

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays on Economic and Behavioral Responses to Constraints

Permalink

<https://escholarship.org/uc/item/0t900473>

Author

Sears, James Matthew

Publication Date

2022

Peer reviewed|Thesis/dissertation

Essays on Economic and Behavioral Responses to Constraints

by

James Sears

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Sofia Villas-Boas, Chair

Associate Professor James Sallee

Professor Hilary Hoynes

Spring 2022

Essays on Economic and Behavioral Responses to Constraints

Copyright 2022

by

James Sears

Abstract

Essays on Economic and Behavioral Responses to Constraints

by

James Sears

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Sofia Villas-Boas, Chair

Throughout life, consumers face constraints that impact all aspects of economic decision-making. Government and agency policies can impose resource constraints, limiting the quantity of a natural resource that a consumer can use or imposing fees that shift individuals away from their preferred consumption levels. Households often face financial constraints, restricting their ability to purchase an optimal consumption basket and preventing families from eating optimal diets. Additionally, social constraints can limit behavior through adherence to social norms, with those who violate these norms suffering disutility due to social pressure or public shame from their peers. In this dissertation, I explore individual and household responses to a range of constraints affecting natural resource use, diet, and mobility patterns.

In the first chapter, I explore the responses of urban water customers to price and non-price water conservation measures. Growing urban water demand combined with a shifting global climate present urban water districts with an acute need for policy approaches that can induce both long- and short-run conservation behavior. Using quasi-experimental long-run variation in prices and exposure to short-run price and non-price policies during severe drought, I estimate the effectiveness of these policy approaches for urban water conservation. This paper is the first to not only identify the water conservation impact of public shame, but also isolate its effect from those due to moral suasion and price-based measures. By following the universe of single family water customers served by a major water district over time, I compare the impacts of different policies on the same households and avoid common sample selection issues. First, I utilize rich administrative panel data to characterize the evolution of causal price elasticities across multiple stages of drought. Next, I leverage quasi-experimental variation in exposure to fees, moral suasion, and public shame to separately identify consumer responses to each during the drought emergency. Demand models yield causal price elasticity estimates of -1 prior to adaptation that attenuate to -0.4 to -0.7 during and -0.2 to -0.5 post-drought, with high-usage households displaying the greatest responsiveness. While

subject to behavioral interventions, top water users display no conservation response to excessive use fees. Moral suasion and public shaming show substantial short-run conservation impacts but display immense sensitivity to emergency messaging and consumer beliefs in crisis. Households called out by name in news stories display sizable reductions in water use even after prior exposure to moral suasion and fees, demonstrating the resilience of public shaming’s conservation effect to crowding out by other concurrent policies. These findings yield important implications for the design of future urban water policies that can balance short and long-run conservation goals.

In the second chapter, I investigate the extent to which Supplemental Nutrition Assistance Program (SNAP) participants utilize food bank resources to smooth expenditures throughout the benefit month. Many recent works have documented the existence of the “SNAP benefit cycle” in which a non-negligible portion of enrolled households exhaust their benefits shortly after receipt. Lack of benefits late in the month carries large consequences, with negative impacts ranging from reduced energy intake and nutritional content, to increased likelihood of hypoglycemia and pregnancy-related emergency room admissions, along with lowered test performance and increased disciplinary action for students. These effects diminish or are otherwise nonexistent for SNAP households that exhibit consumption smoothing. The ability for low-income households to complement resources from federal poverty programs with a local public good has major implications for household welfare and the value of food bank networks. Using the USDA Food Acquisition and Purchasing Survey (FoodAPS) dataset, I provide evidence suggesting differential behavior patterns immediately upon receipt of benefits for SNAP households that utilize food banks. After summarizing prior studies documenting the presence of the SNAP benefit cycle and the expenditure patterns within the FoodAPS sample, I discuss the observed differences in household characteristics and acquisition behavior for those using and forgoing food banks. Next, I present visual evidence of variation in food bank utilization over time among surveyed households and of stark differences in day-zero SNAP benefit utilization for households that visited food banks during the survey period. Following the presentation of visual evidence, I leverage variation in timing since SNAP benefit receipt and food bank use among surveyed households to estimate empirical models of daily food-at-home expenditures. I find that, while typical SNAP households spend \$91-92 more on the day of benefit receipt than on a typical day in the final week of the benefit month, households that use food banks spend \$50-54 less that day. This decrease in day-zero expenditure by food bank clients is offset by higher average spending on the next six days, indicating that the expenditure smoothing exhibited by these households primarily occurs in the first week of the benefit month. Subsequent analyses exploring alternate timing specifications and payment methods confirm that the observed expenditure patterns are driven by payments out of SNAP benefits, and that the expenditure smoothing observed by food bank households is largely isolated to the first week following benefit issuance.

In the the third chapter, co-authored with Sofia Villas-Boas, J. Miguel Villas-Boas, and

Vasco Villas-Boas, we estimate the mobility responses prior and in response to COVID -19 stay-at-home mandates. The recent spread of COVID-19 across the U.S. led to concerted efforts by states to “flatten the curve” through the adoption of stay-at-home mandates that encouraged individuals to reduce travel and maintain social distance. Combining data on changes in travel activity and human encounter rates with state policy adoption timing, we first characterize the overall changes in mobility patterns that accompanied COVID-19’s spread. We find evidence of dramatic nationwide declines in mobility and human encounters prior to adoption of any statewide mandates. Then, using difference-in-differences along with weighted and unweighted event study methods, we isolate the portion of those reductions directly attributable to statewide mandates. Once states adopt a mandate, we estimate further mandate-induced declines between 2.1 and 7.0 percentage points relative to pre-COVID-19 baseline levels. While residents of mandate states soon returned to prior business visitation patterns, the impacts on distances traveled and human encounter rates persisted throughout the observed mandate periods. Our estimates of early mobility reductions and the responses to statewide stay-at-home policies convey important policy implications for the persistence of mobility behavior changes and states’ future re-openings.

To Molly
for today, and every day

To Emily
for filling our lives with so much joy (and for ensuring the brevity of this dissertation)

Contents

Contents	ii
List of Figures	iv
List of Tables	v
1 Culpable Consumption: Residential Response to Price and Non-Price Water Conservation Measures	1
1.1 Introduction	1
1.2 Background	5
1.3 Data	8
1.4 Empirical Framework	12
1.5 Results	18
1.6 Conclusion	28
2 Food for Thought: Food Banks and the SNAP Benefit Cycle	48
2.1 Introduction	48
2.2 Background	50
2.3 The USDA Food Acquisition and Purchasing Survey (FoodAPS)	56
2.4 Food Bank Utilization and Food-At-Home Expenditures, Visualized	60
2.5 Empirical Strategy	62
2.6 Results	63
2.7 Conclusion	67
3 Are we #Stayinghome to Flatten the Curve?	86
3.1 Introduction	86
3.2 Data	89
3.3 Empirical Strategy	92
3.4 Results	95
3.5 Conclusion	100
Bibliography	109

A Appendix to Chapter 1

List of Figures

1.1	East Bay Municipal Utility District (EBMUD) Service Area	31
1.2	East Bay Municipal Utility District (EBMUD) Excessive Use Penalty Bill	32
1.3	Excessive Use Penalty Ordinance (EUP) Timeline	33
1.4	Public Shaming of Excessive Water Users	34
1.5	Variation in Tier 1 Marginal Prices	35
1.6	Average Water Use Over Time by EUP Violation Status	36
1.7	Elevation Band Threshold Example	37
1.8	Price Elasticity Heterogeneity by Billing Cycle Weather	38
1.9	Moral Suasion and Public Shame Event Studies	39
1.10	First-Year Effects of Moral Suasion	40
1.11	First-Year Effects of Public Shame	41
2.1	Food Pantry Use for Food-insecure Households in 2017	69
2.2	CalFresh and CFAP Participation	70
2.3	FoodAPS Primary Respondent Workbook, Food-at-Home Entry	71
2.4	FoodAPS Primary Respondent Workbook, Food-at-Home Source Barcodes	72
2.5	Variation in FoodAPS Instrument	73
2.6	Food Bank Utilization by SNAP Participation	74
2.7	Daily Food-at-Home Expenditures	75
2.8	Daily Food-at-Home Expenditures by Food Bank Use	76
3.1	Changes in Travel Activity and Social Distancing	103
3.2	Unweighted and Weighted Event Studies for Changes in Average Distances Traveled	104
3.3	Unweighted and Weighted Event Studies for Changes in Visits to Non-Essential Businesses	105
3.4	Unweighted and Weighted Event Studies for Changes in the Unique Human En- counter Rate	106
A.1	East Bay Municipal Utility District (EBMUD) Warning Letter	119
A.2	Average and Marginal Water Prices over Time.	120
A.3	Price Elasticity Heterogeneity by Household Water Usage	121

List of Tables

1.1	California Water District Policies During Drought Crisis	42
1.2	Summary Statistics, Billing Cycle Characteristics	43
1.3	Pre-Drought Policy Water Demand Estimates	44
1.4	Pre-Drought Policy Price Elasticity by Season and Location	45
1.5	Full Period Price Elasticity Estimates	46
1.6	Responses to Behavioral Policies and Comparable Price Changes	47
2.1	Evidence of the SNAP Benefit Cycle	77
2.2	Evidence of the SNAP Benefit Cycle in FoodAPS	78
2.3	Household Characteristics by Food Bank Use	79
2.4	Characteristics of Food-at-Home Acquisition Events	80
2.5	Characteristics by Survey-Day	81
2.6	Food-at-Home Expenditures and the SNAP Cycle	82
2.7	Food-at-Home Expenditures, the SNAP Cycle, and Food Bank Use	83
2.8	The SNAP Cycle and Food Banks, Alternate Timing	84
2.9	The SNAP Cycle and Food Banks, Payment Types	85
3.1	Summary Statistics on Travel Behavior and Social Distancing	107
3.2	Statewide Stay-at-Home Mandates, Travel Activity, and Social Distancing	108
A.1	Residential Water Price Schedules over Time	122
A.2	Drought Surcharges, Fiscal Year 2016	122
A.3	EBMUD Elevation Bands	123
A.4	Comparison Across Elevation Band Cutoffs, Alameda County	124

Acknowledgments

I would like to extend my sincere thanks to my dissertation chair, Sofia Villas-Boas, for years of encouragement and unwavering support. I would not be the researcher I am today without you as a collaborator and mentor. I thank Hilary Hoynes for her inspiration and guidance on conducting credible social safety net research. I thank Jim Sallee for his invaluable advising and for taking the time to work through both the nitty-gritty of the finest empirical details and the broadest details to find the compelling narrative. I thank Max Auffhammer for his support from my very first day on campus and for providing me the knowledge and guidance to develop into the instructor I have become. I thank Mark Anderson for his continued mentorship and willingness to help take my work to the next level. I thank Molly Sears for helping solve all of my empirical and theoretical dilemmas while on dog walks and supporting me throughout this entire process.

I would also like to express my gratitude to Katherine Wagner, Casey Wichman, Bryan Pratt, and seminar participants at UC Berkeley, Camp Resources, OSWEET, GARESC, and the AAEA Annual Meetings for helpful suggestions and comments. I also thank Carmen and Diana, without whom I would have surely failed to navigate the many logistical hurdles thrown our way.

Finally, I would like to thank my friends and family for their constant love and support, and for the many hours of childcare – without which this dissertation would not have been possible.

Chapter 1

Culpable Consumption: Residential Response to Price and Non-Price Water Conservation Measures

1.1 Introduction

Climate change and rapid urbanization have established urban water security as a critical concern. Globally, more than 2 billion urban residents of developed and developing countries are projected to reside in water-scarce areas by 2050 (He et al. 2021). In the western United States, extensive reliance on outdoor irrigation places tremendous strain on urban water systems. Years of meager snowpack levels and extreme summer weather situate urban water districts in the difficult position of having to balance reduced long-run supply expectations with rising costs, higher summer demand, and increasingly frequent, severe drought.

In order to manage increased scarcity and growing demand, urban water districts around the globe require extensive knowledge of the policy approaches that can effectively induce water conservation over multiple time horizons. As climate change leads to chronically hotter and drier conditions, water districts must utilize policies that can shift residential consumers to lower long-run consumption paths and balance uncertain supplies (Buck et al. 2020). At the same time, climate change exacerbates short-run climate volatility, increasing both the frequency and severity of droughts (*IPCC* 2021). During these periods of temporary heightened scarcity, districts require policies that can quickly induce rapid, large-scale reductions in water use. By selecting policy approaches appropriate for inducing either long-term adaptive behavior or short-run consumption reduction, water districts can improve the stability and reliability of increasingly-volatile water supplies.

Although policymakers can choose from a wide set of available price and non-price instruments, which policy regimes are the most efficient for achieving water conservation goals remains an area of great debate. Price-based approaches come first to economists' minds for efficiency and cost-effectiveness reasons, but they may lose efficacy as households shift

consumption to more essential indoor uses or may be constrained by government regulations. In recent years, interest has grown in the use of behavioral, non-price methods such as moral suasion and public shame. These information and social pressure-based policies show considerable promise to both reach agents unresponsive to price-based approaches (Johnson 2020) and affect change at dramatically lower cost in comparison to more traditional demand-side management strategies (Kahan and Posner 1999). While behavioral change due to moral suasion, peer comparison, and social pressure has been well-studied (Andreoni 1990; Allcott 2011; Dal Bó and Dal Bó 2014; Brent, Cook, and Olsen 2015; Nemati, Buck, and Soldati 2017), relatively little is known about the potential role of public shame in inducing conservation behavior. This lack of knowledge is driven by the fact that very few existing policies implement public shaming, let alone introduce it quasi-randomly or concurrently with other policy instruments.

In this paper, I take advantage of unique quasi-experimental variation in contemporaneous exposure to price changes, moral suasion, and public shaming in San Francisco’s East Bay Area to explore the relative effectiveness of price versus non-price approaches for short and long-run water conservation. I begin my analyses by estimating credibly causal price elasticities before, during, and after drought policies implemented in the face of historic scarcity. I first estimate long-run elasticities pre and post-drought, taking advantage of marginal price schedules that vary both over time and across difficult-to-observe spatial boundaries. I then leverage the implementation of an additional marginal price tier to identify elasticities for high usage households in response to short-run, policy-induced price changes.

Next, I identify the dynamic conservation effects of moral suasion and public shame separately from price changes during and after drought policies. During the height of the drought crisis, the water district implemented a bundled policy that resulted in quasi-random exposure to fees, moral appeals to conserve, and negative coverage by name in news stories after customers exceeded an arbitrary usage threshold. By nesting event studies for both behavioral instruments into the residential water demand model, I am able to estimate how treated households changed their consumption behavior following direct exposure to public shaming separately from moral suasion or financial penalties. Under parallel trends and treatment effect stationarity assumptions, I isolate the causal effects of each behavioral instrument for inducing conservation among top water users.

Administrative panel data for roughly 300,000 single family homes over eight years in San Francisco’s East Bay Area provide rich variation in water use and prices over time. As the water district charges different marginal prices across three elevation boundaries, households in the same neighborhoods experience different nonlinear pricing schedules in every period. The obscurity of the boundaries’ locations minimizes concerns of sorting, and I confirm balance across boundaries through matching on observable characteristics. Differential price changes over time across consumption tiers and elevation bands contribute a further source of identifying variation. These considerable differences in prices for comparable homes across space and time allow use of an empirical strategy (Blundell, Duncan, and Meghir 1998; Saez, Slemrod, and Giertz 2012; Ito 2014) that isolates causal price responses from mean reversion and distributional shifts – sources of bias that affect typical water price instruments.

Focusing on a water district spanning urban and suburban communities in varying micro-climates provides a study region with extensive heterogeneity in socioeconomic and weather conditions. Using data from 2012-2019, I am able to observe how households responded to both shifting long-run climate and short-run acute scarcity. Differences in outdoor irrigation needs for coastal and inland residents with identical landscaping provides a plausibly exogenous source of variation in the probability of exceeding the usage threshold in a given period. Empirical models provide estimates of behavior in the context of extensive urban outdoor irrigation needs in the world's fifth largest economy.

Price response results yield evidence of large magnitude pre-emergency price elasticities that decrease in magnitude as the drought crisis grew more severe. Prior to the implementation of the district's main drought policies, I estimate unitary price elasticities that are consistent with recent literature and are highly robust to the choice of price measure and sample. Once drought policies are implemented, these elasticities shrink in magnitude to -0.4 for all households and -0.7 for households that ever exceeded the usage threshold. I find no evidence of conservation responses to prices for households that violated the policy and were exposed to moral suasion or public shame, with large, positive price elasticities obtained for these top users while exposed to non-price measures during the policy period. Once the drought emergency ended, price elasticities fall further to -0.1 – -0.2 overall and -0.5 for high volume users. These findings reflect the many available margins of adjustment when outdoor water use is high and the inflexibility of indoor water use following adoption of conservation practices.

Examining the conservation responses to moral suasion and public shame, I estimate large magnitude short-run responses that vanish once the drought emergency is lifted and consumer beliefs in crisis fade. Moral suasion exposure results in a 10-15% reduction in water usage relative to control households for preferred specifications. Being publicly shamed prompts a 2-3 times larger conservation effect that is *in addition to* prior exposure to moral appeals and fees. These findings are highly robust across specifications in the presence of never-treated units and provide evidence of the resilience of public shaming's conservation effect to crowding out by other concurrent policies. Once the district ended its drought emergency roughly one year later, the conservation responses to both behavioral instruments quickly vanished, suggesting that the behavioral responses are not persistent long-term.

