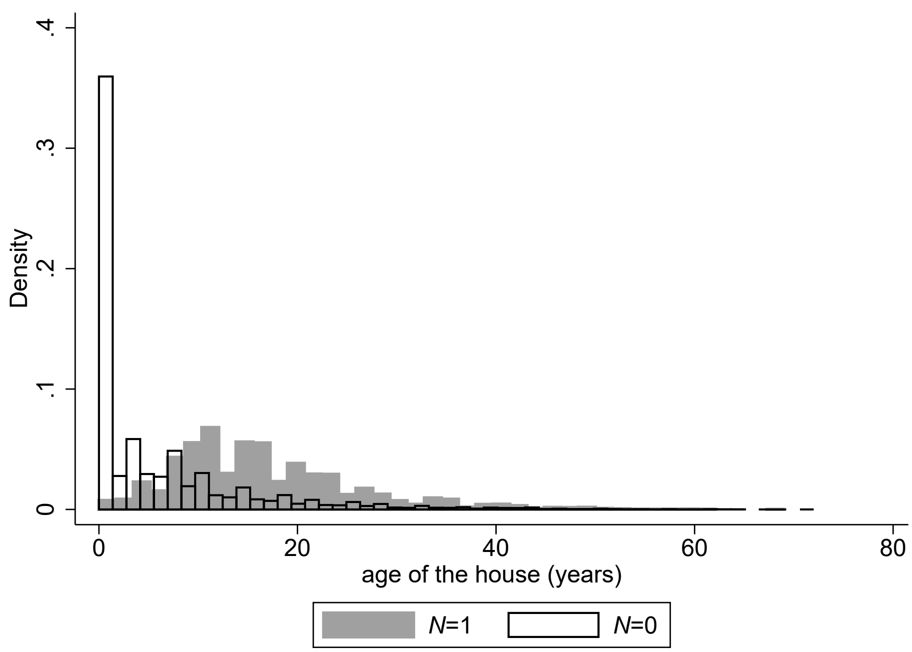

(b) *Neighbors vs. non-neighbors*

Figure A.6: *Density distribution of home age for the model including quarter-block fixed-effects and sale years 1996-2014.*

Lot Area Figure A.7 shows a fairly high degree of overlap between the lot size distributions of participants and non-participants, and between the lot size distributions of neighbors and non-neighbors. This conclusion holds for the repeat sales model and the model that restricts sales to pre-2007.

Price Finally, I compare the distribution of the sale price (in levels) for participants and non-participants, as well as for neighbors and non-neighbors. Figure A.8 illustrates substantial distributional overlap. Though not shown, the same can be said of the repeat sales model and the model that restricts sales to pre-2007.

A.3.4 Robustness to alternative specifications

In this section, I test the sensitivity of my estimates to fixed-effect specifications and other restrictions I placed on the data in section 1.4. For each sensitivity analysis, I hold constant with the main specification all characteristics of the model net of the characteristic under investigation.

Robustness to fixed-effects Quarter-block and quarter-block & parcel fixed-effects control quite flexibly for fixed or varying unobserved neighborhood characteristics, but ask a lot of the data. Figure A.9 illustrates the robustness of the estimates of the direct effect (panel A) and spillover effect (panel B) for the full range of sales (1996-2014) to three sets of fixed-effects. Quarter & block and quarter & parcel fixed-effects absorb average unobserved dynamics in home prices across the LVVWD, and average block or parcel effects. Quarter & block-year and quarter & block-year & parcel fixed-effects absorb average unobserved dynamics across the LVVWD, average block effects in each year of the sample, and average parcel effects for the specification that additionally includes parcel fixed-effects. Finally, quarter-block and quarter-block & parcel fixed-effects, the main specification, control the most flexibly for changes across time and space.

Figure A.9 generally illustrates consistency across fixed-effect specifications. The two exceptions are the quarter & block and quarter & parcel fixed-effect specifications that illus-

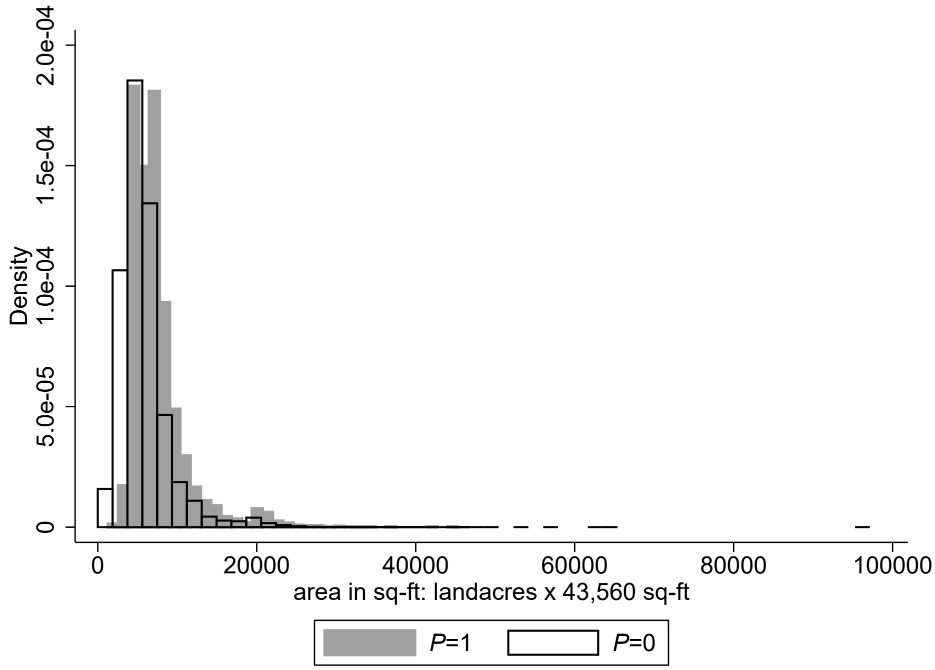

(a) *Participants vs. non-participants*

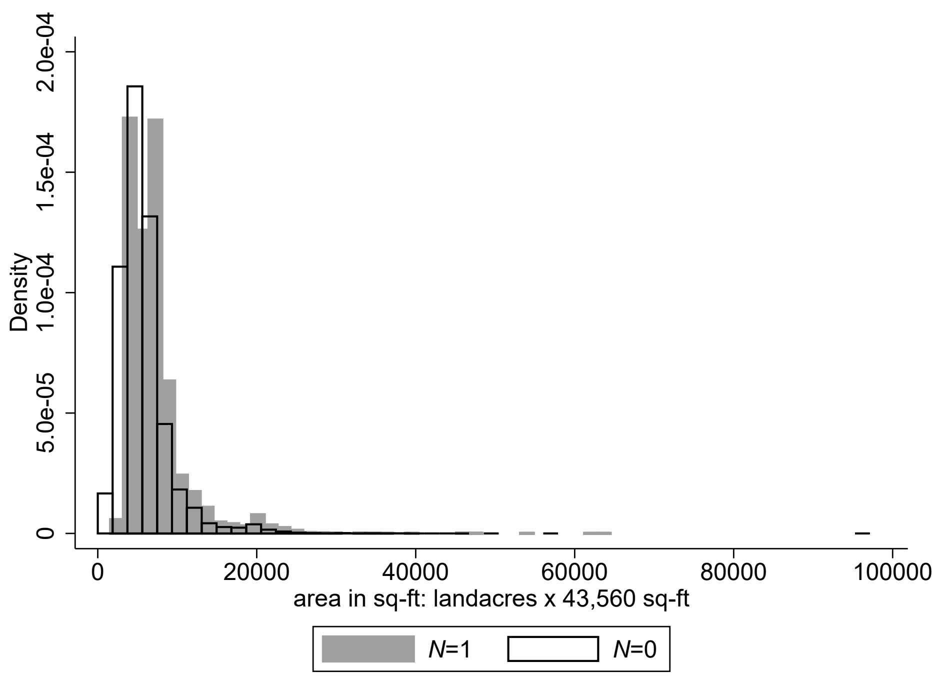

(b) *Neighbors vs. non-neighbors*

Figure A.7: *Density distribution of lot size for the model including quarter-block fixed-effects and sale years 1996-2014.*

