Title

Essays on the Economics of Adapting to and Preventing Climate Change

Author
Winseck, Kevin

Publication Date
2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on the Economics of Adapting to and Preventing Climate Change

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Kevin Winseck

Committee in charge:

Professor Judson Boomhower, Chair
Professor Julian Betts
Professor Prashant Bharadwaj
Professor Joshua Graff Zivin

2023

The Dissertation of Kevin Winseck is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2023

DEDICATION

Dedicated to the dear friends I lost during graduate school, Harris Bunker and Bailee Mulholland.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Judson Boomhower for his support throughout my graduate career, from my first summer research project to field courses, through to my dissertation defense. I also want to thank Professor Joshua Graff Zivin for his insight on all of my work, the opportunity to assist in teaching and research, and especially for the referral that landed me my dream job. Thanks to Professor Prashant Bharadwaj for being the best vice chair for graduate studies that one could ever ask for–always working hard to best serve the needs of students, and support them through the hardest of times (particularly during the early months of Covid). Thanks also to Professor Julian Betts who is likely unaware of his outsized role in my decision to attend UCSD, as my meeting with him during visit day confirmed this as the right place for me, as well as for his excellent teaching of the labor field courses and input on compulsory schooling laws.

I thank all of my committee members for their unconditional support of my career aspirations, carefully balancing advising with listening to my own goals. Finally, I wish to thank all UCSD Economics faculty and graduate students for their continued support, and Joshua Aarons in particular for helping me think through many technical challenges I faced along the way, and many thanks to my fiancée, Julia Weiss, for her infinite support and motivation.

Chapters 1 and 2 are currently being prepared for submission for publication of the material. The dissertation author was the primary investigator and sole author of this material.

Chapter 3, has been given the opportunity to revise and resubmit for publication by the Review of Economics and Statistics. This material is co-authored with Noam Angrist, Joshua Graff Zivin, and Harry Anthony Patrinos. Additional thanks are due to these co-authors, without whom this chapter would have been unattainable.

VITA

2019	B. S. in Economics *summa cum laude with honors*, Arizona State University
2019	B. S. in Mathematics *summa cum laude with honors*, Arizona State University
2019-2020	Reader, University of California San Diego
2020-2022	Graduate Student Researcher, University of California San Diego
2020-2023	Graduate Teaching Assistant, University of California San Diego
2023	M. A. in Economics, University of California San Diego
2023	Ph. D. in Economics, University of California San Diego

FIELDS OF STUDY

Major Field: Economics

Studies in Environmental and Resource Economics, Labor Economics, and Behavioral Economics
Professors Judson Boomhower and Joshua Graff Zivin

ABSTRACT OF THE DISSERTATION

Essays on the Economics of Adapting to and Preventing Climate Change

by

Kevin Winseck

Doctor of Philosophy in Economics

University of California San Diego, 2023

Professor Judson Boomhower, Chair

Climate change is one of the largest challenges of our time. This dissertation assesses the economics of water conservation programs in the face of increasing water scarcity and estimates the causal impact of education on people's climate-related beliefs, behaviors, policy preferences and voting.

Chapter 1 focuses on takeup of water-conserving rebate campaigns in the U.S. Southwest, estimating a demand curve, neighbor spillovers, and heterogeneity for programs replacing thirsty grass lawns with drought-tolerant landscaping. I find strong evidence that households respond to rebate generosity: a 1% increase in generosity leads to a 0.93% increase in takeup. Each conversion leads to an additional 0.61 adoptions the following year. I also find substantial

heterogeneity by race and wealth, with Hispanic rebate applicants 47% more likely than White applicants to have their application rejected, a disparity that completely disappears when using a contractor.

Chapter 2 uses thermal infrared satellite imagery to estimate temperature changes from such landscaping conversions in an event study analysis. Conversions increase summertime parcel temperatures by 0.6°C (1.1°F) with effects twice as large on the hottest 20% of days and for homes with the most removed vegetation. Using existing estimates from the literature, this implies annual heat costs of $1,675 per household, comprised of increased electricity usage ($48) and mortality risk ($306), as well as harder-to-quantify comfort values, diminished cognition, and costly adaptation behavior. In contrast, water savings are only $574-954 for a typical home, depending on water price.

Chapter 3 estimates the causal effect of education on pro-climate survey responses using variation in compulsory schooling laws across 20 European countries. An additional year of education increases pro-climate beliefs, behaviors, and policy preferences with no detectable effect on voting for green parties. Results are robust to the functional form specification of the time trend, the selection of included countries and reforms, as well as an alternative specification of the instrument.

Chapter 1

Takeup, Spillovers, and Heterogeneity in Water-Wise Landscaping Incentive Programs

Abstract

In the face of historic drought in the U.S. West, water agencies throughout the region offer rebates to replace thirsty grass lawns, where a majority of residential water is consumed, with more drought-tolerant landscaping. Using administrative data on turf replacement rebates, and variation in program generosity, this paper estimates a demand elasticity for rebates of 0.93 with respect to generosity and finds that each lawn conversion leads to 0.61 additional conversions in the same neighborhood the following year. Rebate takeup generally increases with wealth, while White and Asian households have the highest baseline takeup. Black households have less elastic demand. Hispanic rebate applicants are substantially more likely to be denied than others, an effect that is largely mitigated by contractor use.

1.1 Introduction

Outdoor water usage accounts for a significant portion of Americans' consumption, with up to 75% dedicated to landscaping in the U.S. West (Mayer et al., 1998). As the population of the region grows, the increasing scarcity of water has led many water agencies to adopt policies promoting water-conserving landscaping through turf replacement and xeriscaping[1]

[1] I use the terms 'turf replacement', 'turf conversion', and 'xeriscaping' interchangeably to refer to the conversion of turf grass to less water-intensive landscaping like mulch interspersed with drought-tolerant plants. Conversions to hardscape such as decks, patios, driveways, or even swimming pools are generally not included in this definition.

rebate programs, replacing water-intensive turf grass with a combination of less thirsty desert-appropriate plants and mulch or rocks. I find an elasticity of demand for xeriscape rebates with respect to generosity of 0.93 along with evidence of spillovers: an additional rebate taker in a given year induces an additional 0.61 takers in their neighborhood the following year. Back-of-the-envelope calculations imply that existing rebate programs would reach California's residential water conservation goal of a 15% reduction in 33 years, while a rebate of $7.38 per square foot (more than three times the current average rebate) would meet the goal in the next decade.

As with other residential rebates, xeriscaping takeup is higher in wealthier neighborhoods, as measured by proxies such as rent, poverty rate, and lot size. Larger and more valuable homes, as well as those owned by Black homeowners are also more likely to be associated with rebates while Hispanic homeowners are least likely to take the rebate. Finally, I show that while Black and Asian families are somewhat more likely than White families to have their rebate application rejected, applications from Hispanic families are 1.5 times more likely to be rejected, even when controlling for wealth proxies. However, this discrepancy shrinks substantially when controlling for whether a contractor was used in the conversion, with Black households almost twice as likely to use a contractor as Hispanic households, while White and Asian households are the least likely. This suggests that from an environmental justice perspective, ensuring equal access to rebate programs might critically depend on access to contractors.

While there is relatively little research outside of the physical sciences about xeriscaping, what exists focuses mainly on program evaluation. Brelsford and Abbott (2021) show that incentivized turf replacement is a cost-effective way to reduce water consumption by nearly one-fifth. Baker (2021) assesses energy and water use along with changes in property values associated with the Las Vegas "Cash for Grass" program, finding an increase in property values (a result which critically depends on neighborhood conformity (Burkhardt et al., 2021)) and energy use, along with the expected decrease in water use. Addink (2005) provides an agronomic analysis of these programs with older data, finding evidence that landscape conversion may actually lead to more water usage across a diverse group of Western cities, a result which may come from the fact that desert native plants may take up to three years to become drought tolerant, and thus need more water during the initial period (Metropolitan Water District of Southern California, 2013). Pincetl et al. (2019) uses Google Street View to categorize biodiversity and network effects. However, none of these are able to examine variation in rebate generosity, nor use spatial and temporal variation for identification.

There is a somewhat larger literature on spillover effects of landscape conversion, with the most thorough being Burkhardt et al. (2021) which takes a behavioral approach by assigning households utility over their level of landscape conformity with their neighbors. Importantly, this paper highlights that peer effects may mute the effectiveness of economic incentives like rebates or water taxes. This paper, along with the earlier Bollinger, Burkhardt and Gillingham (2020) use extremely high resolution aerial imagery and machine learning to classify the greenness of each parcel in their sample of Phoenix, Arizona homes. They show that this greenness measure, meant to represent vegetation, is strongly correlated with water consumption, motivating its use as a measure of outdoor landscaping status. They use new movers, who are more likely to re-landscape their yards, to instrument for neighbors' landscaping decisions; however one may think that the residential flow of a neighborhood might be endogenous,[2] and this IV approach is local to new movers, rather than a broader set of landscape changers. Additionally, Bollinger, Burkhardt and Gillingham (2020) focus more on identifying spillovers in water conservation by using water meter data than on a reduced-form relationship with explicit landscaping choices. Both papers, along with Brelsford and De Bacco (2018) and Pincetl et al. (2019) show evidence of network effects, however, Brelsford and De Bacco (2018) is the only one to use rebate data confirming landscape status rather than imputed values from remote sensing. A policy-relevant question is the extent to which investment in rebates spills over into others adopting the same prosocial behavior, a question which has not been answered with a satisfying level of causal

[2] For example, gentrification may cause a large flux of movers and be correlated with environmental and conformal attitudes, and curb appeal from landscaping may directly affect property values, inherently making the moving decision endogenous.

identification. Furthermore, to my knowledge, none of the related papers actually report a policy-relevant parameter that suggests the expected amount of spillovers resulting from a given level of takeup. A related literature in the solar photovoltaic context, spearheaded by Bollinger and Gillingham (2012) explores similar spillover effects, leveraging first-differencing over time lags in installation to overcome the difficulties in causally identifying peer effects.

There exists a substantial gap in the literature regarding rebate takeup. Most studies focus on a relatively small geography, often one city with one rebate program. To better inform rebate policy, this paper leverages spatial and temporal variation across water districts in Southern California, Arizona, and Southern Nevada, focusing particularly on California where data is especially rich. This is the first paper to estimate a demand curve for these rebates and to examine heterogeneity by race and wealth of the homeowner, in addition to assessing heterogeneity in application denials and contractor use, and provides a novel strategy for causally estimating spillovers.

The paper proceeds as follows: Section 2.2 describes the sources of my data as well as how it is constructed and estimating equations for each analysis, Section 1.3.1 estimates the demand for turf replacement rebates using variation in rebate generosity, Section 1.3.2 identifies the neighborhood spillover effects using cross-boundary rebate variation, Section 1.3.3 analyzes heterogeneity in takeup and application rejection rates, Section 1.4 discusses the results, and Section 2.6 concludes.

1.2 Empirical Strategy

1.2.1 Data Sources

The rebate data in this paper come from public record requests from water agencies, minimally listing the address and date of installation for every home that received a turf replacement rebate. The Metropolitan Water District of Southern California (MWD), Municipal Water

District of Orange County (MWDOC)[3] and Las Vegas Valley Water District (LVVWD from wholesaler Southern Nevada Water Authority, SNWA) provided more detailed data including the area rebated and dollar amount paid, allowing for an empirical calculation of historical rebate generosity (see Section 1.2.2). However, the Arizona jurisdictions, Tempe and Glendale, provided only rebate dates and addresses. Other Phoenix-area water agencies either did not respond to the data request, or were unable to provide the requested information. Data from different jurisdictions are harmonized to have compatible fields and then assigned geographic coordinates from the given street addresses using the Esri geocoder in ArcGIS. This yields a spatial dataset where each point is a house with a rebated landscape conversion. My most recent data is from 2021, and I keep rebates back until 2003 when the Las Vegas area program ramped up. MWDOC rebate data starts in 2011 and MWD in 2013. Parcels are only in the panel data for years in which any parcels in the same water agency had rebates.

Water agency boundaries for Southern California were provided by MWD as part of the records request, and I use census-designated place (CDP) boundaries for the cities of Tempe and Glendale, Arizona to designate the study area of interest. Because the relevant water authority in Nevada extends beyond the city limits of Las Vegas, I allow the region of interest to be CDPs that ever had a rebate, filling in holes in CDP polygons.

In order to acquire additional covariates and to observe non-takers, this rebate dataset is then spatially merged to parcel datasets for each relevant county, paying careful attention to restrict to residential parcels within the study area as the unit of interest. Parcels denote spatial boundaries of land ownership, and databases are generally maintained by county assessors for tax purposes. California parcel data for Los Angeles, Orange, San Bernardino, Ventura, and Riverside counties come from the Southern California Association of Governments' data portal,[4] and San Diego's data is downloaded from the San Diego Association of Governments' data

[3]While the Metropolitan Water District of Southern California sells wholesale water to Orange County member agencies, the Orange County rebate programs were instead administered by the *Municipal* Water District of Orange County, and thus the data is from a separate source.

[4]Southern California Association of Governments (2016, 2019)

warehouse.[5] Parcel data for Clark County, Nevada was provided through a public record request from the county assessor, and that for Maricopa County, Arizona, was received through a request from the Arizona State University library system.[6] As part of my data agreements, I cannot share parcel data from Clark or Maricopa counties, or rebate data from any water agencies.

Due to the varying sources of parcel data, the only covariates that are included in every county are land use codes (which I restrict to residential) and parcel acreage. For some counties, the parcel data has owner name,[7] property valuations,[8] and improvement size,[9] which I leverage in the heterogeneity analysis. I proxy for demographics using Census Block Group (CBG) level statistics from the 2010 Decennial Census[10] and 2017 American Communities Survey 5-year estimates,[11] treating demographics as temporally constant throughout the timeframe of interest. It is important to note that the environmental justice literature recognizes the sensitivity of heterogeneity effects to the level of aggregation (see Banzhaf, Ma and Timmins (2019) and Baden, Noonan and Turaga (2007), for example). For this reason, my preferred specifications use race imputed from homeowner name to give a household-level measure. This uses the Bayesian methodology of Imai and Khanna (2016) to assign a probability of being White, Black, Hispanic, Asian, or Other based on the incidence of a given surname in the census data from a given census tract. This same strategy is used to assign race based on applicant name to assess differential application rejection rates. To avoid falsely classifying the race of parcels not owned by individuals, I exclude from classification parcels with owner names that include terms common to institutions, such as "LLC", "Trust", and "Property". The imputed race is the race predicted by the algorithm to be most likely.

Finally, given that a push factor for tearing out grass is the irrigation cost, I include

[5] San Diego Association of Governments (2022)
[6] Maricopa County Assessor, accessed via https://lib.asu.edu/geo/services
[7] Available for San Diego CA, Clark NV, and Maricopa AZ.
[8] Available for San Diego CA, Clark NV, and Maricopa AZ.
[9] Available for Maricopa AZ and California counties except for Ventura (San Diego, Orange, Los Angeles, Riverside, and San Bernardino).
[10] US Census Bureau (2010)
[11] US Census Bureau (2017)

residential water prices for 2013-2021 from the California State Water Resources Control Board's electronic Annual Report (eAR)[12] as a robustness check. Base usage rates per hundred cubic feet (HCF) of municipal water for single family residential households are matched to the retail agencies in my data by year, leaving missing values when missing in eAR. Because water prices cannot be recovered for a nontrivial portion of my data, I include prices only as a robustness check rather than as a covariate included in primary specifications.

The datasets are merged to provide a parcel by year dataset spanning 2003-2021 for all residential parcels within the water boundaries for which I have rebate data. Parcels with conversions are dropped in all years following a rebate, as they are no longer considered eligible for rebates. It should be noted that this analysis does not observe conversions that do not involve rebates, and so should be considered an analysis of *incentivized* landscape conversion.

1.2.2 Empirical Rebate Generosity Rates

The variation for identifying rebate demand comes from spatial and temporal variation in the generosity of rebate programs. Due to the fact that the data encompasses nearly 200 water retail agencies over 19 years, I calculate the rebate generosity empirically when the data allows it. This is the case for Nevada and Southern California, but not for Arizona. For Tempe and Glendale, Arizona, I use internet archives from the Wayback Machine[13] to characterize a time series of rebate generosity,[14] although these agencies are only used for the heterogeneity analysis, since the primary focus is on California rebates because of their high quality data.

For the jurisdictions in which the generosity is empirically calculable (the vast majority of my rebate data, and that used in the primary analysis), I do so by dividing total rebate amount

[12] California Stater Water Resources Control Board (2023)

[13] Wayback Machine Internet Archive

[14] One difficulty in identifying the rebate generosity at any given point of time, is the variation in rate structure: some rebates are a flat rate (e.g. $100 for converting, or for each front yard and back), most are per square foot (e.g. $2.00 per square foot of turf replaced, most have upper bounds, frequently at 1000 square feet), and some nonlinear per-square-foot rates ($150 for 500-1500 sq. ft., $300 for 1500-2500 sq. ft. etc.). The nonlinear rates tend to track a constant per-square-foot rate, so I use the average per-square-foot rate implied by the rate structure. For flat rates, I divide the lump sum payment by the median replacement area from regions in which this is available (California and Nevada). The median amount of replaced turf is around 1,000 square feet.

issued by the square footage converted. Then, I assign each retail agency a generosity rate in each year by finding the modal (nonzero) generosity[15] among the rebated households within its boundaries in the given year. Because houses within MWD's service area potentially face rebates from three tiers of water agencies (retail, member, wholesaler), I find the modal total rate.[16] This hierarchical structure helps to explain the variation in rebate generosity seen in Figure 1.1, as the three tiers adjust their rebate rate independently of the others, leading to substantial variation with little systematic pattern. A 2019 report on the effect of the Los Angeles rebate program noted, in explaining the rise in rebate rates in 2015:

> One of the likely reasons for this large increase in participation since March, 2015, was the temporary increase in the available rebate. In 2014, [Los Angeles Department of Water and Power] raised the turf rebate from $1.00 per sq. ft. of turf replaced to $1.75 per sq. ft. MWD also offered rebates, making the total amount LA City residents received for lawn replacement $3.75 per sq. ft. (MWD offered $2.00 while DWP offered $1.75). As of July 2015, MWD had exhausted their funds. This lowered the rebate total for Los Angeles residents from $3.75 to the current level of $1.75 per sq. ft. (Jessup, DeShazo and Panjwani, 2019)

This quote highlights the capricious nature of the rebate generosity faced at the retail agency level. To drive home the point that these rates are not systematic and are unpredictable, the report adds: "The rebates have changed rapidly over the past few years and it is unclear what the future holds." (Jessup, DeShazo and Panjwani, 2019)

Coupled with variability seen in the data, this report lends credibility to my identifying assumption throughout this paper that rebate generosity faced by residential customers is effectively random, especially so when conditioning on observables, as well as time and location fixed effects.

Nevertheless, I take great care to control for potential confounding factors. The largest threat to identification is a positive correlation between rebate generosity and localized drought, since an ongoing drought is likely to increase demand for turf replacement regardless of rebate

[15] Dollar per square foot rebate rates are grouped into 5-cent bins centered on multiples of 5 cents to avoid problems with rounding since the mode statistic is sensitive to how data are rounded.

[16] Empirically, the mode of the sum of the three tier rates is more often the actual rate received than is the sum of the modes of each tier.

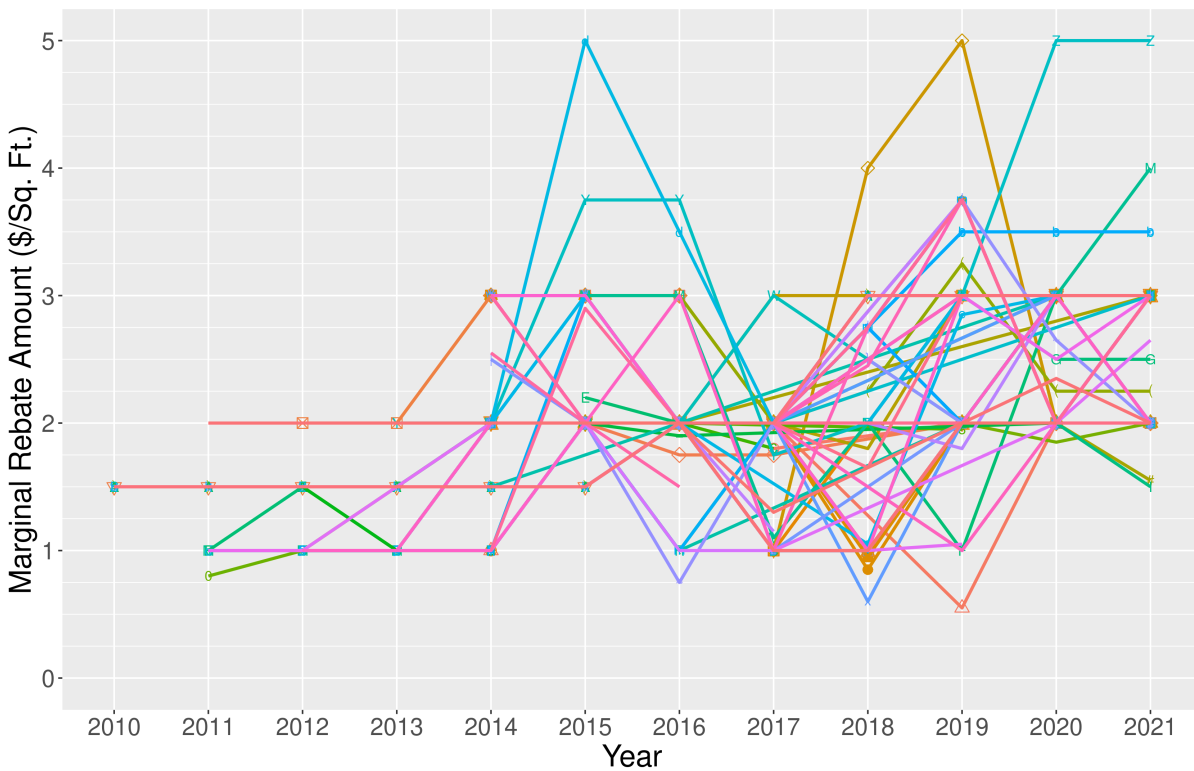

Figure 1.1. Rebate Generosity by Retailer and Year for California
Notes: Imputed rebate generosity by water retailer by year for California data. Each line represents a water retailer. Note the substantial variation in generosity both in space and time.

generosity. Moreover, a drought may make water agencies offer more generous rebates in an attempt to curb water consumption. I use spatially and temporally resolute drought status data from the U.S. Drought Monitor[17] (USDM) to determine the drought status.[18] USDM reports drought boundaries on a weekly basis, and I use the boundary for the first week of each month of the prior year. The drought status variables used as control variables are the number of months of the previous year in each of the five drought categories. I do not use drought status for the current year because I am not able to identify precisely when during the year turf replacements take place and future drought status should not be allowed to affect the rebate takeup decision.

1.2.3 Demand for Rebates

Estimating the demand for landscaping rebates relies on the identifying assumption that conditional on observables, the dollar per square foot rebate generosity is effectively random. I utilize this variation to estimate a linear probability model of takeup as a function of rebate generosity. My primary design is a spatial discontinuity design leveraging highly local variation in rebate generosity at the borders between water agencies. I also show a two-way fixed effects (TWFE) model, and specifications controlling for various combinations of covariates such as demographics and drought status, including different levels of fixed effects. It is worth noting that this section estimates a bundled effect of the rebate generosity and any spillovers from past takeup. In Section 1.2.4, I hone in on the effects of spillovers in isolation from generosity effects. This bundled rebate program effect should also be considered to include advertising effort, which may be correlated with generosity, but for which I am unable to control. The spatial discontinuity design controls for broad advertisements like signs, billboards, and regional television/radio advertisements, but cannot control for targeted mailings, emails, phone calls, or door-to-door

[17]National Drought Mitigation Center, U.S. Department of Agriculture and National Oceanic and Atmospheric Administration (2023)

[18]0-4 with 0 representing moderately dry, and 1-4 the official D1-D4 classification statuses, ranging from moderate to exceptional drought. I allow the omitted category to be no drought.

recruitment efforts.[19]

Let $Takeup_{p,y,a}$ be a binary indicator for whether parcel p in water agency a adopts the rebate in year y and $generosity_{a,y}$ be the rebate generosity in dollars per square foot offered by agency a in year y. For all specifications, $\mathbf{X}_{p,y,a,cbg}$ is a vector of controls that may include CBG level demographics,[20] drought status indicators, parcel acreage, and fixed effects for year, water retail agency, and/or CBG. Then the estimating equation for rebate demand is:

$$Takeup_{p,y,a,cbg} = \alpha + \beta generosity_{a,y} + \Gamma \mathbf{X}_{p,y,a,cbg} + \varepsilon_{p,y,a} \tag{1.1}$$

where β is interpreted as the effect on the probability of a household's takeup in a given year of an increase in rebate generosity of \$1/sf. The fact that baseline takeup is low (averaging about 1.5% of eligible households per year), suggests that the linear approximation of the true demand curve is likely reasonable in the relevant range. Because there has never been a more widespread push for turf replacement in the United States, it is impossible to know what the demand curve looks like at a higher generosities.[21] All regressions cluster standard errors at the water retail agency by year level, the level at which rebate generosity treatment is assigned.

The primary specification, with the strongest identification, uses the spatial discontinuity design by letting $\mathbf{X_{p,y,a}}$ be year by water retailer border fixed effects. In the spirit of Ito (2014), I effectively compare households on one side of a water agency boundary with households just on the other side. To minimize concerns about sorting, I restrict to water boundaries that do not coincide with city boundaries, noting that to the extent that residential sorting occurs, it is unlikely to be motivated by strategically sorting into water districts within a given neighborhood. This is

[19] Bollinger, Burkhardt and Gillingham (2020) also provide evidence through an informal survey that landscape contractors do not generally engage in door-to-door canvassing, so it is unlikely to be a factor in the household landscaping decision.

[20] Race, age, gender, median rent, education

[21] A Nevada law passed in 2021 (AB356) has been the strictest measure yet to limit irrigation of turf grass, banning the use of Colorado River water to irrigate nonfunctional grass. However, even this broad law excludes single family lawns and grass used for recreational activities from the requirement. New construction in Nevada prohibits the installation of grass, but there is no mandate to remove grass from existing single family residential lawns.

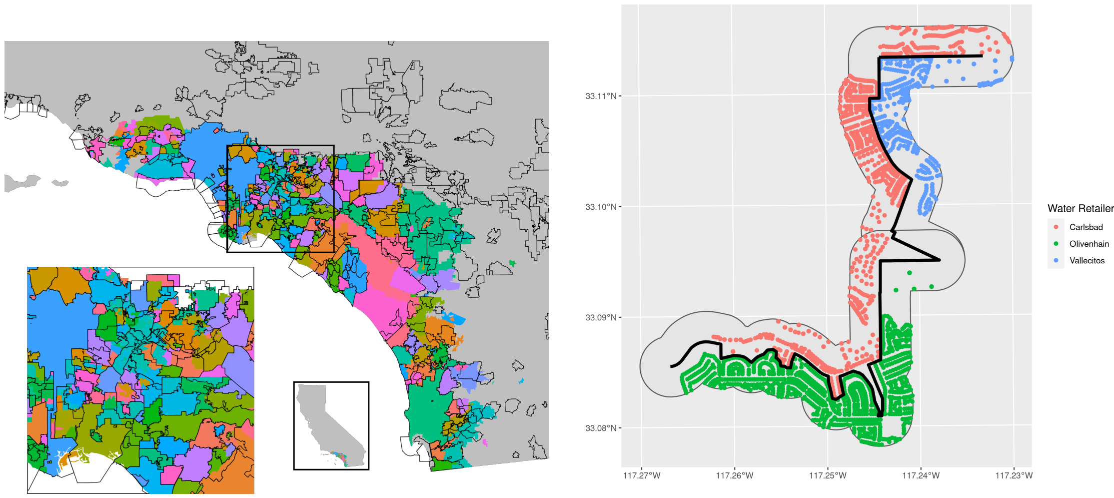

Figure 1.2. Spatial Discontinuity Design
Notes: **Left:** Water retail agency borders (colored polygons) mapped against city boundaries (black lines). Insets show a zoomed in portion of the Los Angeles metropolitan area and a zoomed out map of the state of California for reference. **Right:** An example strip of water agency border with a 1,000ft buffer that is not within 1,000ft of a city boundary.

supported by the fact that fine scale water district boundaries are difficult to find publicly (those used in this paper are the result of a public records request) and anecdotally, that Zillow.com, one of the most popular search tools for home buyers does not list the water agency in home listings, so it is extremely unlikely to be salient in the purchase decision. Moreover, it is not the case for most borders that one side has consistently more generous rebate programs.

For this spatial discontinuity design, I restrict to households within a certain buffer of a water agency boundary (1,000 feet in the primary specification, 500 and 1,500ft in robustness checks) and restrict to water agency boundaries that are not within this same distance from city boundaries.[22] Each contiguous section of water agency border (after removing city border overlaps) receives a border ID number and is interacted with the year fixed effects in Equation (1.1) to compare households on either side of a border for a given border strip in a given year, and which are in the same city. Because this flexibly and nonparametrically controls for time trends, highly local characteristics (such as amenities and demographics), as well as drought

[22]For this paper, city boundaries are the boundaries of Census-designated Places, which are periodically defined and reported by each county to the Census Bureau (US Census Bureau, 2022).

status (through the location×year fixed effect), this is considered to have the cleanest possible identification short of a randomized control trial. Furthermore, even though the estimate is local to households near the boundaries of water agencies, this does not imply that they are local to households on the fringe of a city: by construction, the border sections of interest are not on the edge of a city, and water boundaries cut through urban areas throughout Southern California.