These results carry important implications for the future design of short and long-run water conservation policies. The large-scale responses to moral suasion and public shame under short-run drought emergency provide strong evidence that these behavioral policy instruments can be effective tools for achieving discrete reductions during periods of acute shortage. In contrast, the calculated price increases necessary to match either non-price policy response in the short-run are infeasibly large, matching prior evidence on small-magnitude responses to excessive use fees (Pratt, in press). The relative ineffectiveness of prices during temporary scarcity is unsurprising given the degree of observed adaptation behavior, with households already cutting non-essential water usage and shifting to less elastic regions of the demand curve prior to the implementation of drought emergencies. Further, the estimated positive price elasticities for households while exposed to behavioral

instruments suggests that these non-price mechanisms dominated water users' focus during the drought crisis.

In the long-run, prices remain effective at reaching high-usage households while moral suasion and public shame prove highly sensitive to messaging. Although average households display post-period price elasticities under one-fifth the size of those pre-drought emergency, top users remain relatively more elastic. The rapid disappearance of behavioral policy effects once the signaling of crisis ends mirrors prior evidence of habituation to moral suasion (Ferraro, Miranda, and Price 2011; Ito, Ida, and Tanaka 2018) suggests that conservation responses to behavioral instruments are unlikely to persist in the long-run when individuals no longer believe that such behavior is morally required to ensure availability for essential needs.

My findings contribute to three strands of literature. First, I contribute to the literature on causally-identified water price elasticities. While most prior studies are limited by a lack of quasi-experimental variation in prices, several recent studies have leveraged price variation due to seasonal changes in marginal prices (Klaiber et al. 2014; Yoo et al. 2014) or changes to the pricing structure itself (Nataraj and Hanemann 2011; Wichman 2014). I extend on these works by utilizing an empirical strategy that isolates price responses from mean reversion and distributional shifts – sources of bias common in prior urban water demand studies. My study setting provides unprecedented degrees of price changes and follows the largest-yet set of single family homes through multiple stages of water scarcity and drought policy, allowing use of a rigorous control approach not possible in previous study settings.

Next, I contribute to two distinct literatures on residential responses to urban drought policies and on general responses to behavioral policies, respectively. While a limited set of prior studies in the first literature have examined settings that simultaneously employ price and non-price drought policy instruments (Asci and Borisova 2014; Wichman, Taylor, and von Haefen 2016), this paper is the first to not only identify the water conservation impact of public shame, but also isolate its effect from those due to moral suasion and price-based measures. Prior evidence on behavioral responses to public shaming from the second literature are limited to firms' standards violations (Blackman, Afsah, and Ratananda 2004; Schlenker and Scorse 2011; Johnson 2020) or income tax avoidance (Perez-Truglia and Troiano 2014; Dwenger and Treber 2018; Tsikas 2021); this paper presents the first evidence of individual responses in the context of natural resource conservation.

This paper proceeds as follows. Section 1.2 provides background on California's historic 2011-2017 drought and provides detail on the focus district's policies. Section 1.3 introduces the data sources and discusses water use trends over time. Section 1.4 develops the residential water demand empirical model and discusses the identification strategies for obtaining price elasticities and behavioral responses to drought policies. Section 1.5 presents results, and Section 1.6 concludes.

1.2 Background

2011-2017 California Drought and Policy Landscape

From 2011 to 2017, California experienced historic water shortages. 2012-2014 marked the driest and hottest three-year period yet on record for the world’s fifth-largest economy (Mann and Gleick 2015). Although Governor Jerry Brown proclaimed a state of drought emergency in January 2014, statewide snowpack water content stood at only 5% of typical levels by April 1, 2015 (California Department of Water Resources 2015). As a result of this depleted snowpack and reduced snow melt – a primary source of California’s potable water supply – many urban water districts around the state soon faced critically low reservoir storage levels well below historic averages.

To combat impending water shortfalls, the state took unprecedented action to curtail urban water demand. In April 2015, Governor Brown issued Executive Order B-29-15 which mandated a 25% statewide reduction in urban water deliveries relative to 2013 levels.¹ In addition, the order directed water districts “to develop rate structures and other pricing mechanisms, including but not limited to surcharges, fees, and penalties, to maximize water conservation consistent with statewide water restrictions” (California Executive Order B-29-15 2015).

Water districts across the state now faced both an immediate need to develop conservation policies to bring about the necessary reductions and a newfound flexibility in the policy instruments available to bring about those reductions. Historically, water districts in California have been limited in their ability to set prices beyond levels that recover costs of provision (CA Proposition 218). While districts may include fixed charges or implement increasing block pricing schedules, their rate structures must be approved and ultimately set by the California Public Utilities Commission to ensure that they are not operating as a natural monopoly. However, as a result of state drought emergencies and executive orders, water districts could now design price-base policies for the specific purpose of promoting conservation behavior. Further, districts were free to implement a broad set of non-price policies instead of or in combination with price-based measures.

As a result of this newfound policy freedom, water districts around the state adopted a diverse range of price and non-price based policies to meet state-mandated reductions. Table 1.1 summarises the policy instruments employed by a subset of urban water districts across California. Many water districts – including El Dorado, El Centro, Imperial, and Placer water districts – designed online or mail education campaigns that taught residential customers about ways to reduce indoor and outdoor water usage. Districts often implemented rebate programs, providing financial incentives for customers to adopt water-efficient outdoor irrigation technology (i.e. smart irrigation controllers or in-line drip systems), low-flow

¹While the state sought a 25% overall reduction in urban water deliveries, individual districts faced potentially smaller or larger mandated cuts. For example, while districts whose residents averaged under 65 gallons per capita-day in summer 2014 faced necessary reductions of only 8%, districts with average usage above 215 gallons per capita-day were subject to reduce deliveries by 36%.

showerheads or toilets, or even replace lawns with drought-resistant landscaping. Common price-based policy instruments include the introduction of drought surcharges or replacing flat fee structures with increasing block pricing. In many cases, water districts employed a combination of listed policy instruments, launching mail or web-based information campaigns, introducing rebate programs for high-efficiency landscaping equipment, or modifying pricing structures to reduce demand overall or among top water users.

A common theme among water districts was the targeting of non-essential water use among high-usage customers. While the lion's share of California households utilize moderate amounts of water every month, a disproportionate share of water is utilized by consumers in the top few percent of the water use distribution. As a result, many districts identified top users as the richest source of urban water conservation. The least-invasive policies included contacting outlier users – either over the phone or by placing physical placards on their front doors. Several districts, including Apple Valley Ranchos, San Jose, and Santa Cruz, introduced excessive use fees that increased the marginal price of water for volumes used beyond a set threshold.

East Bay Municipal Utility District and the Excessive Use Penalty Ordinance

While several other California water districts targeted excessive users in various ways, the most unique policy solution was introduced by the East Bay Municipal Utility District (henceforth EBMUD). Located in San Francisco's East Bay Area, EBMUD delivers over 150 million gallons daily to 1.4 million water customers in Alameda and Contra Costa counties. Figure 1.1 provides a map of the EBMUD service area and shows its location within California. The EBMUD service region spans 332 square miles, bounded to the west by the Richmond and Bay Bridges and to the north by the Benicia Bridge. The region extends as far south as San Lorenzo and covers communities along the Highway 24 and 680 corridors through to San Ramon.

EBMUD serves a highly socio-economically and climatically diverse customer base due to its extensive geographic coverage. EBMUD delivers nearly two-thirds of its water to residential customers from urban downtown Oakland to wealthy suburban communities high atop the Berkeley hills or on large lots nestled in the foothills of Mount Diablo to the east. As the district is bisected by a portion of the Pacific Coast mountain range, residents to the west experience cooler, moister conditions relative to customers in the eastern foothills. These factors lead to substantial heterogeneity in usage across communities: while the average residential EBMUD customer consumed 73 gallons per capita-day in 2014, residents of Oakland averaged a mere 57 gallons per capita-day while residents of the affluent Diablo suburb averaged 345 gallons.

As its primary means of meeting their state-mandated 20% delivery reductions, the district designed and adopted the Excessive Water Use Penalty Ordinance (henceforth EUP). While EBMUD implemented drought surcharges for all customers, the district identified

“discretionary, nonessential use” by single-family homes as the greatest source of potential savings. On average, over 50% of annual residential deliveries in the district goes to irrigate lawns or other outdoor landscaping – uses not deemed essential for the common good by state legislation. Further, the distribution of historical water use in the district is top-heavy: from 2012 to 2019, 36% of residential water was consumed by the top 10% of bills.

To induce conservation behavior among top users, the EUP introduced a mix of three price and non-price policy instruments. First, the policy introduced a limit of 80 hundred cubic feet (80 CCF, or 59,840 gallons) of water per two-month billing cycle for all single-family residential water customers, regardless of location or past usage. Households that exceeded this limit faced a \$2 penalty in addition to the typical marginal price for each CCF of water used above the threshold. This price mechanism functioned akin to an additional increasing block price tier, raising marginal prices by 36-45% for only the highest usage households. During the sample period, only 2% of bills exceeded the usage threshold. However, this small subset of bills represented a disproportionate share of total water deliveries, and consumed 13.6% of total residential water.

This seemingly arbitrary threshold value was chosen for two main reasons. First, it represented a level well beyond that of a typical residential customer. On average, a single-family home in the EBMUD service area consumed approximately 250 gallons per day (Cuff and Rogers 2015).² The 80 CCF threshold meant that a household would need to consume over four times the average level in order to violate. Second, the threshold had a clear interpretation in terms of daily flow rates: a water volume of 80 CCF equals approximately 1,000 gallons per day for a standard billing cycle.

The violation messaging and framing of penalties resulted in moral suasion exposure for all violating households. During the drought emergency, EBMUD sent letters to all high-usage households urging them to check for leaks and informing them that exceeding the threshold “is prohibited and has consequences.”³ District messaging referred to the fees as a “civil administrative penalty” for violating an ordinance needed to ensure “sufficient water for human consumption, sanitation, and fire protection.” Indeed, the EUP legislation itself made it unlawful for EBMUD customers to “willfully violate any provisions” of the policy. The moral violation messaging was also reflected on violating customers’ bills. Figure 1.2 provides a copy of an EBMUD bill for a household that violated the EUP and shows how the penalty amount was provided as a line item separate from typical usage volumes, elevation charges, and broad drought surcharges. This aspect of the policy mirrored broader messaging from the district that conveyed a moral need to eliminate nonessential irrigation and modify indoor water use behavior during drought.

²Note that substantial heterogeneity in water use exists between the district’s western cities and drier, affluent eastern suburbs. In 2014 residents in Berkeley and Oakland used 52 and 57 gallons per person per day, respectively, with household averages under 200 gallons per day. Residents of Alamo and Danville – two eastern towns with large homes and expansive lots – averaged 250 and 345 gallons per person per day respectively during the same period, with household averages approaching or exceeding the excessive use level of 1,000 gallons per day.

³See Appendix Figure A.1 for a copy of the full letter.

Additionally, the EUP resulted in public shaming of a subset of top water users who violated the policy. While no charges were ever pressed, violation of the EUP technically resulted in a misdemeanor offense for all customers who exceeded the usage threshold. Due to the particular language of California’s Public Records Act, EBMUD was legally required to make available the names, addresses, and water usage of all violating customers if ever requested through a Public Records Act request. As a result, many regional and local newspapers and news stations placed records requests with the districts and obtained the identities of all violators. While EBMUD’s warning letter to high-usage households in the months preceding the EUP communicated the possibility of this information being released, it did not make the connection to the potential of news media coverage singling out individual water users by name.

Figure 1.3 works through the specific timing of the Excessive Use Penalty Ordinance. The EUP entered into effect for all residential bills beginning July 1, 2015 or later. On October 15, EBMUD created the first Excessive Use Report that contained information on all households that had exceeded the usage threshold by that time. Over the next eight months, EBMUD created eight additional reports, the last of which was created several weeks after the EUP ended and EBMUD lifted the district-wide drought emergency due to increased reservoir levels and a sufficiently wet winter. Soon after each of the reports was created, regional news outlets obtained the lists and began publishing news articles that called individual violators out by name. These articles and television news segments often referred to violators as water “hogs,” “wasters,” or “guzzlers” and frequently lambasted their excessive water use while mentioning their neighborhood of residence or address itself.

Publicly shamed violators largely fell into one of two camps. The first type of article focused on “locally relevant” violators: retired or current professional athletes, musicians, business executives, as well as prominent university faculty and even news anchors themselves. Newspapers published stories about these individuals not because their usage was particularly different from other violators, but because their names were known and could likely generate page views. The second type consisted of “high-shock” users: households who experienced a particularly large water shock relative to other violators in a given period. These households often had typical water use levels in line with other violators, but happened to have abnormally high use in one billing cycle that pushed them either to the top of a particular report or to the top for a given city. Figures 1.4a-1.4b provide examples of each type of article.

1.3 Data

To investigate the price responsiveness of EBMUD customers over time and top users’ responses to moral suasion and public shame, I obtain administrative panel data directly from the water district. Water use data cover the universe of billing cycles for residential customers during the period of April 1, 2012 through May 31, 2019. Each observation includes the billing cycle start and end dates, water volume consumed, and location information.

Water usage is provided in hundred cubic feet units (CCF), which I convert to gallons per day (GPD) to account for slight differences in billing cycle lengths and to match messaging on consumer bills.⁴ The overwhelming majority of EBMUD customers are billed every two months, for a total of six billing cycles per year.⁵

I obtain pricing schedules for each fiscal year from the district website. All residential EBMUD customers pay a common fixed charge in each billing cycle. Residential customers also face increasing block-tier pricing for all units of water consumed across three usage tiers. Tier 1 marginal prices are charged for all CCF of water consumed between 0 and 14 (0-172 GPD), which is set to “the average indoor water year-round usage for single-family residential customers” and represents permanent demand across all seasons. Tier 2 (15-32 CCF/172-393 GPD) and tier 3 (over 32 CCF/393 GPD) marginal prices are set progressively higher to recoup greater portions of the costs associated with fulfilling water demand during peak periods. Each year, planned rate increases to fixed charges and marginal prices in each tier come into effect at the start of the district’s fiscal year on July 1 to account for changes in costs. Using this information I compute marginal and average prices for every billing cycle based on each customers’ consumed water volume. Throughout analyses, I deflate all prices relative to 2012 dollars using the Gross Domestic Product Implicit Price Deflator (U.S. Bureau of Economic Analysis 2021).

As a result of variable price changes each year across the marginal price schedule, drought surcharges, and elevation fees, EBMUD’s price structure provides unprecedented variation over space and time. Figure 1.5 illustrates this variation, plotting the marginal prices over time of consuming a median volume of water (in marginal price tier 1) for households in each of the three elevation bands. The differences in marginal prices across elevation bands reflect per-unit charges that recoup the costs of pumping water to homes in higher pressure zones (roughly 200-600 feet above sea level for band 2 and above 600 feet for band 3).⁶ While marginal prices within a given elevation band change each year, the relative growth between bands varies over time with the gap growing on average throughout the sample period. Further, at any given point in time, households consuming identical volumes face different marginal prices due solely to differences in pumping costs. Identical patterns hold for consumption in tiers 2 and 3, with marginal prices increasing by 83% and 97%, respectively, during the sample period. Further, the addition of per-unit drought surcharges during the height of drought conditions resulted in a period of large price increases followed by decreases to both marginal and average prices.⁷ This variation provides an ideal setting in which to

⁴1 CCF, or unit, equals 748.052 gallons.

⁵As meters are read manually, the number of days in a billing cycle will be slightly above or below an exact two month interval. Monthly billing occurs for a small percentage of cases that include “reasonable and justifiable customer requests” and customers for whom the average monthly bill exceeds \$1,500 (typically commercial customers), or in cases of past credit/collection issues.

⁶See Section 1.4 for a more detailed discussion of a home’s assignment into elevation bands and the resulting identifying price variation. Appendix Table A.1 also reports the exact elevation surcharge amounts in each fiscal year.

⁷See Appendix Table A.2 for exact drought surcharge values. EBMUD enacted Stage 4 (critical) drought in fiscal year 2016. While proposed, surcharges for Stages 1-3 were never implemented.

investigate price responsiveness of urban water users.

Cases of excessive water usage during the drought policy period are obtained under a public records request with EBMUD. These data consist of all instances of household water usage that exceeded the 1,000 gallon per day threshold while the Excessive Use Penalty was in effect. Each entry includes the name, address, city, water usage (in CCF and in gallons per day), statement date, and cycle start/end dates for each violation. These data are a corrected version of those previously obtained by regional news media, omitting instances revised due to leaks or other issues. After merging with the main water use sample, 6,403 violations by 4,951 households remain.

I obtain incidences of public shaming by news media using the NewsBank repository and direct searches of regional newspapers and TV stations. NewsBank indexes articles from 189 regional and statewide print and web-only news sources in California. I performed a targeted search for references to East Bay Municipal Utility District, EBMUD, excessive use, water guzzler, and water hog. Title and search term matches for the 744 returned articles were then manually inspected, noting down any by-name mention of Excessive Use Penalty violators. I perform an identical search procedure for several Bay Area newspapers not available through NewsBank (including the Daily Californian and the Alameda Sun) and all stories (print and video) from the local TV news stations NBC Bay Area, CBS KPIX, Fox KTVU 2, KRON 4, and ABC 7. These two search procedures yield 51 articles with names that match directly to instances of EUP violation. In total, 61 households are called out by name across 202 unique violator/article combinations.

I use PRISM Climate Group’s AN81d gridded daily weather dataset to construct climatic variables (PRISM Climate Group). PRISM inputs weather station data into a climate model that allows realized weather to vary with elevation, coastal atmospheres, and terrain barriers – features that are all present in the San Francisco East Bay – and outputs temperature and precipitation in 4 kilometer square grid cells (PRISM Climate Group). Using these data, I construct household by billing cycle-level measures of experienced rainfall and precipitation that precisely match weather conditions that occurred between a particular billing cycle’s start and end dates. Covered grid cells are obtained by matching customers’ addresses to parcel shape files obtained from the Alameda and Contra Costa County Assessor property files. I further use the property files to limit the sample to only detached, single family homes.