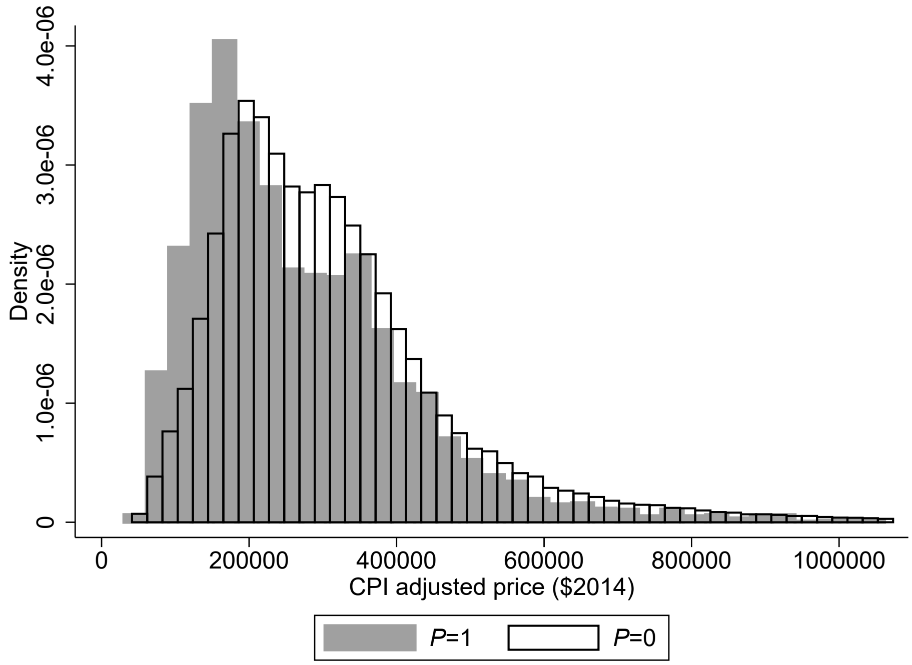

(a) *Participants vs. non-participants*

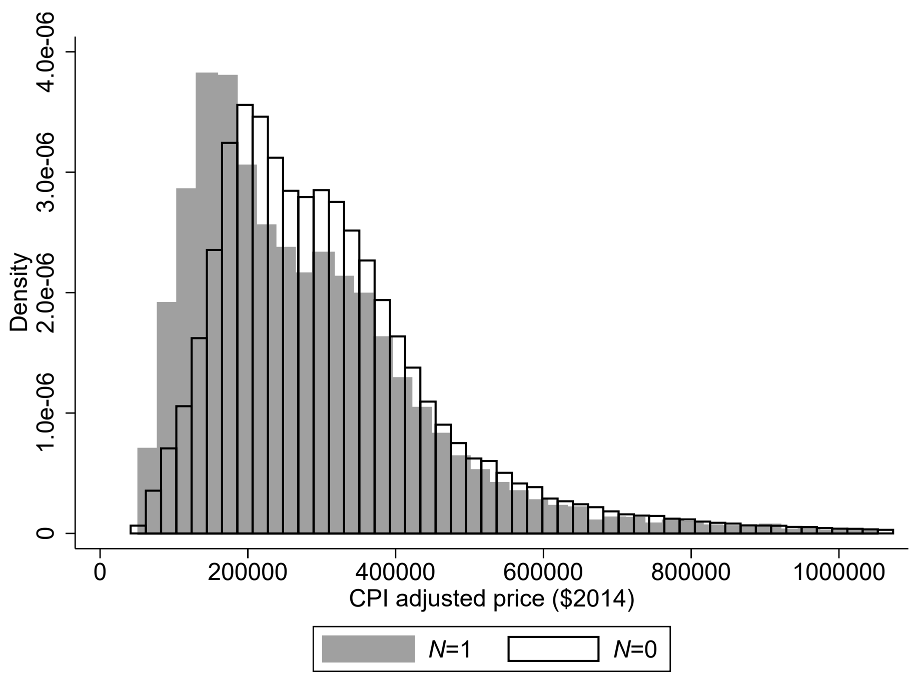

(b) *Neighbors vs. non-neighbors*

Figure A.8: *Density distribution of sale price for the model including quarter-block fixed-effects and sale years 1996-2014.*

trate a positive estimate for the spillover effect. As the least flexible fixed-effect specification, however, these estimates could be biased by unobserved changes across time at the block level. The weight of the evidence still suggests no net spillovers.

Robustness to additions Figure A.10 shows the point estimates and 95 % confidence intervals for estimates with and without the parcels that meet my criteria for undertaking an addition. The estimates do not appear sensitive to the inclusion or exclusion of parcels undertaking additions. To the extent that my addition criteria misses parcels undergoing major structural changes, Figure A.10 suggests any such missed parcels will have only a small impact on my results.

Vacant parcel sales Beginning in 2005, the assessor's office distinguishes sales by the vacancy status of the parcel. I test the sensitivity of my estimates to removing all sales not indicated as "improved" (i.e. not vacant). Dropping such parcels removes all sales of vacant properties after 2005 and nearly all sales of all properties prior to 2005. Since the assessor data designate about 88 percent of post-2004 sales as improved, it is likely that many of my pre-2005 dropped sales are sales of non-vacant properties. However, I have no way of determining with certainty the vacancy status of these pre-2005 sales. Figure A.11 shows the results from dropping all sales not designated as "improved". While I lose some precision when I keep only sales of non-vacant properties, the results appear consistent across vacancy status.

Robustness to data trimming Figure A.12 illustrates the point estimates and 95 % confidence intervals for estimates with no variables trimmed, all variables trimmed (price and all control variables) at their 1^{st} and 99^{th} percentiles, and only the price variable trimmed at the 1^{st} and 99^{th} percentile (the main specification). Though I lose some precision when I do not trim the distribution of any variables, the figure illustrates that estimates are robust to the choice of data trimming. Trimming attempts to prevent outliers from driving results. My estimates here suggest that results are generally robust to outliers.