Table 1.1 shows the balance tests for covariates between these "border homes" and those not near one of these borders but in the sample ("nonborder homes"), as well as for homes on the high generosity sides of borders versus those on the low generosity side. The F-test for Columns (1)-(4) indicate that the listed covariates are able to predict whether a parcel is in a border region, and suggests that border homes may not be completely representative of the population of residential parcels. Columns (5)-(8) regress an indicator for being on the high generosity side of the border in a given year[23] on the listed covariates, with the F-statistic from Column (5) suggesting that these covariates do *not* predict whether a home is on the high generosity side of the border, conditional on being in the border strip. Column (6) includes an indicator for having nonmissing water price while Column (7) restricts the sample to border homes with nonmissing water price (for comparability to Column (8)) and Column (8) includes water price as a covariate. Relative to Column (5), these specifications indicate that there is selection on the availability of water prices, as restricting to the sample where this is available (a loss of 83% of the sample) leads to the covariates being able to successfully predict to which side of the border a household belongs. Additionally, as evidenced by Column (6), having nonmissing water price is associated with 19.9 percentage point increase in the likelihood of being on the more generous side of a border. This is plausibly due to more well-funded water agencies being both more likely to correctly report their prices and to have more generous rebates, reinforcing the omission of water prices from primary analysis since including prices necessarily restricts the data to this selected sample. Nevertheless, robustness checks include water price and show that

[23]Note that while parcels do not change location over time, the side of the border with the more generous rebate frequently changes due to volatility in rebate generosity.

the results do not substantially change. Appendix Tables 1.10 and 1.11 further break down the balance of covariates by showing t-tests for each each covariate. Where statistically significant differences between high- and low-generosity sides occur, they tend not to be consistent between years, suggesting that significant differences are not due to systematic patterns in generosity.

Table 1.1. Balance Test: Border vs. Nonborder and High Generosity vs. Low Generosity Side

Notes: Columns (1)-(4) show regressions of an indicator for being within 1,000 feet of a border on the listed covariates, while Columns (5)-(8) use an indicator for being on the high generosity side of the border as the outcome, restricted to parcels within 1,000 feet of a water agency border (but at least 1,000 feet from a city boundary). 'p' denotes parcel data, 'DEC' is from the Decennial Census, and 'ACS' is the American Communities Survey.

	(1) Border	(2) Border	(3) border	(4) Border	(5) High Side	(6) High Side	(7) High Side	(8) High Side
p acres	-0.001	-0.001	-0.001	-0.000	0.017	0.011	0.005	-0.023
	(0.000)	(0.000)	(0.000)	(0.000)	(0.012)	(0.011)	(0.013)	(0.010)
	[0.000]	[0.000]	[0.000]	[0.104]	[0.172]	[0.314]	[0.716]	[0.031]
DEC race White rate	0.033	0.028	0.174	0.165	0.266	0.199	-0.135	-0.098
	(0.167)	(0.166)	(0.092)	(0.096)	(0.627)	(0.610)	(0.864)	(0.729)
	[0.844]	[0.867]	[0.060]	[0.090]	[0.673]	[0.745]	[0.877]	[0.894]
DEC race Black rate	-0.017	-0.022	0.133	0.148	0.106	-0.111	0.238	0.365
	(0.164)	(0.163)	(0.098)	(0.100)	(0.740)	(0.695)	(0.791)	(0.716)
	[0.916]	[0.892]	[0.180]	[0.140]	[0.886]	[0.873]	[0.766]	[0.613]
DEC race Asian rate	0.179	0.168	0.164	0.142	0.132	0.111	0.269	0.534
	(0.204)	(0.202)	(0.087)	(0.093)	(0.557)	(0.549)	(0.726)	(0.691)
	[0.381]	[0.407]	[0.062]	[0.133]	[0.814]	[0.840]	[0.713]	[0.444]
DEC male rate	-0.328	-0.330	-0.258	-0.207	-0.126	-0.195	-0.818	-1.743
	(0.119)	(0.120)	(0.152)	(0.159)	(1.283)	(1.246)	(1.663)	(1.423)
	[0.006]	[0.007]	[0.093]	[0.195]	[0.922]	[0.876]	[0.626]	[0.228]
ACS median age	0.001	0.001	0.000	0.000	-0.004	-0.004	-0.000	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)	(0.003)	(0.003)
	[0.259]	[0.248]	[0.987]	[0.679]	[0.214]	[0.158]	[0.975]	[0.629]
ACS median rent	-0.000	-0.000	0.000	0.000	0.000	0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	[0.066]	[0.064]	[0.719]	[0.969]	[0.441]	[0.413]	[0.825]	[0.742]
ACS poverty rate	-0.202	-0.200	-0.114	-0.095	0.239	0.192	0.196	-0.270
	(0.055)	(0.054)	(0.050)	(0.056)	(0.272)	(0.268)	(0.488)	(0.347)
	[0.000]	[0.000]	[0.024]	[0.090]	[0.382]	[0.474]	[0.690]	[0.442]
ACS school HS rate	0.019	0.020	-0.035	-0.055	-0.606	-0.643	-1.120	-1.103
	(0.117)	(0.116)	(0.065)	(0.064)	(0.473)	(0.453)	(0.590)	(0.488)
	[0.868]	[0.866]	[0.584]	[0.391]	[0.203]	[0.158]	[0.065]	[0.030]
ACS school some college rate	-0.004	-0.001	0.046	0.011	-0.669	-0.654	-0.665	-0.661
	(0.129)	(0.127)	(0.073)	(0.069)	(0.397)	(0.362)	(0.504)	(0.528)
	[0.976]	[0.995]	[0.526]	[0.868]	[0.095]	[0.074]	[0.195]	[0.218]
ACS school bachelors rate	-0.188	-0.183	-0.132	-0.107	-0.350	-0.364	0.091	-0.275
	(0.117)	(0.116)	(0.060)	(0.065)	(0.398)	(0.374)	(0.456)	(0.521)
	[0.108]	[0.116]	[0.031]	[0.101]	[0.383]	[0.332]	[0.842]	[0.600]
ACS school graduate rate	-0.217	-0.217	-0.206	-0.201	-0.598	-0.569	-1.073	-1.068
	(0.130)	(0.130)	(0.090)	(0.094)	(0.480)	(0.489)	(0.609)	(0.550)
	[0.098]	[0.096]	[0.024]	[0.034]	[0.215]	[0.248]	[0.086]	[0.059]
Nonmissing Water Price		-0.030				0.199		
		(0.012)				(0.091)		
		[0.017]				[0.032]		
Water Price ($/HCF)				-0.009				0.100
				(0.005)				(0.053)
				[0.077]				[0.066]
Constant	0.281	0.292	0.098	0.118	0.910	0.983	1.687	1.964
	(0.130)	(0.131)	(0.104)	(0.092)	(0.823)	(0.791)	(1.060)	(0.953)
	[0.031]	[0.027]	[0.348]	[0.204]	[0.272]	[0.217]	[0.120]	[0.046]
Observations	27720120	27720120	6269747	6269747	389711	389711	64344	64344
Sample	California	California	CA Nonmiss Price	CA Nonmiss Price	CA Border	CA Border	Border+Nonmiss Price	Border+Nonmiss Price
F-Test P-val	0.000	0.000	0.001	0.000	0.392	0.017	0.003	0.000
R^2	0.03	0.03	0.02	0.02	0.01	0.04	0.09	0.18

Standard errors in parentheses, clustered by retailer. P-values in brackets.

1.2.4 Spillovers

For the purposes of this paper, a spillover is defined as an increased probability of parcel p's rebate takeup in a given year attributable to an increase in the number (or fraction) of p's neighbors having taken the rebate in the previous year. The canonical issues in identifying spillovers include endogenous peer group formation, reflection, and correlated unobservables (Bollinger and Gillingham, 2012). To avoid the issue of reflection, I examine only the effect of previous years' takeup on takeup in the current year, not within-year effects. Moreover, I avoid confounding from peer group formation by looking at neighbors across water agency boundaries, which also helps to assuage concerns about correlated unobservables. To further address endogeneity and correlated unobservables, I also show an instrumental variables (IV) specification where lagged crossborder rebate generosity instruments for lagged takeup.

Because of relatively low baseline takeup rates of turf replacement rebate programs, using border strip fixed effects in a design similar to the spatial discontinuity in Section 1.3.1 is underpowered. Instead, I define a parcel's neighbors as other households in its Census Block Group, and restrict to CBGs that contain multiple water retailers in order to use crossborder identification. There are four different relevant metrics for the intensity of treatment: the count or fraction of crossborder neighbors that took the rebate in the previous year, and the count or fraction of *all* neighbors that took the rebate in the previous year. Regression estimates of the effect of the first two on current takeup should be considered causal because crossborder rebate terms are independent of the terms faced by parcel p.The effect of all neighbors' lagged takeup on current takeup is the policy-relevant parameter, but it is generally not causal, since it still suffers from the reflection issue and one cannot disentangle delays in takeup from the causal effect of spillovers in this OLS regression. However, the IV version is causal, and is the most interpretable because the intuitive metric is takeup among all neighbors, rather than just those who are in a different water agency. With regard to defining the denominator of the fraction of lagged takeup, homes are only considered eligible for the rebate if they never take the rebate or have not yet taken it in the current year, the same eligibility definition as in the primary demand analysis.

The OLS specification for crossborder spillovers takes the form of Equation (1.1) with the additional lagged takeup term:

$$
\begin{aligned}
Takeup_{p,y,a,cbg} = \alpha + \beta Takeup_{y-1,-a,cbg} + \gamma_1 generosity_{a,y} + \gamma_2 generosity_{a,y-1} \\
+ \delta \mathbf{X}_{p,y,a,cbg} + \varepsilon_{p,y,a}
\end{aligned}
\tag{1.2}
$$

where $Takeup_{y-1,-a,cbg}$ is the number (or fraction, depending on the specification) of homes in the same CBG as parcel p but a different water agency, $-a$ ("not a") that took up a rebate in the previous year, $y-1$. When regressing on all lagged takeup, the $Takeup_{y-1,-a,cbg}$ term becomes $Takeup_{y-1,\{a,-a\},cbg}$, reflecting lagged takeup in the CBG both inside and outside agency a. Section 1.3.1 reinforces that current rebate generosity is a critical factor in determining takeup, so both current and lagged generosity are included as controls to account for any spatial or temporal correlations that would otherwise lead to omitted variable bias. The vector $\mathbf{X}_{p,y,a,cbg}$ includes CBG-level demographics from the ACS and Decennial Census as well as parcel-level acreage and year fixed effects, as these controls are shown to most closely match the cross-border demand specification with the cleanest identification.[24]

This spillover analysis is subject to the classic bias versus variance tradeoff: restricting to highly local analysis within small buffers is subject to high variance, as noted above, while larger geographic areas will have more precision but attenuated estimates given that more distant neighbors are less likely to affect a household's landscaping decision. One way to get around the bias is by additionally restricting to parcels with a larger percentage of neighbors who are

[24] Border by year fixed effects are problematic in this specification because higher lagged crossborder rebate generosity essentially proxies for being on the less generous side of the border and creates a mechanically negative correlation between takeup and lagged crossborder neighbor takeup when directly comparing one side of the border to another. Furthermore, more power is needed to estimate spillovers than to estimate demand, and the border by year fixed effects consumer too many degrees of freedom in the spillover specification.

crossborder. Noting that the 500-1,500 foot buffer approach is underpowered, and neighbors beyond the census block group are unlikely to have a major effect, the primary focus in this paper is on spillovers from same-CBG neighbors, showing heterogeneity by the share of neighbors belonging to a different water agency.

The IV specification instruments for lagged all-CBG takeup with lagged crossborder generosity and the number of crossborder eligible homes in the same CBG, controlling for both current and lagged own generosity. The demand section of this paper demonstrates the relevance of this instrument, and its validity is given by the fact that the terms of crossborder rebates in agency $-a$ never apply to parcel p in a. To the extent that crossborder generosity is correlated with own generosity, controlling for own current and lagged generosity removes this correlation and leaves only the residual variation for identification. The number of crossborder homes that are eligible for the rebate is also included as an instrument to capture the notion that the number or share of turf replacements induced by crossborder rebate generosity should be proportional to the the number of neighbors in the CBG who fall outside of agency a.[25] Furthermore, both crossborder lagged rebate generosity and number of crossborder neighbors should only affect parcel p's rebate takeup decision through causal spillovers. Recall that the sample is restricted to CBGs with multiple water agencies, so one need not worry about identification coming from variation in closeness to water boundaries. Letting $N_{-a,cbg,y-1}$ be the number of eligible homes in Census Block Group cbg but not in water agency a, the two stage least squares IV estimating equations are:

[25]Given that CBGs are designed to have an approximately constant population, controlling for the *number* of crossborder neighbors is approximately equivalent to controlling for the *share* of crossborder neighbors.

$$
\begin{aligned}
Takeup_{y-1,\{a,-a\},cbg} = & \alpha + \beta_1 generosity_{a,y-1} + \beta_2 N_{-a,cbg,y-1} \\
& + \gamma_1 generosity_{a,y} + \gamma_2 generosity_{a,y-1} + \delta \mathbf{X}_{p,y,a,cbg} + \varepsilon_{p,y,a}
\end{aligned} \tag{1.3}
$$

$$
\begin{aligned}
Takeup_{p,y,a,cbg} = & \alpha + \beta \widehat{Takeup_{y-1,\{a,-a\},cbg}} \\
& + \gamma_1 generosity_{a,y} + \gamma_2 generosity_{a,y-1} + \delta \mathbf{X}_{p,y,a,cbg} + \varepsilon_{p,y,a}
\end{aligned} \tag{1.4}
$$

1.2.5 Heterogeneity

Heterogeneity analysis of rebate demand follows Equation (1.1), but focuses on the elements of $\mathbf{X_{p,y,a,cbg}}$ to examine which demographics have lower baseline levels of rebate takeup, and interacts demographic characteristics with rebate generosity to show differential demand elasticities. For most of the California parcels, I do not observe parcel-level demographics and rely on CBG-level values to represent the characteristics of the neighborhood rather than the household itself, making the spatial discontinuity design infeasible. Instead, I turn to the OLS model which includes in $\mathbf{X_{p,y,a}}$ year fixed effects, parcel size, and CBG level values: percentage of Black, Asian, and Other races, and males, median age and rent, poverty rate, rates of each level of schooling, and the rate of utilities included in rent. This model is also robust to including yearly drought status indicators.

I examine two main dimensions of heterogeneity that are the classic concerns in the environmental justice literature: race and wealth, measured in a variety of ways. First, I use the CBG level rates of each race as the covariates of interest, then I include instead an indicator for CBGs being above median White, with an additional specification interacting this with rebate generosity. Alternatively, I restrict to counties where I have parcel owner name and can use the Imai and Khanna (2016) algorithm to predict race. It should be noted that this algorithm distinguishes White names from Hispanic and Asian names much better than it distinguishes White names from Black names. Predicted race is the race with the highest probability among

White, Asian, Black, Hispanic, and Other. While I do not directly observe wealth, I do observe several reasonable proxies for it. First, at the CBG level, poverty rates, education levels, and median rent should be correlated with wealth. Second, at the household level I have parcel acreage,[26] and for a selection of counties, I have house value and square footage.

The final element of heterogeneity analyzes differential rates of rebate rejection, since the California data reports all applications and the final application status.[27] This analysis restricts to households with an application, and I use applicant name and Census tract to predict race. Additionally, since the California data includes information about whether a contractor is involved in the conversion, I include an indicator for contractor use and interact it with demographics. This model takes the form:

$$Rejected_{p,y,a,cbg} = \alpha_a + \Gamma \mathbf{X_{p,y,a,cbg}} + \varepsilon_{p,y,a} \tag{1.5}$$

In this model, $\mathbf{X_{p,y,a,cbg}}$ includes demographics and wealth proxies interacted with the contractor indicator, as well as either CBG or retail agency fixed effects.

Applications may be denied for a variety of reasons. Eligibility requirements that may be overlooked often include that the grass being replaced must be alive at the time of application (presumably to avoid inframarginality) and grass must generally be replaced with other vegetation, gravel, or organic mulch. Many agencies require photo submissions of the area to be replaced and a calculation of the square footage of the replacement, while some will even perform pre- or post-conversion inspections. Some programs require detailed drawings of planned layouts and have requirements for a minimum number of (native) plants per square foot, and requirements on the irrigation equipment used. Due to the sometimes complex eligibility requirements, language and education barriers may lead people to submit ineligible applications. It is also possible that application reviewers have some level of bias in their reviews, although I am unable to

[26]Parcel size is likely correlated with being suburban/rural in addition to wealth so it is an imperfect measure.

[27]Rejected applications are excluded from the main analysis since they do not have a rebate paid date, which is the relevant date used.

distinguish this from ineligible applications because I observe neither the reasons for denial nor the true quality of the applications.

1.3 Results

1.3.1 Demand

Table 1.2 shows the primary demand estimation results for the spatial discontinuity version of Equation (1.1). Columns (1)-(4) are the analysis for homes within 500 feet of a water agency border (but more than 500 feet from a city border), (5)-(8) are 1,000ft, and (9)-(12) are 1,500ft. The first column of each quartet is the full border sample, while the second column includes lagged rebate generosity to account for delays in takeup in the face of potentially temporally correlated generosity. The third column restricts to parcel-years for which water price data is nonmissing, and the last column in each quartet controls for water prices. All specifications in this table include border section by year fixed effects. Column (5) is the preferred specification, with the medium buffer distance and the sample that is not reduced by missing water prices. The point estimate in Column (5) is interpreted as: relative to crossborder neighbors in the same city, a \$1 per square foot increase in rebate generosity increases a household's probability of taking up the rebate in a given year by 0.66 percentage points (PP). Off a baseline of 1.48 percent takeup per year and with an average rebate generosity of \$2.20/sf, this implies a point elasticity of 0.93 with respect to generosity. Notably, including lagged rebate generosity barely moves the estimate on current generosity but itself has a magnitude about half as large. This suggests that the magnitude of takeup delayed by one year is approximately half that of same-year takeup, highlighting the need for careful identification of spillovers, since this observation reflects delayed takeup rather than a true spillover. Further note that the inclusion of water prices in the regression does not substantially change the magnitude of the estimate of interest, suggesting that it is not an important confound; however, including it tends to somewhat increase estimated elasticities.

Table 1.3 shows results of Equation (1.1) with various controls instead of the spatial discontinuity design with the border section by year fixed effects. Results are robust to the inclusion of year fixed effects, but are sensitive to spatial fixed effects like those for CBG and water agency. This suggests that the important identifying variation acts through spatial variation, and the deviations from a yearly average within a location relative to the location's average (the two way fixed effect model) do not provide enough variation to reliably estimate demand. The fact that Columns (1), (2), (5), (6), (7), and (10), those without spatial fixed effects, more closely match the preferred specifications in Table 1.2 suggests that removing the spatial variation from the analysis removes valuable identifying variation. For spillovers and heterogeneity analysis, I use the specification in Column (10) because it includes a variety of controls as well as year fixed effects, without spatial fixed effects detracting from the identifying variation. This estimate looks nearly identical when further including the drought status indicators (not shown). Column (10) indicates that a $1/sf increase in rebate generosity increases the probability of rebate takeup by 0.94 PP, for an elasticity of 0.82.

Table 1.2. Border RD Demand Estimation for California

Notes: Border regression discontinuity design estimates of the effect of rebate generosity on the linear probability of takeup. Each column represents a different inclusion threshold: parcel centroids within 500, 1,000, or 1,500ft of a water agency boundary. Agency boundaries that are within the given distance of a Census-designated Place boundary are excluded so that included parcels on either side of an agency boundary are in the same city.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated
Rebate Generosity ($/sf)	0.0074	0.0068	0.0076	0.0074	0.0066	0.0061	0.0088	0.0093	0.0066	0.0048	0.0079	0.0089
	(0.0021)	(0.0022)	(0.0042)	(0.0035)	(0.0016)	(0.0019)	(0.0051)	(0.0053)	(0.0012)	(0.0012)	(0.0054)	(0.0059)
	[0.001]	[0.002]	[0.071]	[0.035]	[0.000]	[0.001]	[0.085]	[0.082]	[0.000]	[0.000]	[0.146]	[0.135]
Lagged Generosity		0.0037				0.0045				0.0047		
		(0.0017)				(0.0015)				(0.0012)		
		[0.030]				[0.003]				[0.000]		
Water Price ($/HCF)				-0.0017				0.0027				0.0020
				(0.0007)				(0.0009)				(0.0008)
				[0.019]				[0.004]				[0.017]
Observations	1587726	1185864	237438	237438	2565019	1931254	394743	394743	3060287	2324541	498649	498649
YearXBorderFE	X	X	X	X	X	X	X	X	X	X	X	X
NonmissingWaterPrice			X	X			X	X			X	X
BorderDist	500ft	500ft	500ft	500ft	1000ft	1000ft	1000ft	1000ft	1500ft	1500ft	1500ft	1500ft
DepVarMean	0.0146	0.0146	0.0122	0.0122	0.0148	0.0148	0.0123	0.0123	0.0148	0.0148	0.0120	0.0120
IndepVarMean	2.0682	2.0682	2.2307	2.2307	2.0701	2.0701	2.2048	2.2048	2.0742	2.0742	2.1917	2.1917
Elasticity	1.05	0.96	1.39	1.35	0.93	0.86	1.58	1.67	0.93	0.67	1.44	1.63

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

Table 1.3. OLS Demand Estimation for California

Notes: Linear probability models of a the probability of takeup in a given year at the parcel level for California parcels as a function of the offered rebate generosity rate in dollars per square foot replaced. Demographic covariates include the CBG-level racial makeup and gender ratio from the Decennial Census, and the median age, median rent payment, poverty rate, maximum education level attained, and rate of utilities included in rent payment from the American Community Survey, as well as the acreage of the parcel. Drought status includes parcel indicators for number of months in the previous year in each of the US Drought Monitor's five drought status categories, with the omitted category being no drought. Variance(Generosity) shows the residual variance in rebate generosity after regressing it on the controls in each specification.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated
Rebate Generosity ($/SF)	0.0081	0.0135	-0.0021	-0.0019	0.0053	0.0087	0.0062	0.0004	0.0003	0.0094
	(0.0045)	(0.0043)	(0.0023)	(0.0019)	(0.0031)	(0.0049)	(0.0034)	(0.0011)	(0.0008)	(0.0028)
	[0.072]	[0.002]	[0.363]	[0.303]	[0.089]	[0.076]	[0.069]	[0.698]	[0.708]	[0.001]
Observations	31554035	31554035	31554035	31554035	27701580	31554035	27701580	31554035	31554035	27701580
YearFE		X						X	X	X
RetailerFE			X					X		
CBGFE				X					X	
Demographics					X		X			X
DroughtStatus						X	X			
DepVarMean	0.0255	0.0255	0.0255	0.0255	0.0259	0.0255	0.0259	0.0255	0.0255	0.0259
IndepVarMean	2.2457	2.2457	2.2457	2.2457	2.2606	2.2457	2.2606	2.2457	2.2457	2.2606
Elasticity	0.72	1.19	-0.19	-0.17	0.46	0.77	0.54	0.04	0.03	0.82
Variance(Generosity)	0.559	0.436	0.431	0.432	0.548	0.537	0.523	0.319	0.321	0.425

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

1.3.2 Spillovers

Using the demographic controls from Column (10) of Table 1.3, Table 1.4 shows the effect of lagged neighbor takeup on the current probability of takeup. Columns (1)-(4) are the OLS specifications corresponding to Equation (1.2), while Columns (5) and (6) are the instrumental variables estimates of Columns (4) and (2), respectively, corresponding to Equations (1.3) and (1.4). Because Columns (2) and (4) regress takeup on all lagged takeup in the CBG, they cannot disentangle the effect of delayed takeup from causal spillovers, they are presented in order to compare to the corresponding causal IV estimates. For convenience, the takeup count variable is divided by 100 so that the coefficient is interpreted as the effect of an additional 100 lagged conversions on the probability of current takeup.

While Columns (1) and (3) of Table 1.4 show average treatment effects that are statistically indistinguishable from zero, Table 1.5 demonstrates that effect sizes grow as the sample is restricted to homes with a larger portion of crossborder neighbors. The only estimates from Table 1.4 that are statistically significant are the two that are noncausal, however, the IV estimates reveal similar magnitudes albeit with much more noise, suggesting that there are indeed positive spillovers.

Turning attention to Table 1.5, one sees that the small negative effect on takeup of an additional 100 crossborder neighbors taking up the rebate turns large, positive, and highly significant as the sample is restricted to houses with a larger share of their neighbors being crossborder. “Cutoff” is the proportion of all neighbors in the CBG that are in a different water agency. Among households with at least 95% of their neighbors in a different water agency, an increase of 100 crossborder takeups in the previous year leads to an increase in the probability of takeup of 7.88 percentage points, or 625%. Equivalently, one additional crossborder turf replacement in year $y-1$ increases the probability of takeup in year y by 6.25%. Similarly for Columns (5)-(8) the effect of the fraction of lagged crossborder takeups increases as a larger share of neighbors are crossborder: A one percentage point increase in the share of crossborder neighbors who took up the rebate in year $y-1$ increases the probability of p taking up the rebate in year y by 0.467 percentage points, or 37%. In theory, the crossborder estimate should converge to a causal “all neighbors” estimate as the cutoff approaches 1.0 since this reflects an exogenous increase in the number of neighbors taking up the rebate without changing the terms of p’s rebate program since p is in a different water agency than its crossborder neighbors. Therefore, one should think of columns (4) and (8) as approximately the parameters of interest, noting that effective sample size decreases rapidly as the cutoff approaches 1.0, leading to noisier estimates past the 95% cutoff.

Table 1.4. CBG Level Crossborder Spillovers

Notes: For a given parcel p, *count lagged crossborder takeup / 100* is the number of homes in the same CBG as p, not in the same water agency as p that took the rebate in the previous year divided by 100. *Count lagged all takeup / 100* is the same but among all parcels in the CBG, both inside and outside p's water agency. *Frac elig crossborder takeup* and *frac lag all takeup* are the respective counts divided by either the total number of eligible crossborder neighbors or all eligible neighbors in the CBG. "Eligible" means that the parcel has not received a rebate by year y. For border analysis, all water agency boundaries are included, including those coincident with city boundaries. Includes only CBGs containing multiple water agencies. Controls include current and lagged own rebate generosity, ACS and Decennial Census CBG-level demographics, parcel acreage, and year fixed effects.

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	IV	IV
	rebated	rebated	rebated	rebated	rebated	rebated
Count Lagged Crossborder Takeup / 100	-0.0003					
	(0.0040)					
	[0.949]					
Count Lagged All Takeup / 100		0.0150				0.0071
		(0.0063)				(0.0119)
		[0.017]				[0.548]
Frac Lag Elig Crossborder Takeup			0.0089			
			(0.0114)			
			[0.432]			
Frac Lag Elig All Takeup				0.5284	1.4543	
				(0.1005)	(1.6387)	
				[0.000]	[0.375]	
Observations	3237308	3237308	3232479	3237308	2535733	2535733

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

Table 1.5. OLS Spillover Estimates by Share of Neighbors who are Crossborder

Notes: Table shows Columns (1) and (4) of Table 1.4 replicated on subsamples of the data based on the share of parcel p's neighbors that are in a different CBG. "Cutoff" determines the sample, so Column (2), for example, restricts to parcels with at least 75% of eligible parcels in their CBG in a different water agency. The limit of these estimates as "cutoff" approaches 1.0 is the causal estimate of the effect of lagged neighbor takeup on the probability of own takeup."Eligible" means that the parcel has not received a rebate by year y. For border analysis, all water agency boundaries are included, including those coincident with city boundaries. Includes only CBGs containing multiple water agencies. Controls include current and lagged own rebate generosity, ACS and Decennial Census CBG-level demographics, parcel acreage, and year fixed effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated
Count Lagged Crossborder / 100	-0.0006	0.0063	0.0721	0.0788				
	(0.0012)	(0.0059)	(0.0267)	(0.0327)				
	[0.635]	[0.284]	[0.007]	[0.016]				
Frac Lag Elig XBorder Takeup					0.0698	0.1213	0.4062	0.4671
					(0.0450)	(0.0807)	(0.2497)	(0.3474)
					[0.121]	[0.133]	[0.104]	[0.179]
Observations	632150	186872	41805	16769	632150	186872	41805	16769
DepVarMean	0.0131	0.0133	0.0123	0.0126	0.0131	0.0133	0.0123	0.0126
Cutoff	0.50	0.75	0.90	0.95	0.50	0.75	0.90	0.95

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

1.3.3 Heterogeneity

Lastly, I turn to heterogeneity in demand for rebates and application denials. Because parcel-level information about demographics is limited, results tables are split into sections that are identified for all parcels at the CBG level, and those that are available at the parcel level for a subset of the counties in the data. Column (1) of Table 1.6 reproduces Column (10) of Table 1.3 using the main California sample, but additionally showing the coefficients on the demographic controls, with the proportion of White households being the omitted category. Column 2 replaces the CBG level race proportions with an indicator for the CBG being above median White, and Column (3) adds to this an interaction of this indicator with rebate generosity. Column (3) suggests that households in Whiter CBGs are more likely to take up the rebates, but have a lower price elasticity than those in less White CBGs.

Column (4) restricts the baseline analysis to the counties for which homeowner names and thus predicted race are available (San Diego CA, Clark NV, and Maricopa AZ). The fact that the estimate of the effect of rebate generosity on takeup is the same for this sample as for the main California analysis suggests that this group is comparable and results are not driven primarily by composition effects. Column (5) shows the coefficients on parcel-level predicted race, and Column (6) further interacts these with rebate generosity, omitting the predicted White variable. Black households are substantially more likely to take up the rebates than White households, and Hispanic households are somewhat less likely than White households. However, Black households have a lower responsiveness to changes in rebate generosity while that for Hispanic households is higher than for White households. Asian households have statistically indistinguishable behaviors from their White counterparts.