Table 1.2 summarizes billing cycle characteristics and excessive use during the sample period for the 10,709,816 observations. The first columns report the means and standard deviations across variables for all 301,271 unique households in the sample. The next columns report summary statistics just for the 4,951 “Excessive Users” who violated the EUP by using over 1,000 gallons per day while the policy was in effect. During the sample period the average EBMUD household consumed 18 CCF of water in sixty days, or just over 225 gallons per day. However, the large standard deviation of 280 gallons per day signals considerable heterogeneity, due in part to the presence of high usage households.⁸ On average, customers

⁸Although only 1.7% of observations used over 1,000 gallons per day (surpassing the Excessive Use

face a marginal price of \$3.76 per CCF, less than half that of the typical average price of \$7.59. Including service charges, the average customer paid a total water bill of \$104.67. The EBMUD service area is topographically varied; roughly 36% of billing cycles during the sample period stem from homes in elevation band 2 (approximately 200-600 feet above sea level) and an additional 9.5% from homes in band 3 (roughly more than 600ft). The average single family home in the sample sits on a lot roughly 0.15 acres in size and has a home over 1,800 square feet in size. The mean billing cycle experienced a total of 312 growing degree days accumulated beyond 10°C and received 95.7 millimeters of rain (with considerable variation throughout the year).

Focusing on households that violated the EUP shows the dramatic amounts of water consumed by these users. These high-usage households averaged just over 1,000 gallons per day during the sample period, with water use in 49% of their billing cycles exceeding the EUP usage threshold. Due to this high water use, their average water bill reached a staggering \$502 dollars. These customers are more frequently located at higher elevation bands, reflecting the concentration of affluent suburbs around the Berkeley hills or in the foothills of Mount Diablo. Home and lot sizes reflect the relatively high affluence of violating households, with EUP violations occurring on lots that average over half an acre in size with homes in excess of 3,800 square feet. Broad weather patterns experienced by high-usage households are not systematically different from those seen by all households, with top users also accumulating just over 300 growing degree days and 100 millimeters of precipitation per average billing cycle.

Trends in Water Use

Observing district-wide patterns in water use provides preliminary evidence of extensive conservation prior and in response to drought policies. Figure 1.6 plots the average water use measured in gallons per day for all single family homes in the EBMUD service area from April 2012 through May 2019. The dotted line reports the average usage for customers that violated the EUP while the policy was in effect and were exposed to fees, moral suasion, and potentially public shame (“EUs”). The solid line reports mean water use for all other detached single family home EBMUD customers that were not directly exposed to EUP policy instruments (“Control Households”).⁹

Three main patterns arise from the data. First, consumption by top users is both much higher across the sample period and more variable across seasons. While high-usage households consumed roughly double the average of all other residential customers during winter months, their summer consumption is often four to five times larger. A typical single family home consumed 200 – 315 gallons on average per day during observed summers, whereas households that ultimately violated the EUP averaged from 1220 to over 1800 gallons per day. Second, both groups of customers exhibit conservation trends prior to the height of

Penalty Ordinance’s limit), the total cases of “Excessive” use exceed 200,000 instances. Over 9% of total observations come from households that ever exceed this limit.

⁹See Appendix Figure A.2 for trends in average and marginal prices during the sample period.

the drought crisis. Both annual average and peak summer usage declined for each group from 2012 to 2014, providing initial evidence that households undertook adaptive behavior in response to local and state messaging prior to the implementation of district policies during the height of drought crisis. Third, usage further declines for both groups while the EUP was in effect (indicated by the orange band). Peak consumption during summer 2015 greatly attenuated compared to prior years for both groups, with winter consumption also lower than in 2012-2014. These trends provide preliminary evidence that even while households adapted to changing water supplies in the early drought years, district-wide policies may have induced further conservation.

1.4 Empirical Framework

I next develop an empirical model of residential water demand to investigate how price elasticities evolved under these usage trends and to measure responses to price and non-price policy instruments employed by the EUP. Identification of causal price elasticities in this framework relies on two unique aspects of EBMUD’s water pricing schedules. First, there is cross-sectional variation in marginal prices that is independent of both historical and contemporaneous household water usage. Second, the proportional changes in marginal prices year-to-year differ between groups, providing differences in relative prices faced by each group over time.

Residential Water Demand Model

To obtain estimates of the long-run price elasticity for residential water, I employ differenced models of the form

$$\Delta \ln(Q_{it}) = \beta \Delta \ln(AP_{i,t-1}) + f(Q_{it-3}) + g(W_{it}) + \delta_{zt} + \epsilon_{it} \quad (1.1)$$

where $\ln(Q_{it})$ is the natural log of the gallons of water consumed per day by household i during billing cycle t , $\ln(AP_{i,t-1})$ is the natural log of average prices paid in the previous cycle (two months ago), $f(Q_{it-3})$ is a function of consumption three billing cycles (six months) prior, $g(W_{it})$ is a function of experienced weather, and δ_{zt} are zip code by billing cycle fixed effects.¹⁰ Rather than a one-period difference, I employ a year-long difference: $\Delta \ln(Q_{it}) = \ln(Q_{it}) - \ln(Q_{it-6})$. In this way the idiosyncratic error term $\epsilon_{it} = \mu_{it} - \mu_{it-6}$ can be thought of as the difference in demand shocks between the current period and the same billing cycle in the previous year.

¹⁰While EBMUD technically bills in terms of water volume measured in hundred cubic feet, the price schedules posted online are communicated in flow rates and the “consumption information” section of the monthly bill is provided in terms of gallons used per day. Given that billing cycles can vary in length by several days, my preferred specifications utilize water flow rates in the dependent variable. Results of preferred specifications are robust to the use of water volume in hundred cubic feet.

Considerable attention has been paid of late to *which* price customers respond to; recently consumers have been found to primarily respond to average price, both in terms of household energy consumption (Ito 2014) and residential water behavior (Wichman 2014; Wichman, Taylor, and von Haefen 2016). While I utilize average prices for my preferred specifications in Table 1.3 and later analyses, I also report results utilizing marginal prices that are statistically indistinguishable.¹¹

I employ a simulated instrument approach that targets concerns of both simultaneity and mean reversion (Ito 2014). Typical residential water price instruments seek to characterize the entire price schedule and include service charges in addition to marginal prices at as wide a range of consumption levels as possible (Olmstead 2009; Wichman, Taylor, and von Haefen 2016). While this approach targets the endogeneity of prices, it does not address the possibility of mean reversion – a substantial concern that is further heightened when analyzing policies that target high water usage. I specify the simulated instrument as

$$\Delta \ln(\widehat{AP})_{it} = \ln(AP_t(Q_{it-3})) - \ln(AP_{t-6}(Q_{it-3})) \quad (1.2)$$

where the simulated change in log average price $\Delta \ln(\widehat{AP})_{it}$ is calculated as the difference of two functions of past consumption. The first term, $\ln(AP_t(Q_{it-3}))$, is the log of current average price for a usage level equal to the household’s consumption three billing cycles ago (i.e. a half year prior). The second term, $\ln(AP_{t-6}(Q_{it-3}))$, is the log of average prices when consuming an identical water volume under the price schedule from the same billing period in the previous year. Use of the household’s consumption level six months prior ensures that identification of the simulated instrument stems from typical patterns in household consumption rather than instances of mean reversion (Blomquist and Selin 2010; Saez, Slemrod, and Giertz 2012) and does not introduce the correlation in water demand shocks that occurs when consumption in a fixed base period is used (Ito 2014).

Lagged consumption controls $f(Q_{it-3})$ account for structural shifts in the water demand distribution. A classic example of this issue arises from the domestic income tax schedule: households at the top 1% of the income distribution are likely to have experienced faster income growth than the rest of families in the top 10% (Saez, Slemrod, and Giertz 2012). In the residential water context, outdoor irrigation needs under climate change likely grow at a faster rate for the highest-usage households with more water-intensive landscaping relative to other high users with less irrigated land. Shifts in the underlying residential water use distribution over time will lead to divergent trends between households at different levels of water use, resulting in correlation between the residual ϵ_{it} and past consumption – and ultimately residual correlation with past prices. By including a flexible set of past consumption controls I am able to eliminate this residual correlation channel.

¹¹When both marginal and average prices are included in the same regression, individual elasticity estimates become more volatile. However, the sum of both elasticities is nearly identical to results presented in Table 1.3 for each price measure on its own – reflecting the lack of cross-sectional variation in service charges to distinguish one price measure from the other.

Importantly, this control structure also serves to minimize bias resulting from mean reversion. Households that experience a particularly large water shock in the current period are mechanically much more likely to decrease their water usage in the next period than a comparable household with a smaller shock. Failure to control for a household’s relative position in the water use distribution will then result in price elasticity estimates that conflate actual price responses with differential patterns of mean reversion.

To minimize the amount of imposed structure, my preferred models specify $f(Q_{it-3})$ as a vector of usage percentile-by-time fixed effects (Ito 2014). For each billing cycle, I calculate the usage percentile into which each household’s consumption six months prior (Q_{it-3}) fell. These percentiles are then interacted with each month-year, yielding a dummy variable for each usage percentile in every sample month.¹² As EBMUD customers on either side of elevation band boundaries face different marginal prices at all consumption tiers in any given period, I am able to include such a flexible set of past consumption controls without absorbing all identifying price variation. In Table 1.3 I also report results from models that specify $f(Q_{it-3})$ as a cubic spline with knots at each decile of lagged consumption (as with the weather controls).

To control for differences in weather both over time and across space, I include the set of temperature and precipitation controls $g(W_{it})$. As the sample period includes both wet winters and historical drought periods, substantial variation exists in the temperatures and amount of precipitation experienced by a given household year-to-year. Geographic differences in climate and elevation contribute to additional, cross-sectional variation in weather. My preferred specifications use semi-parametric controls for both temperature and precipitation and include cubic splines of both mean growing degree days and total precipitation with knots included at each decile of the variables’ distribution.

The set of fixed effects δ_{zt} account for remaining unobserved spatial and temporal water demand shifters. In my preferred specifications I model δ_{zt} as the interaction of month-year and nine-digit zipcode (e.g. July 2015 times 94597-0001).¹³ In so doing I control for unobservable determinants of changes to annual water demand that vary by month but are common to homes within the narrow area delineated by postal routes. I also report results that alternately specify δ_{zt} using different combinations of month of year and 5-digit zip codes. Time-invariant household characteristics are removed through differencing.

Quasi-Experimental Price Variation

The nature of the EBMUD service area and the district’s pricing schedule provide a useful source of identifying variation for water demand estimation. First, while sorting around a geographic discontinuity is a first-order concern in cross-sectional analyses, it is less of a concern in the present panel context. Second, residents do not have a choice of water

¹²This process results in a set of 3,800 dummy variables for the pre-drought policy period and 8,600 for analyses utilizing the entire sample.

¹³As the sample includes observations across 54,591 unique nine-digit zipcode, this control approach adds an additional vector of 646,124 fixed effects that absorbs nearly all variation in experienced weather.

provider; once a resident locates within the EBMUD service area they must use the district as their water provider. While residents could theoretically relocate to lower elevation bands within the EBMUD service area to reduce or altogether avoid elevation surcharges, any such re-sorting is complicated by the opacity of the elevation band assignment process along with the financial and hassle costs of moving.

EBMUD elevation surcharges are determined based on the pressure zone in which a home is located. The EBMUD service area is divided into 130 pressure zones that are categorized based on “elements that include elevation and pressure.” Pressure zones are combined into three elevation bands, each with their own elevation surcharges designed to recoup costs associated with pumping. No surcharge exists for homes in elevation band 1, which consists of pressure zones that do not require any pumping for water delivery (the case for most homes within 200 feet of sea level). Homes in elevation band 2 are in pressure zones that require some pumping, and are typically at elevations of 200-600 feet above sea level. Homes in elevation band 2 face marginal prices that are 12-20% higher than homes in elevation band 1, with the exact surcharge varying both over time and across usage tiers. Elevation band 3 consists of homes in pressure bands that require “considerable” pumping, which is needed for most homes above 600 feet above sea level. Homes in band 3 pay between 23 and 42% more per CCF of water than homes that do not require pumping.¹⁴

Figure 1.7 shows the arbitrary nature of the district’s elevation thresholds. This neighborhood in Pinole provides a typical example of how elevation thresholds meander through neighborhoods, with homes in elevation band 1 (orange) often directly next door to or across the street from others in elevation band 2 (pink). In many cases homes that share the same block or even the same elevation face different elevation surcharges, creating substantial variation in the marginal price of water for comparable homes along the band thresholds. For neighborhoods like these around the boundaries, prospective residents would need to both view a home’s past water bill and have prior knowledge of the elevation surcharge system to understand how marginal prices may differ between seemingly-identical homes in the neighborhood. Appendix Table A.4 confirms the similarity of homes near to either side of the elevation thresholds in Alameda County, strengthening the assumption that there are likely no structural differences in observable or unobservable characteristics that impact water demand across the elevation boundaries.

The structure of excessive use fees further allows for isolating responses of top users to short-run, policy-induced price changes. Beginning in July 2015, the EUP added excessive use penalties that functioned as a new, fourth marginal price tier for top users. A high-usage household with a constant consumption level between June and August now faced drastically different costs of consuming a marginal unit of water across adjacent bills. Similarly, the policy introduced a sizable gap in marginal price between a home consuming 1,005 gallons per day and one consuming just below the threshold. As a result, this temporary marginal price tier provides a source of identifying variation for distinguishing excessive users’ short-run

¹⁴See Appendix Table A.3 for a summary of elevation band surcharges and the percent of sampled homes that fall within each band.

responses to temporary fees from annual changes to the usual marginal price schedule.

Behavioral Responses to Non-Price Policy Instruments

To identify the effects of moral suasion and public shame resulting from the Excessive Use Penalty, I extend Eq. 1.1 to allow price responses to vary prior to, during, and after the EUP and include relevant non-price policy elements to capture response dynamics:

$$\begin{aligned} \Delta \ln(Q_{it}) = & \beta_1 \Delta \ln(\widehat{AP})_{it} + \beta_2 \left(\text{EUP} \times \Delta \ln(\widehat{AP})_{it} \right) + \beta_3 \left(\text{Post-EUP} \times \Delta \ln(\widehat{AP})_{it} \right) \\ & + \sum_{k=-18}^{18} \left(\theta^k \text{Suasion}_{it}^k + \tau^k \text{Shamed}_{it}^k \right) + f(Q_{it-3}) + g(W_{it}) + \delta_{zt} + \epsilon_{it}. \end{aligned} \quad (1.3)$$

The term $\text{EUP} \times \Delta \ln(\widehat{AP})_{it-1}$ interacts the simulated average price instrument with a dummy variable equal to one for billing cycles that occurred while the EUP was in effect (July 1, 2015 to May 3, 2016) while the term $\text{Post-EUP} \times \Delta \ln(\widehat{AP})_{it-1}$ allows the average price response to differ after the district’s drought emergency ended. As Figure 1.6 illustrated, EBMUD customers substantially reduced their water use prior to the EUP’s start. If households had already adopted conservation behavior and reduced their propensity to conserve, then they would likely appear less responsive to surcharges and penalties during the drought policy period which would lead to a positive β_2 . Adoption of medium to long-run conservation practices would similarly suggest $\beta_3 > 0$.

I specify concurrent responses to both moral suasion and public shaming using event studies. Each element Suasion_{it}^k is a dummy variable, equal to one for a household k billing cycles since first violating the Excessive Use Penalty, and zero otherwise. $k = 0$ for the billing cycle when a household first violates the policy by using over 1,000 gallons per day and $k = 6$ for billing cycles one year later. As a result, each element of the post-treatment coefficient vector $\{\theta^1, \theta^2, \dots, \theta^{18}\}$ captures any differential annual change in water use and measures the isolated response to the district’s moral conservation appeals two months to three years after first exposure, while controlling for simultaneous fees and public shaming. If violating households responded to the noncompliance language present on their bills, we would expect $\theta^k < 0$ while values close to zero suggest average non-responsiveness.

Similarly, the vector of dummy variables Shamed_{it}^k models household responses to public shaming separate from price or moral suasion effects. For households called out by name in at least one related news article or video, I set Shamed_{it}^0 equal to one during the billing cycle when a household is first mentioned by name.¹⁵ Importantly, shaming coefficients estimate responses *in addition to* prior moral suasion exposure: account holders’ names were only published once they had violated the Excessive Use Penalty and after receiving a bill with moral suasion messaging.¹⁶ To the extent that residential water users experience disutility

¹⁵Findings are robust to specification of the $k = 0$ period as the first full billing cycle after public shaming.

¹⁶On average, first shaming occurs 39.9 days after committing one’s first EUP violation.

from negative news coverage, we would expect a reduction in water use and $\tau^k < 0$ for all post-exposure periods where $k \geq 0$. To ensure that dynamic treatment effects are identified separately from time trends, I bin both endpoints by setting $X_{it}^{-18} = 1 \forall k \leq -18$ and $X_{it}^{18} = 1 \forall k \geq 18$ for $X \in \{Shamed, Suasion\}$ (Schmidheiny and Siegloch 2019). Following convention, I drop the $k = -1$ bin and normalize all event-time effects relative to the billing cycles before being publicly shamed and before first Excessive Use Policy violation. To account for potential differences in pre-treatment conservation behavior among eventually treated households, I residualize $\Delta \ln(Q_{it})$ of treatment cohort-specific pre-trends (Goodman-Bacon, in press).