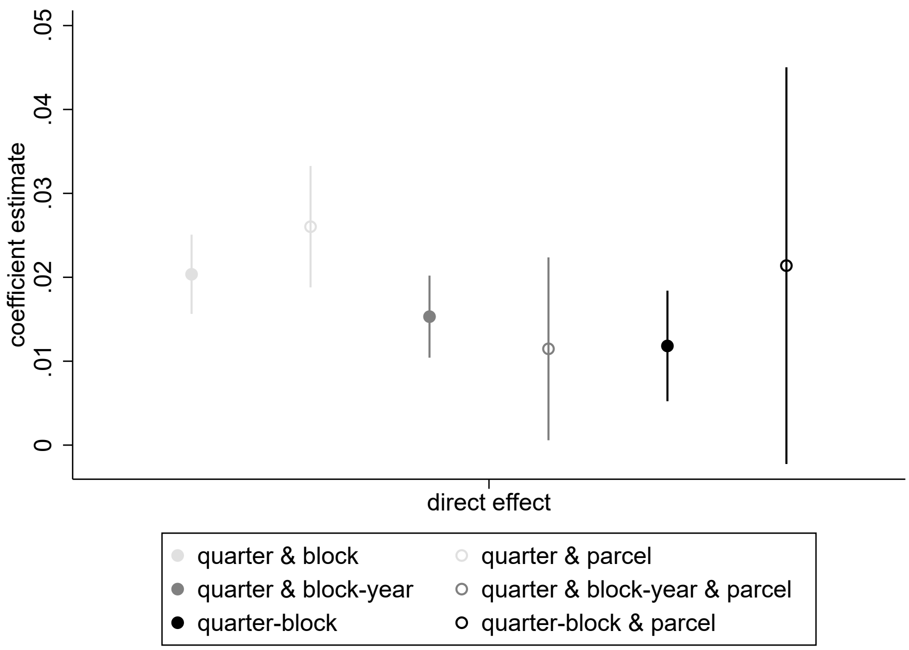

(a) *Direct effect*

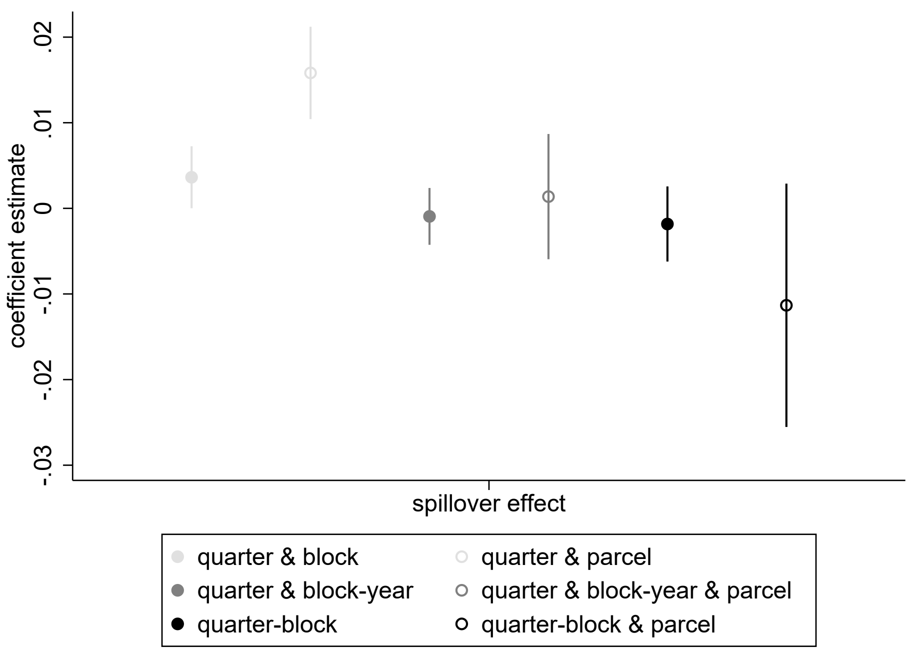

(b) *Spillover effect*

Figure A.9: *Robustness of the direct (A) and spillover (B) effect to various fixed-effect specifications. Point estimates and 95 % confidence intervals shown. Sales range from 1996 through 2014. The rightmost two points (black) represent the main specification.*

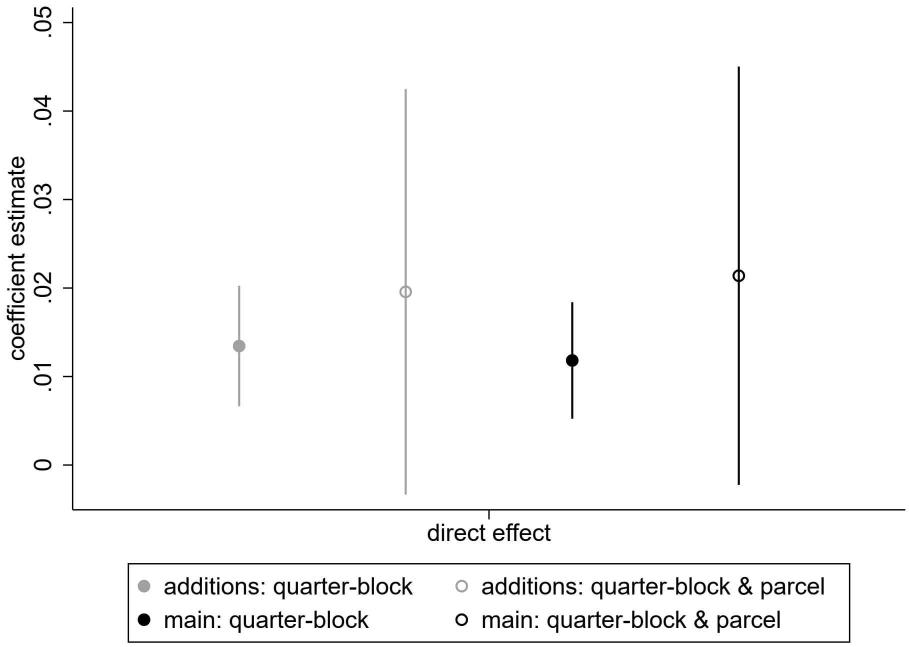

(a) *Direct effect*

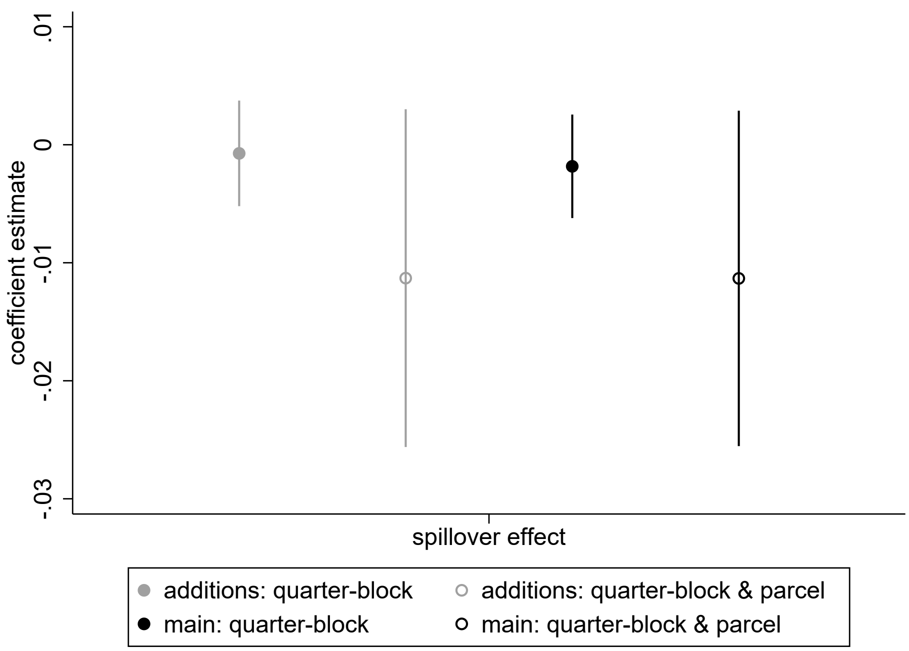

(b) *Spillover effect*

Figure A.10: *Robustness of the direct (A) and spillover (B) effect to including parcels that meet criteria for having undertook an addition. Point estimates and 95 % confidence intervals shown. Sales range from 1996 through 2014. The rightmost two points (black) represent the main specification.*