While I do not directly observe wealth, I observe several proxies that are correlated with wealth such as education, parcel acreage, and CBG level median rent and poverty rate (Table 1.7). For a subset of the counties (San Diego, Clark, and Maricopa), I additionally observe assessed house value and square footage, indicators for which are included in Table 1.8. Households are more likely to take up the rebate if they are in CBGs with lower education, higher rent, lower poverty, and larger lots. Each of these is consistent with higher takeup among wealthy households except for the result on education, since one expects education to be positively correlated with wealth. Similarly, the rent, poverty, and lot size proxies imply households in wealthy areas are more responsive to changes in rebate generosity, while highly educated areas are less responsive. Table 1.8 tells a similar story for takeup rates at the household level, where households with higher home value and larger homes are more likely to take up the rebate. More valuable homes, however, are more responsive to rebate generosity changes, while larger homes are less responsive, all else equal. Nevertheless, because the sample of the data including house value is not the same as that containing square footage, these differences may be driven by compositional effects, highlighted by the differing elasticity in the baseline regression.

The final analysis shows regressions restricted to California applicants, with an indicator

Table 1.6. Heterogeneity in Demand by Race

Notes: Above median white refers to the given CBG having a proportion of White individuals higher than the median in the sample. p whi, p bla, p his, and p asi refer to the household level most likely race among White, Black, Hispanic, and Asian, with Other being the omitted category, based on name and location. Regressions by household race restrict to San Diego CA, Clark NV, and Maricopa AZ due to the availability of homeowner names in the parcel data.

	(1) rebated	(2) rebated	(3) rebated	(4) rebated	(5) rebated	(6) rebated
Rebate Rate	0.0094 (0.0028) [0.001]	0.0099 (0.0029) [0.001]	0.0115 (0.0032) [0.000]	0.0094 (0.0018) [0.000]	0.0072 (0.0016) [0.000]	0.0068 (0.0015) [0.000]
Rebate Rate × above_med_white			-0.0030 (0.0012) [0.009]			
Rebate Rate × p_bla						-0.0249 (0.0097) [0.011]
Rebate Rate × p_his						0.0047 (0.0020) [0.016]
Rebate Rate × p_asi						0.0045 (0.0039) [0.251]
Rebate Rate × p_oth						0.1367 (0.0717) [0.057]
DEC_race_black_rate	0.0231 (0.0036) [0.000]			-0.1088 (0.0125) [0.000]		
DEC_race_asian_rate	-0.0105 (0.0029) [0.000]			0.0033 (0.0146) [0.820]		
DEC_race_other_rate	0.0088 (0.0067) [0.188]			-0.0426 (0.0161) [0.008]		
DEC_male_rate	-0.0210 (0.0088) [0.017]			-0.1323 (0.0311) [0.000]		
above_med_white		-0.0007 (0.0010) [0.490]	0.0062 (0.0028) [0.027]			
p_bla					0.0354 (0.0095) [0.000]	0.0795 (0.0243) [0.001]
p_his					-0.0119 (0.0021) [0.000]	-0.0203 (0.0047) [0.000]
p_asi					0.0072 (0.0049) [0.141]	-0.0013 (0.0083) [0.873]
p_oth					-0.0064 (0.0031) [0.037]	-0.0306 (0.0132) [0.020]
p_acres	0.0114 (0.0004) [0.000]	0.0114 (0.0004) [0.000]	0.0114 (0.0004) [0.000]	-0.0046 (0.0006) [0.000]	-0.0043 (0.0006) [0.000]	-0.0043 (0.0006) [0.000]
Observations	27701580	27701580	27701580	49473	49473	49473
Sample	California	California	California	San Diego+Clark+Maricopa	San Diego+Clark+Maricopa	San Diego+Clark+Maricopa
YearFE	X	X	X	X	X	X
DepVarMean	0.0259	0.0259	0.0259	0.0302	0.0302	0.0302
IndepVarMean	2.2606	2.2606	2.2606	1.7289	1.7289	1.7289

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.∗∗ p< 0.10, ∗∗¡0.05,∗∗∗< 0.01

Table 1.7. Heterogeneity in Demand by Wealth Proxies for Full Sample

Notes: Highedu is an indicator for an above median proportion of people with at least a college degree in each CBG. Above median rent is the same for the median rent of each CBG. Above median acreage is a household level indicator for above median lot size. Sample includes California, Nevada, and Arizona homes.

	(1) rebated	(2) rebated	(3) rebated	(4) rebated	(5) rebated	(6) rebated	(7) rebated	(8) rebated	(9) rebated
Rebate Rate	0.0094 (0.0028) [0.001]	0.0134 (0.0044) [0.002]	0.0188 (0.0065) [0.004]	0.0136 (0.0044) [0.002]	0.0096 (0.0027) [0.000]	0.0135 (0.0043) [0.002]	0.0094 (0.0026) [0.000]	0.0218 (0.0040) [0.000]	0.0160 (0.0025) [0.000]
Rebate Rate × highedu			-0.0086 (0.0041) [0.039]						
Rebate Rate × above_med_rent					0.0062 (0.0033) [0.063]				
Rebate Rate × below_med_poverty							0.0080 (0.0040) [0.043]		
Rebate Rate × above_med_acreage									0.0112 (0.0041) [0.007]
highedu		-0.0043 (0.0033) [0.198]	0.0148 (0.0087) [0.088]						
above_med_rent				0.0193 (0.0049) [0.000]	0.0052 (0.0073) [0.482]				
below_med_poverty						0.0148 (0.0045) [0.001]	-0.0032 (0.0082) [0.693]		
above_med_acreage								0.0185 (0.0032) [0.000]	-0.0061 (0.0093) [0.511]
Observations	27701580	31131883	31131883	27701580	27701580	31131883	31131883	25617025	25617025
Sample	California	California	California	California	California	California	California	California	California
YearFE	X	X	X	X	X	X	X	X	X
DepVarMean	0.0259	0.0256	0.0256	0.0259	0.0259	0.0256	0.0256	0.0259	0.0259
IndepVarMean	2.2606	2.2483	2.2483	2.2606	2.2606	2.2483	2.2483	2.2037	2.2037

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

Table 1.8. Heterogeneity in Demand by Additional Wealth Proxies for Restricted Sample

Notes: Sample restricted to San Diego CA, Clark NV, and Maricopa AZ due to limited availability of home value and square footage in the parcel data. Measures are at the household level.

	(1) rebated	(2) rebated	(3) rebated	(4) rebated	(5) rebated	(6) rebated
Rebate Rate	0.0082 (0.0017) [0.000]	0.0068 (0.0017) [0.000]	0.0114 (0.0017) [0.000]	0.0091 (0.0023) [0.000]	0.0130 (0.0032) [0.000]	0.0084 (0.0019) [0.000]
Rebate Rate × above_med_value			-0.0115 (0.0016) [0.000]			
Rebate Rate × above_med_sqft						0.0117 (0.0032) [0.000]
above_med_value		0.0168 (0.0010) [0.000]	0.0379 (0.0034) [0.000]			
above_med_sqft					0.0140 (0.0030) [0.000]	-0.0111 (0.0069) [0.108]
Observations	19710840	21073759	21073759	28481973	31332510	31332510
Sample	San Diego+Clark+Maricopa	San Diego+Clark+Maricopa	San Diego+Clark+Maricopa	CA-Ventura+Maricopa	CA-Ventura+Maricopa	CA-Ventura+Maricopa
YearFE	X	X	X	X	X	X
DepVarMean	0.0290	0.0299	0.0299	0.0249	0.0250	0.0250
IndepVarMean	1.8193	1.8097	1.8097	2.1308	2.1152	2.1152

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

for whether the application was denied as the outcome of interest. Reviewers have little leeway for rejecting applications other than for not meeting the terms of the rebate program. Often, programs require the grass being replaced to be alive at the time of the conversion. Some also require pictures and/or inspections before any turf removal can take place, and detailed plans including the square footage removed and landscaping to be introduced. For the sample of rebates in the data that include contractor details (the more recent California rebates), more than a quarter of applicants list a contractor.[28] Because contractors frequently engage with the rebate programs, one would expect them to be more familiar with the requirements for successful applications. Indeed, by including an indicator for contractor use in the regressions, I show that households who use contractors are substantially less likely to have their application denied. In fact, most contractors have a 0% denial rate, with only a handful of contractors making up the bulk of the rejected applications that used contractors.

Table 1.9. Heterogeneity in Application Rejection by Contractor Use, Race, and Proxies for Wealth in California

Notes: Restricting the sample to only rebate applicants in California, Asian, Black, Hispanic, and Other represent the applicant's most likely race based on their name and location, with White being the omitted category. Median rent and Proportion at Least College are at the CBG level. Column (8) lets the contractor use indicator be the outcome to reveal differential trends in contractor use.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	App. Denied	App. Denied	App. Denied	App. Denied	App. Denied	App. Denied	App. Denied	Contractor Used
Contractor Used		-0.013 *** (0.002)	-0.017 *** (0.002)	-0.011 *** (0.002)	-0.014 *** (0.002)	-0.045 (0.028)	-0.036 (0.030)	
Asian	0.004 ** (0.002)	-0.003 *** (0.001)	0.004 ** (0.001)	0.005 *** (0.001)	0.006 *** (0.002)	0.006 *** (0.001)	0.005 *** (0.002)	-0.027 *** (0.005)
Black	0.003 (0.003)	0.008 *** (0.003)	0.004 (0.004)	0.005 * (0.003)	0.001 (0.004)	0.005 (0.003)	-0.001 (0.003)	0.378 *** (0.008)
Hispanic	0.056 *** (0.002)	0.006 *** (0.001)	0.004 *** (0.001)	0.008 *** (0.001)	0.005 *** (0.002)	0.006 *** (0.002)	0.005 *** (0.002)	0.219 *** (0.005)
Asian × Contractor Used		0.005 (0.003)	-0.000 (0.004)	-0.003 (0.004)	-0.003 (0.004)	-0.004 (0.004)	-0.004 (0.004)	
Black × Contractor Used		-0.008 ** (0.004)	-0.006 (0.004)	-0.006 (0.004)	-0.004 (0.005)	-0.005 (0.004)	0.001 (0.005)	
Hispanic × Contractor Used		-0.006 *** (0.002)	-0.003 (0.003)	-0.008 *** (0.003)	-0.007 ** (0.003)	-0.007 ** (0.003)	-0.005 * (0.003)	
Log Lot Acres				-0.004 *** (0.000)	-0.003 *** (0.001)	-0.004 *** (0.000)	-0.004 *** (0.000)	-0.013 *** (0.001)
Log Lot Acres × Contractor Used				0.004 *** (0.001)	0.004 *** (0.001)	0.004 *** (0.001)	0.004 *** (0.001)	
Log ACS Median Rent						-0.002 (0.002)	-0.001 (0.002)	-0.051 *** (0.006)
Log ACS Median Rent × Contractor Used						0.004 (0.004)	0.003 (0.004)	
Prop. at Least College						-0.005 (0.003)	-0.009 ** (0.004)	-0.399 *** (0.011)
Prop. at Least College × Contractor Used							0.009 (0.009)	
Observations	82304	78335	78335	59320	59320	52228	52228	52228
CBGFE			X		X			
RetailerFE							X	
DepVarMean	0.038	0.011	0.011	0.012	0.012	0.012	0.012	0.267

Standard errors in parentheses. P-values in brackets.

[28] Regardless of whether homeowners employ a contractor, it is the homeowner's name and not that of the contractor that is listed on the application.

Table 1.9 reveals that while Hispanic applicants are substantially more likely to have their application rejected than others, this effect is largely mediated by contractor use. Listing a contractor on the application is associated with a 1.3 PP lower likelihood of being rejected for White, Black, and Hispanic applicants, respectively (Column (2), sums of contractor use plus contractor$\times$ race minus race estimates). While about 1/4 of applicants report a contractor, those who are predicted to be Asian do so slightly less than White applicants, while Black and Hispanic applicants use contractors substantially more, 37.8 PP and 21.9 PP, respectively (Column (8)). Because contractor use almost guarantees application approval, this large discrepancy between Black and Hispanic applicants likely explains the baseline differences in rejection rates that disappears when controlling for contractor use. Because Hispanic applicants are more likely to primarily speak a language other than English relative to Black applicants, it is plausible that a language barrier keeps these applicants from successfully receiving the rebate. Moreover, because the chances of application rejection are decreased by having a contractor who regularly performs this type of conversion involved in the process, this suggests that familiarity with application requirements (through either a language barrier or an education barrier) may be at play. Nevertheless, this analysis cannot rule out other forms of bias in the application review process.

1.4 Discussion

In July 2021, California governor Gavin Newsom set a statewide goal for water conservation of 15% in the face of historic drought.[29] By fitting a demand curve for turf replacement programs, I can estimate how much of this can be fulfilled by turf replacement in the residential sector, as well as how high the rebate would have to be to induce this level of conservation in a given timeframe. First, I consider how long it would take for 15% of residential water to be conserved under average rebate terms in the data. Sovocool (2005) estimates residential water

[29]https://calmatters.org/environment/2021/07/california-water-use-drought/

savings to be 30% from turf replacement,[30] so 50% of all households would have to convert to reach this goal. Returning to Table 1.2 Column (4), average takeup is 1.48% per year at an average generosity of \$2.0701/sf, so to get 50% takeup it would take $50\%/1.48\%/year = 33.8$ years.[31] Assuming rebate generosity holds steady and demand dynamics are stable, turf replacement could lead to a 15% reduction in average residential water usage in California by 2057.

Suppose instead the exercise is to determine how generous of a rebate is needed to reach this goal in the next 10 years. With the linear demand function of the form: $takeup = \alpha + \beta generosity$, one has to extrapolate from the data, as rates have never been historically high enough to induce such rapid adoption. Setting $\beta = 0.66$, this means the intercept, α, equals 0.114. To have 50% additional takeup in 10 years, 5% is needed in each year. Solving for *generosity* in $5\% = 0.114\% + 0.66 \times generosity$ yields a necessary generosity of \$7.39/square foot, about 3.5 times the current average rate, but less than 50% higher than the highest generosity seen in the data, \$5/sf. Instead of using a linear demand curve, Appendix Section 1.A.2 finds a range of \$5.04-6.37/sf to meet this goal using exponential, quadratic, or cubic demand functions. Of course, it is unlikely that the West's water shortages will be solved exclusively by tearing out grass, but these back-of-the-envelope calculations provide insightful evidence that rebate programs may be one of several effective tools to reduce residential water demand.

As for spillovers, given that the median number of neighboring parcels in a CBG is 768, and assuming that spillovers only act within a CBG,[32] an additional rebate takeup in year $y-1$ is estimated to cause a 0.000788 percentage point increase in the probability of takeup among all 768 neighbors in period y, yielding an expected 0.61 increase in takeup the following year. A

[30]Brelsford and Abbott (2021) find a somewhat lower rate of water savings at 18%, while Baker (2021) finds a decrease of 19-21%.

[31]Strictly speaking, this calculation does not account for the conversions that have already happened. However, implicitly the goal is relative to the current level of usage, so indeed one would need an *additional* 50% of eligible households to convert beyond the current level.

[32]This assumption is unlikely to hold, but it is the level at which spillovers are identified in this analysis. The effect of spillovers likely rapidly decreases with distance as the desire for conformity is unlikely to be as strong for more distant neighbors, and the simple probability of encountering any given home decreases with distance.

rule of thumb used by the rebate programs is that one incentivized takeup causes another 1.0 (unincentivized) conversion. While I cannot account for unincentivized takeup, the estimated spillover of 0.61 does not rule out such a large magnitude for future incentivized conversions, particularly because my estimate is a lower bound: it accounts for only same-CBG spillovers in the following year, it does not account for spillovers that take longer to manifest, those that happen in the same year, or those that cross CBG boundaries.

Turning lastly to heterogeneity, the fact that Hispanic households are significantly less likely to receive these turf replacement rebates and less likely to be approved when they apply raises a question of environmental justice. Even though this group may have lower willingness to pay for participation in water conservation programs (as evidenced by lower takeup), higher application rejection rates indicate that there is a substantial share of Hispanic households who would like to convert their yards, but are denied the rebate. Some may still replace their turf anyway but not receive reimbursement, placing the financial burden of water conservation squarely on the household. Thinking about the inframarginality problem, where programs pay people who would have engaged in the activity without an incentive, there may be less inframarginality among Hispanic households because of higher rejection rates, but this creates an inequity where inframarginal households of other races are much more likely to receive payment.

1.5 Conclusion

As population growth and climate change cause water to become more scarce in arid regions, landscape conversion rebates pose an increasingly attractive option for water agencies to promote conservation. Research in economics and other fields have evaluated the effectiveness of xeriscaping in reducing outdoor water use, with generally positive results. This paper is the first to estimate a demand function for these rebate programs, using a spatial discontinuity design to find a price elasticity of 0.93 with respect to generosity. The magnitude of this result stands up to alternative specifications and after carefully controlling for potential confounds such as residential sorting, drought effects, and water prices.

Existing literature has tried to quantify the neighbor spillover effects and preferences for conformity. However, much of this existing work is in a descriptive rather than causal framework, with the major exception being Bollinger, Burkhardt and Gillingham (2020). This paper supplements the literature by using a novel cross-border approach to estimating spillovers, providing the advantage that this is able to disentangle the effect of delays in takeup from true neighbor spillovers, and supports this estimate with an instrumental variables strategy. I find that an additional neighbor taking up a turf replacement rebate in a given year induces an expected 0.61 additional takeups in their Census Block Group the following year. While I use a somewhat different metric from Bollinger, Burkhardt and Gillingham (2020), Brelsford and De Bacco (2018), and Pincetl et al. (2019), I corroborate their findings of the existence of positive spillovers.

In heterogeneity analysis, I find that Black households are most likely to take up the rebates, followed by Asian, White, then Hispanic households, while Black households have lower responsiveness to changes in generosity than other races. Proxies for wealth generally show increasing takeup and increasing demand elasticities as wealth increases. Finally, I show that while Hispanic rebate applicants are substantially more likely than other races to have their application denied, this effect largely disappears when controlling for contractor use. Since contractors drastically reduce rejection rates, Hispanic families' lower use of contractors than Black families leaves them with a higher rejection rate. Both still have higher rejection rates than Whites, despite lower contractor use, implying a true inequity in access to these pro-social environmental rebate programs.

This chapter is currently being prepared for submission for publication of the material. The dissertation author was the primary investigator and sole author of this material.

1.A Appendix

1.A.1 Additional Balance Tests

This section poses an alternate presentation of the balance test of covariate for border vs. nonborder homes and high generosity vs. low generosity sides of the border.

Border homes have a quarter acre smaller plot than nonborder homes on average, tend to be in less White and more Asian CBGs, and have somewhat lower educational attainment. While there are significant differences in water prices between border and nonborder homes, the direction of this average difference varies from year to year and is not systematic.

Table 1.10. Balance Test: Border Homes vs. Nonborder Homes

Notes: Means and P-value for a t-test of equality between homes within 1,000 feet of a water border (not coinciding with city boundaries) and those outside this buffer among the California sample. The first panel is household level data and has one observation per parcel. The second panel is CBG level data with one observation per CBG per border status (CBGs in a border region also usually have portions that are outside the 1,000ft buffer). Difference is the mean for homes in the border buffer minus the mean outside of the buffer. 'p' denotes parcel data, 'DEC' is from the Decennial Census, and 'ACS' is the American Communities Survey, reported for parcels in the sample in 2021.

Variable	Border Mean (SD)	Nonborder Mean (SD)	Difference	P-value
2021 Water Price ($/HCF)	3.617 (1.518)	3.197 (1.469)	0.420	0.0000
2020 Water Price ($/HCF)	3.658 (0.903)	3.331 (1.243)	0.327	0.0000
2019 Water Price ($/HCF)	3.450 (1.236)	4.574 (1.954)	-1.124	0.0000
2018 Water Price ($/HCF)	3.181 (1.307)	3.243 (1.297)	-0.062	0.0000
2017 Water Price ($/HCF)	2.758 (1.277)	4.367 (1.795)	-1.609	0.0000
2016 Water Price ($/HCF)	4.358 (1.214)	6.036 (1.743)	-1.678	0.0000
2015 Water Price ($/HCF)	2.523 (0.604)	2.435 (0.394)	0.088	0.0000
2014 Water Price ($/HCF)	2.365 (1.026)	3.281 (1.503)	-0.916	0.0000
p acres	0.298 (0.927)	0.561 (4.447)	-0.263	0.0000
p imp sqft	2427.327 (1129.717)	2423.159 (1282.333)	4.168	0.1102
Nonmissing Water Price (any year)	0.226 (0.419)	0.245 (0.430)	-0.019	0.0000
DEC race white rate	0.615 (0.222)	0.614 (0.218)	0.001	0.9416
DEC race black rate	0.065 (0.130)	0.078 (0.141)	-0.013	0.0008
DEC race asian rate	0.141 (0.161)	0.119 (0.131)	0.022	0.0000
DEC male rate	0.491 (0.025)	0.493 (0.031)	-0.002	0.0018
ACS median age	39.539 (8.257)	38.220 (8.562)	1.319	0.0000
ACS median rent	1670.324 (564.195)	1631.780 (571.125)	38.544	0.0319
ACS poverty rate	0.115 (0.101)	0.143 (0.124)	-0.028	0.0000
ACS school HS rate	0.217 (0.096)	0.199 (0.096)	0.018	0.0000
ACS school someCollege rate	0.301 (0.097)	0.281 (0.099)	0.020	0.0000
ACS school bachelors rate	0.201 (0.115)	0.215 (0.129)	-0.014	0.0001
ACS school graduate rate	0.110 (0.098)	0.122 (0.111)	-0.012	0.0002

Table 1.11. Balance Test: High Generosity Side of Border vs. Low Generosity Side

Notes: Means and P-value for a t-test of equality between homes on the generous side of the boundary compared to the less generous side, with the most recent three years of rebate generosity shown. Note that while there are significant differences in given years, these are not systematic between years (owing to the variation in which side of the boundary is the "high" side in a given year). Boundaries between more than two agencies exclude homes with neither the maximum nor minimum rebate. The first panel for each year (p acres and p imp sqft) is household level data and has one observation per parcel. The second panel (DEC and ACS variables) is CBG level data with one observation per CBG. Difference is the mean for homes on the high generosity side minus the mean on the low generosity side. 'p' denotes parcel data, 'DEC' is from the Decennial Census, and 'ACS' is the American Communities Survey, reported for parcels in the sample in 2021.

Variable	High Generosity Mean (SD)	Low Generosity Mean (SD)	Difference	P-value
2021 Water Price ($/HCF)	3.936 (1.097)	1.740 (1.256)	2.196	0.0000
2020 Water Price ($/HCF)	2.790 (0.000)	2.879 (1.071)	-0.089	0.5806
2019 Water Price ($/HCF)	3.274 (1.386)	3.258 (0.765)	0.016	0.5711
2018 Water Price ($/HCF)	5.255 (1.104)	1.843 (1.196)	3.412	0.0000
2017 Water Price ($/HCF)	3.446 (0.759)	5.671 (0.000)	-2.225	0.0000
2016 Water Price ($/HCF)	4.447 (1.936)	5.031 (0.555)	-0.584	0.0000
Nonmissing Water Price (any year)	0.095 (0.294)	0.405 (0.491)	-0.310	0.0000
By 2021 Rebate Generosity				
p acres	0.480 (1.544)	0.356 (1.052)	0.124	0.0000
p imp sqft	2475.760 (1164.180)	2503.666 (1188.406)	-27.906	0.0151
DEC race white rate	0.719 (0.180)	0.800 (0.119)	-0.081	0.0001
DEC race black rate	0.033 (0.034)	0.023 (0.026)	0.010	0.0144
DEC race asian rate	0.132 (0.159)	0.089 (0.077)	0.043	0.0111
DEC male rate	0.489 (0.021)	0.492 (0.028)	-0.003	0.4142
ACS median age	39.960 (9.408)	42.121 (9.483)	-2.161	0.0968
ACS median rent	1819.593 (571.697)	1989.461 (555.514)	-169.868	0.0545
ACS poverty rate	0.092 (0.071)	0.073 (0.063)	0.019	0.0370
ACS school HS rate	0.187 (0.082)	0.153 (0.087)	0.034	0.0043
ACS school someCollege rate	0.306 (0.097)	0.302 (0.100)	0.004	0.7848
ACS school bachelors rate	0.261 (0.123)	0.283 (0.117)	-0.022	0.1832
ACS school graduate rate	0.137 (0.088)	0.187 (0.107)	-0.050	0.0003
By 2020 Rebate Generosity				
p acres	0.306 (0.794)	0.241 (0.555)	0.065	0.0000
p imp sqft	2472.201 (1086.542)	2361.272 (846.559)	110.929	0.0000
DEC race white rate	0.613 (0.194)	0.619 (0.184)	-0.006	0.8727
DEC race black rate	0.073 (0.067)	0.082 (0.080)	-0.009	0.5635
DEC race asian rate	0.094 (0.132)	0.084 (0.087)	0.010	0.6246
DEC male rate	0.494 (0.016)	0.496 (0.022)	-0.002	0.7217
ACS median age	35.296 (7.598)	36.643 (8.437)	-1.347	0.4033
ACS median rent	1680.162 (685.473)	1626.824 (497.851)	53.338	0.6733
ACS poverty rate	0.135 (0.104)	0.115 (0.100)	0.020	0.3255
ACS school HS rate	0.208 (0.079)	0.244 (0.083)	-0.036	0.0292
ACS school someCollege rate	0.326 (0.115)	0.316 (0.090)	0.010	0.6449
ACS school bachelors rate	0.164 (0.123)	0.169 (0.105)	-0.005	0.8236
ACS school graduate rate	0.093 (0.091)	0.082 (0.075)	0.011	0.5184
By 2019 Rebate Generosity				
p acres	0.319 (1.075)	0.267 (0.771)	0.052	0.0000
p imp sqft	2515.333 (1180.685)	2393.994 (1126.804)	121.339	0.0000
DEC race white rate	0.676 (0.197)	0.648 (0.201)	0.028	0.1294
DEC race black rate	0.039 (0.074)	0.037 (0.054)	0.002	0.8099
DEC race asian rate	0.122 (0.144)	0.157 (0.164)	-0.035	0.0132
DEC male rate	0.488 (0.021)	0.493 (0.022)	-0.005	0.0090
ACS median age	40.668 (8.643)	40.318 (7.510)	0.350	0.6424
ACS median rent	1690.638 (546.584)	1780.988 (616.626)	-90.350	0.1247
ACS poverty rate	0.104 (0.088)	0.108 (0.091)	-0.004	0.6319
ACS school HS rate	0.206 (0.101)	0.200 (0.101)	0.006	0.4942
ACS school someCollege rate	0.289 (0.102)	0.290 (0.098)	-0.001	0.9304
ACS school bachelors rate	0.230 (0.127)	0.219 (0.113)	0.011	0.3354
ACS school graduate rate	0.130 (0.107)	0.138 (0.109)	-0.008	0.3958

1.A.2 Nonlinear Demand Estimation

Table 1.12. Nonlinear Crossborder Demand Estimation

Notes: Replication of the primary demand specification, Column 5 of Table 1.2, allowing the relationship between generosity and takeup to take a nonlinear form. Column 1 replicates the primary estimate (reporting more digits for a precise calculation) while Columns 2 and 3 are quadratic and cubic forms. Column 4 lets the independent variable be the exponential function of rebate generosity. Given that the specification is a fixed effects model, there is no single intercept; the reported intercept is the average of the fixed effects, and equals the predicted takeup level when rebate generosity is zero.

	(1)	(2)	(3)	(4)
	rebated	rebated	rebated	rebated
Rebate Generosity ($/sf)	0.006624	0.004425	0.004500	
	(0.0016)	(0.0051)	(0.0139)	
	[0.000]	[0.383]	[0.746]	
Generosity^2		0.000449	0.000415	
		(0.0010)	(0.0055)	
		[0.662]	[0.939]	
Generosity^3			0.000004	
			(0.0006)	
			[0.994]	
exp(Generosity)				0.000244
				(0.0001)
				[0.001]
Constant	0.001064	0.003580	0.003530	0.012520
	(0.0034)	(0.0063)	(0.0111)	(0.0008)
	[0.757]	[0.569]	[0.750]	[0.000]
Observations	2565019	2565019	2565019	2565019
YearXBorderFE	X	X	X	X
BorderDist	1000ft	1000ft	1000ft	1000ft

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

Solving for the rebate generosity that yields a 5% annual takeup for each model:

$$0.05 = 0.001064 + 0.006624 generosity$$

$$\implies generosity = \$7.39/sf$$

$$0.05 = 0.003580 + 0.004425 generosity + 0.000449 generosity^2$$

$$\implies generosity = \$6.37/sf$$

$$0.05 = 0.003530 + 0.004500 generosity + 0.000415 generosity^2 + 0.000004 generosity^3$$

$$\implies generosity = \$6.36/sf$$

$$0.05 = 0.0125201 + 0.0002437 e^{generosity}$$

$$\implies generosity = \$5.04/sf$$

1.A.3 Crossborder Demand with Lagged Generosity

Table 1.13 replicates Table 1.2 while additionally controlling for one year lagged rebate generosity. While lagged rebate generosity has a positive effect on current takeup, indicating that there is some delay in takeup, the magnitude of the coefficient on current generosity is largely unchanged. This indicates that results are not driven by delayed takeup.