Identification of Responses to Behavioral Drought Policy Measures

Interpretation of $\hat{\theta}^k$ and $\hat{\tau}^k$ as unbiased estimates of the causal impacts of moral suasion and public shaming on residential water use rely on two key assumptions. First, unbiasedness relies on a parallels trend assumption: the trends in water use over time among the non-violating households (for moral suasion) and the violating but unshamed households (for public shaming) are what the trends would have been for the violating/shamed households absent the treatment(s). Second, an assumption of stationarity for the average treatment effect on the treated in each treatment cohort (households first treated in the same period) ensures recovery of causally interpretable event-time estimates (Sun and Abraham, in press). That is, after conditioning on the time since treatment and included controls, cohort-level ATTs do not systematically differ from one another. As preferred specifications of Eq. 1.3 account for time-invariant household characteristics, experienced weather, and include Zip-9 by month-year and percentile of past consumption by month-year fixed effects, potential violations of this assumption requires remaining changes in unobservable demand determinants or water shocks over time to be associated with different treatment effect paths.

$Suasion_{it}^k$ and $Shamed_{it}^k$ coefficients reflect intent to treat estimators. Despite observing which households received bills containing a moral suasion message or were featured in a local news story, I cannot guarantee whether the household was actually exposed to each behavioral policy component. Households with paperless billing enabled could use the district’s web portal to pay their water bill without ever directly viewing the bill and fail to see the moral suasion appeal. Use of auto-pay further reduces a customer’s likelihood of observing the bill’s conservation messaging. Even if a homes’ bills were handled manually they may not have been seen by the property’s residents. Compared to their lower-usage neighbors, high-usage households on large lots are more likely to employ landscape services that are responsible for managing water bills. Although media coverage of the policy was widespread and many shamed households were called out across multiple news sites or directly contacted by reporters, a household would miss the story altogether if they relied on national news providers and a friend or neighbor failed to alert them. In this way both sets of treatment effects likely represent lower bounds on the overall impact of the behavioral drought interventions.

1.5 Results

I next report results showing how price elasticities evolved over the sample period in response to changing water scarcity and to what extent high-usage households responded to price and non-price EUP policy instruments.

Pre-Drought Price Elasticities

Table 1.3 reports results for average and marginal price models from estimating Equation 1.1. These models use data covering the period April 1, 2012 through July 28, 2015, prior to the drought policy period. The top panel reports estimates using average price measures while the bottom panel reports comparable estimates where marginal price variables are specified. Column (1) provides estimates from an ordinary least squares regression in differences without the simulated average price instrument and only month-year by 9 digit-zip code fixed effects. Column (2) adds cubic splines for precipitation and growing degree days in each billing cycle. Column (3) adds the simulated price instrument but omits past consumption controls $f(Q_{it-3})$. Columns (4) and (5) add past consumption controls $f(Q_{it-3})$, with Column (4) specifying cubic splines of six months-lagged consumption with notches at each decile and Column (5) using the preferred month-year by percentile of past consumption fixed effects. Column (6) adds to the specification of Column (5) an interaction of the price instrument with an indicator for households that ever used a level of water that would have been considered “excessive” under the Excessive Use Policy (over 1,000 gallons per day). Standard errors in all columns are two-way clustered by 5-digit zip code and fiscal year to account for likely correlation patterns in water shocks (Colin and Miller 2015).

Turning first to the ordinary least squares specifications for average price in Columns (1) and (2) of the top panel, the models estimate price coefficients between -1.62 and -1.66. Once the price change variable is replaced with the simulated instrument, the coefficient estimate falls in magnitude to a nearly unit elastic value of -1.11 in Column (3). This elasticity estimate remains stable in Columns (4) and (5) when either control approach for past consumption is included, with point estimates between the cubic spline and distribution by time fixed effect approaches indistinguishable from one another. Allowing the price elasticity in Column (6) to vary for households who ever used over 1,000 gallons per day yields a negative but statistically insignificant coefficient. This finding suggests that households who had at least one high water use draw in the pre-drought policy period did not respond to price changes in a way that was statistically distinguishable from lower usage households prior to implementation of drought measures.

Price elasticity estimates are unchanged when identified using variation in marginal prices. Despite recent empirical and theoretical findings that average price is likely the more salient measure for both residential water and energy decision-making purposes (Ito 2014; Wichman 2014; Wichman, Taylor, and von Haefen 2016), neoclassical theory remains unchanged in its prediction that utility-maximizing consumers should base water use decisions on marginal prices. Columns (1) and (2) in the bottom panel of Table 1.3 yield

large, positive, and highly statistically significant coefficients that imply a more than 2 and a half percent *increase* in water use for each percent increase in price. Once the marginal price change is replaced with the corresponding simulated instrument in Column (3), the coefficient falls to -1.02 and both remains stable as past consumption controls are added in Columns (4) and (5) and proves statistically indistinguishable from all corresponding average price elasticities. Once again we observe a negative but statistically insignificant coefficient on the interaction term in Column (6), yielding no additional evidence of systematic differences in price responsiveness for households across the arbitrary 1,000 gallons per day threshold. These estimates' equivalency to the average price coefficients is unsurprising given the lack of cross-sectional variation in service charges and that the identifying variation in both variables is driven by the marginal price differences across elevation bands.

The large magnitude of these price coefficients reflects the many, varied margins of adjustment available to households over time. During periods of temporary scarcity, residential water users may choose to make more temporary, intensive-margin adjustments in response to price changes that they expect to be short-lived. As price changes become more frequent and consumer expectations around the time path of water availability adjust, water users can make a wide range of decisions to more permanently shift their water needs. Replacement of old, inefficient toilets with current low-flow versions can save roughly 13,000 gallons per year while installation of efficient showerheads and clothes washers save approximately 8 gallons per day and 30 gallons per load, respectively (EPA 2021; Flamer 2021). Households looking to reduce outdoor irrigation use are often encouraged by water districts – at times with rebate incentives – to install drip irrigation systems or self-adjusting irrigation controllers and replace lawns with drought-resistant landscaping (EBMUD 2021; (BAWSCA) 2021; SoCal WaterSmart 2021; City of Sacramento 2021; Valley Water 2021). In the long-run, homeowners with in-ground pools could empty them or take the more drastic step of removing them altogether.

Elasticity estimates from preferred specifications are in line with recent causal estimates from the residential water demand literature. Nataraj and Hanemann (2011) use a regression discontinuity design around the introduction of an additional block tier for the marginal price schedule in Santa Cruz County, California. This block affected only high-usage households who consumed 40 or more CCF per two-month cycle and was implemented shortly after the easing of price and non-price controls during the state's 1987-1992 drought (Nataraj and Hanemann 2011). Although the authors interpret their findings as a -0.12 short-run marginal price elasticity, Wichman (2014) shows that the change in average price implied by their treatment effects corresponds to an average price elasticity of -1.16 (Wichman 2014).

Klaiber et al. (2014) utilize changes in the distribution of water consumption at the census block group level over a three year period in Phoenix, Arizona to estimate the price responses resulting from seasonal changes in marginal prices. They estimate a wide range of seasonal elasticities by water usage percentiles and obtain estimates from an annually-differenced model ranging between -1.93 and -1.53 for winter usage and -0.99 to -0.45 during summer months for 2000-2003 (Klaiber et al. 2014). This approach relies on interpretation of the intercept as a local approximation to the full price response and is potentially biased

if the intercept term is not stable to alternate control specifications (Yoo et al. 2014). Yoo et al. (2014) employ a similar differenced regression that includes marginal price and obtain a long-run estimate of -1.16 over the period of 2000-2008. While bias arising from instability of the intercept is eliminated (Yoo et al. 2014), the city’s pricing schedule does not exhibit cross-sectional variation in marginal prices, preventing the authors from including controls for past consumption as utilized here that can account for shifts in the underlying water use distribution over time.

Pre-Drought Policy Heterogeneity

Table 1.4 and Figure 1.8 explore variation in price responsiveness across measures of billing cycle weather conditions and typical water usage. These models account for the fact that residential consumers’ sensitivity to price changes may look quite different during hot summer months or for high-usage households in comparison to colder winter periods or in more temperate climates.

Table 1.4 reports results from models estimated using Eq. 1.1 for specific seasonal and geographic sub-samples. Columns (1)-(4) estimate price elasticities and differentials for excessive users in spring, summer, fall, and winter billing cycles respectively. While point estimates for the price instrument coefficient in spring, summer, and winter are highly similar to the overall estimates, the point estimate for fall is positive but statistically insignificant. Of these, only the summer price coefficient is statistically significant, suggesting greater uniformity in price responsiveness when outdoor irrigation needs are at their highest (as reflected by an average usage of 334 gallons per day). Once again no statistically identifiable difference in responsiveness is observed for households who ever use over 1,000 gallons per day during a billing cycle.

Columns (5) and (6) split the sample for 5-digit zipcodes to the east and west of the Berkeley hills. A sub-range of the Pacific Coast Ranges, the Berkeley hills run north-south through the entire EBMUD service area, separating cities along the San Francisco Bay from their inland counterparts. As the range reaches nearly 2,000 feet in elevation, it blocks coastal fog from reaching the eastern cities – leading to drastically different climates and outdoor irrigation needs for comparable homes on either side.¹⁷ Coefficient estimates and mean water usage by region reflect these climatic differences. Homes west of the Berkeley hills use on average 187 gallons per day during the sample period, while homes to the east use over 120% more per day, at 412 gallons. The eastern region yields a -1.45 price coefficient in Column (6) while a much smaller magnitude effect of -0.83 is observed for homes in the western portion (statistically different from each other at the 95% level, $z = 2.06$). Residents to the east being over 75% more responsive to price than their western counterparts likely reflects their greater irrigation needs, higher average consumption levels, and subsequent larger propensity to conserve.

¹⁷These differences in outdoor irrigation needs are reflected in water budget tables provided by the district. Homes east of the hills are expected to require an additional gallon per day for each hundred square feet of either lawn or shrubs (EBMUD 2021).

Figure 1.8 explores heterogeneity in price responsiveness across the distribution of weather exposure. Figure 1.8 plots the coefficients and 95% confidence intervals for the overall price elasticity at each percentile of experienced weather, with estimates for growing degree days reported in the top panel and for total precipitation in the bottom panel.¹⁸ The coefficient at 1% (far left) in the top panel indicates the estimated price elasticity for a home that experienced the first percentile of growing degree days during their billing cycle (no days with average daily temperature above 10°C) with the furthest right estimate reporting the same for the 99th percentile (664 accumulated growing degree days). The bottom panel of Figure 1.8 reports corresponding estimates for the distribution of rainfall (from 0mm at the 1st percentile through 354mm in the 99th).

Figure 1.8 shows the role that weather plays in informing water needs and sensitivity to price. In the top panel, billing cycles under the hottest conditions yield the largest magnitude coefficients, with an estimate of -1.7 for temperatures in the 99th percentile (664+ accumulated growing degree days). At the opposite ends of the spectrum we observe dramatically smaller effects. For households experiencing no or few accumulated growing degree days at the bottom of the temperature distribution, we observe greatly attenuated price elasticity estimates between -0.4 and -0.8. Estimates for the 2nd-8th and 9th-17th percentiles (all larger than -1) are statistically distinguishable from that of the 99th percentile. While point estimates decrease monotonically from the 1st through 99th percentiles, none are statistically different than at the median of the distribution (-1.1).

A similar but more extreme pattern holds across the distribution of rainfall in the bottom panel. Under no rainfall (1st - 7th percentiles of billing cycles), households display a price elasticity of roughly -2.6, statistically distinguishable from the elasticity at the median (-1.1) and all but 13 percentiles in the bottom half of the distribution. Coefficients between the 2nd and 65th percentiles remain highly stable, and while point estimates increase nearly monotonically from the 75th through 95th percentiles they are statistically indistinguishable from one another. At the highest levels of rainfall (291mm for the 95th percentile) where households likely have no need for additional irrigation, estimated elasticities approach one and even become largely positive, with 96th-99th statistically different from the median estimate with estimates between -0.1 and 1.4.

Taken together, the panels in Figure 1.8 illustrate that homes exhibit their highest price responsiveness under climatic conditions that prompt the greatest irrigation need. When temperature is highest and precipitation lowest – times when demand for outdoor irrigation water and total bills would be at their peaks – I estimate that households are the most responsive to changes in price. As temperature falls and precipitation increases, irrigation needs decline and overall expenditures on residential water fall – and so too does sensitivity to price. These patterns suggest that price sensitivity falls as a household’s water usage

¹⁸Estimates are obtained using models that follow Eq. 1.1 where the price instrument is interacted with each percentile of a given variable’s distribution. Each point estimate is calculated as the sum of the percentile’s interaction term and the coefficient on the baseline term (50th specified as the omitted group) with the standard error of the sum calculated from the relevant components of the variance-covariance matrix.

reduces to expected indoor-only consumption levels and that many households are willing to substitute away from outdoor irrigation in cases of both high climatic need and high prices.

To explore the role of chronic usage on price responsiveness, Appendix Figure A.3 estimates price elasticities across the distributions of mean household water usage (top panel) and mean parcel vegetation levels (bottom panel). In both cases, estimates for any given percentile are statistically indistinguishable from any other across the entire distribution. Vegetation index models yield a median price elasticity of -1.3 while households at the median of average water usage (184 gallons per day) present an elasticity of roughly -1 that is indistinguishable from those of preferred specifications in Table 1.3.

Shifting Price Elasticities Under Drought Policy

I next examine how price-responsiveness of EBMUD customers changed following implementation of the Excessive Use Penalty Ordinance during the height of drought-induced water scarcity. Allowing the price elasticity to vary across periods shows how households' responses to water prices varied under drought policy and water storage conditions. To the extent that households engaged in conservation behavior prior to the start of the EUP, they had likely shifted their consumption to a relatively less elastic portion of the demand curve – resulting in a positive estimate for $\Delta \widehat{\ln(AP)} \times \text{EUP}$. If further behavioral adjustments or structural changes to outdoor irrigation needs took place in response to drought policies, then a similarly positive $\Delta \widehat{\ln(AP)} \times \text{Post-EUP}$ would likely arise. However, if households heeded the end of district and state-level drought emergencies as signs that conservation was no longer necessary, than they could have increased consumption – either relative to during the EUP or even to the pre-policy period.

In Table 1.5 I report the estimated price coefficients from models following Equation 1.3. The first row of the top panel reports coefficients for the un-interacted average price simulated instrument, $\Delta \widehat{\ln(AP)}$, reflecting average responsiveness prior to the drought policy period (akin to the pre-drought policy elasticities in Table 1.3). The second row reports the coefficient estimate for $\Delta \widehat{\ln(AP)} \times \text{EUP}$, reflecting differential price responsiveness while the EUP was in effect. The third row similarly reports the interaction of the price instrument with the post-drought emergency period. The bottom panel reports the overall elasticities for each period and their standard errors.

Column (1) reports estimates from the ordinary least squares model omitting any controls. Column (2) adds weather splines and past consumption controls, while Column (3) includes month-year and Zip-5 fixed effects and Column (4) using month-year by Zip-9 fixed effects. Column (5) uses the specification of Column (4) but limits the sample to households that ever use over 1,000 gallons per day. Column (5) restricts the sample to households that violated the Excessive Use Penalty and includes month-year by zip-5 fixed effects to ensure identifying variation is not fully absorbed by controls. As these violating households were the only group subjected to excessive use penalties, their elasticity includes identifying variation around this temporary marginal price tier and provides a valid estimate of plausibly

short-run responses.

Results in Column (1) that fail to account for mean reversion or experienced weather yield large magnitude estimates of -1.48 and 1.73 for the pre-EUP price elasticity and the post-EUP elasticity difference, respectively. Estimates in columns (2) and (3) return elasticities of -0.70 prior to the EUP, which attenuate to nearly zero after the policy ends. Both models estimate large, positive, and statistically significant coefficients during the EUP, implying price elasticities during the drought policy period of roughly 1.7 .

Estimates obtained in Column (4) under the preferred control specification mirror long-run elasticity estimates prior to enactment of drought policies. The coefficient of -0.96 on the un-interacted price instrument matches overall pre-drought policy estimates from Table 1.3. The time period interactions suggest monotonic attenuation over time, with a statistically significant price elasticity of -0.44 during the EUP and an insignificant estimate of -0.16 after its removal. These findings support the notion that after households adopted water conservation practices and district-wide usage fell, they became increasingly less responsive to future price changes.

Columns (5) and (6) illustrate increased pre and post-policy price responsiveness for high usage households, with no conservation impact of price increases for excessive users while drought policies were in effect. When considering only households that ever used an “excessive” volume in Column (5), we observe a pre-EUP price elasticity estimate of -1.24 , larger in magnitude than the overall estimate from Table 1.3 but smaller than that for all Eastern EBMUD customers from Table 1.4. The price elasticity falls by roughly half to -0.68 during while the EUP was in effect, and the -0.45 estimate once again becomes statistically indistinguishable from no response once the policy period ends.

Limiting the sample to only ever-treated units that violated the EUP, estimates in Column (6) provide evidence of increased price sensitivity for top water users – except while drought policies were in effect. Prior to experiencing increased prices and fines during the EUP, households that would eventually violate display a price elasticity of -2.23 , over three times as large as the equivalent district-wide average in Column (3). Once the EUP came into effect and these households experienced moral suasion (and potentially public shaming), their price-responsiveness is weakly statistically significant with a positive overall elasticity of 1.59 . This positive elasticity arises from a 3.82 EUP-period elasticity differential, more than twice the magnitude of policy-period differentials of $1.72 - 1.73$ for all households in Columns (2) and (3). Once the EUP ends, top users’ responsiveness falls in line with other high-usage households, with a statistically insignificant post-EUP elasticity of -0.51 . This pattern suggests that, when already subject to behavioral policy instruments, excessive use fees did not induce further conservation by violating households.¹⁹

¹⁹Equivalent estimates obtained using month-year by Zip-9 fixed effects are highly similar but much noisier and yield coefficients and standard errors of -2.61 (2.19) for $\Delta \ln(\widehat{AP})$, 4.24 (1.74) for $\Delta \ln(\widehat{AP}) \times \text{EUP}$, and 1.52 (2.18) for $\Delta \ln(\widehat{AP}) \times \text{Post-EUP}$.