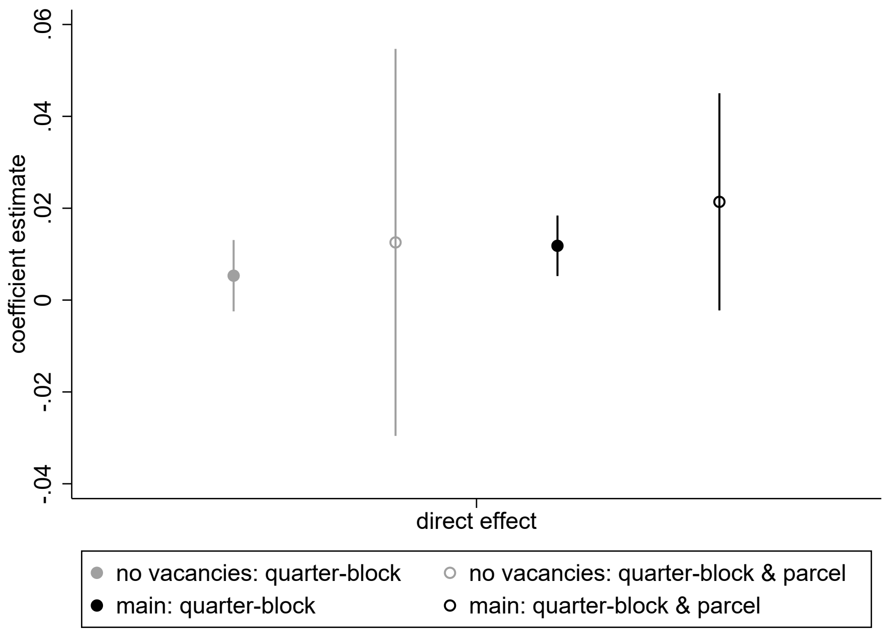

(a) *Direct effect*

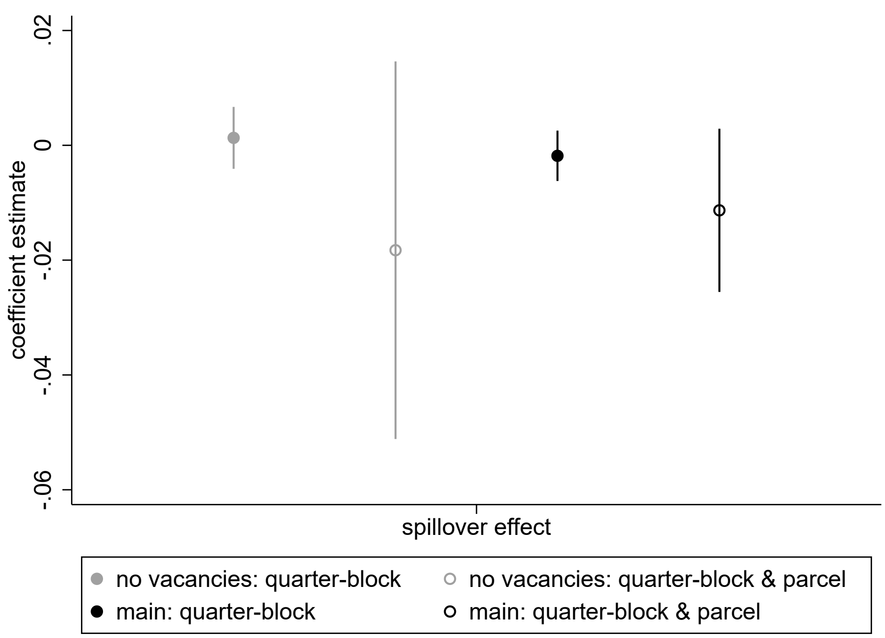

(b) *Spillover effect*

Figure A.11: *Robustness of the direct (A) and spillover (B) effect to the designation of vacant or improved parcel sales. Point estimates and 95 % confidence intervals shown. For the estimates that include only sales of improved properties, dates range from 2005 through 2014 (though a very few sales are designated as improved prior to 2005). For the estimates that include all sales regardless of the sale designation (the main specification), sales range from 1996 through 2014. The rightmost two points (black) represent the main specification.*

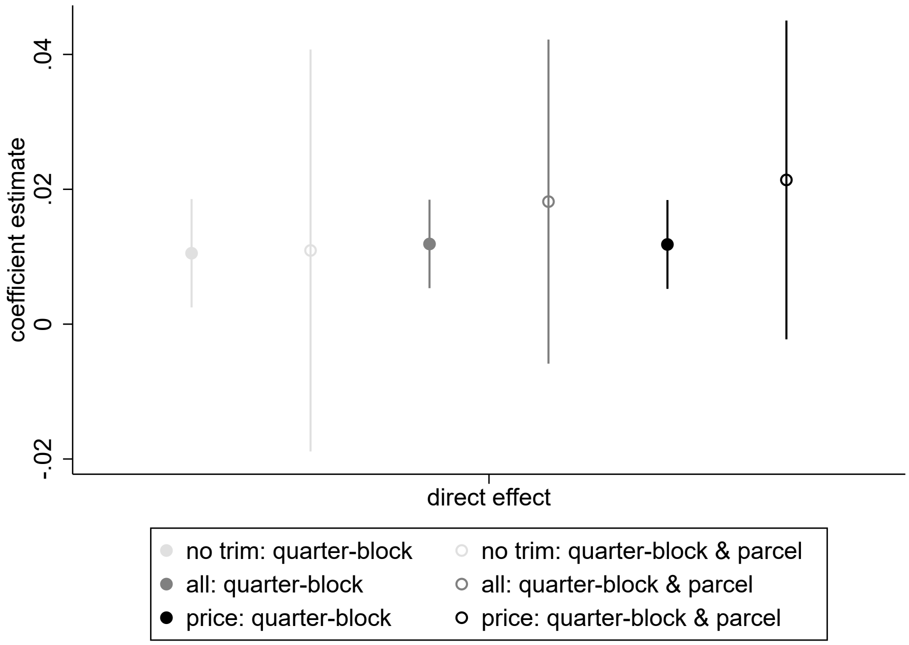

(a) *Direct effect*

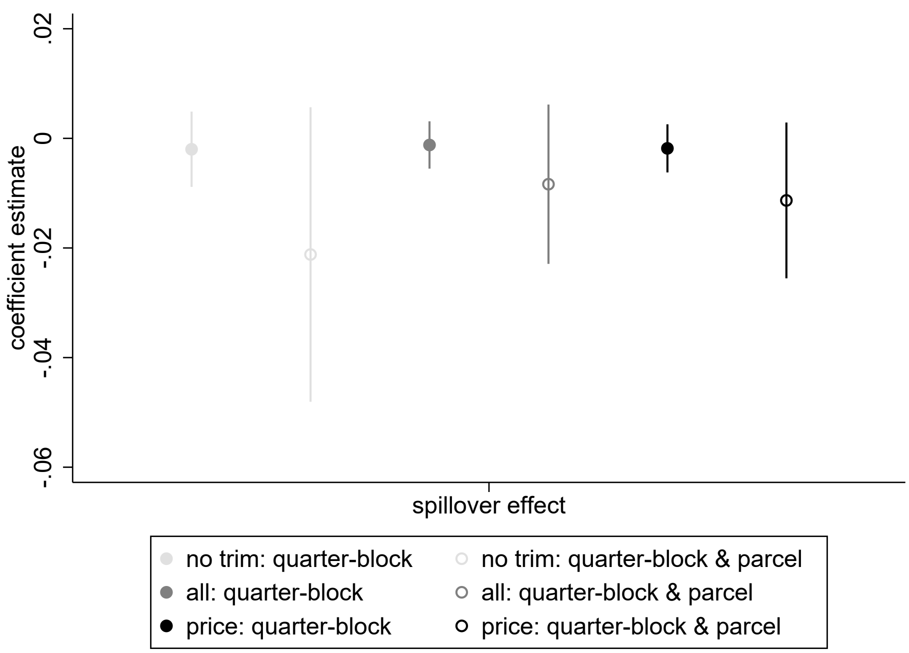

(b) *Spillover effect*

Figure A.12: *Robustness of the direct (A) and spillover (B) effect to how many variables in the model are trimmed at the 1st and 99th percentiles. Point estimates and 95 % confidence intervals shown. Sales range from 1996 through 2014. The rightmost two points (black) represent the main specification.*