Table 1.13. Demand Including Lagged Generosity

Notes: Primary crossborder demand specification additionally including 1 year lagged rebate generosity in the specification to partially disentangle delayed takeup from spillovers.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated	rebated
Rebate Generosity ($/sf)	0.0068	0.0079	0.0043	0.0061	-0.0043	-0.0019	0.0048	0.0009	0.0038
	(0.0022)	(0.0047)	(0.0038)	(0.0019)	(0.0078)	(0.0090)	(0.0012)	(0.0078)	(0.0088)
	[0.002]	[0.092]	[0.253]	[0.001]	[0.585]	[0.834]	[0.000]	[0.912]	[0.665]
Lagged Generosity	0.0037	0.0052	0.0060	0.0045	0.0057	0.0052	0.0047	0.0078	0.0072
	(0.0017)	(0.0021)	(0.0017)	(0.0015)	(0.0019)	(0.0021)	(0.0012)	(0.0024)	(0.0025)
	[0.030]	[0.013]	[0.001]	[0.003]	[0.004]	[0.015]	[0.000]	[0.001]	[0.005]
Water Price ($/HCF)			-0.0020			0.0013			0.0017
			(0.0008)			(0.0009)			(0.0008)
			[0.011]			[0.125]			[0.030]
Observations	1185864	193576	193576	1931254	326745	326745	2324541	418965	418965
YearXBorderFE	X	X	X	X	X	X	X	X	X
NonmissingWaterPrice		X	X		X	X		X	X
BorderDist	500ft	500ft	500ft	1000ft	1000ft	1000ft	1500ft	1500ft	1500ft
DepVarMean	0.0146	0.0122	0.0122	0.0148	0.0123	0.0123	0.0148	0.0120	0.0120
IndepVarMean	2.0682	2.2307	2.2307	2.0701	2.2048	2.2048	2.0742	2.1917	2.1917
Elasticity	0.96	1.45	0.79	0.86	-0.77	-0.34	0.67	0.16	0.70

Standard errors in parentheses, clustered at retailerXyear. P-values in brackets.

Bibliography

Addink, Sylvan. 2005. "'Cash for Grass'-A Cost Effective Method to Conserve Landscape Water?" *Turfgrass Research Facility, University of California, Riverside*, 14.

Baden, Brett M., Douglas S. Noonan, and Rama Mohana R. Turaga. 2007. "Scales of justice: Is there a geographic bias in environmental equity analysis?" *Journal of Environmental Planning and Management*, 50(2): 163–185. Publisher: Routledge _eprint: https://doi.org/10.1080/09640560601156433.

Baker, Jonathan E. 2021. "Subsidies for Succulents: Evaluating the Las Vegas Cash-for-Grass Rebate Program." *Journal of the Association of Environmental and Resource Economists*, 8(3): 475–508.

Banzhaf, Spencer, Lala Ma, and Christopher Timmins. 2019. "Environmental Justice: The Economics of Race, Place, and Pollution." *Journal of Economic Perspectives*, 33(1): 185–208.

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

Bollinger, Bryan, Jesse Burkhardt, and Kenneth T. Gillingham. 2020. "Peer Effects in Residential Water Conservation: Evidence from Migration." *American Economic Journal: Economic Policy*, 12(3): 107–133.

Brelsford, Christa, and Caterina De Bacco. 2018. "Are 'Water Smart Landscapes' Contagious? An Epidemic Approach on Networks to Study Peer Effects." *Networks and Spatial Economics*, 18(3): 577–613.

Brelsford, Christa, and Joshua K. Abbott. 2021. "How smart are 'Water Smart Landscapes'?" *Journal of Environmental Economics and Management*, 106: 102402.

Burkhardt, Jesse, Nathan W. Chan, Bryan Bollinger, and Kenneth T. Gillingham. 2021. "Conformity and Conservation: Evidence from Home Landscaping and Water Conservation." *American Journal of Agricultural Economics*. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajae.12224.

California Stater Water Resources Control Board. 2023. "Electronic Annual Report (eAR)." *https://www.waterboards.ca.gov/drinking*$_{w}$*ater/certlic/drinkingwater/ear.html*.

Imai, Kosuke, and Kabir Khanna. 2016. "Improving Ecological Inference by Predicting Individual Ethnicity from Voter Registration Records." *Political Analysis*, 24(2): 263–272. Publisher: Cambridge University Press.

Internet Archive. n.d.. "Wayback Machine." *https://web.archive.org/.*

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–563.

Jessup, Kelsey, J.R. DeShazo, and Ali Panjwani. 2019. "Turf Replacement Program Impacts on Households and Ratepayers: An Analysis for the City of Los Angeles." *UCLA Luskin Center for Innovation.*

Maricopa County Assessor. n.d.. "Premium Secured Master-Full County - R115." *https://mcassessor.maricopa.gov/page/data*$_s$*ales/.*

Mayer, Peter W, William B DeOreo, Eva M Opitz, Jack C Kiefer, William Y Davis, Benedykt Dziegielewski, and John Olaf Nelson. 1998. "Residential end uses of water." *American Water Works Association.*

Metropolitan Water District of Southern California. 2013. "California Friendly Turf Replacement Incentive Program Southern California Final Project Report."

National Drought Mitigation Center, U.S. Department of Agriculture, and National Oceanic and Atmospheric Administration. 2023. "GIS Data." *https://droughtmonitor.unl.edu/DmData/GISData.aspx.*

Pincetl, Stephanie, Thomas W. Gillespie, Diane E. Pataki, Erik Porse, Shenyue Jia, Erika Kidera, Nick Nobles, Janet Rodriguez, and Dong-ah Choi. 2019. "Evaluating the effects of turf-replacement programs in Los Angeles." *Landscape and Urban Planning*, 185: 210–221.

San Diego Association of Governments. 2022. "SanGIS GIS Data Warehouse."

Southern California Association of Governments. 2016. "2016 Land Use Information for Ventura County." *https://gisdata-scag.opendata.arcgis.com/datasets/SCAG::2016-land-use-information-for-ventura-county/about.*

Southern California Association of Governments. 2019. "2019 Annual Land Use." *https://hub.scag.ca.gov/datasets/ea9fda878c1947d2afac5142fd5cb658*$_0$*/about.*

Sovocool, Kent A. 2005. "Xeriscape Conversion Study." *Southern Nevada Water Authority*, 93.

US Census Bureau. 2010. "Decennial Census of Population and Housing Data." *https://www.census.gov/programs-surveys/decennial-census/data.html*, Section: Government.

US Census Bureau. 2017. "American Community Survey 5-Year Data (2009-2021)." *https://www.census.gov/data/developers/data-sets/acs-5year.html*, Section: Government.

US Census Bureau. 2022. “TIGER/Line Shapefiles.” Section: Government.

Chapter 2

Heat Islands in a Sea of Water Conservation: Heat Costs of Turf Replacement Programs

Abstract

Homeowners face a water versus heat tradeoff when it comes to replacing thirsty grass lawns with water-conserving landscaping, a process touted by water agencies in the U.S. West as a top way to conserve water. Using nearly 200,000 rebate records and remotely-sensed temperature data from Southern California and Southern Nevada, I assess the effect of water-conserving lawn replacement programs on local temperatures and estimate the associated costs of increased heat. Conversions increase summertime parcel temperatures by 0.6°C (1.1°F) on average with substantial heterogeneity throughout the temperature spectrum. Heat effects are twice as large on the hottest 20% of days and for homes with the most removed vegetation. Back-of-the-envelope calculations suggest that the annual costs from increased heat may be up to \$1,675 per household, comprised of increased electricity usage (\$48) and mortality risk (\$306), as well as harder-to-quantify comfort values, diminished cognition, and costly adaptation behavior. These costs exceed the maximum expected water savings for a typical home in Southern California (\$954) and Las Vegas (\$574), where municipal water is cheaper. This suggests that such rebate programs may not be welfare-improving.

2.1 Introduction

Outdoor water usage makes up as much as 3/4 of household water consumption, largely due to irrigated lawns (Mayer et al., 1998). As water becomes more scarce, water agencies throughout the American West have begun offering incentives for households to replace their grassy lawns with more drought-tolerant landscaping, a process known as 'xeriscaping'. Such replacements, or 'turf conversions', relieve strain on scarce water resources in arid regions by converting thirsty non-native turf grass to desert-adapted vegetation, rocks, and mulch. However, removing grass means forfeiting the cooling effects provided by lush vegetation through evapotranspiration. In this paper, I document the heat effects of removing vegetation that plausibly outweigh the private benefits from water savings. The average yard in my sample experiences a 0.6°C (1.1°F) increase in summertime surface temperature after converting from a grass lawn to xeriscaping, an effect that is doubled on the hottest days and for yards with the most vegetation removed. I show that surface temperature is a good proxy for air temperature when measuring differences-in-differences, even though the two metrics may vary in levels. Using Klaiber, Abbott and Smith's (2017) estimate of households' hedonic valuation of localized temperature in Phoenix, Arizona, this temperature change is responsible for a welfare cost of $1,675 per year per household. In contrast, Sovocool's (2005) estimates of the private water savings of these conversions yield up to $954/year for the median home under the highest marginal billing rate for Southern California, and substantially less for lower billing tiers and those in other water districts. Optimal water policy demands a thorough understanding of the costs and benefits of these turf replacement programs, yet the heat costs have not been well documented, despite the programs' highly-touted benefits.[1] Much of the pushback to turf replacement has been on aesthetic grounds, but I show that the heat effects are too large to ignore.

Landscape conversions have been proposed as one of the most impactful ways for households to conserve water. In addition to indoor and outdoor household use, water provides value to society through agriculture, industry, recreation, ecological services, and existence values, many of which are not fully internalized in the market (Young and Loomis, 2014). Because of this market failure and misaligned incentives (Timmins, 2002), overuse of water

[1]Benefits primarily include 20-30% water savings, or 55.8 gallons per square foot per year (Baker, 2021; Brelsford and Abbott, 2021; Sovocool, 2005), but also ecological benefits such as restoring native habitats for animals. In addition, there may be substantial social benefits to water conservation, see Young and Loomis (2014) for a thorough discussion of the private and public valuation of water.

relative to the economic optimum coupled with extended drought has led to increasing scarcity. Regulators can step in to help correct the market by implementing conservation programs like turf replacement.[2] However, due to water's variety of both use and non-use values, assigning a single number to the social value of conserved water is difficult. Thinking of the value of water conservation as the shadow value of water, recent water shutoffs in Arizona suggest that the short term value of an additional hundred cubic feet of water is about $100.[3] However, in the long run there is time to adapt by investing in more water-efficient capital (landscaping, appliances, etc.) so the long run shadow price of water is likely lower. This is in contrast to typical prices for piped municipal water of $2-12/HCF. Nevertheless, these costs are not borne by the vast majority of households who receive water at a price lower than its marginal social cost.

This paper is concerned with the household decision problem and focuses on the private costs of xeriscape rebates, while acknowledging the existence of unaccounted-for social benefits of water conservation[4] and abstracting away from up-front costs of replacement, which I do not observe and can vary substantially with the type and quality of new landscaping used.

To estimate the increase in temperature from xeriscaping, I use geocoded turf replacement rebates from Southern California and Nevada along with thermal infrared (TIR) satellite imagery which records the temperature of the Earth's surface. This panel dataset allows me to document within-household temperature changes over time back to 1999 when the satellite began collecting data. Using an event study with matched controls, I find a precise 0.6°C (1.1 °F) increase in the average summertime surface temperature of a property immediately following a landscape conversion, after correcting for measurement error. Consistent with Braun and Schlenker's

[2] Correcting prices through Pigouvian taxes is difficult for multiple reasons. First, many view water as a necessity that should be cheaply available to all. Second, pinning down the true social marginal cost of water is extremely difficult given the variety of uses that depend on location, quality, and timing. Third, complex historical legal claims to water rights complicate the design of an efficient market.

[3] $13.40/10,000gal = $100.2/HCF. "Arizona city cuts off a neighborhood's water supply amid drought" *The Washington Post*. January 16, 2023. https://www.washingtonpost.com/climate-environment/2023/01/16/rio-verde-foothills-water-scottsdale-arizona/

[4] Estimates for the recreation value of water, one component of the social benefit, range from $0.01 to $1.37 per hundred cubic feet (HCF) (Loomis, 2012; Ward, Roach and Henderson, 1996; Young and Loomis, 2014). The upper bound is non-trivial relative to typical municipal water rates of $2-12/HCF, but still small compared to the highest rates.

(2023) observations of temperature heterogeneity from agricultural irrigation, I show that the heat effect from xeriscaping is largest on the hottest days when the health risks of heat are already elevated, and is larger for parcels with greater vegetation losses. Examining related outcomes, I estimate the increase in household electricity usage to be 57.6 kWh/year in Southern California and use the heat mortality estimates from Deschênes and Greenstone (2011) to predict the annual cost of increased heat mortality risk to be $76.42 per person exposed to increased heat, or $306 for a family of four. Under a scenario in which all Southern Californians have a xeriscaped yard, I predict an additional 142 heat deaths per year, leading to $1.5 billion in annual losses. This predicted increase in mortality suggests that the rebates in my data are currently responsible for 3.75-5 additional heat deaths per year.

To the best of my knowledge, the heat versus water tradeoff of xeriscaping has not been rigorously studied. This analysis ties together three distinct literatures: a small body of economic work on xeriscape rebates, a larger interdisciplinary collection of evaluations of these water conservation programs, and a vast corpus of work by climatologists, urban planners, physicists, and sustainability scientists (among others) on urban heat island effects, a subset of which examines xeriscaping with small sample sizes. The urban heat island (UHI) is a phenomenon in which urban and suburban areas have increased air and surface temperatures relative to their surrounding rural environment (Buyantuyev and Wu, 2010). This temperature difference is largely due to two factors characterizing urban areas: decreased vegetation and increased hardscape (buildings and pavement), highlighting the role of plants in urban microclimates.

There is little mention of temperature among the program evaluation literature on landscape rebates. Baker (2021) is the only study to mention temperature effects of turf replacement rebates, finding a 3% increase in energy use following conversion without directly measuring temperature. In their analysis of the Las Vegas "Cash for Grass" program (a subset of the rebates in my data), Brelsford and Abbott (2021) show that xeriscaping lowers household water usage by about 20%, and that rebate programs are a cost-effective method of water conservation relative to other options faced by water suppliers. In a field experiment, Sovocool (2005) finds larger water savings of 30%, with 55.8 gallons saved annually per square foot replaced. Moreover, in an agronomic analysis of landscape conversion programs, Addink (2005) notes the cooling effect of grass as a potential downside of turf replacement. Chow and Brazel (2012), however, suggest that even desert-friendly trees can help cool, making the ex ante effect of conversion ambiguous. Bollinger, Burkhardt and Gillingham (2020) and Burkhardt et al. (2021) use remote sensing to characterize household landscape changes and peer effects, but do not assess temperature or the related costs. Both Bollinger, Burkhardt and Gillingham (2020) and Klaiber, Abbott and Smith (2017) perform hedonic analyses related to xeriscaping, finding opposite effects on home values. However, in this paper I utilize Klaiber, Abbott and Smith's analysis which estimates willingness to pay for a one degree temperature change holding all else (including landscaping) constant. In contrast, Bollinger, Burkhardt and Gillingham's estimates are the bundled effects of a landscaping change, which make it hard to disentangle pure heat effects from the effects of a newly renovated yard. Finally, on a county scale, Braun and Schlenker (2023) show substantial cooling effects of agricultural irrigation.

The UHI literature on the effect of landscaping on temperature is more developed, although the relevant work is overwhelmingly cross-sectional, which ignores the potential bias of endogenous selection into landscape conversions. Of the few papers that look at within-location landscape changes,[5] landcover is measured coarsely using changes in the Normalized Difference Vegetation Index (NDVI) rather than a ground-truthed measure of a xeriscape conversion, and not calibrated to reflect household-level conversions. This elucidates the relationship between temperature and remotely-sensed landscaping indices without providing the causal effects of xeriscaping at the household level. These studies fall into the category of papers in which remote sensing allows for a large sample size, but the spatial resolution is low. Coarse resolution is not an issue in these papers, as the analysis is not meant to be at the household level.

In the other category of related UHI work are the small scale studies which perform analysis at the household level. Many of these studies use simulations and models of the physical

[5]See Benz, Davis and Burney (2021), Lo, Quattrochi and Luvall (1997), and Stabler, Martin and Brazel (2005).

world using climatology and fluid- and thermo-dynamics to simulate counterfactual landscaping arrangements in neighborhoods from which data is measured.[6] Such modelling critically assumes a knowledge of the true relationship between landscaping choices and temperature; in contrast, this paper identifies this relationship completely empirically. Shashua-Bar, Pearlmutter and Erell (2009) and Shashua-Bar, Pearlmutter and Erell (2011) empirically show the cooling effects of vegetation in Israel, albeit with a very small sample size of two. To my knowledge, this is the first paper to examine within-household temperature effects of actual xeriscape conversions on a large scale.

The remainder of this paper is organized as follows. Section 2.2 describes the data used and construction of key variables and Section 2.3 defines the estimator used in the main analysis. Section 2.4 shows the main results as well as heterogeneity by baseline temperature and remotely sensed vegetation change. Section 2.5 quantifies the heat costs of xeriscaping, and Section 2.6 concludes.

2.2 Data

2.2.1 Turf Replacement Rebates

Over the past two decades, many water providers in the U.S. Southwest have offered incentives for replacing lush turf grass with a less water-intensive alternative. This rebate is typically either a flat payment or a rebate per square foot replaced, with a cap. The typical per-square-foot rate varies substantially from \$0.25 to \$6.00, with significant variation over both time and space. Most rebates subsidize but do not fully cover the cost of a landscape conversion.

I use data from public record requests of three large water wholesalers: Southern Nevada Water Authority (covering Clark County, NV, home of Las Vegas), Municipal Water District of Orange County (a densely suburban county between Los Angeles and San Diego, California) and Metropolitan Water District of Southern California (the rest of coastal Southern California).

[6] See Chow and Brazel (2012), Gober et al. (2009), (Middel et al., 2007), Middel et al. (2014) Pearlmutter, Berliner and Shaviv (2007), and Shashua-Bar and Hoffman (2004).

Collectively, these agencies cover a population of about 20 million, and with rebates back to 2011 in California and 1996 in Nevada, they account for approximately 175,000 residential landscape conversions.

Rebate data includes the address, date issued, square footage removed, and amount paid. I use ESRI's geocoder in ArcGIS to assign addresses to a latitude and longitude in order to merge with other spatial datasets (temperature and parcel boundaries). The time variable I use, the date of rebate *issuance*, necessarily lags the actual date of conversion, motivating the primary analysis at the yearly (rather than monthly) level. Relative to the date households pre-registered for the rebate (which must be done before any grass is removed) in the California data, the rebate is typically issued at least a month, or as much as 1.5 years after the rebate process is initiated, with an average of about 6 months. Thus, it should not be surprising to see some effects in the calendar year prior to the issuance of the rebate.

2.2.2 Parcels

Parcels represent the legal boundary of a property owner's land, data for which is commonly maintained by county assessors' offices for tax purposes. I use publicly available GIS shapefiles of parcels in Los Angeles, Orange, San Diego, Riverside and San Bernardino counties in California,[7] and restricted data from Clark County, Nevada, which was made available by the assessor for research purposes. While most of these files include additional information about the parcels (ownership, zoning type, number of rooms, etc.), the varying sources of the data means that these values are not consistently available in every county so I do not include them as covariates.

2.2.3 Remotely Sensed Temperature and Vegetation

While the true outcome of interest in this study is air temperature, surface temperature is the only temperature data available at a fine enough resolution that has not been modeled

[7]County of Los Angeles (2022); County of Orange California (2020); Riverside County (2019); San Bernardino County (2022); San Diego Association of Governments (2022).

from a sparse network of temperature sensors.[8] Surface temperature is measured by detecting the infrared radiation emitted by the surface of the Earth, which is proportional to temperature. This measure comprises the thermal infrared (TIR) band of several publicly accessible satellite datasets; I use Landsat 7 for its relatively high spatial resolution, and time series back to 1999.[9] The source data is carefully calibrated by remote sensing experts to reflect true temperatures, and daily maximum air temperatures closely track surface temperatures empirically (see Appendix 2.A.1). Nevertheless, a literature in the physical sciences has shown that the harmony between air temperature and surface temperature can vary when using temperature *levels*.[10] However, I show in Appendix 2.A.1 that for measuring *differences-in-differences* (as the event study in this paper does), there is no significant difference between air temperature and surface temperature. In fact, the discrepancy when comparing surface temperature to air temperature is no larger than the discrepancy when comparing one air temperature reading to that from another air temperature sensor of the same type when measuring differences-in-differences.

In order to prevent missing data resulting from clouds and Landsat 7's scan line correcter failure,[11] I take a composite of summertime cloud-free images. Specifically, my primary temperature measure is the median (across images) surface temperature over nonmissing pixels in a given location from June 1 to September 1 of a given year after removing pixels with a cloudiness likelihood of at least 10%. Landsat 7 has a 16-day revisit time, so this median is taken over up to six observations in this three-month window. In addition to minimizing missing data, this strategy helps to control for daily variation in weather, creating a smoothed metric of typical

[8] Any gridded air temperature product inherently has some degree of interpolation, which is likely agnostic to vegetation and certainly ignorant of household-level year-over-year changes in landscaping. Thus it provides no value as an outcome measure in this context, as it mechanically does not respond to changes in household landscaping decisions.

[9] U.S. Geological Survey (2022).

[10] See for example Schneider et al. (2023) Hu and Brunsell (2015) Schwarz et al. (2012) Pichierri, Bonafoni and Biondi (2012) Vancutsem et al. (2010) Fung et al. (2009) Hartz et al. (2006), and Dousset and Gourmelon (2003).

[11] The 2003 failure of Landsat 7's scan line corrector means that no single image is complete but does not bias the measurement. Nevertheless, I use several images in each year to compensate for the image gaps. According to the USGS: "Although these scenes only have 78 percent of their pixels remaining after the duplicated areas are removed, these data are still some of the most geometrically and radiometrically accurate of all civilian satellite data in the world" (U.S. Geological Survey, https://www.usgs.gov/landsat-missions/landsat-7).

summertime temperature. While Landsat 7's visible light bands yield pixels 30 meters wide, the TIR band is coarser at 60m. However, Google Earth Engine and the U.S. Geological Survey provide all Landsat data at 30m using a cubic spline algorithm to resample the 60m TIR band to 30m. This means that the actual variation in measured temperature happens at the 60m level, and is interpolated to 30m.

For the temperature heterogeneity analysis, I do not aggregate temperatures over time. Instead, I select the first Landsat observation of a given area in each month. Since the cloud removal algorithm above relies on having a time series of images (for example, the three-month summer) to determine cloudiness, I take a different approach when using single-day observations. Any pixels for which the USGS quality assurance pixel indicates the presence of clouds are removed, leaving missing data for that pixel$\times$month.

The three-month summer metric and the disaggregated monthly temperature measure are mapped to parcels by taking the spatial average of the Landsat pixels covering a given parcel, weighted by the percentage of the pixel that intersects the parcel. This creates a panel dataset with one temperature observation for each parcel in each time period. I use data from 1999 through 2022 and restrict to a panel balanced in event time for all analyses.

2.3 Empirical Strategy

Relative to the existing literature on the relationship between heat and vegetation, this study is better suited to identify the causal effect of turf replacement programs for several reasons. First, it is identified off of a large sample of actual rebate-takers, making these results more reflective of real-world conversion programs compared to studies examining broader vegetation changes or those looking at specific small-scale case studies. The other advantages come from my large panel data, since most other studies either measure cross-sectional differences which ignore selection or very small sample studies that likely do not represent a typical rebate taker. By comparing to a pure control that is very near the converted household, I am able to flexibly account for time-constant within-household factors as well as local microclimates that may vary over time. This helps to control for omitted variable bias from highly localized weather patterns on the day of measurement, as well as selection into landscape conversion programs. It is plausible that households in hotter, drier areas (even within a given city) are more likely to xeriscape because of the increased water cost of maintaining a grass lawn. Thus a naïve cross-sectional regression of temperature on landscaping status would overstate the heat effects of turf replacement. Alternatively, homeowners with cooler baseline temperatures might be more willing to remove vegetation, as the heat effects would be more tolerable at a lower starting temperature, leading to an underestimate. Thus the cross-sectional regression is potentially biased, but the sign of the bias is ambiguous. In contrast, my difference-in-differences strategy overcomes this bias by comparing the change in temperature of a xeriscaped home to the changes in a nearby control home which is measured by the satellite at the same time, and thus under the same weather conditions.

2.3.1 Matching Estimator

Given the novelty of my large panel structure in identifying heat effects of xeriscaping, my primary specification assesses within-parcel temperature changes as a function of the time relative to conversion. The natural estimator given the data is a two-way fixed effects (TWFE) event study model with time and location fixed effects. While this TWFE model yields similar results, to account for the effects of highly localized (e.g. neighborhood level) weather patterns which might endogenously relate to landscaping choices or differences in the date of satellite measurement or geography,[12] I employ a matching strategy where each converted parcel is

[12]Baseline temperatures may be positively or negatively correlated with the decision to xeriscape. On one hand, hotter temperatures demand higher water expeditures, but on the other hand, cooler locations would be more able to comfortably absorb a temperature increase from landscape conversion. These temperatures likely correlate with elevation and how far from the city center a home is. Furthermore, the farther apart two points are, the higher the chance that the Landsat satellite measures them on different dates, given that one pass of the satellite measures a finite band of the Earth, and farther apart locations are more likely to have differing weather, potentially removing one due to clouds but not the other. These considerations motivate the use of a control home from the same neighborhood.

matched to a nearby parcel that is not in the turf replacement rebate data. Specifically, controls are drawn from parcels intersecting the donut more than 275ft (84m) but less than 350ft (107m) from the treated parcel[13] to ensure that the control parcel is very close to the converted parcel without risking contaminating the same Landsat pixel by being too close. From this set of candidate control parcels, I exclude any that are less than half or more than twice the area of the rebated parcel[14] and select one from the remainder at random.

For the primary specification, I use one-to-one exact matching, where the match is based simply on location: converted parcel i is matched to the untreated parcel that was selected as described above. Letting $T_{i,t}^{treated}$ be the temperature of treated parcel i at time t and $T_{i,t}^{control}$ be the temperature of its match, the matching estimator is the sequence of βs arising from:

$$T_{i,t}^{treated} - T_{i,t}^{control} = \alpha_i + \sum_{\tau=-5,\ \tau\neq-2}^{5} \beta_\tau \mathbb{1}\{t - conversionYear_i = \tau\} + \varepsilon_{i,t} \tag{2.1}$$

where α_i is a household pair fixed effect (allowing for one parcel to be hotter at baseline), $\mathbb{1}$ is the indicator function, and τ represents event time, letting $\tau = -2$ be the omitted reference year, since some conversions have already begun in $\tau = -1$. There is no need for time or location fixed effects because matching on the control parcel in a given location and in the same time period inherently captures and differences out this variation. This model is equivalent to a model where the left hand side is temperature, and the right hand side includes event time and parcel pair by year fixed effects. Standard errors are clustered at the parcel pair level to allow for serial correlation in the errors and other potential correlated errors. This model is identified off of the assumption that other than xeriscaping, neither the treated nor the control parcels change their lots in a way that is correlated with both event time and temperature; there are few ways other than landscaping to manipulate the surface temperature of one's lot.[15]

[13] A TIR pixel is 60m on each side, or $60\sqrt{2} \approx 84.9$m$\approx$ 278ft across the diagonal.

[14] Conditioning on similar sized parcels helps to confirm that the parcels are of a similar type in the absence of other covariates on which to condition. Furthermore, because parcel datasets do not all include zoning, it is possible that large parcels are commercial or industrial rather than residential and thus make for a poor control.

[15] It is plausible that households take up multiple incentive programs at once. For example, a home doing a

Lastly, in order to prevent concerns about the interpretation of the regression estimates due to the staggered rollout design (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Sun and Abraham, 2020), the primary specification restricts to a panel balanced in event time and includes pure controls (never-treated units). By comparing each treated unit only to its matched control in any given year (effectively a household pair by year fixed effect rather than a TWFE), this estimator avoids the negative weighting problem arising from the "forbidden comparison" of treated households to already-treated households that arises from a TWFE model (de Chaisemartin and D'Haultfoeuille, 2022; Goodman-Bacon, 2021).

2.3.2 Parcel Orientation Correction

Equation (2.1) yields unbiased estimates of the temperature change caused by landscape conversion in the absence of measurement error. However, due to the fact that the typical parcel is smaller than the 60m Landsat pixel, the temperature change in the pixels understates the true effect. Take for example, a parcel that occupies $1/4$ of a pixel and whose temperature increased by 1°C after the conversion (while the other $3/4$ of the pixel remains unchanged). Because the Landsat sensor aggregates the temperature change over the whole 60m pixel, rather than seeing a 1°C change on a quarter of the pixel, it sees a 0.25°C change on the whole pixel, biasing estimates downward. Moreover, if a parcel spans multiple pixels, its measured temperature change is diluted by the untreated portions of all the pixels it touches. This paper's primary analysis rescales the measured effect, given data on the shape, size and orientation of each parcel.