Price and Non-Price Responses to Drought Policies

I next present results from Eq. 1.3 in the form of event study figures and aggregate post-treatment effect figures. Event study figures are a common, useful tool for reporting treatment effects and illustrating their evolution over time. A further benefit of event study methods is the avoidance of weighting concerns inherent in typical difference-in-differences estimators (Goodman-Bacon, in press) and the ability to aggregate effects to time periods relevant to the research setting. I present the main event study estimates in Figure 1.9, and in Figures 1.10-1.11 I show the robustness of estimated effects to alternate control structures and sub-samples.

Figure 1.9 plots the moral suasion and public shame event studies for the preferred specification following Eq. 1.3. The blue points and band in the top panel correspond to the event-time coefficient estimates and 95% two-way clustered confidence intervals for moral suasion exposure. In the bottom panel, the red points and band report the public shaming event-time effects from the same model. As the event-time coefficient $\hat{\tau}^k$ measures the annual change in log gallons per day due to being k billing cycles since treatment, it can be interpreted as a semi-elasticity and reports the average percentage annual change in water use relative to the billing cycle prior to treatment ($k = -1$). A coefficient equal to zero conveys that treated units experienced no difference in their annual change in water use relative to their pre-treatment cycle while a coefficient of -20% indicates that water use fell by 20% on average among treated homes k periods after first exposure relative to control households.

Turning first to moral suasion in the top panel, we observe delayed conservation responses that die out within a year and a half. Annual changes in water use prior to treatment for violating households consistently appear slightly higher than in control households, ranging from 0 and 15% . When a violation first occurs the gap between violating and control households grows to 52.5% at event time 0, suggesting that violation is correlated with a particularly large water shock or outlier use period. This effect falls to 26% and 4% over the following two billing cycles before turning negative in event time 3. Treatment effects grow in magnitude from -9% to -16% between event times 3 and 5 (6 and 10 months post-treatment) before reaching a maximum change of -67% after 12 months ($k = 6$), with all four coefficients statistically significant. The moral suasion effect reduces in magnitude to -43% at event-time 7 with estimates for all future periods even smaller in size and often statistically indistinguishable from zero.

Public shaming displays a larger magnitude but similarly short-run conservation response. Pre-treatment event time indicators in the bottom panel are largely statistically indistinguishable from zero, with point estimates oscillating around zero with little clear pattern. Water use spikes by 69% at event time $k = -2$, suggesting that eventually-shamed households engaged in trend-deviating activity in the months shortly before their first shaming (and likely prior to their first EUP violation). Once public shaming occurs, we observe a negative but statistically insignificant point estimate for the billing cycle during which shaming occurs ($k = 0$), reflecting the partially-treated nature of that cycle. Estimates

remain negative but statistically indistinguishable through the next six months (through event time $k = 3$). Water usage declines further over the next six months post-shaming, with statistically significant reductions of 36%, 71%, and 60% estimated for billing cycles 8, 10, and 12 months after shaming, respectively. Once again, additional conservation effects of public shaming disappear shortly after the EUP ends, with the coefficient for 14 months after first exposure ($k = 7$) losing all significance and jumping to a small, positive value. Coefficients for future periods are similarly indistinguishable from no difference relative to control households.

The disappearance of treatment effects shortly after district-wide conservation policies end confirms the observed patterns in overall consumption changes. Figure 1.6 showed that average residential water consumption fell substantially as drought severity grew, even before the EUP came into effect. During the EUP period, consumption further fell among both violating and non-violating households. In contrast, consumption patterns largely stabilized once the drought emergency was lifted. At that point, consumers were no longer inundated with messaging to conserve – both from the district and at the state-level (California’s drought emergency ended in April 2017). The signaling of crisis by the water district and the risk of future public shaming for high-usage households both ended once drought policies were abolished. As a result, the welfare costs to these top users of maintaining greater reductions in water use relative to their neighbors largely disappeared and prompted a return to more typical consumption patterns.

Robustness of Behavioral Responses

To explore the robustness of estimated treatment effects for moral suasion and public shame, I next present estimates of average first-year treatment effects across a range of specifications and sub-samples. Figures 1.10 and 1.11 plot average first-year treatment effects on the treated and 95% confidence intervals clustered by 9-digit zipcode and fiscal year for the full analysis sample (“Full,” $N = 8.8M$) and four progressively smaller samples of control households: those in 5-digit zip codes bisected by an elevation band (“Band Split,” $N = 7.0M$), that ever used over 1,000 gallons per day in a billing cycle (“Ever 1K+,” $N = 1.1M$), that average at least 1,000 gallons per day (“Avg 1K+,” $N = 163K$) and only those that violated the Excessive Use Penalty (“EUs,” $N = 149K$) without never-treated units.

For each set of control households I report estimates from left to right for four sets of increasingly conservative model specifications. The first column omits weather controls and all fixed effects from Eq. 1.3 and is akin to a household-level fixed effects model that does not control for past consumption. The next specification includes weather splines and past consumption controls, and specifies δ_{zt} as month-year fixed effects. The third column adds Zip-5 fixed effects, while the fourth column utilizes month-year by Zip-9 fixed effects and matches the specification of Figure 1.9.

Moral suasion results in Figure 1.10 show stability of first-year ATTs across the first four sub-samples with the largest effects estimated when considering only the households that violated the EUP. Ordinary least squares estimates with no controls (left-most estimates) are

highly stable across samples, ranging between -16% and -21% with no ability to statistically distinguish any of the point estimates from one another. The three increasingly-conservative control structures also return highly stable estimates for samples that include never-treated units: estimated average first-year effects fall between -10% and -28% across all specifications and between -10% and -17% for the preferred month-year by Zip-9 fixed effects model. These estimates correspond to average reductions 10-28% greater for households exposed to moral suasion in the year following EUP violation in comparison to never or not yet-treated households.

First-year conservation effects of moral suasion grow substantially when the sample is limited to EUP violators. When only identifying off comparisons between currently and not yet-treated households, moral suasion ATTs increase in magnitude to -49% to -58% when models include controls. The -58% effect for the preferred month-year by zip-9 fixed effects specification (farthest right) is statistically distinguishable from equivalent models run on the full, band split, and ever over 1,000 gallons per day samples (but is indistinguishable from the average 1,000 gallons per day sample estimate of -17%). This increase in effect magnitude suggests that never-treated households decreased their water usage during the EUP at faster rates than households who eventually violate the ordinance, resulting in attenuated first-year treatment effect estimates in models that include pure control units.

Public shame results in Figure 1.11 provide evidence of increased stability across samples and confirm the effectiveness of public shaming on inducing additional conservation beyond those elicited by fees or moral suasion. Estimated first-year treatment effects on the treated are statistically indistinguishable across samples and specifications with estimated changes in water use ranging from -14% to -32% . These estimates indicate that publicly shamed households further increased the magnitude of their water reductions in comparison to control households for the twelve months following being called out in the local news.

While estimates are within 3 percentage points of each other for the full and band split samples, estimates from the most conservative control specification become smaller in magnitude and more noisy as the sample sizes decrease. This attenuation and increased variance likely arise from the low number of households now identifying each fixed effect. When the sample is reduced to only EUP violators, 2,300 out of 3,189 9-digit zip codes only ever contain one household – with 2,817 containing fewer than three households.²⁰ As a result, all identifying variation for 17 out of 56 shamed households (and 2,283 unshamed EUP violators) will be fully absorbed by the month-year by Zip-9 fixed effects.

Comparison of Responses to Price and Non-Price Drought Policy Instruments

To better understand which policy instruments are well-suited to short or long-run conservation policies, I next explore how the conservation effects of non-price behavioral drought

²⁰In contrast, 39,845 out of 54,588 9-digit zip codes in the full analysis sample contain 3 or more households. Only two of these singleton-household zip codes contain a publicly shamed account.

policies compared to consumers' responses to price changes. Combining price elasticity estimates from Tables 1.3 and 1.5 with first-year responses to moral suasion and public shame from Figures 1.10-1.11, I calculate the additional short-run price changes that would have been necessary to match behavioral responses to non-price drought policy instruments.

Table 1.6 reports the implied price increases necessary to match households' average first-year response to moral suasion and public shame exposure across a range of samples and price elasticities. Models (1)-(3) utilize behavioral first-year ATTs from the fourth "Full" sample estimate in Figure 1.10 that includes weather splines, past consumption controls, and month-year by 9-digit zipcode fixed effects. These models use price elasticities from Column (4) of Table 1.3 and calculate necessary price increases using price elasticities before, during, and after the EUP policy period, respectively. In this way Model (1) reflects the price increase necessary if households were to return to consumption behavior observed before extensive drought reductions, while Model (3) reflects typical post-drought behavior.

Models (4)-(5) focus on the responsiveness of higher usage households during the drought policy period, using estimates for households that ever used over 1,000 gallons per day ("Ever 1K+") or violated the Excessive Use Penalty ("EUs"), respectively.²¹ Calculations for moral suasion are reported in the top panel while the bottom panel repeats the calculations for responses to public shame.

Turning first to the top panel of Table 1.6, we see increasingly large price changes necessary to match responses to moral suasion as households adapted to drought. Model (1) calculations show that a relatively modest average price increase of 11% could have induced similar water use reductions as the first-year effect of moral suasion if households returned to their pre-drought price elasticity of -1.09 . However, this necessary increase grows substantially in magnitude as households adapt to drought and shift to less elastic regions of the demand curve. The average price increase of 30% for all single family home customers in Model (2) reflects the much-attenuated price elasticity of -0.435 observed during the height of drought. Once the EUP period ended and price elasticities fell further to -0.159 , I find that prices would have had to more than double, with an increase of 107% needed to match moral suasion's impact among top water users.

Turning to Models (3)-(4), we see a stark difference in necessary price changes when considering potential versus actual EUP violators. Back-of-the-envelope calculations for households that ever used above the EUP threshold in Model (3) suggest a price increase of 15% to match price responsiveness while policies were in effect. Shifting to the set of households that actually violated the policy in Model (4), we see that no price increase could induce short-run conservation behavior given the large and positive short-run price elasticity (1.587). This finding suggests that excessive use fees primarily functioned as a short-term revenue source rather than as an instigator of conservation behavior among the highest-use households.

²¹While Model (3) follows the specification of Models (1)-(2), calculations for excessive use violators in Model (4) include month-year and 5-digit zipcode fixed effects in place of the interacted fixed effects to ensure sufficient identifying variation. See the discussion at the end of Section 1.5 for more details.

Looking to the bottom panel for price increases needed to match responses to public shame, we observe similar patterns at generally higher magnitudes. Full sample models imply necessary average price increases of 32% under pre-drought price sensitivity, 101% for conditions while drought policies were in effect, and 574% after drought emergencies ended. Models limiting control households to those that ever used “excessive” volumes during the entire sample period suggest average price hikes of 41% during the EUP. Once again, considering only ever-treated households who violated the policy yields no conservation ability of short-run price changes during the EUP.

These findings indicate that feasible price changes during drought crisis likely would not have resulted in comparable conservation effects for high-usage household as moral suasion and public shame combined. When considering observed price responsiveness for all households during the policy period, these non-price instruments induced conservation on par with a short-run sextupling of average prices – an infeasible change in nearly all conceivable settings and policy contexts. Although considering elasticities for general high-usage households reduces this total price change to 56%, the top users actually exposed to excessive use fees showed no price-based conservation response while simultaneously exposed to both non-price policy instruments.

This result of limited power for feasible short-run price increases to induce conservation during extreme drought mirrors existing findings across a broad set of prior price instruments and drought conditions. Doubling the marginal price faced by high-usage households during a prior drought in Santa Cruz, CA resulted in a 12% reduction in water used by treated households (Nataraj and Hanemann 2011). Similarly, a 52% increase in average price was needed to match conservation impacts of mandatory use restrictions in North Carolina (Wichman, Taylor, and von Haefen 2016). Perhaps the most comparable policy occurred a year before the EUP in Santa Cruz, CA. In 2014, Santa Cruz Municipal Utilities imposed excessive use fees that increased average prices by 54-299% for consumption over 125 gallons per day. While nearly 14% of residential customers received at least one fine, aggregate water use fell by only 5-10% as a result of these fines and broader social pressure (Pratt, in press). Evidence from these studies suggests that price-based conservation would likely not have matched responses to non-price instruments even if EBMUD had set the excessive use threshold at a much lower point in the water distribution.

1.6 Conclusion

In this paper I take advantage of unique variation in price to estimate causal price elasticities and separately identify excessive users’ responses to price and non-price water conservation policies during recent extreme drought in the San Francisco East Bay Area. Using administrative data for the universe of residential water users in a major California water district, I estimate medium to long-run price elasticities of -0.97 to -1.09 prior to the imposition of conservation-oriented drought regulation. I find that price elasticities are largest under conditions prompting the highest levels of outdoor irrigation and for high-usage households.

However, residential price responsiveness attenuates substantially over time as households adopt conservation behavior and reduce non-essential outdoor use, shifting to less elastic regions of the demand curve. The employed simulated instrument approach builds upon previous causal water demand estimation strategies by eliminating mean reversion and distributional shifts as potential sources of bias. Further, estimates prove highly robust to the choice of price measures and control approaches.

Models focused on responses to price and non-price drought policies provide the first empirical evidence of public shame's effectiveness at inducing resource conservation. Moreover, I separately identify responses to public shame from changes due to excessive use fees or moral suasion messaging. I find evidence of sizable short-run conservation due to behavioral instruments but no comparable response to price changes. In the year following exposure to moral appeals to conserve, treated households reduced water consumption by 10-50% relative to control units. Households that were publicly shamed after moral suasion exposure demonstrate additional conservation, further reducing water use by 23-30%. In contrast, I estimate a large, positive price elasticity for excessive users exposed to short-run increases to marginal price. Back-of-the-envelope calculations imply that drought policies would have had to increase average prices by an additional 11-32% to match conservation responses to either moral suasion or public shame alone given overall pre-drought elasticities. This necessary price change grows to 30-574% when considering overall price responsiveness during or following drought policies. Importantly, no further price increase could have prompted additional conservation from top water users who violated the drought usage threshold.

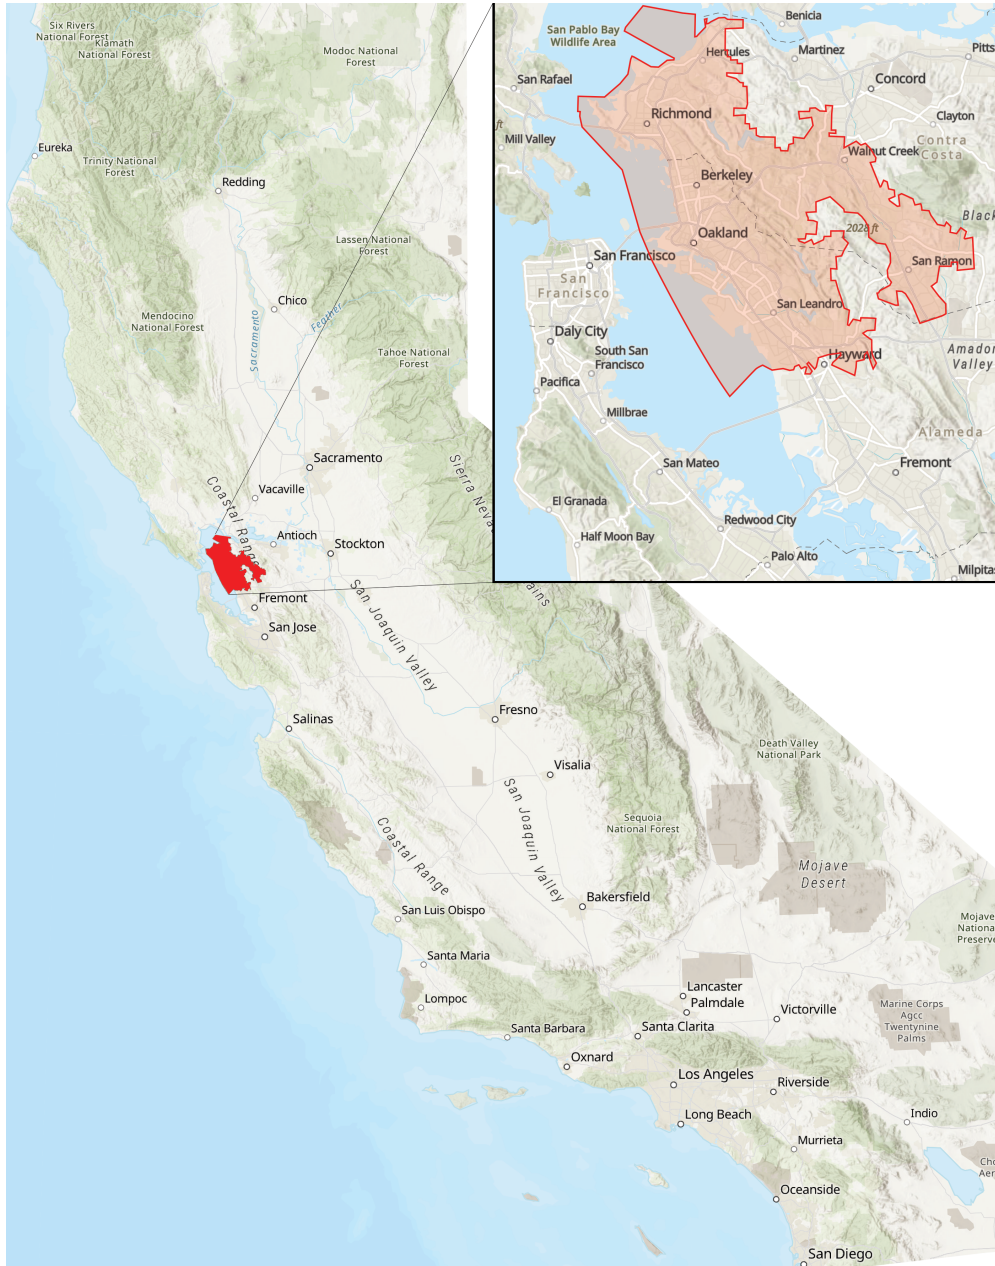
Taken together, these findings convey important policy implications for the design of future urban water conservation policies. In the study setting, price-based policies proved most effective at inducing long-run adaptation behavior and had the least power at affecting short-run, discrete reductions among top water users. Early on in the sample, price-based policies proved effective at driving adaptation behavior as drought conditions grew more severe. Further, prices showed the greatest ability to reduce non-essential, outdoor irrigation over several years of increasing water scarcity. However, financial incentives displayed substantially lower effectiveness at driving short-run, discrete reductions in response to acute shortages, as households had already eliminated many non-essential water uses and adjusted along the most promising conservation margins. This is especially true among top water users who displayed no conservation response to excessive use fees while already exposed to behavioral policy measures. However, these same users showed the greatest price-responsiveness post-drought emergency – suggesting that price-based policies targeting gradual reduction of outdoor water uses could still prove valuable for adjusting to new equilibrium water supply levels in the future.