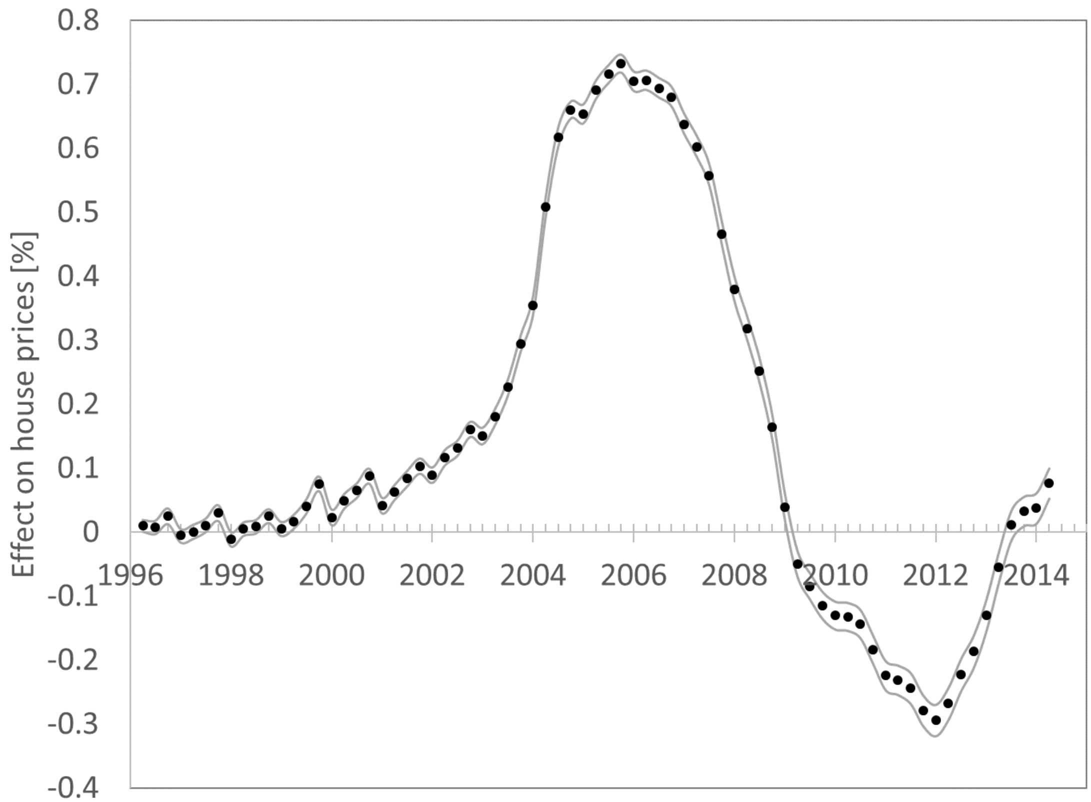

Figure A.13: *Plot of point estimates and 95 % confidence intervals for quarter of sample fixed-effects (qrt) estimated from the following model (where blk refers to census block fixed-effects):* $\ln p_{it} = \alpha_1 DP_i + \beta_1 P_{it} + \alpha_2 DN_i + \beta_2 N_{it} + \delta Z_i + blk + qrt + \epsilon_{it}$ *(which I estimate using `areg`; all other hedonic models I estimate using `reghdfe` (Correia, 2016)). Point estimates are relative to the first quarter (i.e. quarter 1 in 1996).*

Robustness to choice of pre-Crash date The housing market crash hit Las Vegas especially hard. Figure A.13 illustrates point estimates and 95 % confidence intervals of quarterly fixed-effects for a model akin to Eq. (1.6) that includes census block and quarter of sample fixed-effects. The figure clearly illustrates the housing bubble in the Las Vegas valley. Out of a concern that my quarter-block fixed-effects do not completely absorb all the effects of the housing crisis, in my main specification I additionally include models that limit sales to pre-housing crisis years. I choose pre-2007 sales since the most precipitous drop illustrated in Figure A.13 occurs after 2006. In Figure A.14, I illustrate the robustness of this selection by further estimating models that limit sales to pre-2006 and pre-2008. Specifically, Figure A.14 shows the point estimates and 95 % confidence intervals for estimates derived from sale years 1996 to 2005, 1996 to 2006 (my main specification), and 1996 to 2007. I only estimate models with quarter-block fixed-effects; including parcel fixed-effects severely reduces the precision of the estimates since so few data exist in the samples limited to pre-crash years. Results appear robust to the choice of the beginning of the housing crisis.

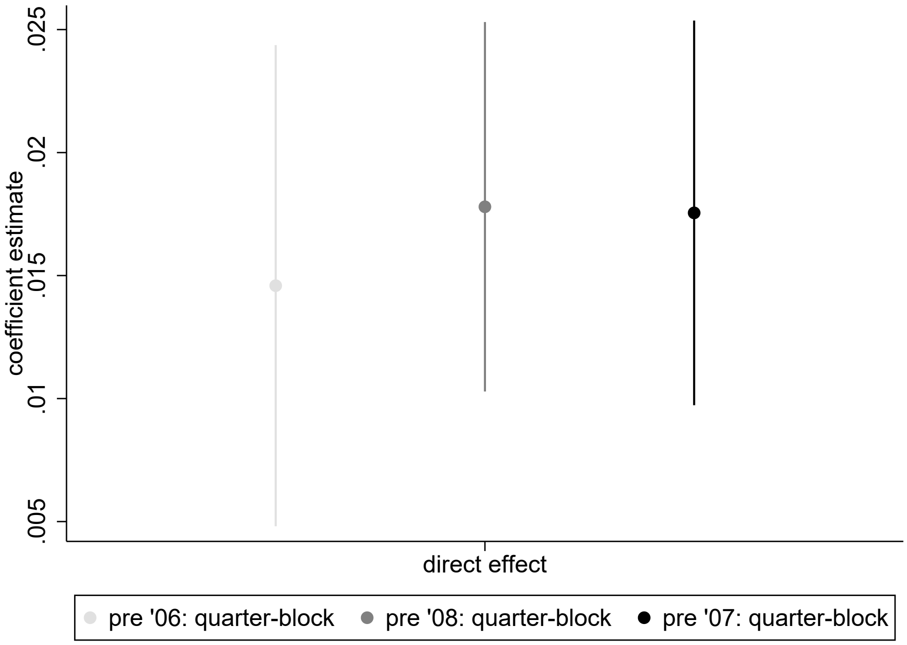

(a) *Direct effect*

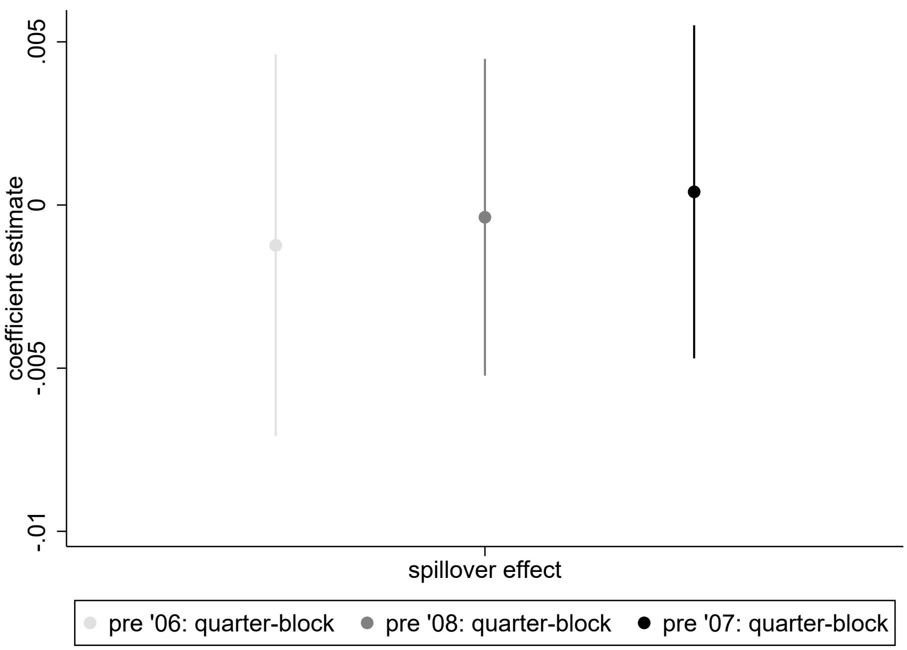

(b) *Spillover effect*

Figure A.14: *Robustness of the direct (A) and spillover (B) effect to the choice of the pre-housing market crash period. Point estimates and 95 % confidence intervals shown. Sales range from 1996 through 2005, 2006 or 2007. The rightmost two points (black) represent the main specification.*

A.4 Calculation details of $/kgal-saved and net benefits

In this appendix, I provide further discussion regarding the details of my calculation of the annual cost per gallon saved and estimate of net benefits.

Additional details regarding annual cost per gallon saved

- **Total rebate outlays:** I assume rebate totals reported by the water authority are nominal to the year the rebate was granted (i.e. a rebate administered in 2007 would be recorded in 2007 dollars). I therefore adjust rebates to reflect 2014 dollars using the CPI index for all urban consumers (Bureau of Labor Statistics Series Id: CUUR0000SA0).