A multiplicative correction can be defined assuming that the temperature of the entire parcel increases by the same amount, and under the mild assumptions that the temperature of

turf replacement might also install solar panels and/or a new air conditioning unit at the same time. Both of these might affect the heat characteristics of the parcel, as solar panels are designed to absorb solar radiation, while air conditioners transfer heat from the interior of a home to the exterior (perhaps with more efficiency after an upgrade). However, both of these examples occupy substantially smaller space than the typical landscaped area of a yard, suggesting the heat effect will be small relative to that of removing vegetation. Furthermore, I assume that the percentage of households for which such other upgrades coincide with turf replacement is small. A landscape redesign might also include the installation of a swimming pool. While not eligible for a turf replacement rebate, the pool should be considered a bundled amenity (with respect to heat) to the extent that it is also induced by the rebate.

unconverted parcels does not change relative to the treated parcel's counterfactual, and that a pixel measurement is the average surface temperature of the area it spans. Letting $\mathscr{P}$ be the set of pixels that a given parcel i intersects, ΔT be the measured temperature change equivalent to Equation (2.1), and $\Delta\tau$ be the actual treatment effect of landscape conversion. Then, letting s_p be the share of pixel $p \in \mathscr{P}$ covered by the parcel, the correction multiplier m_i is such that:

$$\begin{aligned} \Delta\tau &= m_i \Delta T \\ \text{for } m_i &= \frac{\sum_{p \in \mathscr{P}} s_p}{\sum_{p \in \mathscr{P}} s_p^2} \end{aligned} \tag{2.2}$$

So the estimated actual treatment effect is the observed temperature difference (after partialing out the parcel pair fixed effect) times the ratio of summed pixel coverage to the sum of squared pixel coverage. See Appendix 2.A.3 for the proof. In practice, I partial out the fixed effects by regressing the temperature difference, $T_{i,t}^{treated} - T_{i,t}^{control}$ on the α_is, then multiply this residual by the scale multiplier for the respective parcel and regress on the event time indicators to retrieve the series of β_τs:

$$T_{i,t}^{treated} - T_{i,t}^{control} = \alpha_i + \eta_{i,t} \tag{2.3}$$

$$m_i \widehat{\eta_{i,t}} = \sum_{\tau=-5,\ \tau \neq -2}^{5} \beta_\tau \mathbb{1}\{t - conversionYear_i = \tau\} + \varepsilon_{i,t} \tag{2.4}$$

2.4 Results

2.4.1 Main Results

Relative to the periods before the turf replacement, the periods after have increased surface temperatures of 0.6°C (1.1°F), as shown in Figure (2.1). There is a slight uptick in temperature in the year before the rebate is issued, which can likely be explained by two factors. First, the rebate issue date lags the actual conversion so some conversions actually take place in period -1, and second, households may limit watering and let their yards die off in the months leading up to conversion, increasing the temperature.[16]

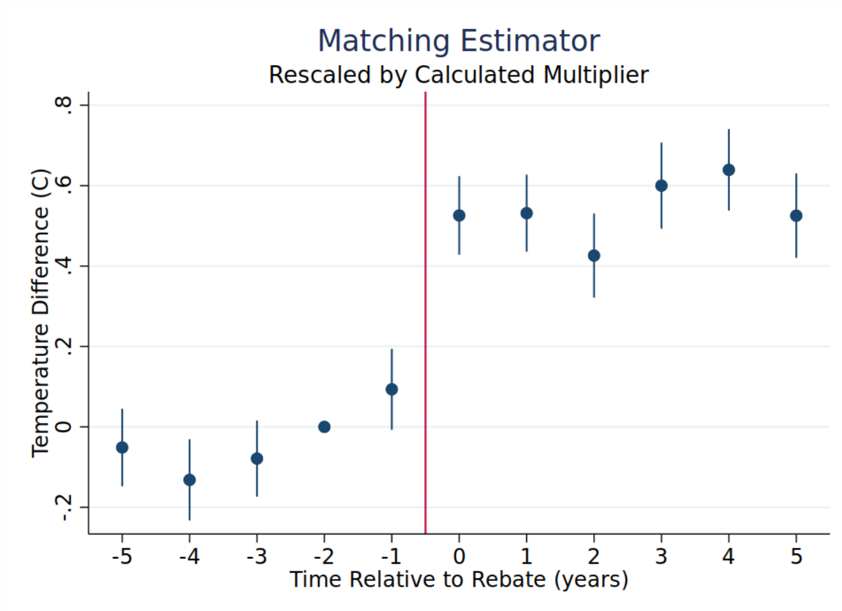

Figure 2.1. Temperature Effect of Landscape Conversion

Notes: Figure shows the matching difference-in-differences estimates for summertime surface temperature of converted parcels relative to their matched control. 95% confidence intervals shown with standard errors clustered at the parcel pair level. The outcome, temperature differential, is scaled by the calculated multiplier as described in Section 2.3.2.

2.4.2 Temperature Heterogeneity

In this section, I focus on the heterogeneity of heat effects across the temperature spectrum by interacting the event time indicators with temperature bins. This analysis requires a refinement of the Landsat temperature measurement used in the primary analysis: the aggregate summer measure will not reveal underlying heterogeneity. Therefore I use one Landsat image from each month of the year without aggregation for this analysis.

Recall that Landsat 7 has a 16-day cycle, so I select imagery for each month by choosing the first image of the month. Since the USGS algorithm for removing clouds relies on a collection of images for each location, I instead use the Quality Assurance (QA) flag to remove any pixels that are identified to be any type of cloud. Without additional imagery to replace these

[16]Most xeriscape rebate programs explicitly require that the grass being replaced is alive at the time of application, but households may still water it less or simply let it die as soon as the application is submitted (or even as soon as pictures are taken).

pixels, this leaves holes and thus missing observations in the data; this is the price to pay for disaggregated data. Furthermore, due to the added computational burden of disaggregation, I restrict heterogeneity analysis to two counties with the most conversions: Clark County, Nevada, and Los Angeles County, California.

To avoid conditioning on the outcome of interest, I condition on the temperature of the matched *untreated* parcel, binning these temperatures into quantiles $Q_{i,t}$ for $b = \{1, \ldots, B\}$ bins, constructing quantiles and running regressions separately by county. Then the heterogeneous treatment effects of landscape conversion are the $\beta_{\tau,b}$ sequences for temperature bin b. Equation (2.4) becomes:

$$m_i\widehat{\eta_{i,t}} = \sum_{b\in\{1,\ldots,B\}} \sum_{\tau=-5}^{5} \beta_{\tau,b} \mathbb{1}\{t - conversionYear_i = \tau\} \times \mathbb{1}\{Q_{i,t} = b\} + \varepsilon_{i,t} \tag{2.5}$$

where $m_i\widehat{\eta}_{i,t}$ is the rescaled and residualized temperature difference as in Equation (2.3). Figure (2.2) shows the sequence of temperature heterogeneity estimators with five temperature bins. Clearly, the heat effect of xeriscaping is concentrated in hot days: on the hottest 20% of days, the average temperature differential between a xeriscaped and non-xeriscaped yard is about 0.8°C (1.4°F) in Los Angeles and exceeds 1°C (1.8°F) in Clark County.[17] On cooler days, the temperature effect is more moderate but still positive and significant except for Clark's second quintile which shows no discernible difference from the pre-conversion period. This may be the effect of the interaction between humidity and temperature in the early spring and late fall when second quintile temperatures are likely to occur, such that the increased humidity from vegetation does not provide a cooling effect. Furthermore, since the negative consequences of heat are particularly acute at high temperatures, it is important to note that these are precisely the temperatures at which turf replacement increases heat the most. See Appendix 2.A.4 for point

[17]Note that quintiles are calculated separately by county, so because Clark has higher high temperatures than Los Angeles, Clark's top quintile represents hotter days than that for Los Angeles.

estimates by 5°C bins rather than by quintile.

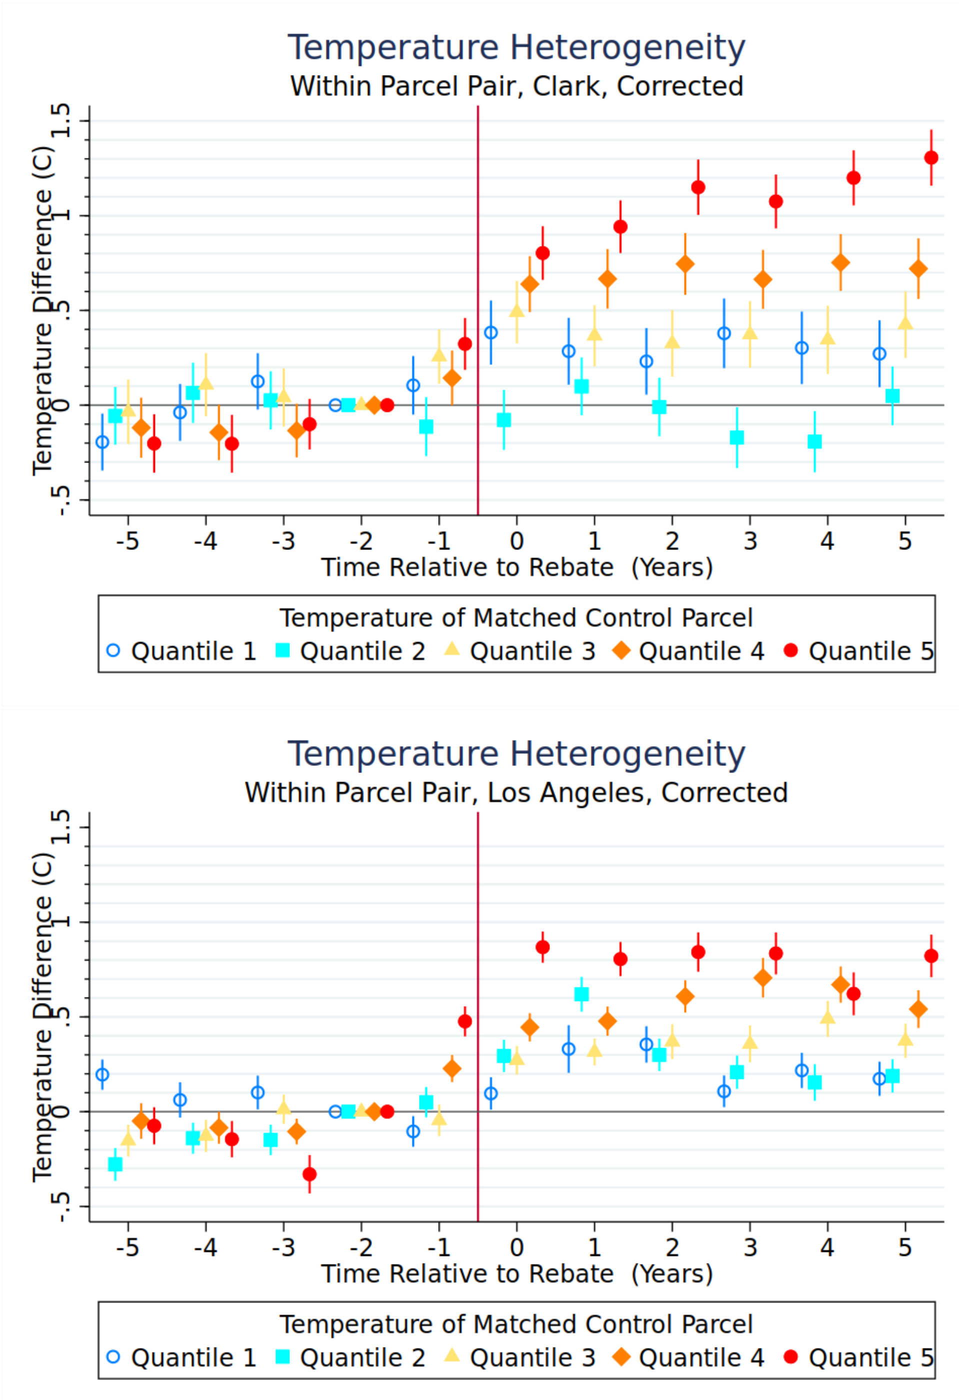

Figure 2.2. Temperature Heterogeneity for Clark and Los Angeles Counties

Notes: Figure shows temperature effects of landscape conversion by quantile of temperature of the matched control parcel, with pair fixed effects and rescaled using the correction multiplier as in Section 2.3.2. Clark and Los Angeles counties are run separately. Uses disaggregated temperature observations with one measurement from each month of the year.

I additionally use this disaggregated monthly data to estimate the year-round average temperature effect. Regressing the rescaled temperature difference on on an indicator for $\tau \geq 0$ and an indicator for $\tau = -1$ (since this is an in-between period, letting $\tau \leq -2$ be the reference period) reveals that the average post-conversion temperature is 0.46°C (0.83°F) hotter than the pre-period. The year-round estimate is attenuated relative to the summertime number due to the smaller heat effects of cooler days.

2.4.3 Heterogeneity by Vegetation Change

The main analysis of this paper explores the extensive margin of the heat effects of xeriscaping: the increased temperature as a result of any xeriscaping. Here, I examine the intensive margin: whether larger vegetation loss induces larger temperature gains. To do so, I turn to Landsat 7's Normalized Difference Vegetation Index (NDVI), the satellite's near infrared band minus the red band divided by the sum of the two values. NDVI ranges from -1 to 1 and proxies for vegetativeness, with higher values representing more vegetation. The more complex Enhanced Vegetation Index (EVI) has also been proposed and yields results nearly indistinguishable from those using NDVI.

While my primary analysis has the advantage of being able to measure signal out of relatively noisy household-level satellite data, assessing within-household NDVI change at the individual level does not benefit from the large sample size. To model within-household changes, I assign the vegetation change for each household to be its average post-conversion NDVI relative to the control minus the average pre-conversion NDVI differential, or the household level difference-in-difference:[18]

[18] Event time $\tau = -1$ is excluded from this calculation because the primary analysis suggests that the year prior to the rebate being issued has a mix of homes that have already converted and those that have not.

$$\Delta NDVI_i \equiv \tag{2.6}$$
$$\frac{1}{|\{\tau : \tau \geq 0\}|} \sum_{\{\tau:\tau\geq 0\}} NDVI_{i,\tau}^{treated} - NDVI_{i,\tau}^{control}$$
$$- \frac{1}{|\{\tau : \tau \leq -2\}|} \sum_{\{\tau:\tau\leq -2\}} NDVI_{i,\tau}^{treated} - NDVI_{i,\tau}^{control}$$

The distribution of $\Delta NDVI_i$ surprisingly is centered just slightly below zero, indicating that it is a relatively noisy measure of xeriscaping, but shows some aggregate loss of vegetation.

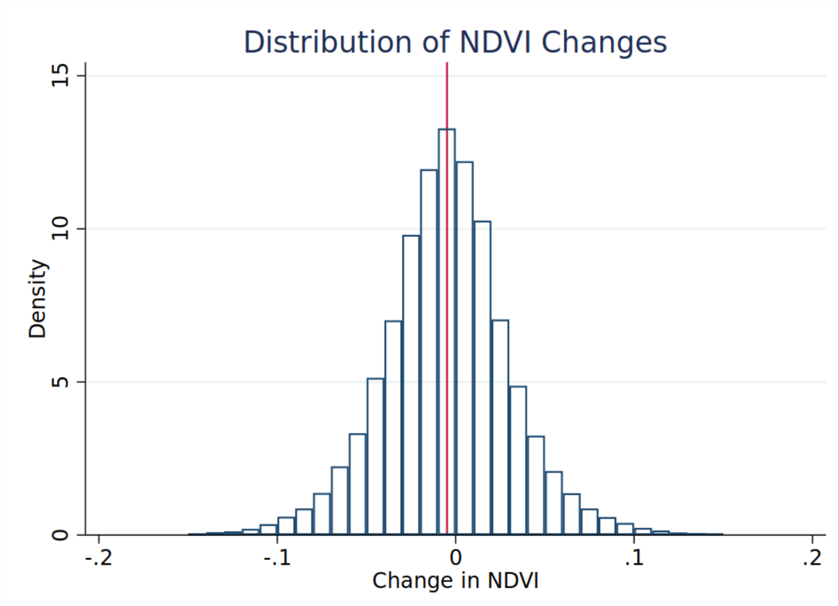

Figure 2.3. Estimated Household-level NDVI Change

Notes: Figure shows the distribution of within-household NDVI change following the landscape conversion, as calculated by Equation (2.6). Vertical line shown is the mean change in NDVI. As an index, NDVI ranges from -1 to 1.

Conditioning on quartiles of $\Delta NDVI_i$, Figure (2.4) shows substantial heterogeneity by vegetation loss. Homes with the top 25% of vegetation loss following a conversion see summertime temperature increases of 1.0-1.5°C (1.8-2.7°F), approximately doubling the average estimate from the primary analysis. The quartile with the largest NDVI *gain* sees no discernible temperature effects, while the middle 50% have modest increases in temperature.

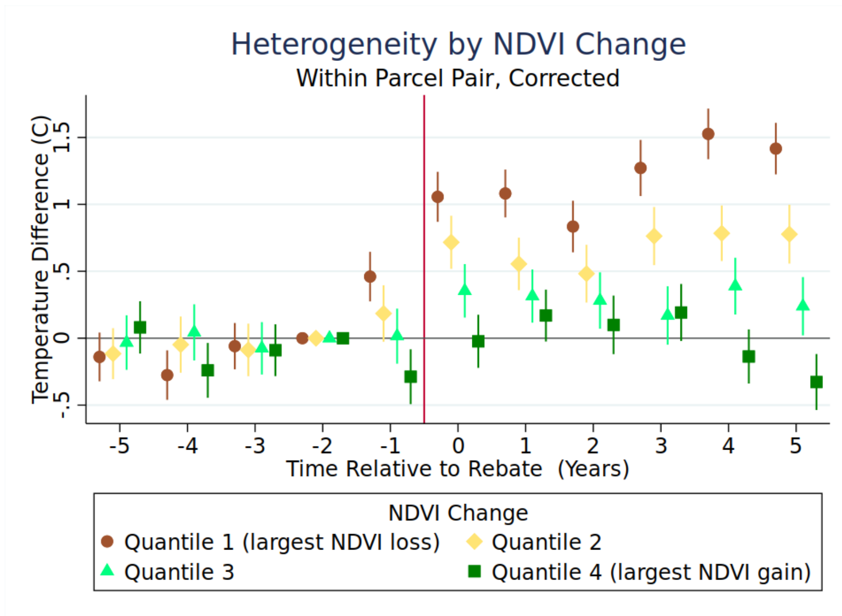

Figure 2.4. Heterogeneity by Vegetation Change

Notes: Temperature effects by quartile of measured vegetation change. Because the change in NDVI is centered nearly around zero, it is helpful to think of the top two quartiles as parcels with observed vegetation gain and the bottom two as those with vegetation loss.

This analysis reinforces the hypothesis that temperature changes from xeriscaping are from the removal of vegetation, as homes with the most removed vegetation have the largest increase in surface temperatures. However, it is also possible that parcels with the largest measured NDVI loss are those that are simply better picked up by the satellite, either due to the orientation of the parcel relative to the Landsat pixel grid or are those with a more drastic change.

2.5 Discussion

A vast climate literature enumerates the effects of temperature on various outcomes, from health to cognition and energy use. In this section, I provide estimates of some of the economic costs of the increased heat due to landscape conversions.

First, I estimate the effect on electricity usage and expenditure of the increased temperature due to turf replacement. I use publicly available ZIP code by month electricity usage data from San Diego Gas and Electric (SDGE) and Southern California Edison (SCE)[19] matched to ZIP code daily temperatures from PRISM[20] to estimate the effects on electricity demand of increasing temperature.[21] I first map daily maximum temperatures[22] to the actual 5 degree Celsius bins used in Table (2.2) of the appendix and also assign the bins under the counterfactual heterogeneous temperature change scenario implied by Table (2.2). Then, I regress ZIP code level monthly temperature usage on the number of days in each temperature bin[23] with calendar month and ZIP code fixed effects to predict the average household electricity usage in each month both before and after the landscape conversion-induced temperature change.

Finally, I average the monthly difference between predicted converted and unconverted usage within each ZIP code and show the distribution of ZIP code-level change in per-household electricity usage in Figure (2.5), weighted by the number of residential electricity customers. The average predicted increase in electricity usage as a result of landscape conversion is 4.8 kWh per month, or 57.6 kWh per year, representing a 0.9% increase. This is smaller than the 3% increase found by Baker (2021) using household-level data in Las Vegas. The discrepancy may be due to Baker's focuses on Las Vegas which is hotter than Southern California (recall that hotter days see larger heat effects) and has near universal air conditioning penetration, so electricity demand is expected to be more responsive to turf replacement. Furthermore, estimates using household-level electricity data may also be higher if households engage in other electricity-intensive upgrades at the same time as the landscape conversion. Even if the upgrades are to more efficient appliances, the "rebound effect" literature has shown that this can lead to

[19]San Diego Gas and Electric (2022); Southern California Edison (2022).

[20]PRISM Climate Group, Oregon State University (2022).

[21]Plenty of studies have already estimated this so-called "temperature response function", but few report the actual parameterization of the function and instead display a graph. In order to quantify the effect of heterogeneous (through the temperature spectrum) changes in temperature, I estimate and record the parameters of my own temperature response function, constraining to Southern California which is my main study area of interest, and where electricity data is publicly available.

[22]Recall that the surface temperature measure from Landsat 7 most closely maps to daily maximum air temperatures.

[23]Because there are some days missing from PRISM temperature data, I scale up each observed count by the proportion of all observed days they account for. Consider a month with 31 days but only 20 observations, 10 of which were in the 10-15 C bin. Because half of the observed days were 10-15 C, I assign half of the month's days to this bin, that is, 15.5 days.

more energy consumption; see Fowlie, Greenstone and Wolfram (2018), for example. On the other hand, my ZIP code level temperature response function should not respond to efficiency upgrades in this way. Nevertheless, the electricity response may well be larger in Las Vegas than in Southern California.

Under SDGE's highest standard time-of-use rate of $0.83325/kWh[24], this 57.6 kWh is an increase in annual expenditure of $48. Note that marginal air conditioning (A/C) use plausibly occurs in summer on-peak hours when electricity is most expensive. Since A/C usage is the primary mechanism for electricity's response to temperature, this estimate is expected to be substantially higher in Clark County, where A/C adoption is near universal, and baseline temperatures are hotter. Furthermore, as A/C penetration increases, the effect on California's electricity usage is likely to increase, although this may be partially offset by efficiency improvements.

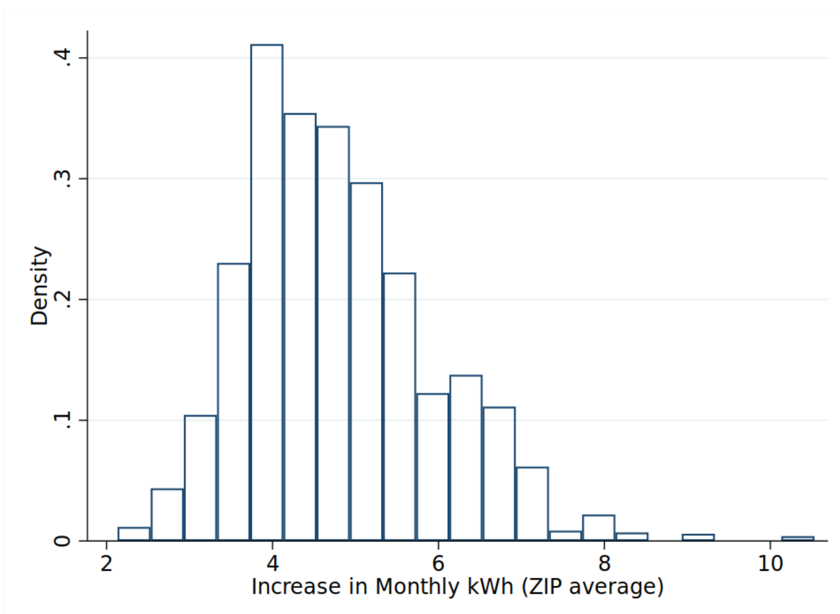

Figure 2.5. Estimated Increase in Monthly Electricity Usage

Notes: Figure shows the distribution of the estimated increase in per-household electricity usage as a result of a turf replacement. Each observation is the within-ZIP code average of per-household estimated monthly electricity usage changes.

[24] See https://www.sdge.com/total-electric-rates plan "Schedule TOU-DR1" beginning on 1/1/23.

Another factor contributing to the disutility of heat is the increased risk of heat mortality. I use Deschênes and Greenstone's (2011) estimates of heat mortality as a function of temperature by binning daily average temperatures into the 10°F (5.6°C) bins used in their study, and assigning my heterogeneous temperature change estimates based on the days' maximum temperatures, following the same aggregation procedure as for electricity. Note that the methodology used by Deschênes and Greenstone (2011) captures the mortality effect of heat exposure *net of avoidance behavior*, so these estimates represent additional costs on top of costly avoidance behavior (like increased air conditioning and decreased time outdoors). Comparing the population-weighted distribution of temperatures currently experienced by Southern Californians with the additional heat deaths from shifting the temperature distribution, I estimate that universal adoption of xeriscaping in the region would lead to 142 more heat deaths per year. Using the EPA's $10.7 million value of a statistical life (Environmental Protection Agency, 2022),[25] universal adoption of xeriscaping in Southern California would lead to an additional $1.5 billion of loss from heat mortality each year. This $306 per year for a household of four is a cost that should be factored into the household's decision problem when considering landscape conversion. Importantly, this cost is faced primarily by the household occupying the xeriscaped lot regardless of the number of other households who convert their yards. Using this estimate, the 175,000 conversions in my data are responsible for 3.75-5 additional heat deaths each year, assuming 3-4 individuals exposed to each converted property.

Lastly, I turn to holistic estimates of the marginal willingness to pay (MWTP) to avoid heat, as revealed by hedonic pricing. To the extent that homebuyers are aware of the effects of heat (whether implicitly or explicitly), increasing temperatures should become capitalized into housing values: all else equal, a house with a hotter lot should sell for less. Using a hedonic pricing approach of home values in Phoenix, Arizona, Klaiber, Abbott and Smith (2017) estimate a marginal willingness to pay for a highly localized 1°F (0.56°C) decrease in summertime

[25]$7.4 million in 2006 dollars, converted to 2022 dollars. The consumer price index for all items in U.S. cities in 2006 was 201.6 and 292.6 in 2022 using the default 1982-1984 index (U.S. Bureau of Labor Statistics, 2022).

temperature of \$72 (\$129 in 2022 dollars[26]) per month holding landscaping and other amenities constant.[27] Converting this to an annual cost for a 0.6°C change, this implies a \$1,675 (2022 dollars) per year loss in welfare due to heat effects. Using Sovocool's (2005) estimate of 55.8 gallons of water saved per square foot replaced, a median replacement of 1,000 square feet, and Los Angeles' top marginal rate for municipal water of \$12.79/HCF,[28] the typical turf replacement saves as much as \$954 per year in water,[29] failing to break even compared to the estimated disutility costs. Baker (2021) and Brelsford and Abbott's (2021) smaller estimates of water savings from xeriscaping would imply an even smaller value of water savings.

Using the MWTP estimate of \$1,675 in annual costs, the marginal cost of water would have to be \$22.45/HCF in order to break even, well above even the highest rates in the region. Meanwhile, the current top rate for Las Vegas single family residences is just \$7.70/HCF, substantially below the breakeven rate, and San Diego's top rate of \$12.49/HCF fails to break even as well.[30] The majority of the water saved for the median xeriscaped home in Las Vegas falls below the top tier billing rate (Sovocool, 2005), so this is likely an upper bound on the value of private water savings. Lower tier marginal rates range from \$1.96-5.18/HCF in Las Vegas and \$5.50-8.88/HCF in San Diego. Benefits of xeriscaping are indeed likely lower than the costs imposed from increased heat for many households.

Accounting for just heat mortality and electricity costs, a typical turf replacement imposes \$354 in costs annually, but saves up to \$954 in water for Los Angeles or \$574 in Las Vegas

[26] Klaiber, Abbott and Smith (2017) normalize prices to 1998. The consumer price index for all items in U.S. cities in 1998 was 163.0 and 292.6 in 2022 using the default 1982-1984 index (U.S. Bureau of Labor Statistics, 2022).

[27] Strictly speaking, Klaiber, Abbott and Smith (2017) use July minimum air temperatures. I show in Appendix 2.A.1 that minimum air temperatures move at the same rate as maximum temperatures throughout the temperature spectrum, and that maximum temperatures are well approximated by surface temperatures, the outcome used in this paper.

[28] Los Angeles Department of Water and Power (2021)

[29]

$$\frac{1{,}000 sq.\ ft.}{1} \times \frac{55.8 gal}{1 sq.\ ft.} \times \frac{1 HCF}{748 gal} \times \frac{\$12.79}{1 HCF} = \$954$$

[30] See "Water rates effective Jan. 1, 2023" https://www.sandiego.gov/public-utilities/customer-service/water-and-sewer-rates/water for San Diego and https://www.lvvwd.com/customer-service/pay-bill/water-rates.html for Las Vegas (multiplied by 1,000 gal / 748 HCF).

with the difference due to varying marginal costs of municipal water. These cost estimates fall short of Klaiber, Abbott and Smith (2017)'s estimated MWTP for temperature, implying large unaccounted-for comfort values, other omitted costs, and/or an overestimate of MWTP. While difficult to quantify in this context, other literature suggests that additional costs may operate through channels such as reduced cognition (Dai et al., 2016; Park et al., 2020), in utero and postnatal development (Isen, Rossin-Slater and Walker, 2017), or costly adaptation through changes in time use (Graff Zivin and Neidell, 2014).

2.6 Conclusion

Using household level rebate data coupled with remotely sensed surface temperatures, I find a highly statistically significant and pervasive 0.6°C (1.1°F) increase in summertime temperatures after replacing grassy lawns with drought-tolerant landscaping. Effects are twice as large among homes with the largest quartile of vegetation loss, and on the hottest 20% of days. This additional heat imposes costs on rebate adopters which they may not factor into their landscaping decision. I estimate these costs at up to $1,675 per household per year, composed of about $48 in increased electricity and $306 from elevated mortality risk, with the remainder likely composed largely of comfort value, but also including diminished cognition and development and a costly shift in time use.

Estimates of the electricity and mortality risk costs of landscape conversion approach the private value of water savings from such programs, particularly in lower billing tiers and where municipal water is cheaper. Meanwhile, the more comprehensive hedonic estimates of this cost eclipse the private benefits. Therefore, while xeriscape rebate programs may be a cost-efficient strategy for water providers to mitigate water shortages, I provide evidence that they are not necessarily welfare-increasing for the typical household (even those who select into takeup), particularly given the fact that the rebates rarely cover the full up-front monetary cost of the replacement.

This chapter is currently being prepared for submission for publication of the material. The dissertation author was the primary investigator and sole author of this material.

2.A Appendix

2.A.1 Air Temperature versus Surface Temperature

This section motivates the use of surface temperature of Landsat as a proxy for experienced air temperature by comparing the two. First, I use the Landsat surface temperatures in the main analysis to compare to ZIP code level air temperatures. In case one is concerned that this is too coarse of a measurement, Section 2.A.1 shows that when measuring differential changes in temperatures between land surface types, there is not a significant difference in the measure resulting from air temperature compared to surface temperature.