In contrast, moral suasion and public shamed proved both highly effective at driving discrete, short-run conservation and fully reliant on consumer belief in crisis. Together, moral suasion and public shame reduced consumption among the highest-usage households by 40-70% for nearly a full year at the peak of the drought emergency. However, once the district ended its drought emergency, the messaging these households faced shifted to one of a wet winter and sufficient supplies to more than meet essential needs. At that point, all

conservation effects from prior exposure to moral suasion and public shame disappeared – indicating that households stopped responding as soon as they lost belief in the moral need to conserve. This lack of long-run response mirrors prior findings of residential water and energy users’ habituation to moral conservation appeals (Ferraro, Miranda, and Price 2011; Ito, Ida, and Tanaka 2018) and extends the conclusion to the case of public shaming.


Future research can continue to improve our estimation of residential water demand and to identify the impacts of non-price policy mechanisms for eliciting conservation behavior. Uncovering plausibly causal price elasticities in additional settings will provide a foundation for improving the targeting of price-based policies and designing rate structures that can balance long-run conservation goals with existing constraints. Studies that explore the extent to which behavioral measures crowd out price responses in the short and long-run will help water districts improve the design of future conservation policies. Further, understanding how pricing strategies interact with various short-run drought policies can both guide the choice of optimal rate structures and increase our understanding of the equity implications of various conservation policy approaches. Research providing such insights has the potential to improve the reliability and availability of urban water around the globe and help reduce a major obstacle to the continued success of modern urban societies.

Figure 1.1: East Bay Municipal Utility District (EBMUD) Service Area



This figure plots the location of the EBMUD service area in California. The inset map in the upper-right corner provides a zoomed-in view of San Francisco's East Bay Area and shows the EBMUD water service area in red (shape files obtained from EBMUD).

Figure 1.2: East Bay Municipal Utility District (EBMUD) Excessive Use Penalty Bill



EAST BAY MUNICIPAL UTILITY DISTRICT
 EBMUD 1-866-40-EBMUD YOUR ACCOUNT NO. IS: 12345678001

Effective July 1, 2015 new rates are in effect to pay for replacement of aging pipelines and long-term infrastructure improvements. To fund increasing drought costs, a temporary 25% drought surcharge is in effect. Excessive use penalties will apply to residential customers averaging more than 984 gallons per day. For more information visit www.ebmud.com

EBMUD
 Florence Waters
 1234 PIPELINE ST
 OAKLAND CA 94607-1234

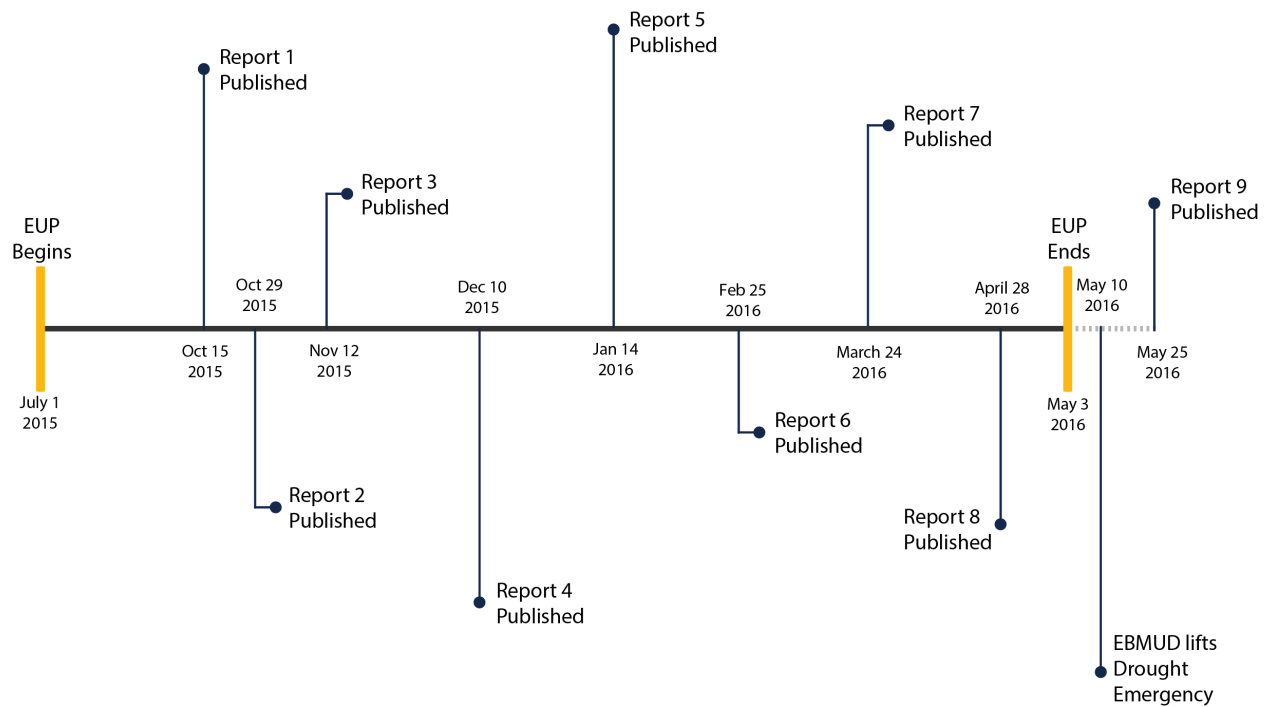
Bill Date: 09/03/15
 Billing Period
 From: 07/01/15 To: 09/01/15

For:	AMOUNT	TOTAL
1234 Pipeline St Private Residence		
WATER CHARGES - EBMUD		
WATER SERVICE CHARGE	38.88	
WATER FLOW CHARGE	44.25	
15 UNITS @ 2.95	73.08	
18 UNITS @ 4.06	359.12	
87 UNITS @ 5.38	60.00	
WATER ELEVATION CHARGE	115.55	
100 UNITS @ 0.80		
DROUGHT SURCHARGE		
EXCESSIVE USE PENALTY	40.00	730.68
PLEASE SEE REVERSE SIDE FOR BILLING EXPLANATION	Please Pay This Amount Now Due	882.47

METER SIZE	ELEV. Band	METER READINGS		UNITS	CONSUMPTION INFORMATION		
		Current	Previous		Gallons	Days	Ga/Day
5/8 inch	2	2,197	2,097	100	74,800	82	1,208
		2013	USAGE	125	93,500	63	1,484


Source: EBMUD. This figure shows a copy of a residential water bill for an EBMUD household that violated the Excessive Use Penalty and faced penalty fees in addition to drought surcharges and typical marginal prices and service charges.

Figure 1.3: Excessive Use Penalty Ordinance (EUP) Timeline



This figure provides a timeline of the East Bay Municipal Utility District’s Excessive Use Penalty Ordinance (EUP) during the height of the 2011-2017 California drought.

Figure 1.4: Public Shaming of Excessive Water Users



(a) Source: ABC Channel 7 News

EBMUD Provides Records Identifying 1,108 Excessive Water Users

Seventeen Pleasant Hill residents made the list. Are any your neighbors?

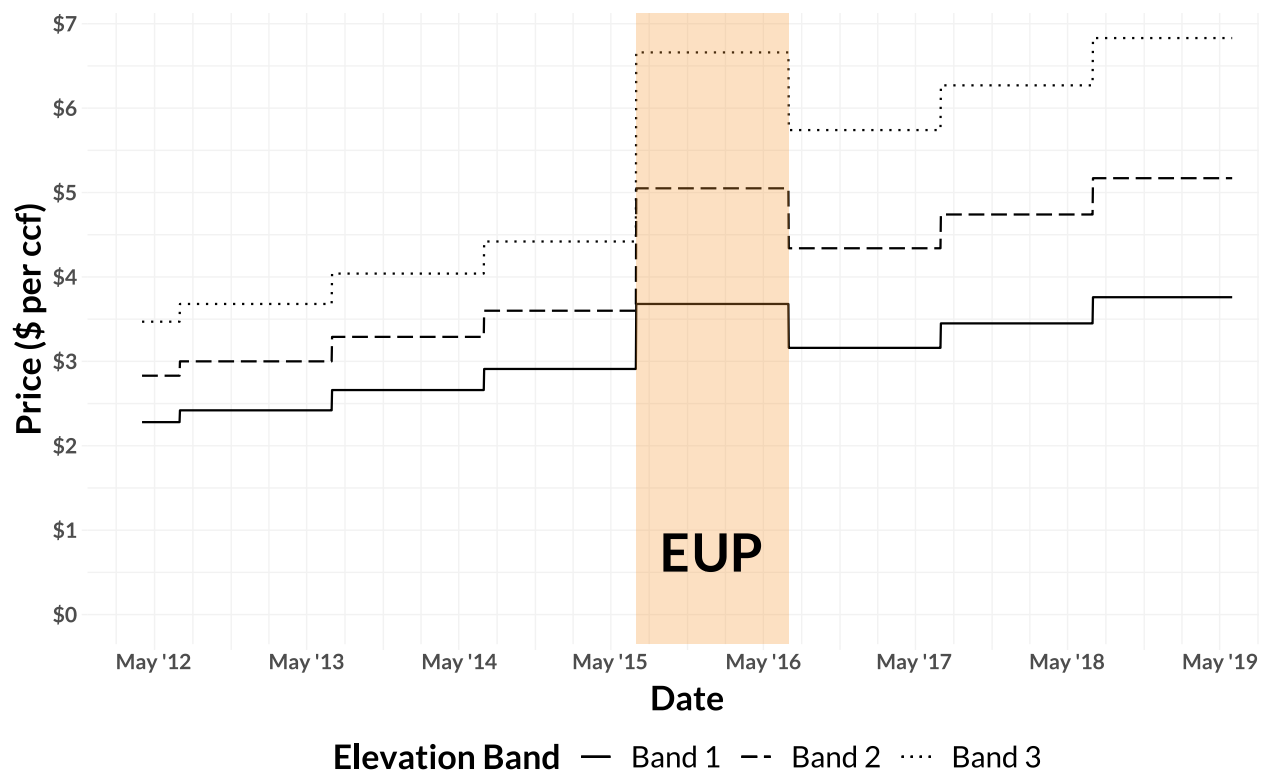
And at the top of the list was [REDACTED] of Danville, a former [REDACTED] executive who used 12,579 gallons per day. Many of the households that violated the ordinance were in East Bay communities with larger properties, Figueroa said, such as Danville and Alamo. EBMUD expected that would be the case.

The top Pleasant Hill water user is [REDACTED]

(b) Source: Patch.com

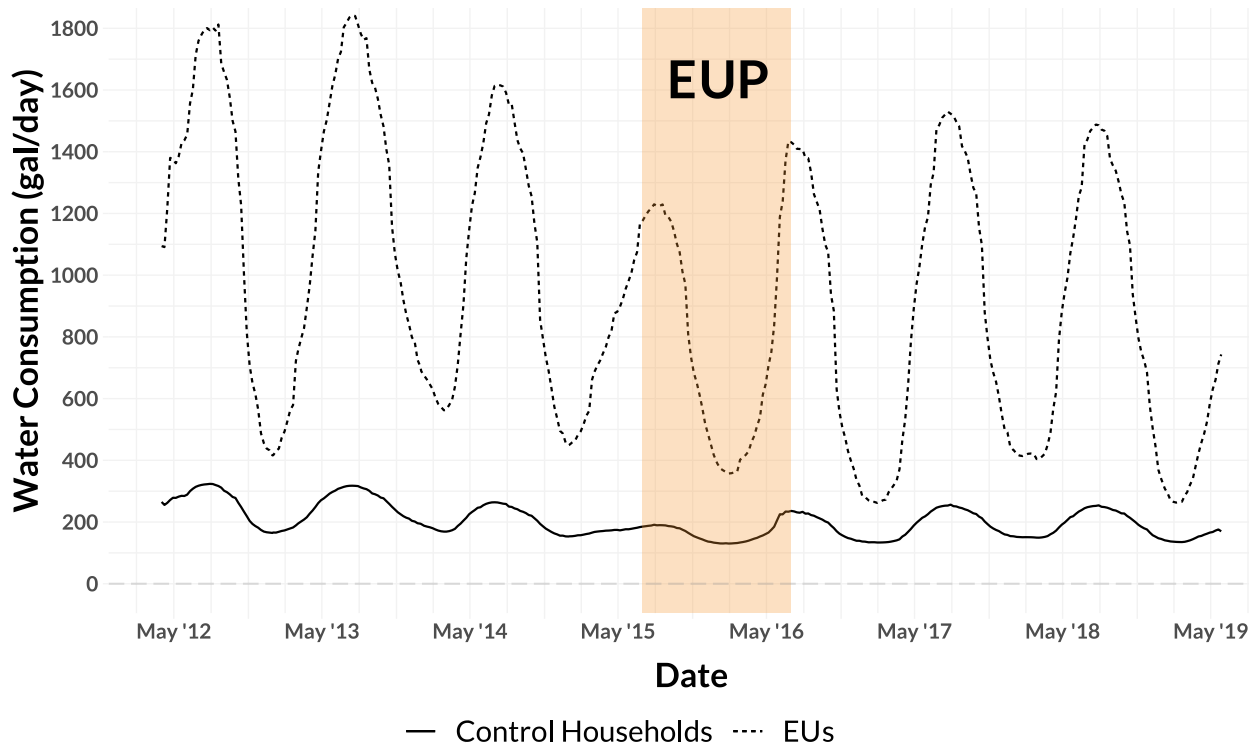
These figures provide examples of news articles calling individual water users out by name after violating the Excessive Use Penalty Ordinance. Figure 1.4a shows a story calling out an individual of local relevance, while Figure 1.4b shows users being shamed for using the most water on a given monthly report or for a given city.

Figure 1.5: Variation in Tier 1 Marginal Prices



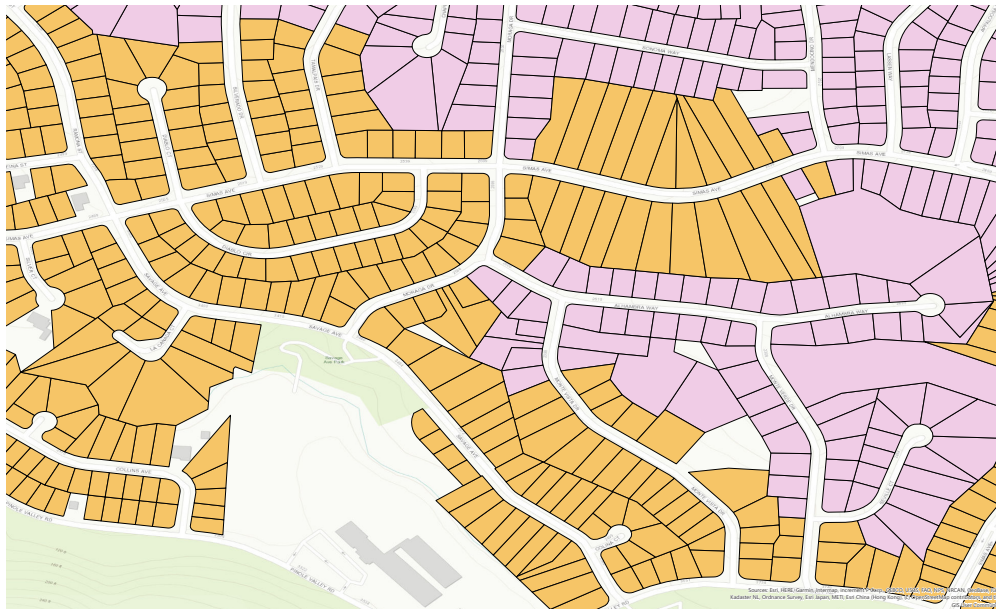
This figure plots variation in marginal prices across time and space for a median consumption household (in marginal price tier 1). The solid line reports the marginal price of tier 1 consumption over time for households in the base elevation band who face no elevation surcharges (band 1). The dashed line reports the marginal price for a household located in elevation band 2 (approximately 200-600 feet above sea level) consuming an identical volume of water. The dotted line reports equivalent marginal prices for homes in elevation band 3 (approximately 600+ feet above sea level).