- **Financing costs:**[158] I ignore financing costs. Up until 2009, the rebate was funded through one-time connection charges applied to new service meters. These new connection fees would introduce little market distortion, and therefore negligible additional cost. Starting in 2009, the rebate was funded through bond measures. But the bond issue for the 2015/2016 fiscal year could be paid off entirely with only a small percentage increase in water rates. Since the water authority would pay off a bond over many years, I consider costs due to paying off bonds to be small. For these reasons, I ignore financing costs in my analysis (both in the estimate of cost per gallon saved and my estimate of net benefits).

- **Average water bill:** I calculate the water bill per 1000 gallons for an average LVVWD customer in 2013. Water charges depend upon the meter size and in 2013, over 99 percent of single-family LVVWD customers in my panel have a 1 inch, 3/4 inch or 5/8 inch meter. Using a bill calculator provided by the water authority I estimate the annual water bill per 1000 gallons for each meter size (1", 3/4" and 5/8") and calculate a weighted average water bill, with weights defined by the share of customers associated with each meter size (see Appendix A.1). The weighted average annual

[158] I am grateful to Joe Aldy for a helpful discussion on this point.

water bill equals \$3.54/kgal. Note that I assume 2013 dollars are equivalent to 2014 dollars.

- **Out-of-pocket conversion costs:** The average rebate from the 26,488 conversions that make up my water savings panel falls just under \$1,996. Since the average conversion size equals 1,348 ft^2, the average rebate per square-foot equals about \$1.50/ft^2. I assume total conversion costs equal \$3/ft^2, implying that the out-of-pocket expenditure for an average rebate recipient equals about \$1.50/ft^2. Using the average conversion size (1,348 ft^2) and the total number of conversions (26,488), I estimate that total out-of-pocket expenditures equal about \$54M. I add this to total costs to the utility reported in section 1.5 (\$65M), re-compute annualized cost using a 30-year time horizon and a 5 percent discount rate, and divide by total annual savings (1.6M kgal/year). The resulting program costs come to \$4.84/kgal-saved.

- **Estimate of the opportunity cost of scarce water:** Edwards and Libecap (2015) report that in the Truckee river basin, "the median price of 1,025 agriculture-to-urban water rights sales between 2002 and 2009 (2008 prices) was \$17,685/Acre Foot (AF)". I adjust this value to 2014 dollars,[159] and then divide by 325,851 gallons per acre-foot, resulting in the estimate of \$0.06/gal reported in section 1.5. Using sales occurring in Nevada would seem to best approximate the value of water for a Nevada water utility, and for a municipal water utility, agriculture to urban sales is a more relevant proxy for the value of water than what would be reflected in intra-agricultural sales.

Additional details regarding net benefits To be precise,[160] benefits associated with converting to desert landscape include the private benefits to the household reflected in the hedonic estimates, the scarcity value of water, as well as reduced operating and maintenance costs associated with lower delivery requirements (and any positive externalities which I ignore for the purposes of this discussion). Costs include conversion costs, and reduced

[159] http://data.bls.gov/cgi-bin/cpicalc.pl

[160] Many thanks to Nick Hagerty for the insights into net benefits contained in this discussion.

revenue for the utility (which should equal the benefits consumers receive from lower water bills). However, most water utilities endeavor to price water such that revenues cover costs. Thus, the benefits from reduced operating and maintenance costs should be approximately equal to lost revenue. In my analysis in section 1.5, I implicitly make this assumption. But because the SNWA has other sources of revenue, it may be that operating costs exceed water bill revenue, leading me to understate net benefits.[161]

[161]To make matters more complicated, the SNWA does not actually distribute water to customers. The LVVWD and other Las Vegas area water districts supply tap water to residents' homes, buying treated water wholesale from the SNWA. In my analysis, I have assumed that the LVVWD and the SNWA are essentially one financial entity.

Appendix B

Appendix to Chapter 2

B.1 Regression discontinuity design

As a robustness check on our differences-in-differences approach, we apply a regression discontinuity design (RD) to test for effects of the publishing, mailing, and online posting requirements. Because the policy affects water systems differently based on size, and because the number of customers served by water systems varies more or less continuously, the information disclosure policy shock provides an ideal setting for applying an RD.

To motivate our RD analysis, consider average violations as a function of system size in 1998 illustrated in Figure 2.2. A clear drop in violations on the right side of the threshold (blue vertical lines in the figure) compared to the left would provide evidence in the raw data that an impact due to the information disclosure policy exists. Figure 2.2 does not illustrate such a drop in violations. In fact, average violations increase immediately after each threshold. And while average violations in 1998 fall sharply in the second bin to the right of each disclosure threshold, the same appears to be true for average violations in 1997. At the very least, Figure 2.2 weakens the expectation we might have of finding an effect from an RD, and due to the noise in the data illustrated by the figure, also implies that results of any RD may be highly sensitive to functional form.

In each year of our sample we implement a linear and quadratic parametric RD, corresponding to Eq. (B.1) and Eq. (B.2) respectively.[162] In each model, v_i describes health-based water quality violations from water system i for a given year in our sample, T_i refers to the information disclosure treatment category, sp_i corresponds to the service population of water system i, and k refers to information disclosure service population cutoff corresponding to the publishing, mailing, or online posting requirement.

$$v_i = \beta T_i + \alpha \left(sp_i - k\right) + \gamma \left(sp_i - k\right) T_i + \epsilon_i \tag{B.1}$$

$$\begin{aligned} v_i = \beta T_i + \alpha_1 \left(sp_i - k\right) + \alpha_2 \left(sp_i - k\right)^2 \\ + \gamma_1 \left(sp_i - k\right) T_i + \gamma_2 \left(sp_i - k\right)^2 T_i + \epsilon_i \end{aligned} \tag{B.2}$$

In the two models above, the estimate on T describes the average effect of the information disclosure requirement at the cutoff k. If $k = 501$, the estimate on T describes the effect of the publishing requirement. For $k = 501$, $T_i = 1$ for water systems serving more than 500 customers. If $k = 10,000$, the estimate on T describes the effect of the mailing requirement. For $k = 10,000$, $T_i = 1$ for water systems serving at least 10,000 customers. Finally, if $k = 100,000$, the estimate on T describes the effect of the online posting requirement. For $k = 100,000$, $T_i = 1$ for water systems serving at least 100,000 customers.

Figure B.1 illustrates results from estimating the linear RD model, Eq. (B.1), in each year of our sample. There appears to be little impact of being on either side of the publishing threshold prior to the advent of the information disclosure policy (pre-1998),[163] and a noticeable drop in violations after the policy came into effect in 1998. And while the magnitude of the violation reductions shrinks in each year after 1998, we cannot rule out the possibility that reductions remain stable given the confidence intervals around our estimates. Furthermore, our preferred differences-in-differences estimate of the reduction in violations at the publishing threshold (indicated by the black dashed line in Figure B.1) falls within the confidence intervals of our RD results in each year after 1998. Our RD results therefore

[162] Gelman and Imbens (2014) recommend against using higher order polynomials.

[163] The decrease in violations in 1997 may suggest some anticipatory behavior on the part of water systems, though the impact in 1997 is nearly indistinguishable from zero.