Landsat Surface vs. PRISM Air Temperatures

Using the ZIP code level PRISM air temperature from Section 2.5, I fit Landsat surface temperature against maximum daily air temperature with a local linear regression in Figure (2.6). I use Clark County because of its similar climate to Phoenix, Arizona, where Klaiber, Abbott and Smith's (2017) study took place, and because the disaggregated surface temperature data is readily available from my temperature heterogeneity analysis. For reference, I show a 45° line through the origin reflecting a one-for-one mapping of daily maximum air temperature to surface temperature. Figure (2.6) shows that the actual relationship closely follows the 45° line, suggesting that there is little need for concern that surface temperatures do not reflect sensed air temperatures. The histogram highlights the fact that where the fit line deviates from the 45° line, there are relatively few observations.

The right panel of Figure (2.6) demonstrates that maximum and minimum temperatures tend to move in tandem. Fitting a global regression line yields an intercept (the difference between maximum and minimum temperatures) of 11.44°C and a slope of 1.08, indicating that on average, a one degree increase in the maximum temperature corresponds to an approximately one degree increase in the minimum temperature.

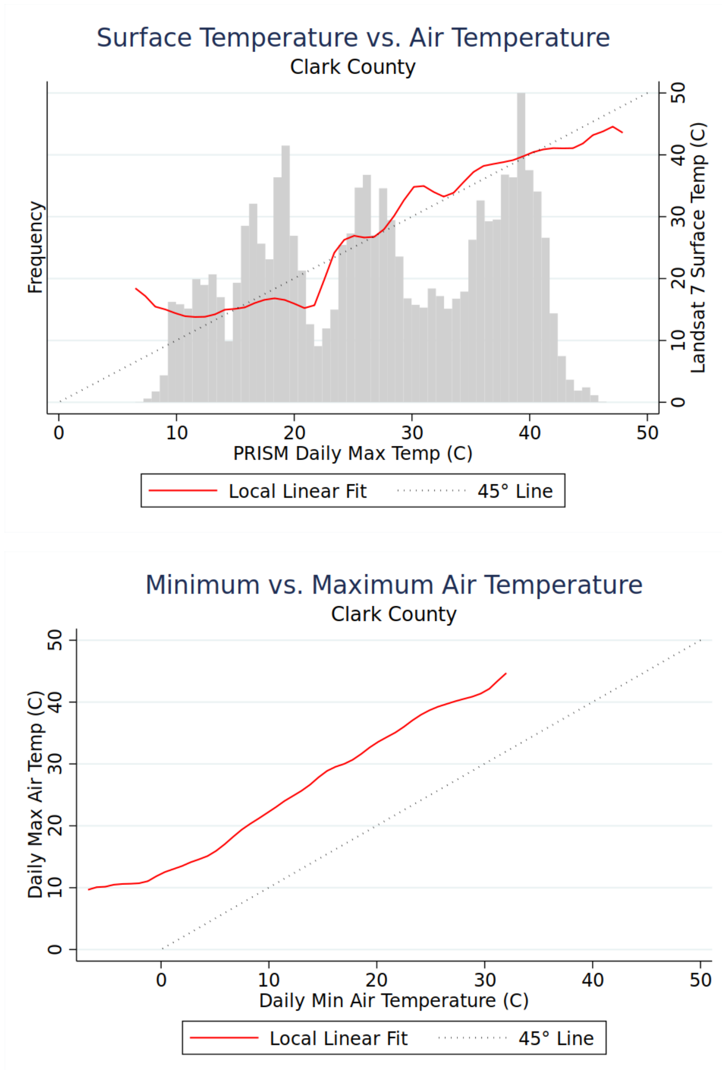

Figure 2.6. Surface Temperature vs. Air Temperature

Notes: Top: Landsat surface temperature versus PRISM maximum daily air temperature for parcels in Clark County, with histogram of air temperature frequency and local linear fit with bandwidth 1. The fit line closely tracks the 45° line through the origin, suggesting that surface temperature is a reasonable approximation of daily maximum air temperatures. *Bottom:* Maximum daily air temperature plotted against minimum daily air temperature, with local linear fit of bandwidth 1°C. The fit line is approximately parallel to the 45° line, indicating that a 1° increase in maximum temperature coincides with an approximately 1° increase in minimum temperature.

Measured Surface vs. Air Temperature from Schneider et al. (2023)

In this section I use ground measurements of surface and air temperature to demonstrate that surface temperatures perform no worse than a secondary air temperature sensor at measuring differential changes in air temperature in the urban environment. For this, I use publicly accessible data from Schneider et al. (2023) which attempts to characterize differences in heat effects from reflective pavements, comparing air temperature, surface temperature, and mean radiant temperature. Data used in this analysis is collected by driving a car packed with sensors through two routes in each of three Phoenix, Arizona neighborhoods four times over the course of a day (pre-sunrise, noon, afternoon, and post-sunset; each neighborhood was measured on a different day in August and September 2020). The car regularly records air temperature from multiple thermocouple sensors at 2 meters above the ground, in addition to surface temperature from an infrared radiometer. Because Landsat images are always taken during the day, I use the noon and afternoon traverses for this analysis, omitting the data from after sunset and before sunrise.

The Schneider et al. (2023) data is not a perfect analog to the current paper because it lacks repeated observation in a given location over different years (and inherently does not capture within-location landscaping changes). However, by repeatedly measuring temperatures across various surfaces over the course of the day, I can assess how temperatures differentially change across different surfaces. For this analysis, I compare 'cool pavement' (CP), a product designed to absorb less heat, to regular asphalt, dropping concrete surfaces from the analysis due to a much smaller sample of concrete surfaces. I compare each change in temperature by the sensor used to measure it. The data has a primary air temperature sensor, TH06, that is active for all included observations (T_{air}^{prim}), and three others that are each active for a subset of traverses; I combine these other three into a single variable, 'alternate air temperature' (T_{air}^{alt}).[31] Surface temperature is always measured by exactly one sensor ($T_{surface}$). Figure 2.7 reveals an idiosyncrasy in data collection for the Westcliff neighborhood, where in the data from the afternoon, the primary and alternate air temperature sensors are nearly uncorrelated. In contrast, the measurements from the Garfield and Maryvale neighborhoods neatly follow the 45°line,

[31] I use TH08 when available, and if missing I use TH09. Results are not sensitive to instead using TH10 when TH08 is missing. TH09 and TH10 are always missing when TH08 is nonmissing (active sensors are [TH08] xor [TH09 and TH10]).

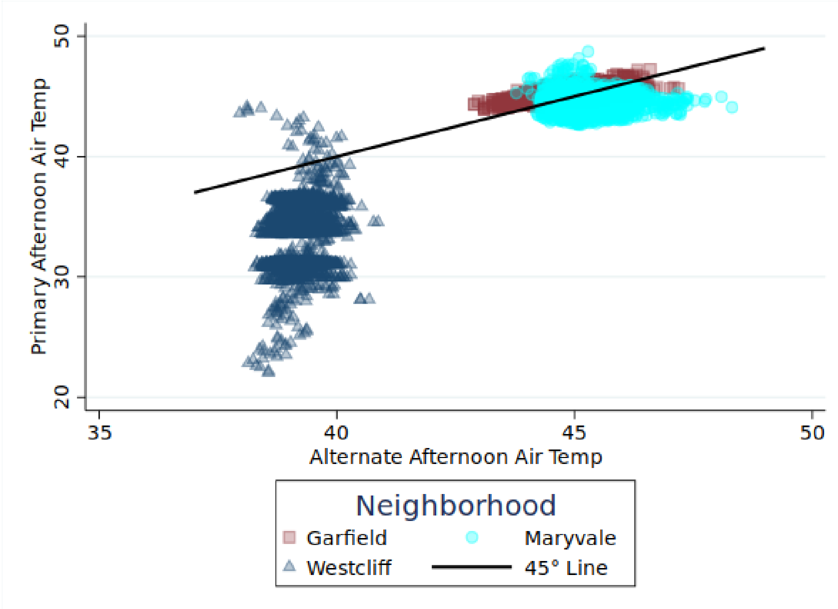

Figure 2.7. Afternoon Primary vs. Alternate Air Temperature by Neighborhood
Notes: Primary versus alternate afternoon air temperature measurements by neighborhood. Because both air temperature measures are captured using the same type of equipment but with different sensors, one should expect the observations to closely follow a 45°line, with any deviations due to noise. This is indeed the case for Garfield and Maryvale, while Westcliff deviates significantly from the 45°line, suggesting an issue with the data collection. For this reason, Westcliff is omitted from the analysis.

as one would expect for two different measurements of the same value using the same type of sensor. Because of this data irregularity, I omit the Westcliff neighborhood from this analysis.

Table 2.1 shows two triple difference models and a quadruple difference. Column 1 shows the difference between the air temperature and surface temperature sensors in the change of temperature from the noon reading to the afternoon reading for the cool pavement relative to asphalt. Column 2 shows the same but for the difference between the primary air temperature sensor and the alternate air temperature reading. Finally, Column 3 compares the differences in discrepancies between the sensors by doing a quadruple difference. Neighborhood by traversal fixed effects are included in all regressions. [32] For data point i in neighborhood n and traversal t,

[32]Day fixed effects are subsumed by the neighborhood fixed effects because each neighborhood is only measured on one day. Time of day (noon versus afternoon) is not considered a "control" because this is one level over which one of the differences is taken, in order to emulate the time element of the panel data in the current paper, and is instead a main coefficient of interest.

the estimating equations are as follows:

$$\text{Column 1:} \tag{2.7}$$
$$T_{surf,i} - T^{prim}_{air,i} = \alpha_{n,t} + \beta_0 Afternoon_i + \beta_1 CP_i + \beta_3 Afternoon_i \times CP_i + \varepsilon_i$$
$$\text{Column 2:} \tag{2.8}$$
$$T^{prim}_{air,i} - T^{alt}_{air,i} = \alpha_{n,t} + \beta_0 Afternoon_i + \beta_1 CP_i + \beta_3 Afternoon_i \times CP_i + \varepsilon_i$$
$$\text{Column 3:} \tag{2.9}$$
$$(T_{surf,i} - T^{prim}_{air,i}) - (T^{prim}_{air,i} - T^{alt}_{air,i}) =$$
$$\alpha_{n,t} + \beta_0 Afternoon_i + \beta_1 CP_i + \beta_3 Afternoon_i \times CP_i + \varepsilon_i$$

The parameter of interest for confirming the validity of using surface temperatures is β_3, or the coefficient on the interaction between afternoon and cool pavement. This is the sensor discrepancy in the differential effect of cool pavement from noon to afternoon. Table 2.1 shows the results of these regressions. The β_3 estimate of -0.0893 (°C) in Column (1) reveals a small and nonsignificant[33] discrepancy between using an air temperature sensor and a surface temperature sensor to measure differential changes in temperature across land surface types at different times of measurement. Taking the point estimate as given, this represents a less than 15% error resulting from using surface temperature rather than air temperature, relative to this paper's headline result of a 0.6°change. Moreover, Column (1) reveals that if anything, the surface temperature sensor measures *smaller* differential changes than those measured by the air temperature sensor, implying that the surface temperature may understate the differences-in-differences, meaning that the true effect is even larger.

For comparison, Column (2) shows compares the difference-in-differences measurement for the primary air temperature sensor against the alternate air temperature sensor, and reveals a

[33] These regressions use unclustered OLS standard errors in order to avoid failing to reject the null hypothesis when it should have been rejected. More conservative standard errors would make these observed differences even more nonsignificant, furthering the argument that surface temperature is an acceptable proxy for air temperature when measuring differences-in-differences.

larger and more statistically significant difference than for the surface vs. primary air temperature regression (Column 1). Column (3) performs the quadruple difference comparing Column (1) to Column (2) and fails to reject that the discrepancy between air temperature sensors is the same as that between the surface and air temperature sensor when measuring differential changes in temperature by land surface type.

Overall, Table (2.1) provides strong evidence that using surface temperatures to measure changes over time between land surface types is an adequate proxy for air temperature. In fact, using the air temperature sensor common to all routes in the data as “ground truth”, I find that the surface temperature sensor *more* accurately captures differences-in-differences than an alternate air temperature sensor of the same type, suggesting that the discrepancy due to sensor type (air vs. surface) is smaller than that due simply to noise.

Table 2.1. Discrepancies between Air and Surface Temperatures when Measuring Differential Changes in Temperature

Notes: *Afternoon* is an indicator for the afternoon measurement rather than noon and *Cool Pavement* is an indicator for the reflective pavement type as opposed to regular asphalt. Column (1) reveals that there is not a significant difference between the surface temperature measurement and the primary air temperature measurement in capturing the difference-in-differences of change in temperature for different land surface types. Column (2) suggests that comparing two different air temperature sensors has a larger and more significant discrepancy than for the surface temperature vs. primary air temperature. Column (3) fails to reject that Columns (1) and (2) are different.

	(1)	(2)	(3)
	Surface - Air Temp	Air Temp - Alt Air Temp	(Surface - Air) - (Air - Alt Air) Temp
	b/se/p	b/se/p	b/se/p
Afternoon	4.6528	0.7119	3.9409
	(0.1176)	(0.0192)	(0.1257)
	[0.000]	[0.000]	[0.000]
Cool Pavement	-3.7128	-0.3651	-3.3476
	(0.1331)	(0.0217)	(0.1422)
	[0.000]	[0.000]	[0.000]
Afternoon × Cool Pavement	-0.0893	0.1992	-0.2885
	(0.1965)	(0.0321)	(0.2101)
	[0.650]	[0.000]	[0.170]
Constant	10.5342	0.1149	10.4193
	(0.1205)	(0.0197)	(0.1289)
	[0.000]	[0.000]	[0.000]
Observations	11379	11379	11379
r2	0.426	0.569	0.439

Standard errors in parentheses. P-values in brackets.

2.A.2 Matching Estimator by County

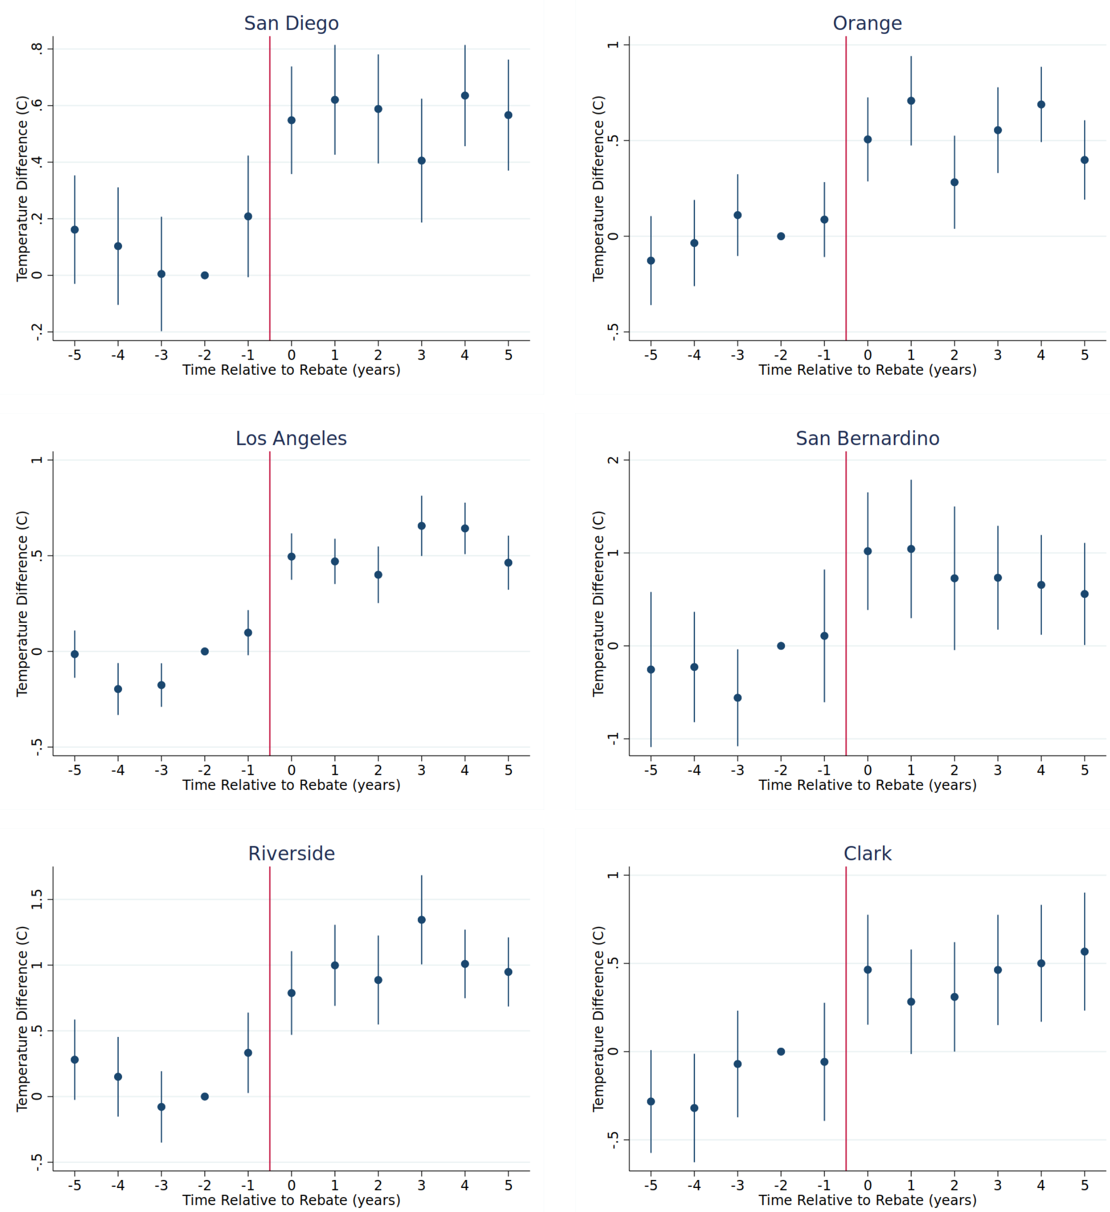

Figure 2.8. Main Results by County

Notes: Figure shows the main summertime temperature analysis by county. Includes parcel pair fixed effects with standard errors clustered by parcel pair.

2.A.3 Proof of Equation (2.2)

This section derives the formula for the multiplier used in the main text. For a given parcel, let:

- $\mathscr{P} = \{\text{pixels touched by parcel}\}$
- s_p =share of pixel p covered by parcel
- Δt_p = measured temperature change of pixel p
- $\Delta\tau_p$ = actual change in temperature of the lot in pixel p. Assume this is constant across all pixel regions covered by the parcel, $\Delta\tau_p = \Delta\tau \ \forall p \in P$.
- ΔT is the overall measured change in temperature. By construction, this is $\frac{\sum s_p \Delta t_p}{\sum s_p}$, or the area-weighted average temperature of the pixels intersected by the parcel.

For notational simplicity, let all sums be over all $p \in \mathscr{P}$, the pixels touched by the parcel.

$$
\begin{aligned}
\Delta T &= \frac{\sum s_p \Delta t_p}{\sum s_p} \\
&\quad \text{noting that } \Delta t_p = \underbrace{(1-s_p)0}_{\text{no change in unconverted region}} + \underbrace{s_p \Delta\tau_p}_{\Delta\tau_p \text{ change in converted region}} = s_p \Delta\tau_p, \\
\Delta T &= \frac{\sum s_p^2 \Delta\tau_p}{\sum s_p} \text{ and assuming } \Delta\tau_p = \Delta\tau \ \forall p \in \mathscr{P}, \\
&= \Delta\tau \frac{\sum s_p^2}{\sum s_p} \text{ so,} \\
\Delta\tau &= \Delta T \frac{\sum s_p}{\sum s_p^2}
\end{aligned}
$$

So multiplier $m = \frac{\sum s_p}{\sum s_p^2}$. Note that $\sum s_p$ is the total parcel area in units of pixels.

2.A.4 Temperature Heterogeneity Tabulation

Table 2.2. Temperature Heterogeneity Point Estimates

Notes: Table shows regression coefficients of the treatment effect of landscape conversion by 5°C temperature bin of the control parcel. Data is the disaggregated monthly temperature used in Section 2.4.2. Column (1) shows estimates for Los Angeles and Clark counties combined, while Columns (2) and (3) separate each county. The *post* indicator equals one for event times greater than or equal to 0. Regression also includes an indicator for event time -1 due to its ambiguous treatment status, leaving event times less than -1 as the omitted baseline.

	(1) Pooled	(2) Los Angeles	(3) Clark
Post × 0-5 C	0.991	-1.399	1.035
	(0.112)	(0.740)	(0.115)
Post × 5-10 C	-0.212	-2.607	-0.00905
	(0.0724)	(0.178)	(0.0844)
Post × 10-15 C	0.421	0.170	0.417
	(0.0388)	(0.0513)	(0.0557)
Post × 15-20 C	-0.186	-0.305	0.185
	(0.0296)	(0.0322)	(0.0613)
Post × 20-25 C	0.475	0.768	-0.118
	(0.0285)	(0.0309)	(0.0598)
Post × 25-30 C	0.0480	0.0873	-0.0598
	(0.0265)	(0.0294)	(0.0576)
Post × 30-35 C	0.637	0.714	0.371
	(0.0261)	(0.0270)	(0.0685)
Post × 35-40 C	0.884	0.935	0.851
	(0.0296)	(0.0354)	(0.0503)
Post × 40-45 C	1.177	1.099	1.206
	(0.0458)	(0.0758)	(0.0571)
Post × 45-50 C	1.621	10.23	1.423
	(0.0752)	(0.707)	(0.0756)
Observations	3345281	2041361	1303920

Standard errors in parentheses

Bibliography

Addink, Sylvan. 2005. "'Cash for Grass'-A Cost Effective Method to Conserve Landscape Water?" *Turfgrass Research Facility, University of California, Riverside*, 14.

Baker, Jonathan E. 2021. "Subsidies for Succulents: Evaluating the Las Vegas Cash-for-Grass Rebate Program." *Journal of the Association of Environmental and Resource Economists*, 8(3): 475–508.

Benz, Susanne A., Steven J. Davis, and Jennifer A. Burney. 2021. "Drivers and projections of global surface temperature anomalies at the local scale." *Environmental Research Letters*, 16(6): 064093. Publisher: IOP Publishing.

Bollinger, Bryan, Jesse Burkhardt, and Kenneth T. Gillingham. 2020. "Peer Effects in Residential Water Conservation: Evidence from Migration." *American Economic Journal: Economic Policy*, 12(3): 107–133.

Braun, Thomas, and Wolfram Schlenker. 2023. "Cooling Externality of Large-Scale Irrigation." National Bureau of Economic Research Working Paper 30966. Series: Working Paper Series.

Brelsford, Christa, and Joshua K. Abbott. 2021. "How smart are 'Water Smart Landscapes'?" *Journal of Environmental Economics and Management*, 106: 102402.

Burkhardt, Jesse, Nathan W. Chan, Bryan Bollinger, and Kenneth T. Gillingham. 2021. "Conformity and Conservation: Evidence from Home Landscaping and Water Conservation." *American Journal of Agricultural Economics*. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajae.12224.

Buyantuyev, Alexander, and Jianguo Wu. 2010. "Urban heat islands and landscape heterogeneity: linking spatiotemporal variations in surface temperatures to land-cover and socioeconomic patterns." *Landscape Ecology*, 25(1): 17–33.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Chow, Winston T. L., and Anthony J. Brazel. 2012. "Assessing xeriscaping as a sustainable heat island mitigation approach for a desert city." *Building and Environment*, 47: 170–181.

County of Los Angeles. 2022. "Parcels."

County of Orange California. 2020. "Parcel Polygons."

Dai, Lingzhen, Itai Kloog, Brent A. Coull, David Sparrow, Avron Spiro, Pantel S. Vokonas,

and Joel D. Schwartz. 2016. "Cognitive function and short-term exposure to residential air temperature: A repeated measures study based on spatiotemporal estimates of temperature." *Environmental Research*, 150: 446–451.

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2022. "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey."

Deschênes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics*, 3(4): 152–185.

Dousset, B., and F. Gourmelon. 2003. "Satellite multi-sensor data analysis of urban surface temperatures and landcover." *ISPRS Journal of Photogrammetry and Remote Sensing*, 58(1-2): 43–54.

Environmental Protection Agency. 2022. "Mortality Risk Valuation."

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2018. "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program*." *The Quarterly Journal of Economics*, 133(3): 1597–1644.

Fung, W. Y., K. S. Lam, Janet Nichol, and Man Sing Wong. 2009. "Derivation of Nighttime Urban Air Temperatures Using a Satellite Thermal Image." *Journal of Applied Meteorology and Climatology*, 48(4): 863–872. Publisher: American Meteorological Society Section: Journal of Applied Meteorology and Climatology.

Gober, Patricia, Anthony Brazel, Ray Quay, Soe Myint, Susanne Grossman-Clarke, Adam Miller, and Steve Rossi. 2009. "Using Watered Landscapes to Manipulate Urban Heat Island Effects: How Much Water Will It Take to Cool Phoenix?" *Journal of the American Planning Association*, 76(1): 109–121. Publisher: Routledge _eprint: https://doi.org/10.1080/01944360903433113.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Graff Zivin, Joshua, and Matthew Neidell. 2014. "Temperature and the Allocation of Time: Implications for Climate Change." *Journal of Labor Economics*, 32(1): 1–26. Publisher: The University of Chicago Press.

Hartz, D. A., L. Prashad, B. C. Hedquist, J. Golden, and A. J. Brazel. 2006. "Linking satellite images and hand-held infrared thermography to observed neighborhood climate conditions." *Remote Sensing of Environment*, 104(2): 190–200.

Hu, Leiqiu, and Nathaniel A. Brunsell. 2015. "A new perspective to assess the urban heat island through remotely sensed atmospheric profiles." *Remote Sensing of Environment*, 158: 393–406.

Isen, Adam, Maya Rossin-Slater, and Reed Walker. 2017. "Relationship between season of birth, temperature exposure, and later life wellbeing." *Proceedings of the National Academy of Sciences*, 114(51): 13447–13452. Publisher: National Academy of Sciences Section: Social Sciences.

Klaiber, H. Allen, Joshua K. Abbott, and V. Kerry Smith. 2017. "Some Like It (Less) Hot: Extracting Trade-Off Measures for Physically Coupled Amenities." *Journal of the Association of Environmental and Resource Economists*, 4(4): 1053–1079.

Lo, C. P., D. A. Quattrochi, and J. C. Luvall. 1997. "Application of high-resolution thermal infrared remote sensing and GIS to assess the urban heat island effect." *International Journal of Remote Sensing*, 18(2): 287–304. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/014311697219079.

Loomis, John B. 2012. "Comparing households' total economic values and recreation value of instream flow in an urban river." ISSN: 2160-6544.

Los Angeles Department of Water and Power. 2021. "LADWP Water Rate Changes Now in Effect."

Mayer, Peter W, William B DeOreo, Eva M Opitz, Jack C Kiefer, William Y Davis, Benedykt Dziegielewski, and John Olaf Nelson. 1998. "Residential end uses of water." *American Water Works Association*.

Middel, Ariane, Björn Hagen, Anthony Brazel, and Soe Myint. 2007. "Simulation of possible scenarios for local scale energy balances in residential neighborhoods in Phoenix, AZ, USA." 17.

Middel, Ariane, Kathrin Häb, Anthony J. Brazel, Chris A. Martin, and Subhrajit Guhathakurta. 2014. "Impact of urban form and design on mid-afternoon microclimate in Phoenix Local Climate Zones." *Landscape and Urban Planning*, 122: 16–28.

Park, R. Jisung, Joshua Goodman, Michael Hurwitz, and Jonathan Smith. 2020. "Heat and Learning." *American Economic Journal: Economic Policy*, 12(2): 306–339.

Pearlmutter, D., P. Berliner, and E. Shaviv. 2007. "Urban climatology in arid regions: current research in the Negev desert." *International Journal of Climatology*, 27(14): 1875–1885. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/joc.1523.

Pichierri, Manuele, Stefania Bonafoni, and Riccardo Biondi. 2012. "Satellite air temperature estimation for monitoring the canopy layer heat island of Milan." *Remote Sensing of Environment*, 127: 130–138.

PRISM Climate Group, Oregon State University. 2022. "PRISM Climate Data."

Riverside County. 2019. "Parcels CREST."

San Bernardino County. 2022. "GIS Data Layer Catalog."

San Diego Association of Governments. 2022. "SanGIS GIS Data Warehouse."

San Diego Gas and Electric. 2022. "Energy Data Access."

Schneider, Florian A., Johny Cordova Ortiz, Jennifer K. Vanos, David J. Sailor, and Ariane Middel. 2023. "Evidence-based guidance on reflective pavement for urban heat mitigation in Arizona." *Nature Communications*, 14(1): 1467. Number: 1 Publisher: Nature Publishing Group.

Schwarz, Nina, Uwe Schlink, Ulrich Franck, and Katrin Großmann. 2012. "Relationship of land surface and air temperatures and its implications for quantifying urban heat island indicators—An application for the city of Leipzig (Germany)." *Ecological Indicators*, 18: 693–704.

Shashua-Bar, Limor, and Milo E. Hoffman. 2004. "Quantitative evaluation of passive cooling of the UCL microclimate in hot regions in summer, case study: urban streets and courtyards with trees." *Building and Environment*, 39(9): 1087–1099.

Shashua-Bar, Limor, David Pearlmutter, and Evyatar Erell. 2009. "The cooling efficiency of urban landscape strategies in a hot dry climate." *Landscape and Urban Planning*, 92(3): 179–186.

Shashua-Bar, Limor, David Pearlmutter, and Evyatar Erell. 2011. "The influence of trees and grass on outdoor thermal comfort in a hot-arid environment." *International Journal of Climatology*, 31(10): 1498–1506. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/joc.2177.

Southern California Edison. 2022. "Quarterly Customer Data Reports."

Sovocool, Kent A. 2005. "Xeriscape Conversion Study." *Southern Nevada Water Authority*, 93.

Stabler, Linda B., Chris A. Martin, and Anthony J. Brazel. 2005. "Microclimates in a desert city were related to land use and vegetation index." *Urban Forestry & Urban Greening*, 3(3): 137–147.

Sun, Liyang, and Sarah Abraham. 2020. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." 53.