Figure 1.6: Average Water Use Over Time by EUP Violation Status



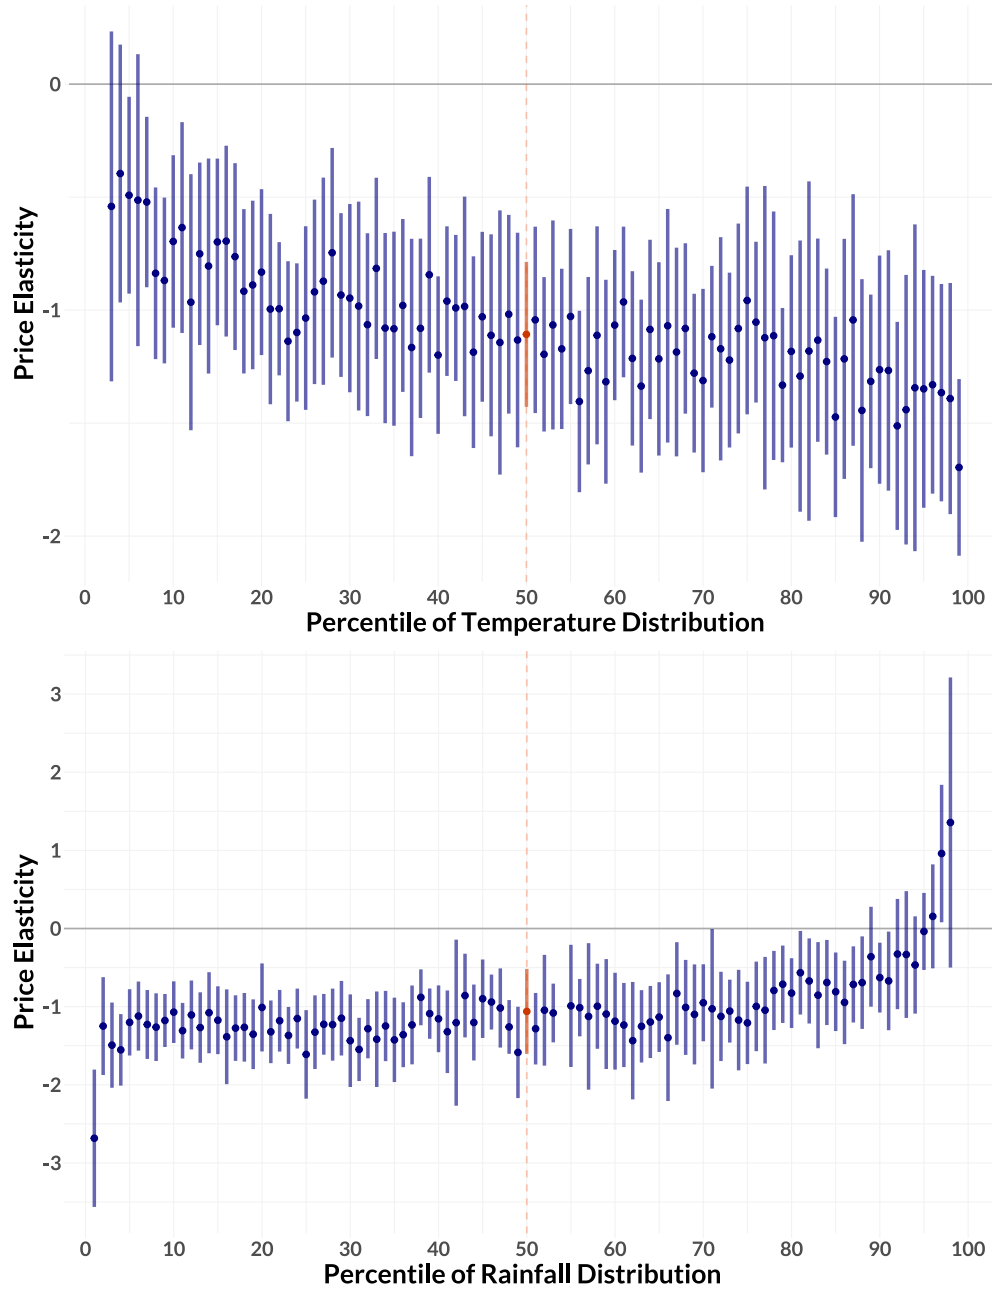
This figure plots the average water use measured in gallons per day for all single family homes in the EBMUD service area from April 2012 through May 2019. The dotted line reports the average usage for customers that violated the EUP while the policy was in effect and were exposed to fees, moral suasion, and potentially public shame (“EUs”). The solid line reports mean water use for all other detached single family home EBMUD customers that were not directly exposed to EUP policy instruments (“Control Households”). The orange band indicates the period that the EUP was in effect.

Figure 1.7: Elevation Band Threshold Example



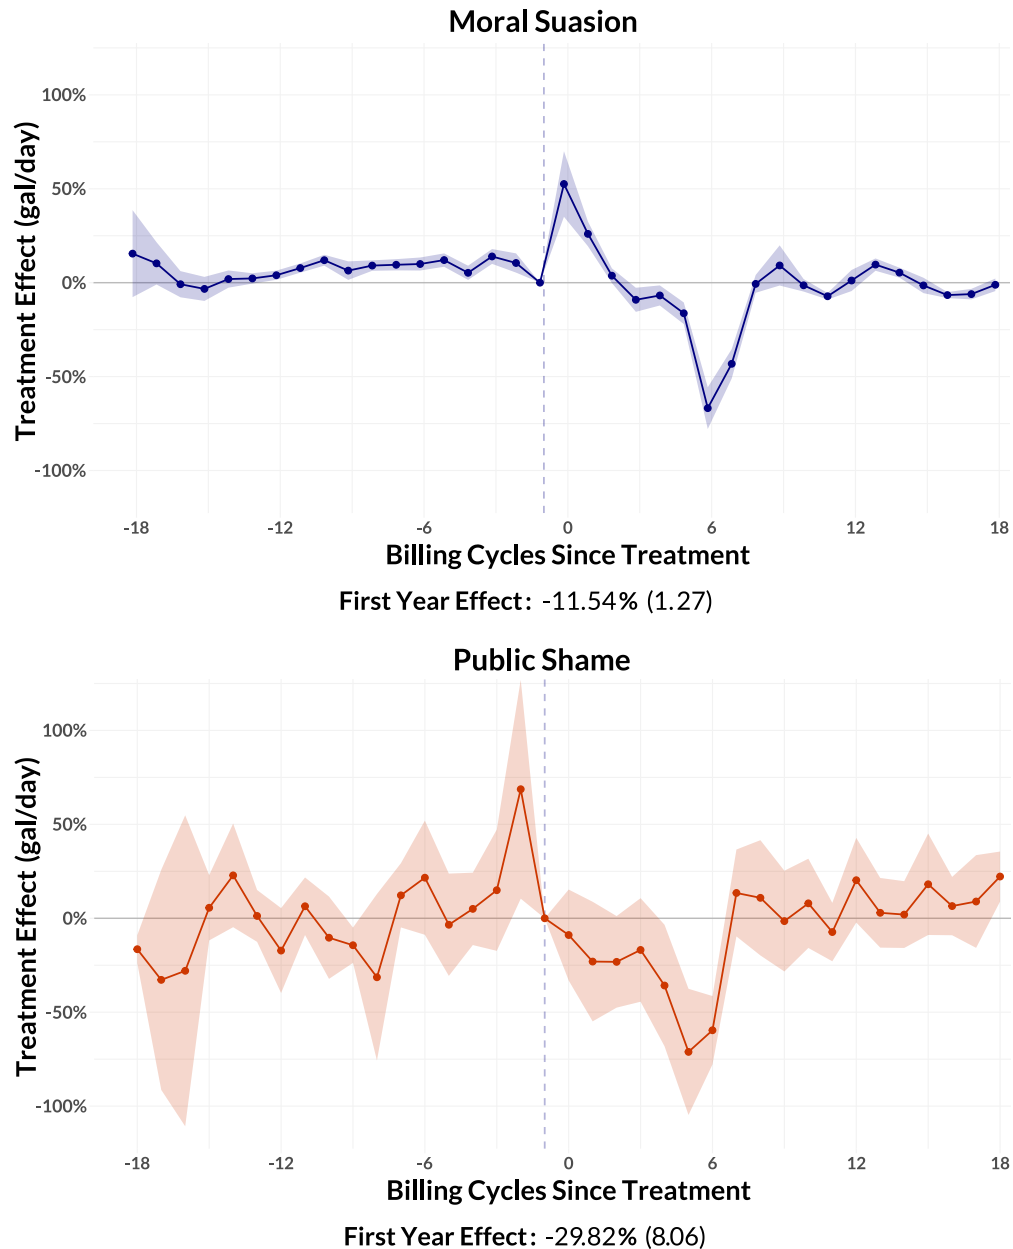
This figure provides an example in Pinole, CA of the elevation band threshold between homes in band 1 (Orange) who face no elevation surcharges and in band 2 (Pink) who pay higher marginal prices at all consumption levels.

Figure 1.8: Price Elasticity Heterogeneity by Billing Cycle Weather



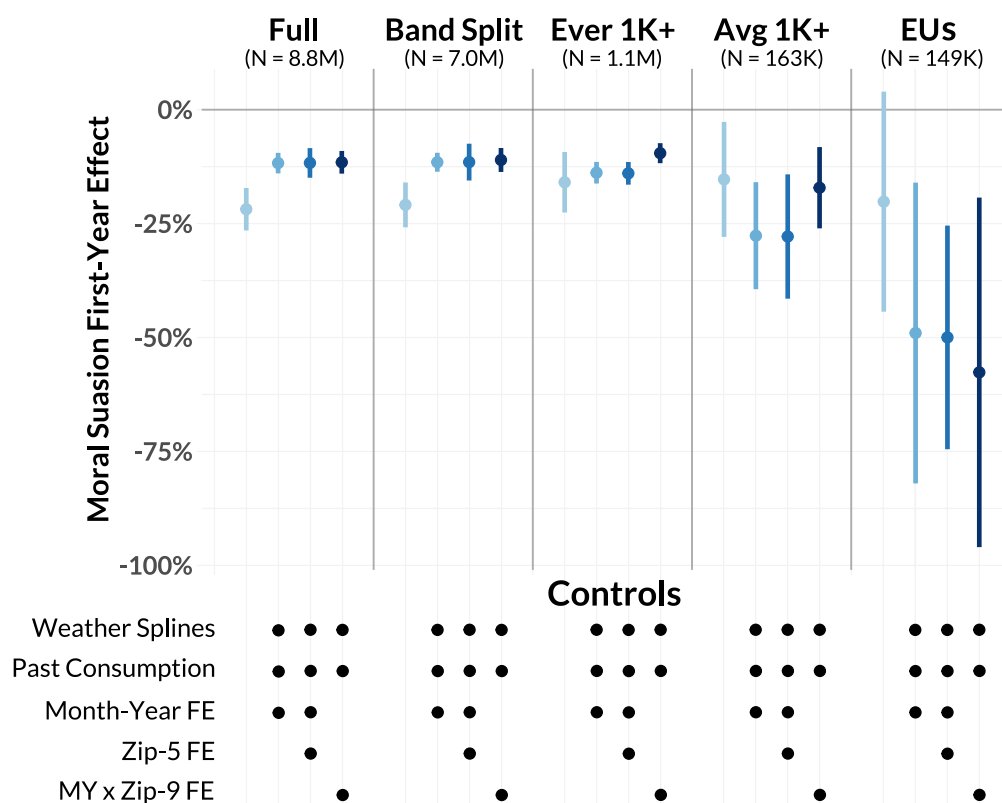
This figure plots the coefficients and 95% confidence intervals for the overall price elasticity at each percentile of the growing degree day (top panel) and precipitation (bottom panel) distributions. Estimates are obtained using models that follow Eq. 1.1 where the simulated price instrument is interacted with each percentile of a given variable's distribution. Each point estimate is calculated as the sum of the percentile's interaction term and the coefficient on the baseline term (50th) specified as the omitted group) with the standard error of the sum calculated from the relevant components of the variance-covariance matrix.

Figure 1.9: Moral Suasion and Public Shame Event Studies



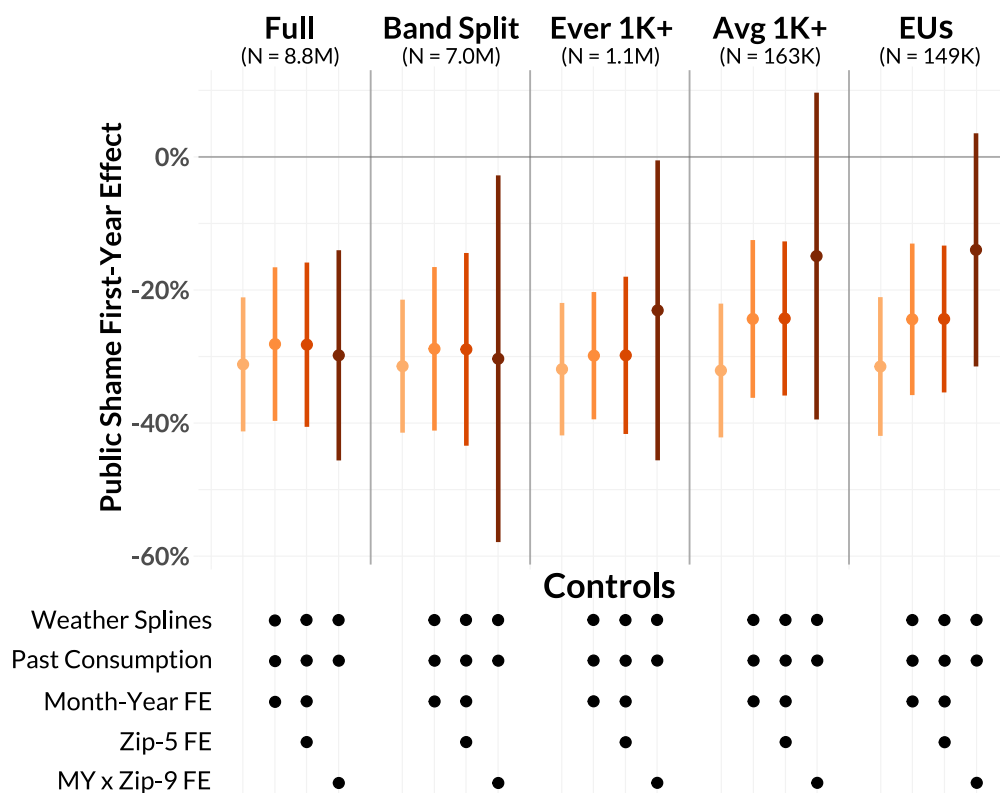
This figure reports event-time coefficient estimates and 95% standard errors clustered by Zip-9 and fiscal year for exposure to moral suasion (top panel, blue) and public shame (bottom panel, red) obtained using Eq. 1.3. The model is estimated in annual differences and includes month-year by Zip-9 fixed effects and controls for underlying shifts in the water use distribution with interactions of the month-year with each percentile of the water use distribution six months' prior. Each point estimate measures the change in water usage k billing cycles since exposure to moral suasion or public shame. "First Year Effect" reports the average treatment effect in the first year post-treatment.

Figure 1.10: First-Year Effects of Moral Suasion



This figure plots point estimates and 95% confidence intervals clustered by Zip-9 and fiscal year for the average response to moral suasion for each billing cycle in the year after first exposure ($1 \leq k \leq 6$). “Full” denotes the entire analysis dataset ($N = 8.8M$) while “Band Split” limits control households to those living in 5-digit zipcodes bisected by an elevation band ($N = 7.0M$), “Ever 1K+” to households who use over 1,000 gallons per day in at least one billing cycle ($N = 1.1M$), and “Avg 1K+” to accounts that average at least 1,000 gallons per day ($N = 163K$). “EUs” includes only households who violated the EUP (and are included in all other sub-samples).

Figure 1.11: First-Year Effects of Public Shame



This figure plots point estimates and 95% confidence intervals clustered by Zip-9 and fiscal year for the average response to public shame for each billing cycle in the year after first exposure ($0 \leq k \leq 5$). “Full” denotes the entire analysis dataset ($N = 8.8M$) while “Band Split” limits control households to those living in 5-digit zipcodes bisected by an elevation band ($N = 7.0M$), “Ever 1K+” to households who use over 1,000 gallons per day in at least one billing cycle ($N = 1.1M$), and “Avg 1K+” to accounts that average at least 1,000 gallons per day ($N = 163K$). “EUs” includes only households who violated the EUP (and are included in all other sub-samples).

Table 1.1: California Water District Policies During Drought Crisis

Policy Instrument	Water Districts
Education campaigns	El Dorado, El Centro, Imperial, Placer
Rebate programs	Calaveras, Folsom, Sacramento, Yuba City
Drought surcharges	Contra Costa, Irvine, LA
Increasing block fees	Citrus Heights, Coachella
Contact high users	Anaheim, LA, Milpitas, Morro Bay, Turlock
“Excessive use” fines	Apple Valley, San Jose, Santa Cruz

This table reports examples of the main policy instruments implemented during the height of California’s 2011-2017 drought for a subset of the states’ water districts.

Table 1.2: Summary Statistics, Billing Cycle Characteristics

	<u>Full Sample</u>		<u>Excessive Users</u>	
	Mean	Std. Dev.	Mean	Std. Dev.
Water Volume (CCF)	18.010	21.804	82.932	79.273
Days per Billing Cycle	60.172	5.737	60.072	6.171
Gallons per Day (GPD)	225.385	280.344	1,041.979	1,001.662
“Excessive” Use (1,000+ GPD)	0.019	0.136	0.409	0.492
Ever Used 1,000+ GPD	0.124	0.329	1	0
Marginal Price (\$/CCF)	3.761	1.069	5.956	1.680
Average Price (\$/CCF)	7.593	4.927	6.108	2.292
Total Bill Amount (\$)	104.674	126.582	501.627	550.458
Elevation Band 2 (\approx 200-600ft)	0.363	0.481	0.615	0.486
Elevation Band 3 (\approx 600ft)	0.095	0.293	0.297	0.457
Lot Size ¹ (acre)	0.147	0.226	0.520	1.412
Home Size ¹ (ft ²)	1,842.858	894.296	3,856.909	1,642.458
Growing Degree Days	311.669	177.409	306.397	216.236
Total Precipitation (mm)	95.657	119.410	102.662	129.834
Observations	10,709,816		184,637	
Households	301,271		4,951	

This table reports summary statistics of billing cycle characteristics for the analysis sample. “Full Sample” columns report the mean and standard deviations of billing cycle characteristics for all households while “Excessive Users” reports values only for households that violated the Excessive Use Penalty Ordinance by using over 1,000 gallons per day in a billing cycle while the policy was in effect. “Excessive Use” is a dummy variable equal to one for any billing cycle with water use over 1,000 gallons per day and “Ever Used (1,000+ GPD)” is a dummy variable equal to one in all billing cycles for customers that ever use over 1,000 gallons per day in a single period. All prices are deflated using the U.S. Bureau of Economic Analysis’ Implicit Price Deflator. ¹Home characteristics are matched for only 7,454,864 observations.