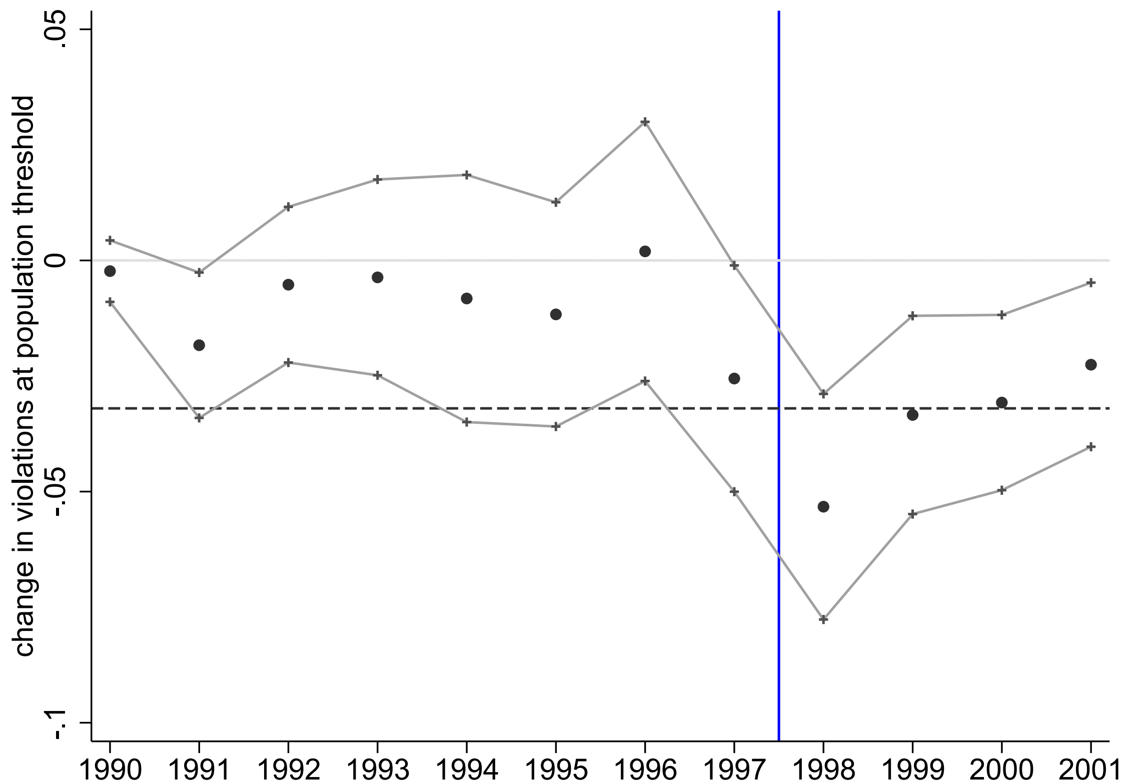

Figure B.1: *The effect of the publishing requirement considering a linear regression discontinuity. The figure shows estimates and associated 95 percent confidence intervals of T from Eq. (B.1) with* $k = 501$*. We estimate Eq. (B.1) separately for each year in our sample. The solid horizontal light gray line indicates zero effect, the dashed black line illustrates our preferred point estimate from our differences-in-differences model, and the blue vertical line delineates the post-information disclosure policy period.*

validate the reductions we estimate in our differences-in-differences models, and at least weakly support our conclusion that reductions in violations remain stable over time. We also estimate (but do not show) a quadratic RD model, Eq. (B.2), and observe a nearly identical pattern to what we observe in Figure B.1. The stability of the estimates across the linear and quadratic specifications demonstrates the robustness of the RD results, which further validates our differences-in-differences estimates of the effect of the publishing requirement.

We also test the robustness of the RD results at the publishing threshold by limiting the sample to those water systems serving between 25 and 9,999 customers and re-estimating Eq. (B.1) and Eq. (B.2). This limited set of water systems faces only the publishing information disclosure threshold and therefore avoids any bias stemming from including water systems facing a different disclosure threshold. While not shown, the results from re-estimating Eq. (B.1) with the limited sample demonstrate the same general pattern illustrated in Figure B.1. However, estimating Eq. (B.1) and Eq. (B.2) with the limited sample produces imprecise results, perhaps due to the smaller sample size.

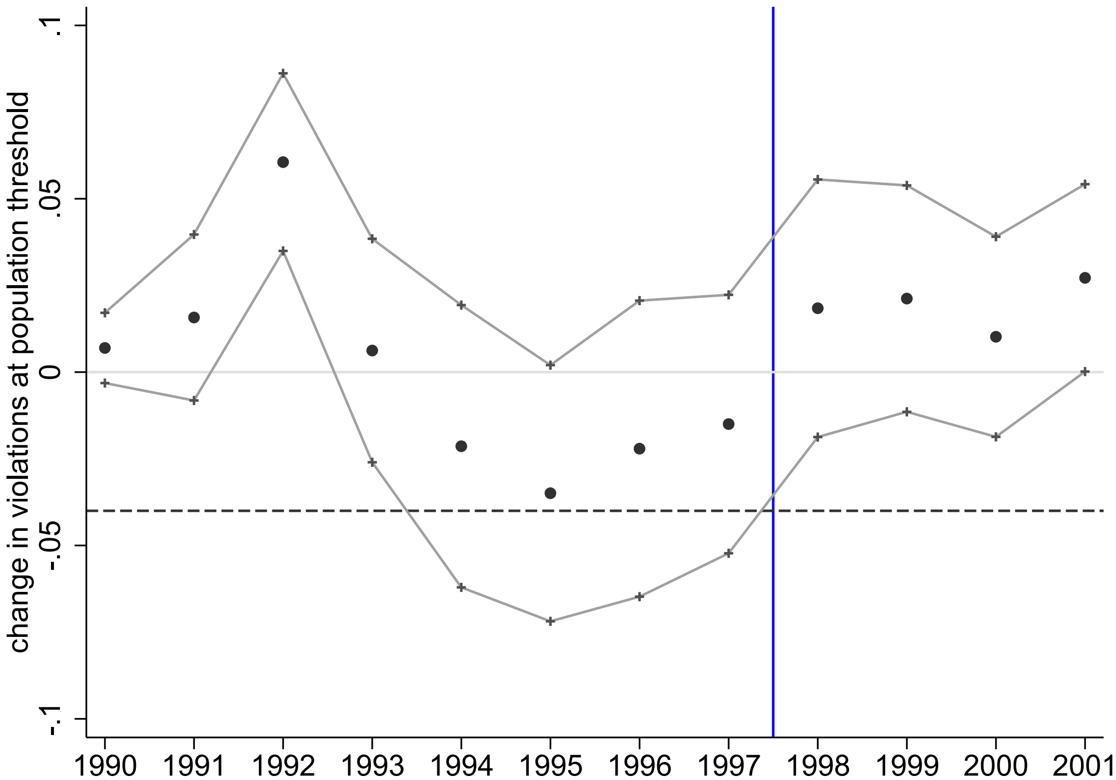

Figure B.2: *The effect of the mailing requirement considering a linear regression discontinuity. The figure shows estimates and associated 95 percent confidence intervals of T from Eq. (B.1) with* $k = 10,000$. *We estimate Eq. (B.1) separately for each year in our sample. The solid horizontal light gray line indicates zero effect, the dashed black line illustrates our preferred point estimate from our differences-in-differences model, and the blue vertical line delineates the post-information disclosure policy period.*

Figure B.2 and Figure B.3 provide no evidence of any effect of the mailing or online threshold. While not shown, we also generate quadratic RD estimates, as well as linear and quadratic RD estimates with a limited sample, and find no effect due to the mailing or online thresholds. These results validate our differences-in-differences estimates of the effect of the online threshold, but not of the mailing threshold. Overall then, our RD results confirm the effect we estimate with the differences-in-differences models at the publishing and online posting requirements, but do not support the results we find for the mailing requirement. Resolving these discrepancies is the subject of ongoing investigation.