Timmins, Christopher. 2002. "Does the Median Voter Consume Too Much Water? Analyzing the Redistributive Role of Residential Water Bills."

U.S. Bureau of Labor Statistics. 2022. "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average." Publisher: FRED, Federal Reserve Bank of St. Louis.

U.S. Geological Survey. 2022. "Landsat 7."

Vancutsem, Christelle, Pietro Ceccato, Tufa Dinku, and Stephen J. Connor. 2010. "Evaluation of MODIS land surface temperature data to estimate air temperature in different ecosystems over Africa." *Remote Sensing of Environment*, 114(2): 449–465.

Ward, Frank A., Brian A. Roach, and Jim E. Henderson. 1996. "The Economic Value of Water in Recreation: Evidence from the California Drought." *Water Resources Research*, 32(4): 1075–1081. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/96WR00076.

Young, Robert A., and John B. Loomis. 2014. "Determining the Economic Value of Water, Concepts and Methods."

Chapter 3

Human Capital and Climate Change

Abstract

Addressing climate change requires individual behavior change and voter support for pro-climate policies, yet surprisingly little is known about how to achieve these outcomes. In this paper, we estimate causal effects of additional education on pro-climate outcomes using new compulsory schooling law data across 20 European countries. We analyze effects on pro-climate beliefs, behaviors, policy preferences, and novel data on voting for green parties. Results show that a year of education substantially increases pro-climate beliefs, behaviors, and policy preferences.

3.1 Introduction

The costs and consequences of climate change are enormous and multifaceted (Carleton and Hsiang, 2016; Graff Zivin and Neidell, 2013; Intergovernmental Panel on Climate Change, 2022; Isen, Rossin-Slater and Walker, 2017; Park, Behrer and Goodman, 2021), with monetized impacts estimated to be as large as 20% of annual global GDP within a generation (Nordhaus, 2007). On current trajectories, the world is on track to experience 2.7°C warming above pre-industrial levels within the next century, far above the global goal of 1.5°C (Tracker, 2022). Individual behavior change and government policy are needed to dramatically alter the trajectory of emissions. Despite the urgency and scale of the challenge, current efforts are underwhelming, in part because sizable populations around the globe remain skeptical about climate change and policies to tackle it (Bechtel, Scheve and van Lieshout, 2020; Dechezleprêtre et al., 2022;

Sunstein et al., 2017). Surprisingly little is known about how to overcome such resistance.

One promising approach is the accumulation of human capital through increased educational attainment.[1] More educated individuals may be better equipped to understand the complexities of climate science and have more awareness of the risks of climate change, as well as increasing their trust in science. Descriptive correlations suggests this might be true: a global survey found people with more education were more likely to see climate change as a major threat (Pew Research Center, 2019). More education might also yield transferable skills across occupations, encouraging voting for policies which promote less-polluting industries, such as renewable energy subsidies. Yet determining the causal effect of human capital accumulation on pro-climate beliefs and behaviors is challenging. People who choose to pursue more education are, by revealed preference, forward looking and thus more concerned with the future consequences of climate change. It might not be education that is causing pro-climate beliefs and actions, but rather time preferences. Reverse causality is another challenge: individuals who believe in climate change might choose to pursue more education to better adapt to a changing world. Furthermore, family background is an unobservable factor that may confound this relationship.

In this paper, we overcome causal inference challenges by assembling a new database on compulsory schooling laws (CSLs) to estimate the causal effect of human capital accumulation on a series of climate outcomes in Europe. The use of CSLs as a plausibly exogenous shift in educational attainment has a rich tradition in labor and health economics (Angrist and Krueger, 1991; Black, Devereux and Salvanes, 2008; Brunello, Fort and Weber, 2009; Gathmann, Jürges and Reinhold, 2015; Goldin and Katz, 1997; Lleras-Muney, 2005; Oreopoulos, 2006), but is much more limited on climate.[2] Moreover, due to data limitations, studies have been largely limited to single countries. We build on this nascent climate literature leveraging 41 CSL reforms in 20 countries, identified via a new reforms database. In addition, studies to date analyze limited

[1] Human capital captures an individual's knowledge and skills (Becker, 1962) and is typically measured by education metrics including years of schooling (Barro, 2001) and learning (Angrist et al., 2021).

[2] A small set of studies explore environmental outcomes (Meyer, 2015; Powdthavee, 2021).

outcomes. We study new climate outcomes which extend well beyond standard measures of beliefs and behaviors, also examining the highly consequential domains of policy preferences and voting.

Europe is an ideal setting for this study. Countries in Europe enacted dozens of education reforms in the twentieth century, expanding the number of years of education legally mandated through compulsory schooling laws. At the same time, Europe has large, harmonized multi-country surveys, enabling credible within- and cross- country analyses, with recent climate modules added to the European Social Survey (ESS) which we analyze in this study. Moreover, Europe has a robust green party movement, which has an explicit environmental agenda.[3] We codify a novel dataset of green party voting outcomes, enabling identification of pro-climate voting behavior.

Our analysis focuses on outcome indices as well as on specific indicators within each index, including comparisons between correlations and causal estimates. We find significant impacts on nearly all pro-climate measures. Our headline results show that an additional year of education leads to an increase of 1.9 percentage points (PP) in pro-climate beliefs, 3.0 PP in behaviors, 0.8 PP in policy preferences, and 0.3 PP in green voting. Relative to status quo rates, these impacts are non-trivial, translating into 2.9% increase for beliefs, 4.3% for behaviors, 1.3% for policy preferences, and a 4.3% increase for green party voting.

These results are notable since education has been conspicuously absent from most major climate change discussions.[4] Our results show that human capital accumulation can play an important role in shaping beliefs about the costs and benefits of policies to reduce emissions (Dechezleprêtre et al., 2022) and extend directly to consequential outcomes such as policy preferences and voting. This motivates renewed focus on policies expanding access to general education as part of the menu of approaches considered in tackling climate change.

The rest of the paper is organized as follows. The next section describes our data. Section

[3]Green political parties' environmental focus includes climate change, pollution, and industrial agriculture.

[4]A recent analysis showed that only 24% of countries mention youth education in the context of the Paris Agreement (Kwauk, 2021) – a historic international treaty on climate change.

III details our empirical strategy and Section IV presents our results. Some brief concluding remarks are offered in Section V.

3.2 Data

Data on pro-climate outcomes – including beliefs, behaviors, policy preferences, and voting outcomes – come from the European Social Survey (ESS).[5] The ESS is conducted biennially across dozens of European countries using stratified random sampling with a total sample size ranging from 20,000 to 40,000 individuals per round. The ESS is a large microdata set capturing information on a host of social issues and is harmonized over time and across countries. In 2016, the ESS introduced novel questions on climate outcomes, such as "how often do you do things to reduce energy use?" and "how likely are you to buy energy efficient appliances?" Moreover, the ESS collected data on policy preferences such as "to what extent are you in favour or against using public money to subsidise renewable energy such as wind and solar power?" Finally, we codify data on voting for green parties since 2002. Europe has a thriving green party movement in 32 countries. We codify a novel dataset of "green voting" across Europe based on party platforms. Many political parties around the world have broad mandates, and are thus too general to explore specific climate voting patterns. In contrast, green parties have an explicit environmental agenda, enabling identification of pro-climate voting. Green voting can be derived from all ESS rounds, whereas all other climate outcomes are only included in 2016.

Table 3.1 shows the climate outcomes we consider in our analysis and Table A2 in the Online Appendix includes the parties we classify as "green" in each country. Each climate outcome is transformed into a binary 'pro-climate' indicator if the response is equal to or above the median. For example, a response is 'pro-climate' if the respondent answered "strongly in favor" or "somewhat in favor" when asked about policies to subsidize renewable energy, since the

[5]European Social Survey European Research Infrastructure (ESS ERIC). (2020). ESS8 - integrated file, edition 2.2 [Data set]. Sikt - Norwegian Agency for Shared Services in Education and Research. https://doi.org/10.21338/ESS8E02_2

median response is "somewhat in favor". Alternatively, we also consider a continuous outcome, where 1 is the most pro-climate response and 0 is the least.

In addition to analyzing individual outcomes, we aggregate climate outcomes into three indices: beliefs, behaviors, and policy preferences. Table 3.1 lists each question and denotes the index to which it belongs; indices are simple within-individual averages. Our main results also include an indicator for whether respondents voted for a member of a green party in the last election for countries where such a party exists.

Table 3.1. Climate Outcomes – Beliefs, Behaviors, Policy Preferences, and Voting

Notes: Each outcome is grouped by index category. Each index is computed as an average for each individual across the indicated questions. The final outcome, green voting, is a stand-alone binary outcome not aggregated with others into an index. For beliefs about the source of electricity, we create a sub-index: the ESS has questions about individuals' opinions on electricity generation from coal, gas, hydroelectric, nuclear, solar, wind, and biofuel. Given these outcomes are highly inter-related, we average pro-hydroelectric, pro-solar, and anti-coal beliefs. We exclude indicators which might be collinear with renewables captured by solar and hydro-electric, such as wind, as well as indicators with more ambiguous climate interpretations, such as nuclear.

Question	Beliefs	Behaviors	Policy	Voting
Do you think the world's climate is changing	✓			
Climate change good or bad impact across world	✓			
How worried about climate change	✓			
How much electricity should be generated from coal/hydro/solar	✓			
How worried too dependent on fossil fuels	✓			
How much thought about climate change before today		✓		
How likely to buy most energy efficient home appliance		✓		
How often do things to reduce energy use		✓		
Favor increase taxes on fossil fuels to reduce climate change			✓	
Favor subsidize renewable energy to reduce climate change			✓	
Favor ban of inefficient household appliances to reduce CC			✓	
Voted for green party in last national election				✓

We restrict our analysis to respondents at least 25 years old at the time they were surveyed to capture effects for those who have completed their schooling. In particular, we analyze outcomes for cohorts who received schooling and were affected by education reforms in the 1960s through the 1980s and were adults being surveyed in the ESS from 2002 to 2018. In addition to climate and voting outcomes, the ESS data contains birth year and years of education for every individual, which are critical to mapping climate outcomes to cohorts of students affected by compulsory schooling laws, and who in turn experienced exogenous shocks to their educational attainment.

To examine the causal effect of education on climate outcomes, we leverage a new World Bank dataset on compulsory schooling laws (CSLs) in Europe. Europe has had dozens of education reforms throughout the twentieth century expanding the number of years of education legally mandated through compulsory schooling laws. Figure 3.3 in the Appendix includes a map of the number of compulsory schooling law reforms which can be mapped to the ESS data over this time period. For each CSL, we have information on the year it was passed, the year it came into effect, and the new minimum schooling requirement under the law. For most CSLs, we also have the school starting age, and assume this to be 6 years – the most common school starting age – for CSLs for which it is missing; this lets us calculate the birth year of the first affected cohort. We identify the CSL which applies to each respondent by finding the CSL that is applicable to their birth year cohort in the country in which they were surveyed.

Together, these two unique datasets yield exogenous shocks to education which can be mapped directly onto climate outcomes including beliefs, behaviors, policy preferences, and voting.

3.3 Empirical Strategy

3.3.1 Compulsory Schooling Laws as an Instrument

Compulsory schooling laws are commonly used in the economics literature as an instrument for educational attainment. We briefly review the necessary conditions for their use in our context. First, compulsory schooling must affect educational attainment. While this may seem obvious, we show in Section 3.3.2 that this relationship holds for many reforms, but does not necessarily hold for all. Thus, as an additional specification, we follow (Oreopoulos, 2006), to carefully identify reforms which bind – that is, reforms which affect a large enough share of students to have a detectable increase in educational attainment. Our primary specification includes all reforms to alleviate concerns about restricting the analysis to a selected sample.

Second, compulsory schooling must affect climate outcomes through the educational attainment channel, and not be confounded by other factors. Given that the passing of compulsory schooling laws is a national, exogenous shock, resulting gains in education are largely orthogonal to other factors that would otherwise make the individual schooling decision endogenous. For example, a potential confounding variable in the education-climate relationship is individuals' valuation of the future (e.g. their discount rates or degree of present bias), which can simultaneously motivate them to pursue education as an investment in their future, as well as be concerned about the future costs of climate change. Compulsory schooling laws overcome this confounder by mandating individuals to obtain greater educational attainment, regardless of these factors.

The plausibility of the assumption that CSLs affect climate outcomes only through the education channel is further bolstered by the fact that most of the possible effects of CSLs on other mediating factors, such as income, likely increase as a direct result of the education channel. This means our estimate is the bundled effect of education, including changes in income and other mediators, that come with an exogenous increase in schooling. In line with both of these points, Table A4 in the Online Appendix shows a strong first stage on education, while no statistically significant effect on other variables which should not be affected by CSL changes and would not operate through the education channel, such as gender or country of birth.

Our estimation strategy instruments for years of education using a series of indicators for whether each compulsory schooling law binds for a given cohort of individuals. We construct these indicators cumulatively, that is, the estimated effect of the current law is the marginal effect of the law relative to the prior law. We run a two-stage least squares regression where the second stage regresses our climate outcomes on predicted education based on the applicable compulsory schooling laws, controlling for time trends and country fixed effects.[6] For a given individual i

[6] We further Winsorize educational attainment at the 1 percent level, to minimize outlier bias and address spurious coding in the ESS data of extreme values. With Winsorization, we have a minimum of two years of schooling and a maximum of 22 years. Without Winsorization, 414 respondents or 0.11% of our sample report at least 30 years of education, which clearly does not map to our standard notion of years of full-time education, even for individuals with a PhD, motivating Winsorization. Otherwise, the maximum reported education is 60 years of schooling which exerts undue leverage on the rest of the data. Nevertheless, our results also hold when using raw years of education or topcoding at 20 years of education instead.

we estimate:

$$E_{icy} = \alpha_c + \beta_r \mathbf{CSL}_{icyr} + \delta T_y + \varepsilon_{icy} \tag{3.1}$$

$$Y_{icy} = \alpha_c + \beta_r \widehat{\mathbf{E_{icyr}}} + \delta T_y + \varepsilon_{icy} \tag{3.2}$$

where $\mathbf{CSL_{icyr}}$ is a binary indicator of whether an individual i in country c is a member of a cohort y affected by the reform r, and is therefore in the treatment group.[7] We estimate effects across multiple countries and reforms, with $\mathbf{CSL_{icyr}}$ representing a vector of binary indicators across all included reforms r.[8] In Equation (3.1) we estimate the first stage of the effect of CSLs on educational attainment E_{icy}. Since educational attainment has trended upward over time, we also condition on a time trend T_y.[9] We include country fixed effects δ_c given that we analyze results in a unified cross-country framework. We interact time trends and country fixed effects to produce country-specific time trends. Standard errors are clustered at the country-law (e.g., the CSL) level, which is the level of treatment assignment. We estimate Equation (3.2), the causal effect of additional education on a given climate outcome Y_{icy}, with two-stage least squares, where the first stage is estimated from Equation (3.1) with educational attainment instrumented by CSL reforms.

This specification mirrors those most common in the economics literature (Acemoglu and Angrist, 2000; Lleras-Muney, 2005; Oreopoulos, 2006). It is important to note that these strategies all identify local treatment effects of education that are applicable to individuals on the margin of dropping out in the absence of the CSL. This is the policy-relevant estimate if the policy in question is to increase minimum schooling requirements.

[7] This is defined based on each respondents' birth year and starting school age to derive when the reform would first take effect for a given individual.

[8] Omitted indicators are the earliest laws in each country, such that the earliest laws take the value of the country fixed effect, and each subsequent law has a positive β estimate as long as the reform i increased education relative to the country's time trend.

[9] Our estimates are robust to using linear, quadratic, and cubic time trends as well as completely flexible birth cohort fixed effects, implying that the functional form of the trend does not drive results.

Furthermore, we include country-averaged treatment effect estimates to address the negative weighting issue caused by comparing countries increasing treatment to already-treated countries (Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfoeuille, 2022; Goodman-Bacon, 2021; Sun and Abraham, 2020). In the spirit of Callaway and Sant'Anna (2021), this process estimates Equations (3.1) and (3.2) separately for each country, ensuring valid comparisons, then averages estimates across countries to produce a single point estimate for each outcome. In the aggregation step, we weight each country proportionally to the inverse of the variance of its estimate, meaning that more precisely estimated treatment effects receive more weight. These estimates are reported in the results as "country-averaged treatment effects" or "country-avg TE" and are broadly similar to the standard IV estimates (correlation 0.87), implying that cross-country comparisons do not significantly bias the main estimation (see Online Appendix Figure A4 for a scatterplot).

3.3.2 First Stages: the Effect of CSLs on Education

Compulsory Schooling Laws (CSLs) legally mandate an increase in educational attainment, often by raising the minimum school leaving age. For example, in 1963, Italy increased minimum schooling from 5 years of education to 8 (equivalent to increasing the minimum school leaving age from 11 to 14 years old).As an additional robustness check, we carefully identify reforms for which there is a binding first stage – that is, where an increase in required years of schooling by CSLs increases average educational attainment, net of the time trend. While legally enforceable, changes to CSLs will only have a strong first stage if they are enforced, rolled out rapidly, and bind for those who would otherwise not proceed to attain more schooling without the law (e.g., some individuals may attain 8 years of education in Italy even before it was legally required).

Figure 3.1 shows the 20 countries with relevant reforms (and up to 41 country-reforms, with several reforms in some countries). Online Appendix Figure A2 shows that results are not sensitive to restricting to only reforms with positive or positive and statistically significant first stages. Table A3 in the Online Appendix shows all first stages estimates with positive effects, including those that are not statistically significant. Countries in the main analysis include Austria, Belgium, Switzerland, Czechia, Germany, Estonia, Spain, Finland, France, Hungary, Iceland, Ireland, Italy, Lithuania, Netherlands, Norway, Poland, Portugal, Russia, and Sweden. We exclude countries, such as the United Kingdom, where reforms occurred at the sub-national level and do not map cleanly to the ESS data.

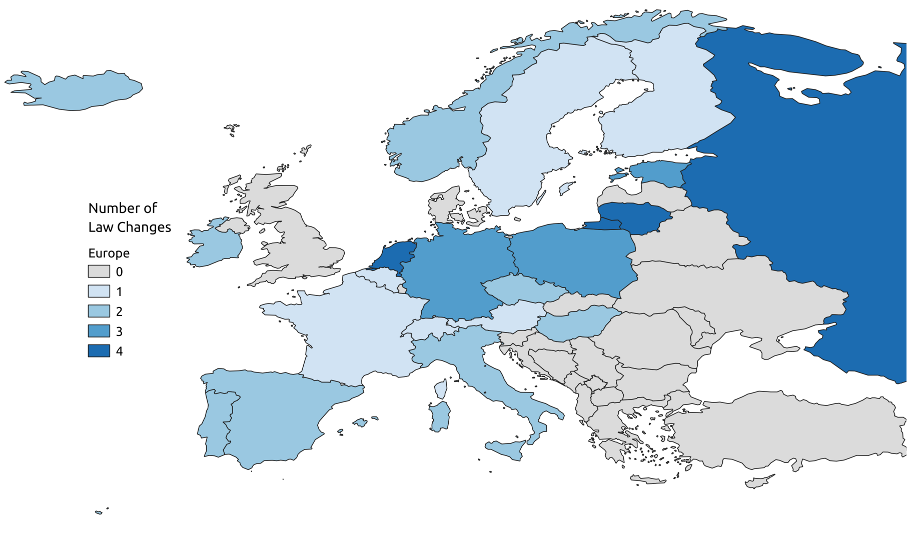

Figure 3.1. Compulsory Schooling Law Changes by Country

Notes: This figure shows the number of compulsory schooling law changes by country. Countries in the main analysis include Austria, Belgium, Switzerland, Czechia, Germany, Estonia, Spain, Finland, France, Hungary, Iceland, Ireland, Italy, Lithuania, Netherlands, Norway, Poland, Portugal, Russia and Sweden.

3.4 Results

Results on our three main pro-climate indices - beliefs, behaviors, and policy preferences - as well as green voting are shown in Table 3.2. An additional year of education leads to highly statistically significant increases of 1.9 percentage points (PP) in pro-climate beliefs, 3.0 PP in behaviors, 0.8 PP in policy preferences, and a nonsignificant 0.3 PP in green voting. These impacts translate into a 2.9% increase for beliefs, 4.3% for behaviors, 1.3% for policy preferences, and 4.3% for green party voting. Panel B of Table 3.2 shows the results with continuous outcomes to ensure results are not driven by binary threshold values defined as being "pro-climate"; results remain consistent. Point estimates are positive and p-values also follow a similar pattern. For example, an additional year of education has large and statistically significant effects on pro-climate beliefs and behaviors, with p-values below < 0.001 in both panels. Of note, while effect directions and statistical significance can be compared, the magnitudes in Panels A and B are not directly comparable.[10] In Online Appendix Figure A2 we include a series of robustness tests, such as various time trends and restrictions to both positive and positive and significant first stages. Results show consistently large and positive effects of education on pro-climate beliefs, behaviors, and policy preferences. Appendix Figure A3 additionally shows the robustness of these estimates to the inclusion of particular countries or reforms by plotting the distribution of leave-one-out β estimates. Benchmarking these results to the effect of education on income, we find a similar sized effect, 1.8%, on respondent income being above median.

In Figure 3.2, we compare the causal effects derived from IV estimates on the three pro-climate indices and green voting to their corresponding OLS correlation estimates (also with country fixed effects and linear time trend) , expressed in terms of standard deviations for comparability between outcomes. In Figure 3.2 and Table 3.3 we analyze outcomes using binary indicators for ease of interpretation. Results are similarly robust whether using binary

[10] In Panel A, a one unit change in the outcome is the difference between being below and above median, whereas in Panel B, a move from 0 to 1 means changing from the most anti-climate response to the most pro-climate.

or continuous outcomes. The gains shown in Table 3.2 translate to 0.072 standard deviation increase for pro-climate beliefs, a 0.096 increase for behaviors, a 0.026 increase for policies, and a 0.013 for green party voting. Moreover, IV causal estimates are substantially larger than OLS estimates for beliefs, and behaviors. One important potential explanation for these larger causal estimates is downward bias in the OLS estimates due to income effects. More educated individuals are often richer, and richer individuals are often more conservative – a standard assumption in political economy models (Meltzer and Richard, 1981) – and thus might be less pro-climate. Indeed in Table A1 in the Online Appendix we see correlations exactly along these lines. The substantial increase in causal IV estimates relative to OLS estimates for these outcomes highlights the importance of credible causal identification of the effects of education on pro-climate outcomes.

Table 3.2. The Effect of Education on Pro-Climate Outcomes.

Notes: This table shows the causal effect of a year of education on each pro-climate outcome index, as in Equation (3.2). The outcome in Panel A denotes effects on being pro-climate defined in binary terms (relative to the median). Panel B shows averages of the continuous outcomes, where 1 is the most pro-climate response to each question and 0 is the least. Standard errors clustered by country×CSL in parentheses. Standard errors are in parentheses and p-values are in brackets. Sample sizes vary due to variation in the availability of outcomes across survey questions. Green voting is available in multiple rounds of the survey, but only for select countries with green parties.

	(1) Pro-climate beliefs	(2) Pro-climate behaviors	(3) Pro-climate policy preferences	(4) Green voting
Panel A: indicators for above-median climate stance				
Years of education	0.019	0.030	0.008	0.003
	(0.005)	(0.004)	(0.003)	(0.008)
	[0.000]	[0.000]	[0.001]	[0.670]
Country-Avg Treatment Effect	0.023	0.043	0.010	0.005
Mean	0.652	0.703	0.632	0.080
Percent Change	2.9 %	4.3 %	1.3 %	4.3 %
Panel B: continuous pro-climate variables				
Years of education	0.012	0.017	0.007	0.003
	(0.003)	(0.003)	(0.002)	(0.008)
	[0.000]	[0.000]	[0.000]	[0.670]
Country-Avg Treatment Effect	0.010	0.024	0.005	0.005
Observations	33238	33238	32698	100474
Clusters	66	66	66	71
Mean	0.642	0.645	0.603	0.080
Percent Change	1.9 %	2.7 %	1.1 %	4.3 %

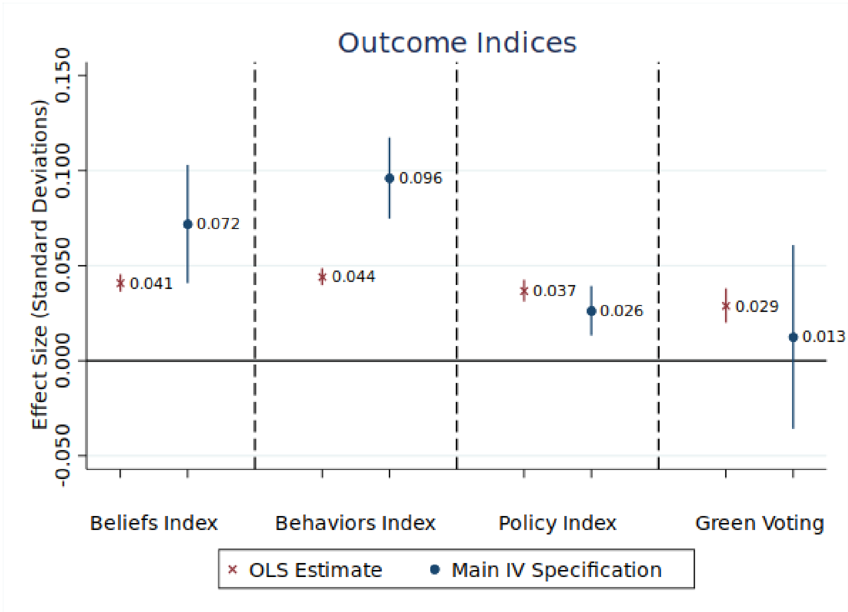

Figure 3.2. Effects of Education on Pro-Climate Outcomes - Standardized Causal Estimates vs. Correlations

Notes: This figure plots estimates from our main IV specification which captures causal estimates compared to the OLS estimate which shows correlational estimates, both with a pooled linear time trend and country fixed effects. The OLS regression is restricted to the same sample as the IV. The indices are standardized and expressed in terms of standard deviations. 90% confidence intervals shown from standard errors clustered at the country×law level.

While Table 3.2 shows our primary results, the panels of Table 3.3 break down each of the indices into their components, showing positive and significant estimates on most sub-outcomes. In terms of specific indicators, on beliefs, we find one year of education causes a 2.8 percentage point increase in thinking the world's climate is changing, with somewhat smaller effects on thinking that climate change has a bad impact, worrying about climate change and worrying about dependency on fossil fuels. In terms of behaviors, we find 2.7 and 2.9 percentage point increases in reducing energy use and buying energy efficient appliances, respectively, with a 3.3 PP increase in having thought about climate change before today. For policy preferences, we find a 1.1 PP increase in favoring bans on the sale of inefficient appliances and a 2.1 PP increase on favoring subsidies for renewable energy. In contrast, we find a null or even slightly negative effect on preferences to increase taxes on fossil fuels, a result that attenuates our policy index despite two of the three components being strongly positive. It is plausible that respondents either fail to see the equivalence between taxes and subsidies or the redistributive effects of taxes are more harmful to those on the margin of receiving additional education than are the effects of subsidies. The impacts on green voting are indistinguishable from zero, possibly because green parties remain a small player on the global political stage, with only 8% of respondents reporting having voted for one, even when restricting to countries with active green parties.

3.5 Conclusion

Climate change poses existential risks to the planet and generates trillions of dollars in annual costs to society. While changing pro-climate beliefs, behaviors, policy preferences, and voting is difficult, one approach that can move the needle is additional education. This paper provides strong causal evidence that education can impact a range of pro-climate outcomes. We find that an additional year of education is linked with increases in pro-climate beliefs, behaviors, and most policy preferences, with little effect on voting for green parties.

While education is often a footnote in climate change agendas, this paper reveals the promise of education as an additional tool to combat climate change. Europe in particular is a context where climate change is receiving substantial attention, including efforts such as the European Green New Deal, yet education remains an underutilized lever. Moreover, while educational attainment has expanded dramatically in recent decades, the median school reform law in 2020 in Europe guaranteed only 10 years of schooling, a full two years below a complete primary and secondary education of 12 years. These gaps are even more dramatic in the developing world; in sub-Saharan Africa educational reform laws only guarantee 8 years of schooling on average. Expanding access to education has traditionally been believed to play a transformative role in the economic and social well-being of societies – it now also appears to play a vital role in the battle against climate change.

This dissertation chapter was co-authored with Noam Angrist, Joshua Graff Zivin, and Harry Anthony Patrinos and is currently being revised for resubmission to the Review of

Table 3.3. Effect of Education on Each Element of Pro-Climate Outcome Indices.

Notes: This table shows point estimates for each of the elements of the indices. Panel A shows the elements of the beliefs index, Panel B the behaviors index, and finally Panel C shows both the policy preferences index and green voting. Outcomes are binary, so multiplying the point estimate by 100 yields the percentage point increase in the likelihood of having a pro-climate stance on the given outcome from an additional year of education. Standard errors clustered at the country×law level in parentheses and p-values in brackets. Estimates include country fixed effects and country-specific linear time trends. “CC” means “climate change”.