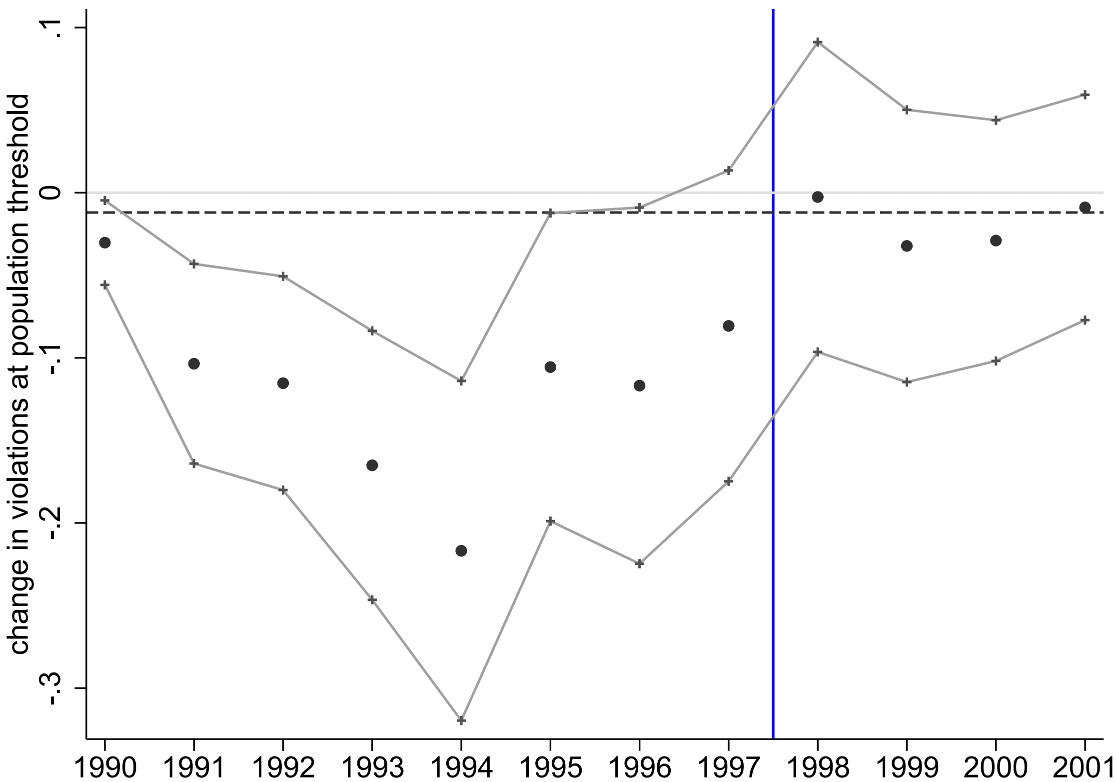

Figure B.3: *The effect of the online posting requirement considering a linear regression discontinuity. The figures shows estimates and associated 95 percent confidence intervals of T from Eq. (B.1) with* $k = 100,000$. *We estimate Eq. (B.1) separately for each year in our sample. The solid horizontal light gray line indicates zero effect, the dashed black line illustrates our preferred point estimate from our differences-in-differences model, and the blue vertical line delineates the post-information disclosure policy period.*

B.2 Additional robustness check

In our main results we include water systems subject to multiple disclosure methods. To address any concern that including the full sample of systems biases our results, we run models that include only those water systems affected by a single information disclosure requirement. To explore the publication requirement, we consider water systems that serve up to 9,999 customers (i.e. we exclude water systems subject to the mailing and online posting requirement). To isolate the effect of the mailing requirement, we consider two samples: water systems serving more than 500 customers, and water systems serving between 501 and 99,999 customers.[164] To isolate the effect of the online posting requirement, we consider systems serving 10,000 or more customers. Thus each sample considers water systems impacted by only one form of information disclosure on either side of the service population threshold.

Table B.1: *Regression results illustrating the impact of the publishing requirement with systems limited to those only impacted by the publishing requirement (i.e. we ignore systems required to mail or post the water quality report online).*

	(1)	(2)	(3)	(4)
$T_{pub} \times Post$	-0.034	-0.028	-0.032	-0.004
	(0.004)***	(0.004)***	(0.004)***	(0.008)
$f(size)$			-0.011	-1.689
			(0.003)***	(0.513)***
$f(size)^2$			7.9e-05	1.4e+01
			(4.0e-05)**	(6.2e+00)**
Service pop. range	full	< 10k	full	< 10k
adj. R^2	0.25	0.25	0.25	0.25
Systems	46,900	42,752	46,900	42,752
Observations	562,800	513,024	562,800	513,024

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$
Water system clustered standard errors (reported in parentheses).
All models include water system and state-by-year fixed effects.
Flexible function of system size: $f(size)^n = Post_t \times size_i^n$
size refers to the water system service population in 100,000s.

We find mixed evidence for our results being robust to sample size limitations. Table B.1 shows that estimates of the publication requirement are consistent between the full and limited samples when excluding the flexible function of system size (columns 1 and

[164]Systems subject to the online requirement must still mail the water quality report. Both samples therefore compare systems that must mail to systems that must publish. The second sample perhaps isolates the mailing requirement more precisely, since it considers systems that only must mail the report.

2), but when including the flexible function of system size (columns 3 and 4), the effect of the publication requirement vanishes under the limited sample. We observe similarly in Table B.2; the estimates of the mailing requirement for the models that exclude the

Table B.2: *Regression results illustrating the impact of the mailing requirement considering systems only impacted by the mailing requirement.*

	(1)	(2)	(3)	(4)	(5)	(6)
$T_{mail} \times Post$	-0.043	-0.023	-0.021	-0.040	-0.021	0.009
	(0.006)***	(0.007)***	(0.007)***	(0.007)***	(0.007)***	(0.015)
$f(size)$				-0.006	-0.004	-0.202
				(0.003)**	(0.003)	(0.097)**
$f(size)^2$				1.4e-05	-1.6e-05	1.7e-01
				(3.7e-05)	(3.6e-05)	(1.1e-01)
Service pop. range	full	> 500	500-100k	full	> 500	500-100k
adj. R^2	0.25	0.27	0.27	0.25	0.27	0.27
Systems	46,900	21,052	20,638	46,900	21,052	20,638
Observations	562,800	252,624	247,656	562,800	252,624	247,656

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$
Water system clustered standard errors (reported in parentheses).
All models include water system and state-by-year fixed effects.
Flexible function of system size: $f(size)^n = Post_t \times size_i^n$
size refers to the water system service population in 100,000s.

flexible function of system size are consistent across sample size (columns 1, 2, and 3), but with the flexible function of system size, the effect of the mailing requirement becomes statistically indistinguishable from zero under increasingly limited samples (columns 4, 5, and 6). We argue in section 2.4 that the flexible function of systems size isolates the effect at the disclosure threshold. If this is true, then why does the presence of the flexible function of systems size seem to render the results so sensitive to sample size? We suggest that it could be a result of non-linearity in our data. The coefficient estimates of $f(size)$ and $f(size)^2$ change by several orders of magnitude between the full sample and limited sample in both Table B.1 and Table B.2. This suggests that non-linearity in our data from systems serving between 500 and 100,000 customers drastically influence the shape of the flexible function of system size that then ceases to influence results in the full sample.

Finally, in Table B.3 we estimate the robustness of our online posting estimates to sample limitations. Unlike the publishing and mailing requirements, estimates of the impact of the online posting requirement are consistent across sample size for the model including the flexible function of system size (columns 3 and 4). Estimates are not consistent across sample size when considering the models that exclude the flexible function of system size (columns 1 and 2). This suggests that the small systems drive the result in column 1. But when we are able to isolate the effect at the online posting requirement, either by limiting the sample size in column 2 or by including the flexible function of system size in columns 3 or 4, we find that the true effect of the online posting requirement is minimal.

Table B.3: *Regression results illustrating the impact of the online posting requirement considering systems only impacted by the online posting requirement.*

	(1)	(2)	(3)	(4)
$T_{web} \times Post$	-0.056	-0.017	-0.012	-0.010
	(0.014)***	(0.016)	(0.021)	(0.021)
$f(size)$			-0.016	-0.002
			(0.005)***	(0.004)
$f(size)^2$			1.5e-04	-5.0e-05
			(6.3e-05)**	(4.6e-05)
Service pop. range	full	$\geq$ 10k	full	$\geq$ 10k
adj. R^2	0.25	0.24	0.25	0.24
Systems	46,900	4,146	46,900	4,146
Observations	562,800	49,752	562,800	49,752

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Water system clustered standard errors (reported in parentheses).

All models include water system and state-by-year fixed effects.

Flexible function of system size: $f(size)^n = Post_t \times size_i^n$

size refers to the water system service population in 100,000s.