	Climate Outcomes				
	(1)	(2)	(3)	(4)	(5)
Panel A: pro-climate beliefs	Think the world’s climate is changing	CC has bad (not good) impact across world	Worried about CC	Pro-clean energy beliefs	Too dependent on fossil fuels
Years of education	0.028	0.021	0.016	0.005	0.016
	(0.008)	(0.011)	(0.007)	(0.003)	(0.006)
	[0.000]	[0.049]	[0.022]	[0.196]	[0.008]
Country-Avg TE	0.026	0.024	0.030	0.012	0.025
Observations	32631	31139	32059	32460	31971
Clusters	66	65	66	66	66
Mean	0.559	0.580	0.747	0.729	0.676
Percent Change	4.9 %	3.6 %	2.2 %	0.6 %	2.4 %
Panel B: pro-climate behaviors	Thought about CC before today	Likely to buy most efficient appliance	How often do things to reduce energy use		
Years of education	0.033	0.029	0.027		
	(0.007)	(0.009)	(0.007)		
	[0.000]	[0.002]	[0.000]		
Country-Avg TE	0.049	0.046	0.030		
Observations	33093	32496	32869		
Clusters	66	66	66		
Mean	0.717	0.676	0.724		
Percent Change	4.6 %	4.2 %	3.7 %		
Panel C: pro-climate policy preferences & voting	Favor increase taxes on fossil fuels to reduce CC	Favor subsidise renewable energy	Favour ban sale of inefficient household appliances	Green Voting	
Years of education	-0.006	0.021	0.011	0.003	
	(0.006)	(0.005)	(0.006)	(0.008)	
	[0.314]	[0.000]	[0.079]	[0.670]	
Country-Avg TE	-0.018	0.024	0.024	0.005	
Observations	31913	32308	32089	100474	
Clusters	66	66	66	71	
Mean	0.546	0.761	0.591	0.080	
Percent Change	-1.1 %	2.8 %	1.8 %	4.3 %	

Economics and Statistics.

3.A Appendix

3.A.1 Correlations between Education, Income, and Conservatism

Table 3.4 shows the correlations between education, income and schooling. As expected, income and schooling are positivley correlated, as are income and conservatism. However, the negative correlation between conservatism and schooling mediates the schooling-income relationship, making the effect of schooling on pro-climate outcome ex ante ambiguous.

Table 3.4. Correlations Between Education, Income, and Conservatism.

Notes: This table shows correlation coefficients between income, schooling, and conservatism. Conservatism reflects where respondents self-report falling on a 0-1 scale where 1 is most right-leaning and 0 is most left-leaning on the political spectrum. Years of schooling is the Winsorized years of education attained, as in the main text. Lastly, income is the self-reported household income decile, normalized to fall on the 0-1 range. Raw correlations are simply the correlation coefficients in our main analysis sample. Residualized coefficients are the result of first residualizing income, schooling, and conservatism on country fixed effects and the pooled linear time trend as in the main analysis.

	Raw		Residualized	
	Schooling	Conservatism	Schooling	Conservatism
Income	0.388	0.071	0.313	0.077
Schooling		-0.023		-0.016

3.A.2 Green Party Coding

Table 3.5. Green Party Coding.

Notes: An individual is coded as voting green if they reported voting for one of the listed parties in the last election. Missing responses and those from countries with no green parties in the relevant election are coded as missing. Those who voted for a different party in countries with green parties at the time are coded as not voting green.

Country	Abbr.	Green Parties
Austria	AT	Grüne
Belgium	BE	Groen!, Ecolo
Switzerland	CH	Green Party
Cyprus	CY	The Cyprus Green Party
Czechia	CZ	Česká pirátská strana
Germany	DE	Alliance 90/The Greens
Denmark	DK	SF Socialistisk Folkeparti, Alternativet
Estonia	EE	Erakond Eestimaa Rohelised
Spain	ES	En Comú Podem, Iniciativa per Catalunya-Verds
Finland	FI	Green League
France	FR	EELV (Europe Ecologie Les Verts)
Hungary	HU	LMP (Lehet Más A Politika)
Ireland	IE	Green Party
Iceland	IS	Vinstri hreyfinguna - grænt framboõ
Lithuania	LT	Lithuanian Green Party (LZP)
Latvia	LV	Zaļo un Zemnieku savienība
Netherlands	NL	Green Left
Norway	NO	Miljøpartiet De Grønne
Portugal	PT	PAN - Pessoas-Animais-Natureza
Sweden	SE	Miljöpartiet de gröna

3.A.3 First Stage Estimates

In this paper, we leverage a new dataset on compulsory schooling laws in Europe from the World Bank, which is one of the largest databases on CSLs to date. Figure 3.3 shows the number of compulsory schooling law reforms by country.

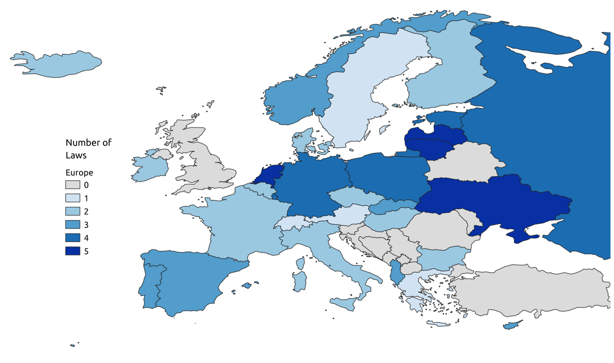

Figure 3.3. Number of compulsory schooling laws (CSL) by country.
Notes: The map shows all CSLs that can be mapped to the ESS data. Note that a British reform commonly used in literature is excluded from our analysis, because this reform is region-specific and the ESS data does not have enough geographic granularity to accurately assign regional laws to respondent's individual level climate outcomes.

Table 3.6. Share of CSL Compliers and Green Votes and Average Education by Country.

Notes: This table shows the share of respondents in the sample who report having attained at least the minimum level of schooling required of the CSL assigned by their year of birth and survey country (i.e. they are "compliers"). 95% of respondents in our sample are compliers. Noncompliers may have truly not achieved the legally mandated level of schooling despite the requirement. Alternatively, they may have been incorrectly assigned to a CSL, potentially because of migration out of their country of birth, exceptions to the compulsory schooling laws, or mischaracterization of the CSL rules. Average Winsorized educational attainment and share of respondents who voted for a green party are also shown. Missing green vote share means there was no active green party in the given country during our sample.

Country	Share Compliers	Avg Education (Years)	Green Vote Share
AT	0.969	12.639	0.110
BE	0.934	13.816	0.066
CH	0.940	11.393	0.093
CZ	0.909	12.548	0.085
DE	0.960	14.479	0.163
EE	0.978	13.246	0.028
ES	0.907	12.675	0.008
FI	0.990	13.946	0.109
FR	0.848	12.428	0.064
HU	0.880	11.929	0.046
IE	0.995	14.550	0.014
IS	0.968	15.128	0.197
IT	0.941	11.322	.
LT	0.970	12.917	0.019
NL	0.927	13.790	0.062
NO	0.983	14.313	0.030
PL	0.989	12.553	.
PT	0.933	10.196	0.015
RU	0.989	13.138	.
SE	0.976	13.441	0.081

Table 3.7. CSL Changes with Any Education Effect

Notes: This table shows first stage estimates for Equation (1) for each CSL that positively affects educational attainment. The point estimate is the effect on educational attainment following each CSL's implementation, controlling for country-specific linear time trends and country fixed effects. The numbers following each country code indicate the years of schooling required by each law (AL8 requires 8 years of schooling in Albania). The listed year is the first birthyear affected by the law. CSL changes not included in this table have nonpositive first stage estimates. Standard errors are in parentheses.

Reform & Year	Estimate	Positive	Positive+Significant
AL8	2.079	✓	✓
1963	(0.271) [0.000]		
AT9	0.595	✓	✓
1966	(0.201) [0.004]		
BE8	1.217	✓	✓
1919	(0.122) [0.000]		
BG8	0.718	✓	✓
1960	(0.349) [0.042]		
CH9	0.069	✓	
1970	(0.209) [0.741]		
CY6	0.194	✓	
1962	(0.480) [0.687]		
CZ9	0.363	✓	✓
1948	(0.013) [0.000]		
DE13	0.502	✓	✓
1992	(0.044) [0.000]		
DE4	0.315	✓	✓

Table 3.7. CSL Changes with Any Education Effect Continued

1920	(0.015) [0.000]		
DE8	0.981	✓	✓
1946	(0.026) [0.000]		
DK7	0.893	✓	✓
1958	(0.247) [0.000]		
DK9	0.109	✓	
1972	(0.310) [0.726]		
EE6	0.850	✓	
1920	(0.776) [0.276]		
EE8	0.729	✓	
1958	(0.666) [0.276]		
ES8	0.331	✓	
1970	(0.355) [0.353]		
FI6	1.136	✓	
1921	(1.399) [0.419]		
FR10	0.116	✓	
1967	(0.061) [0.059]		
HU10	0.568	✓	✓
1961	(0.173) [0.001]		
HU8	1.077	✓	✓
1945	(0.095) [0.000]		
IE9	0.182	✓	
1972	(0.203) [0.371]		
IS7	1.361	✓	✓

Table 3.7. CSL Changes with Any Education Effect Continued

1936	(0.555) [0.016]			
IT8	1.040	✓		✓
1963	(0.516) [0.046]			
LT11	0.023	✓		
1980	(0.045) [0.607]			
LT5	1.894	✓		✓
1937	(0.051) [0.000]			
LT7	0.766	✓		✓
1953	(0.030) [0.000]			
LT8	1.477	✓		✓
1958	(0.082) [0.000]			
LT9	0.135	✓		
1980	(0.124) [0.276]			
LU10	0.717	✓		✓
1977	(0.098) [0.000]			
LU11	0.785	✓		✓
1993	(0.019) [0.000]			
LV5	1.176	✓		✓
1937	(0.055) [0.000]			
LV7	0.188	✓		✓
1953	(0.032) [0.000]			
LV8	0.736	✓		✓
1958	(0.088) [0.000]			
NL10	0.142	✓		

Table 3.7. CSL Changes with Any Education Effect Continued

1973	(0.162) [0.382]		
NL7	0.169	✓	
1928	(0.292) [0.565]		
NL8	0.570	✓	
1950	(0.289) [0.051]		
NL9	0.220	✓	
1969	(0.173) [0.207]		
NO7	0.761	✓	
1936	(0.772) [0.327]		
NO9	0.105	✓	
1969	(0.809) [0.897]		
PL8	0.309	✓	
1966	(0.463) [0.506]		
PT6	1.170	✓	
1964	(0.804) [0.148]		
PT9	0.756	✓	
1986	(0.750) [0.316]		
RU5	1.482	✓	✓
1937	(0.039) [0.000]		
RU7	0.855	✓	✓
1953	(0.024) [0.000]		
RU8	0.651	✓	✓
1958	(0.086) [0.000]		
RU9	0.069	✓	

Table 3.7. CSL Changes with Any Education Effect Continued

2004	(0.096) [0.472]		
SE9	0.357	✓	
1963	(0.864) [0.680]		
SK8	1.338	✓	✓
1948	(0.153) [0.000]		
SK9	0.709	✓	✓
1948	(0.145) [0.000]		
UA12	0.714	✓	✓
2002	(0.024) [0.000]		
UA5	1.575	✓	✓
1937	(0.077) [0.000]		
UA7	1.144	✓	✓
1953	(0.043) [0.000]		
UA8	0.474	✓	✓
1958	(0.141) [0.001]		
UA9	0.102	✓	
1996	(0.156) [0.517]		
Observations	315927		

Table 3.8. First Stage Estimates for the Pro-climate Beliefs Index

Notes: First stage estimates for each of the reform indicator variables, with the omitted variable in each country being the first CSL in the sample. Includes country fixed effects and a pooled linear time trend.

Table 3.8. First Stage Estimates for the Pro-climate Beliefs Index Continued

	(1)		
	Years of education		
	Estimate	Standard Error	P-Value
AT9	-0.094	(0.186)	[0.614]
BE8	0.949	(0.191)	[0.000]
CZ9	-1.207	(0.878)	[0.169]
DE13	-0.326	(0.710)	[0.646]
DE4	1.023	(0.520)	[0.049]
DE8	0.936	(0.507)	[0.065]
HU10	0.055	(0.238)	[0.816]
HU8	0.311	(0.650)	[0.633]
IS7	3.828	(2.390)	[0.109]
IT8	2.121	(0.174)	[0.000]
LT5	-0.015	(1.102)	[0.989]
LT7	0.425	(0.487)	[0.383]
LT8	0.926	(0.415)	[0.026]
NL8	0.840	(0.371)	[0.023]
RU5	0.110	(0.941)	[0.907]
RU7	0.142	(0.504)	[0.778]
RU8	0.052	(0.435)	[0.906]
CH9	-0.658	(0.209)	[0.002]
EE6	0.275	(0.250)	[0.272]
EE8	0.594	(0.231)	[0.010]
ES8	2.031	(0.198)	[0.000]
FI6	0.481	(0.193)	[0.013]

Table 3.8. First Stage Estimates for the Pro-climate Beliefs Index Continued

FR10	0.950	(0.175)	[0.000]
IE9	0.607	(0.155)	[0.000]
LT11	0.066	(0.449)	[0.882]
LT9	-0.277	(0.191)	[0.148]
NL10	-0.152	(0.350)	[0.664]
NL7	0.289	(0.369)	[0.433]
NL9	0.602	(0.341)	[0.077]
NO7	2.935	(1.518)	[0.053]
NO9	0.636	(0.212)	[0.003]
PL8	0.884	(0.216)	[0.000]
PT6	2.196	(0.263)	[0.000]
PT9	1.204	(0.289)	[0.000]
RU9	-0.404	(0.323)	[0.212]
SE9	0.443	(0.202)	[0.028]
BE12	0.338	(0.188)	[0.072]
CZ10	-0.492	(0.214)	[0.021]
DE9	-0.248	(0.694)	[0.721]
EE11	-0.169	(0.386)	[0.661]
EE9	-0.595	(0.195)	[0.002]
ES10	1.218	(0.222)	[0.000]
ES6	-1.752	(0.214)	[0.000]
FI9	1.585	(0.174)	[0.000]
FR8	0.951	(1.952)	[0.626]
IE10	-0.861	(0.229)	[0.000]
IS9	0.981	(0.255)	[0.000]

Table 3.8. First Stage Estimates for the Pro-climate Beliefs Index Continued

IT12	0.173	(0.365)	[0.636]
NL12	-0.505	(0.241)	[0.036]
PL10	0.299	(0.742)	[0.687]
PL7	-0.913	(0.234)	[0.000]
PL9	0.030	(0.331)	[0.928]
PT4	0.986	(0.315)	[0.002]
Observations	33238		
Fstatistic	147.8		

3.A.4 CSL Validity Test

Table 3.9. Validity Test

Notes: This table shows the coefficient on the indicator for being after the first CSL change in a country while additionally controlling for country fixed effects and linear pooled time trends. “any reform” is zero for respondents in the sample under the first compulsory schooling law in the analysis window and one for all others. The outcomes are (1) an indicator for the respondent being male, (2) an indicator for being born in the country they are surveyed in, and (3) Winsorized years of education. The small and nonsignificant estimates in Columns (1) and (2) suggest that the instrument is not predictive of other factors like gender and whether the respondent was born in the country in which they are surveyed,supporting the validity of the instrument, as CSL changes have no discernible effect on predetermined outcomes like gender and birth country. The effect in Column (3) is not large because this only compares respondents under the first law to all other laws, losing important variation which could otherwise not be captured in a single concise statistic. Note that while the ESS has plenty of other outcomes that could be tested in this manner, gender and birth location are the primary ones that we do not expect to be influenced by education, as these are determined before the amount of schooling is realized.

	(1) Male	(2) Born in Country	(3) Years of Education
Any reform	-0.003	0.013	0.070
	(0.012)	(0.014)	(0.252)
	[0.820]	[0.370]	[0.782]
Observations	314017	313907	314125

Standard errors in parentheses. P-values in brackets.

Table 3.10. Placebo Results for Main Indices.

Notes: This table shows the IV estimates when assigning placebo instruments and countries for the four main outcomes.

	(1) Pro-climate beliefs	(2) Pro-climate behaviors	(3) Pro-climate policy preferences	(4) Green voting
Years of Education	0.007	0.005	0.001	0.006
	(0.007)	(0.009)	(0.007)	(0.007)
	[0.363]	[0.595]	[0.907]	[0.393]
Observations	33238	33238	32698	100474
Clusters	104	104	104	107
Mean	0.652	0.703	0.632	0.080

3.A.5 Placebo Tests

To demonstrate that the instruments are truly the driving force in identifying results, we include placebo test results where each observation is assigned the country and accompanying CSL assignments of another randomly drawn respondent in the sample. Then, the main index regressions are rerun using these placebo instruments, and country assignment, with the original birth year and outcome. Each outcome has a highly nonsignificant placebo point estimate under one percentage point, suggesting that identifying variation is indeed coming from the correctly assigned CSLs.

3.A.6 Robustness and Alternate Specifications

In this section, we consider the robustness of our estimates to several modeling decisions. We analyze results with all positive first stages and all positive and significant first stages. To establish the first stage, we estimate Equation 3.1 on all rounds of the ESS with standard errors clustered by country×law[11]. In addition, we analyze results with alternative time trends such as quadratic and cubic time trends as well as completely flexible birth cohort fixed effects. Finally, rather than using indicators for compulsory schooling laws as the instrument for educational

[11]By using all rounds of the ESS to determine strong first stages, we have more power to estimate the true effect of compulsory schooling laws beyond the country's time trend.

attainment, we use the current level of the minimum schooling requirement rather than a binary indicator, controlling for the upward time trends and country fixed effects. Figure 3.4 shows a plot of estimates across these robustness tests, showing broadly similar patterns and robustness. We see broadly consistent results across specifications; the positive and large effects of education on pro-climate outcomes persist. Figure 3.5 includes a set of additional robustness tests, assessing robustness to leaving out any one country or reform, also showing highly consistent results.

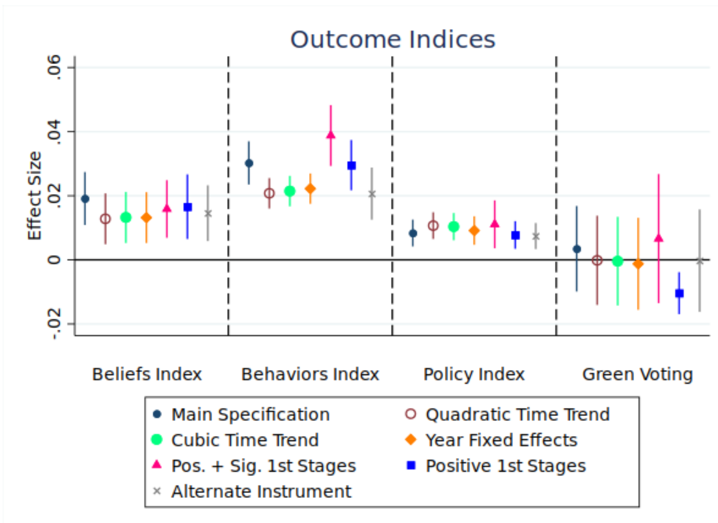

Figure 3.4. Robustness Checks

Notes: Figure shows IV estimates for the four outcome indices under alternative time trend specifications and inclusion criteria for the first stage. The main specification is as in Section 3.3 (linear pooled time trend and country fixed effects for all CSLs in the sample). *Quadratic Time Trend* is the same as *Main* adds in a squared birth year term, while *Cubic Time Trend* also adds a cubic term. *Year Fixed Effects* uses year fixed effects as the time trend along with country fixed effects, pooling birth years earlier than 1920 with the 1920 cohort (less than 0.5% of the sample). *Pos. + Sig. 1st Stages* restricts the analysis to reforms with a positive and statistically significant first stage, and *Positive 1st Stages* uses all positive first stages regardless of p-values. *Alternate* is the secondary IV specification where the instrument is the number of years of schooling interacted with country. 90% confidence intervals shown from standard errors clustered at the country$\times$law.

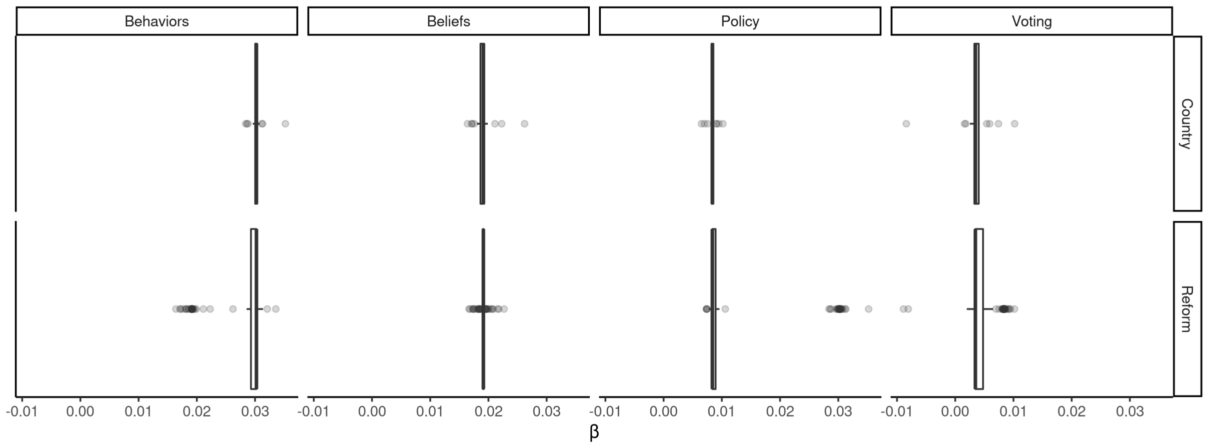

Figure 3.5. Robustness to Leaving Out Countries or Reforms

Notes: This figure shows boxplots of the distributions of β estimates for each of the four main indices when leaving out one country (top row) or all reforms from one country (bottom row). Nearly all estimates from this jackknife procedure remain positive, with the 25th-75th percentile clustered tightly around the main estimate, indicating that these results are not sensitive to the inclusion of any particular reform or country.

Finally, figure 3.6 shows the correlation between the standard IV estimates and the Callaway and Sant'Anna (2021)-inspired country-averaged treatment effect estimates. The correlation is 0.87, with both estimators closely aligned on all outcomes, suggesting that cross-country comparisons are not significant sources of bias in the main analysis.

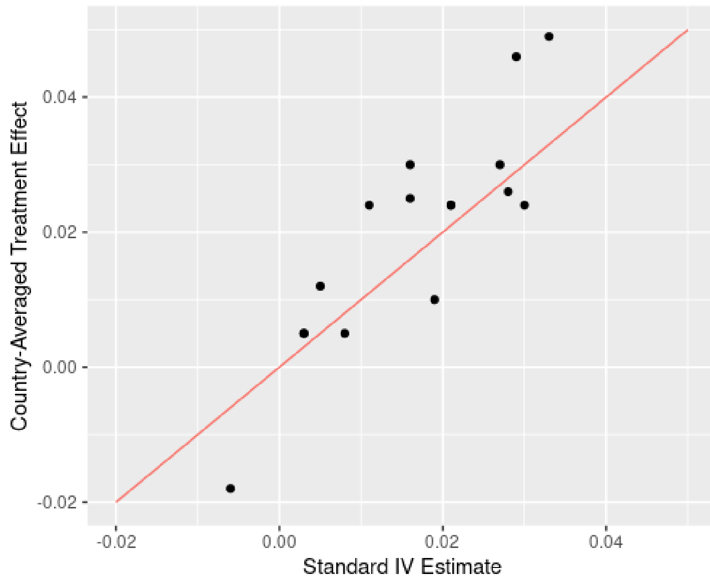

Figure 3.6. Correlation between Estimators

Notes: This figure demonstrates the correlation across all outcomes between the two main estimators used in this analysis: the standard IV estimate and the country-averaged treatment effect estimate. Each point represents one outcome, either one of the main indices or one of the constituent questions. 45°line through the origin shown for reference.

3.A.7 ESS Question Text and Pro Environmental Beliefs Definitions

We include exact question wording and coding for our main pro-climate outcomes.

- **How likely to buy most energy efficient home appliance**: If you were to buy a large electrical appliance for your home, how likely is it that you would buy one of the most energy efficient ones?
  0 Not at all likely - 10 Extremely likely

- **How often do things to reduce energy use**: There are some things that can be done to reduce energy use, such as switching off appliances that are not being used, walking for short journeys, or only using the heating or air conditioning when really needed. In your daily life, how often do you do things to reduce your energy use?

- **How much electricity should be generated from [energy source]**: The highlighted box at the top of this card shows a number of energy sources that can be used to generate electricity. Please take a moment to look over them. How much of the electricity used in [country] should be generated from each energy source? First, how much of the electricity used in [country] should be generated from [energy source]?
  Note: pro-clean energy beliefs outcome is an average of being pro-hydro and solar, and anti-coal.

- **How worried too dependent on fossil fuels**: How worried are you about [country] being too dependent on using energy generated by fossil fuels such as oil, gas and coal?

- **Do you think the world's climate is changing**: You may have heard the idea that the world's climate is changing due to increases in temperature over the past 100 years. What is your personal opinion on this? Do you think the world's climate is changing?

- **How much thought about climate change before today**: How much have you thought about climate change before today?

- **How worried about climate change**: How worried are you about climate change?
- **Climate change good or bad impact across world**: How good or bad do you think the impact of climate change will be on people across the world? Please choose a number from 0 to 10, where 0 is extremely bad and 10 is extremely good.
  0 Extremely bad - 10 Extremely good
- **Favour increase taxes on fossil fuels to reduce climate change**: To what extent are you in favour or against the following policies in [country] to reduce climate change? Increasing taxes on fossil fuels, such as oil, gas and coal.
- **Favour subsidise renewable energy to reduce climate change**: To what extent are you in favour or against the following policies in [country] to reduce climate change? Using public money to subsidise renewable energy such as wind and solar power.
- **Favour ban of least energy efficient household appliances to reduce climate change**: To what extent are you in favour or against the following policies in [country] to reduce climate change? A law banning the sale of the least energy efficient household appliances.

Bibliography

Acemoglu, Daron, and Joshua Angrist. 2000. "How Large Are Human-Capital Externalities? Evidence from Compulsory Schooling Laws." *NBER Macroeconomics Annual*, 15: 9–59. Publisher: The University of Chicago Press.

Angrist, Joshua D., and Alan B. Krueger. 1991. "Does Compulsory School Attendance Affect Schooling and Earnings?*." *The Quarterly Journal of Economics*, 106(4): 979–1014.

Angrist, Noam, Simeon Djankov, Pinelopi K. Goldberg, and Harry A. Patrinos. 2021. "Measuring human capital using global learning data." *Nature*, 592(7854): 403–408. Number: 7854 Publisher: Nature Publishing Group.

Barro, Robert J. 2001. "Human Capital and Growth." *American Economic Review*, 91(2): 12–17.

Bechtel, Michael M., Kenneth F. Scheve, and Elisabeth van Lieshout. 2020. "Constant carbon pricing increases support for climate action compared to ramping up costs over time." *Nature Climate Change*, 10(11): 1004–1009. Number: 11 Publisher: Nature Publishing Group.

Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy*, 70(5): 9–49. Publisher: University of Chicago Press.

Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. 2008. "Staying in the Classroom and Out of the Maternity Ward? The Effect of Compulsory Schooling Laws on Teenage Births." *The Economic Journal*, 118(530): 1025–1054.

Brunello, Giorgio, Margherita Fort, and Guglielmo Weber. 2009. "Changes in Compulsory Schooling, Education and the Distribution of Wages in Europe." *The Economic Journal*, 119(536): 516–539.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Carleton, Tamma A., and Solomon M. Hsiang. 2016. "Social and economic impacts of climate." *Science*, 353(6304): aad9837. Publisher: American Association for the Advancement of Science.

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2022. "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey."

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. 2022. "Fighting Climate Change: International

Attitudes Toward Climate Policies." *NBER Working Paper Series*.

Gathmann, Christina, Hendrik Jürges, and Steffen Reinhold. 2015. "Compulsory schooling reforms, education and mortality in twentieth century Europe." *Social Science & Medicine*, 127: 74–82.

Goldin, Claudia, and Lawrence Katz. 1997. "Why the United States Led in Education: Lessons from Secondary School Expansion, 1910 to 1940." National Bureau of Economic Research, Inc NBER Working Paper 6144.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics*, 225(2): 254–277.

Graff Zivin, Joshua, and Matthew Neidell. 2013. "Environment, Health, and Human Capital." *Journal of Economic Literature*, 51(3): 689–730.

Intergovernmental Panel on Climate Change. 2022. "Climate Change 2022: Impacts, Adaptation and Vulnerability." *IPCC Report 2022*.

Isen, Adam, Maya Rossin-Slater, and Reed Walker. 2017. "Relationship between season of birth, temperature exposure, and later life wellbeing." *Proceedings of the National Academy of Sciences*, 114(51): 13447–13452. Publisher: Proceedings of the National Academy of Sciences.

Kwauk, Christina. 2021. "Who's making the grade on climate change education ambition?"

Lleras-Muney, Adriana. 2005. "The Relationship Between Education and Adult Mortality in the United States." *The Review of Economic Studies*, 72(1): 189–221.

Meltzer, Allan H., and Scott F. Richard. 1981. "A Rational Theory of the Size of Government." *Journal of Political Economy*, 89(5): 914–927. Publisher: The University of Chicago Press.

Meyer, Andrew. 2015. "Does education increase pro-environmental behavior? Evidence from Europe." *Ecological Economics*, 116: 108–121.

Nordhaus, William D. 2007. "A Review of the Stern Review on the Economics of Climate Change." *Journal of Economic Literature*.

Oreopoulos, Philip. 2006. "Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter." *American Economic Review*, 96(1): 152–175.

Park, R. Jisung, A. Patrick Behrer, and Joshua Goodman. 2021. "Learning is inhibited by heat exposure, both internationally and within the United States." *Nature Human Behaviour*, 5(1): 19–27. Number: 1 Publisher: Nature Publishing Group.

Pew Research Center. 2019. “A look at how people around the world view climate change.”

Powdthavee, Nattavudh. 2021. “Education and pro-environmental attitudes and behaviours: A nonparametric regression discontinuity analysis of a major schooling reform in England and Wales.” *Ecological Economics*, 181: 106931.

Sun, Liyang, and Sarah Abraham. 2020. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” 53.

Sunstein, Cass R., Sebastian Bobadilla-Suarez, Stephanie C. Lazzaro, and Tali Sharot. 2017. “How People Update Beliefs about Climate Change: Good News and Bad News.” *Cornell Law Review*, 102(6): 1431–1444.

Tracker, Climate Action. 2022. “Warming Projections Global Update November 2022.” Publisher: Climate Action Tracker.