Table 1.3: Pre-Drought Policy Water Demand Estimates

<i>Average Price</i>	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(AP)$	-1.622*** (0.021)	-1.660*** (0.017)				
$\Delta \ln(\widehat{AP})$			-1.110** (0.148)	-1.093*** (0.106)	-1.090*** (0.152)	-1.046*** (0.152)
$\Delta \ln(\widehat{AP}) \times \text{Exc.}$						-0.342 (0.242)
<i>Marginal Price</i>	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(MP)$	2.521*** (0.099)	2.546*** (0.106)				
$\Delta \ln(\widehat{MP})$			-1.024* (0.323)	-0.967** (0.111)	-1.074*** (0.151)	-1.024*** (0.152)
$\Delta \ln(\widehat{MP}) \times \text{Exc.}$						-0.345 (0.243)
MY \times Zip-9 FE	Yes	Yes	Yes	Yes	Yes	Yes
Weather Controls	No	Yes	Yes	Yes	Yes	Yes
Past Q	No	No	No	Splines	Pctl	Pctl
Observations	2.9M	2.9M	2.9M	2.9M	2.9M	2.9M
Adjusted R ²	0.680	0.699	0.117	0.118	0.121	0.121

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are two-way clustered by fiscal year and 5-digit zipcode. These models estimate the pre-drought average and marginal price elasticities of residential water demand for detached single-family homes in the San Francisco East Bay Area using data from April 1, 2012 through July 28, 2015. The dependent variable $\Delta \ln(Q)$ measures the change in the log of gallons per day by a household between the current billing cycle and the same period in the previous year. In the top panel $\Delta \ln(AP)$ is the equivalent change in the log of average price, and $\Delta \ln(\widehat{AP})$ is the simulated price instrument using consumption levels from six months earlier. $\Delta \ln(\widehat{AP}) \times \text{Exc.}$ interacts the simulated instrument with an indicator for whether the household ever used over 1,000 gallons per day in a billing cycle. All models include MY \times Zip-9 fixed effects, interacting the month-year with each postal route (646,124 fixed effects). The bottom panel reports estimates for identical specifications that use the corresponding marginal price measure in place of average price. *Weather Controls* include cubic splines in growing degree days and total precipitation per billing cycle, with knots at each decile of the variables' distributions. *Past Q* controls of "Splines" employ a cubic spline approach for household consumption six months prior akin to weather controls, while "Pctl" indicates percentiles of consumption six months prior interacted with the month-year (2,905 fixed effects).

Table 1.4: Pre-Drought Policy Price Elasticity by Season and Location

	Season				Hills Location	
	Spring (1)	Summer (2)	Fall (3)	Winter (4)	West (5)	East (6)
$\Delta \ln(\widehat{AP})$	-1.063 (1.226)	-1.087*** (0.115)	1.213* (0.919)	-1.047 (7.837)	-0.821*** (0.221)	-1.446*** (0.208)
$\Delta \ln(\widehat{AP}) \times \text{Exc.}$	-0.491*** (0.211)	-0.359*** (0.147)	-0.290 (0.412)	-0.326 (0.430)	-0.254 (0.296)	-0.398* (0.254)
Weather Splines	Yes	Yes	Yes	Yes	Yes	Yes
MY \times Zip-9 FE	Yes	Yes	Yes	Yes	Yes	Yes
Past Q	Yes	Yes	Yes	Yes	Yes	Yes
Mean GPD	244	334	269	181	187	412
Observations	727K	727K	730K	682K	2.02M	907K
Adjusted R ²	0.100	0.071	0.139	0.149	0.082	0.192

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are two-way clustered by fiscal year and 5-digit zipcode. These models estimate the pre-drought policy price responsiveness of residential water demand for detached single-family homes in the San Francisco East Bay Area using data from April 1, 2012 through July 28, 2015. The dependent variable $\Delta \ln(Q)$ measures the change in the log of gallons per day by a household between the current billing cycle and the same period in the previous year. $\Delta \ln(AP)$ is the equivalent change in the log of average price, and $\Delta \ln(\widehat{AP})$ is the simulated price instrument using consumption levels from six months earlier. $\Delta \ln(\widehat{AP}) \times \text{Exc.}$ interacts the simulated instrument with an indicator for whether the household ever used over 1,000 gallons per day in a billing cycle. All models include MY \times Zip-9 fixed effects, interacting the month-year with each postal route (646,124 fixed effects), “Weather Splines” in both growing degree days and precipitation with knots at each decile of the variable’s distribution, and “Past Q” controls where the percentiles of consumption six months prior interacted with the month-year (2,905 fixed effects).

Table 1.5: Full Period Price Elasticity Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Coefficients</i>						
$\Delta \ln(\widehat{AP})$	-1.479** (0.6842)	-0.695*** (0.2064)	-0.698*** (0.2294)	-0.960*** (0.1098)	-1.235*** (0.2903)	-2.231*** (0.5724)
$\Delta \ln(\widehat{AP}) \times \text{EUP}$	0.622 (0.658)	2.423*** (0.295)	2.419*** (0.536)	0.525*** (0.132)	0.558 (0.452)	3.818*** (0.819)
$\Delta \ln(\widehat{AP}) \times \text{Post}$	1.733** (0.736)	0.653*** (0.220)	0.632** (0.246)	0.801*** (0.120)	0.782*** (0.242)	1.717*** (0.450)
<i>Elasticities</i>						
Pre-EUP	-1.479** (0.684)	-0.695*** (0.206)	-0.698*** (0.229)	-0.960*** (0.110)	-1.235*** (0.290)	-2.231*** (0.572)
EUP	-0.857*** (0.106)	1.728*** (0.215)	1.721*** (0.541)	-0.435*** (0.151)	-0.677*** (0.286)	1.587*** (1.096)
Post	0.254 (0.199)	-0.042 (0.060)	-0.066 (0.063)	-0.159*** (0.106)	-0.454 (0.368)	-0.514 (0.418)
Weather Cntrl.	No	Yes	Yes	Yes	Yes	Yes
Past Q	No	Yes	Yes	Yes	Yes	Yes
MY + Zip-5 FE	No	No	Yes	No	No	Yes
MY \times Zip-9 FE	No	No	No	Yes	Yes	No
Sample	Full	Full	Full	Full	Ever 1K+	EUs
Observations	8.9M	8.9M	8.9M	8.9M	1.1M	149K
Adjusted R ²	0.031	0.092	0.092	0.133	0.316	0.297

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are two-way clustered by fiscal year and 5-digit zipcode. These models report price coefficients (top panel) and combined elasticities (bottom panel) for detached single-family homes in the San Francisco East Bay Area using data from April 1, 2012 through May 31, 2019. “EUP” denotes the period of time while the Excessive Use Penalty Ordinance (EUP) was in effect from July 1, 2015 through May 3, 2016 while “Post” denotes the period after the policy ended. The dependent variable $\Delta \ln(Q)$ measures the change in the log of gallons per day by a household between the current billing cycle and the same period in the previous year. $\Delta \ln(\widehat{AP})$ is the simulated log annual change in prices using consumption levels from six months earlier. *Weather Cntrl.* denotes cubic splines in growing degree days and total precipitation per billing cycle, with knots at each decile of the variables’ distributions. *Past Q* indicates percentiles of consumption six months prior interacted with the month-year (6,012 fixed effects). models either include month-year and Zip-5 or month-year times Zip-9 fixed effects, The “Ever 1K+” sample restricts the sample to households who ever use over 1,000 gallons per day while “EUs” only uses households that violated the Excessive Use Penalty Ordinance.

Table 1.6: Responses to Behavioral Policies and Comparable Price Changes

Moral Suasion					
Model	First-Year Effect	Period	Elasticity Estimate	Median AP	Necessary AP %↑
(1) Full	-0.115***	Pre-EUP	-1.090***	\$5.05	11.2%
(2) Full	-0.115***	EUP	-0.435***	\$7.05	30.3%
(3) Full	-0.115***	Post-EUP	-0.159***	\$7.09	107%
(4) Ever 1K+	-0.095***	EUP	-0.677**	\$6.73	15.1%
(5) EUs	-0.500***	EUP	1.587*	\$6.73	NA
Public Shame					
Model	First-Year Effect	Period	Elasticity Estimate	Median AP	Necessary AP %↑
(1) Full	-0.303***	Pre-EUP	-1.090***	\$5.05	32.1%
(2) Full	-0.303***	EUP	-0.435***	\$7.05	100.7%
(3) Full	-0.303***	Post-EUP	-0.159***	\$7.09	574%
(4) Ever 1K+	-0.231***	EUP	-0.677**	\$6.73	40.6%
(5) EUs	-0.244***	EUP	1.587*	\$6.73	NA

This table calculates the average price changes necessary to have induced conservation behavior equivalent to behavioral policy responses. All first-year ATTs and price elasticities are obtained from corresponding models reported in Figures 1.10-1.11 and Table 1.5, with median average prices calculated from appropriate samples. “Full” denotes estimates from Column (4) in Table 1.5 that utilizes the full analysis sample and includes all preferred controls. “Ever 1K+” uses observations for households that ever used over 1,000 gallons per day (Column 5) while “EUs” restrict the sample to households that violated the Excessive Use Penalty (Column 6). “Pre-EUP” elasticities utilize the coefficient on the un-interacted price instrument while “EUP” and “Post-EUP” add to it the interaction for the differences during and after the EUP, respectively. The necessary percent increases in average prices is calculated as $\% \Delta AP = \left(\frac{AP' - AP}{AP} \right) \times 100$, where $AP' = \exp(ATT/\eta + \ln AP)$ is the necessary average price to achieve the same conservation given a behavioral instrument first-year effect ATT and a price elasticity η .

Chapter 2

Food for Thought: Food Banks and the SNAP Benefit Cycle

2.1 Introduction

According to a 2021 USDA report, approximately 10.5% of households were food-insecure during the previous year and at some point did not have “enough food for an active, healthy life for all household members” (Coleman-Jensen et al. 2021). Nearly 60% of food-insecure households participate in one or more of the three federal nutrition assistance programs. Despite receiving federal benefits designed to provide access to a healthy diet, many households still do not have enough to eat: 45.4% of Supplemental Nutrition Assistance Program (SNAP) participants in 2017 still characterized themselves as food-insecure. Every month, millions of households must find alternative sources of food to supplement federal program benefits (Coleman-Jensen et al. 2021).

Food banks stand as a potentially key resource for food-insecure households. Even the USDA, the agency responsible for operating federal nutrition programs, recognizes the enormous role that food banks play in reinforcing the federal safety net. Figure 2.1, produced by the USDA Economic Research Service, plots food pantry use by food insecurity status and shows that, since 2012, over 25% of food-insecure households have used food pantries. For households on SNAP, who likely represent a more food-insecure portion of the distribution, food bank utilization is even higher. In 2003, nearly 90% of food pantry clients were eligible to participate in SNAP and half of all pantry clients had used SNAP benefits in the prior year (Briefel et al. 2003). In total, two-thirds of pantry clients received some form of federal food assistance.

The insufficiency of SNAP benefits to support a nutritious diet varies geographically. As SNAP is a federal program, its monthly payments are set consistently across states, varying primarily as a function of the federal poverty level (FPL). To be eligible for SNAP, households must generally be below 130% FPL. Although SNAP benefits are adjusted for national cost of living at the beginning of each fiscal year, no adjustments are permitted at

the state level.¹ Food prices differ from state to state, altering the *real* value of benefits. For a set level of nutritional quality, residents of regions with higher local food prices can afford a smaller amount of food out of benefits than program participants living in less expensive areas. The need to make trade-offs is heightened for participants with lower real benefits; healthcare utilization and children’s health have been shown to be margins that suffer (Hoynes, Bronchetti, and Christensen 2017).

This paper aims to detail the extent to which impoverished households use local nutritional assistance to supplement resources obtained from federal programs throughout the month. In particular, I focus on the ability for food bank access to counteract the negative impacts of cyclicity in federal benefit utilization. Over the last two decades, researchers have detailed the existence of a benefit utilization cycle among SNAP participants. This cycle constitutes a pattern of early benefit redemption, wherein a large share of households rapidly draw down or completely exhaust benefits soon after they are received. While many papers have confirmed the existence and persistence of the pattern, frequently referred to as the “SNAP benefit cycle,” little is known about the methods households use to mitigate the cycle’s negative impacts. This paper extends our understanding of potential coping strategies by providing rigorous analysis of a major potential mechanism: food bank utilization.

Using data from USDA’s Food Acquisition and Purchasing Survey (FoodAPS), I provide motivating evidence in favor of the important role food banks play in the national nutritional assistance system. I find that SNAP participants visit food banks less often at the start of the month and increase visitation in the second half of the month relative to non-SNAP households. Exploiting plausibly exogenous variation in timing since receipt of SNAP benefits, I show that sampled households who visit food banks exhibit smoother consumption trajectories for food-at-home purchases in comparison to non-food bank households, especially at the start of the benefit month. On average, SNAP households who used food banks during the survey period retained an additional \$50-54 in SNAP benefits (18-19% of the monthly total) beyond the day of benefit receipt relative to typical spending in the fourth week of the benefit month. While daily spending by non-food bank households falls by \$78 for the rest of the first benefit week, food bank users’ only decrease average spending by \$38-41. These differences are driven entirely by spending out of SNAP benefits, with minimal patterns observed in out of pocket spending for food-at-home.

The rest of this paper is organized as follows. Section 2.2 describes the state of federal and local nutritional assistance, and discusses relevant literature. Section 2.3 discusses the USDA FoodAPS instrument in more detail and summarizes survey data. Section 2.4 presents graphical evidence on food bank utilization and daily food-at-home expenditures. In Section 2.5 I present the empirical strategy. I report econometric results in Section 2.6. Section 2.7 concludes.

¹While maximum allotments are consistent for the 48 continental states, unique benefit limits exist for Alaska, Hawaii, Guam, and the US Virgin Islands. Alaska benefit levels even differentiate between urban and two levels of rural conditions.

2.2 Background

SNAP and CalFresh

Formerly known as the Food Stamp Program, the Supplemental Nutrition Assistance Program (SNAP) is the U.S. government’s primary tool for reducing food insecurity. SNAP has formed the core of the social safety net since its nationwide implementation in the 1970’s. SNAP is the largest USDA program, serving 40.3 million people in 2018 with an average monthly benefit of \$125.79 per person (USDA). SNAP operates at a much larger scale than other USDA nutrition programs; in FY 2014, the National School Lunch Program (\$9.6 billion), Special Supplemental Nutrition Program for Women, Infants, and Children (\$3.6 billion) and the School Breakfast Program (\$3.7 billion) combined for \$16.9 billion, four times less than the \$74.2 billion in program costs for SNAP (Hoynes and Schanzenbach 2015). SNAP benefits do not vary across states, and few program reforms have occurred over time.

CalFresh is the name of the California implementation of SNAP. As a federal program, SNAP eligibility and benefits are set at the national level. CalFresh receives USDA funds in order to manage applications, confirm eligibility, and distribute benefits. Enrollment is handled independently at the county level; the availability of offices and effort required to register varies geographically. All CalFresh applicants must complete an interview with their county office (generally conducted over the phone). Further, all applicants must provide “proof of your income, expenses, and other circumstances to see if you are eligible” (California Department of Social Services 2022). Required documents include identification (driver’s license, passport), proof of address, social security numbers, recent bank statements, and pay stubs for all working household members. Applicants are suggested to also bring documentation of housing, utility, and childcare costs, phone bills, and child support documentation. Further, all participants must re-certify eligibility with the county at the end of each certification period (typically 12 months, with 24 month periods for disabled or elderly members). Re-certification requires submission of Re-certification Application CF37, the SAR7 Eligibility Status Report, and generally completion of another interview. Difficulties in enrollment or re-certification could contribute to low participation in California relative to national levels – in 2016 California ranked 47th out of 50 states in SNAP participation among eligible residents (Cunnyngham 2020).

Figure 2.2 plots household participation by month for Fiscal Years 2012-2017 for CalFresh and the California Food Stamps Program (the state-provided food stamp program for non-citizens). Beginning in 2012, participation in the nutritional assistance programs grew substantially through 2015, peaking in December of that year at 2,142,921 households. Since 2015, participation in the state programs fell consistently over time, reaching 1,970,777 households in July 2017. Substantial variation in participation exists across counties as well; in 2013, 3.8% of the population in Marin County participated in CalFresh while 25.1% of all Tulare County residents were enrolled (California Budget and Policy Center 2014).

Although benefit levels are constant across households of the same size, variation does

exist in the timing of benefits. CalFresh benefits are distributed by electronic benefits transfer (EBT) card. The date of benefit disbursement is determined by the final digit of the CalFresh case number, beginning on the first of the month and ending on the tenth. As a result, by the time one SNAP household in Alameda County receives its benefits, another could have had access to benefits for a third of the month.

Food Banks

For those in need of nutrition assistance, food banks represent an alternative for or supplement to federal nutrition programs. Broadly, food banks are non-profit organizations designed to fight hunger in their communities. While food banks can distribute food directly to the hungry (giving either groceries or prepared meals), they more typically fill a distributor role. These food banks engage with a wide variety of partner agencies – food pantries, hot meal programs, senior centers, religious organizations, and other non-profits – and provide food resources to the agencies themselves, who then handle preparation and distribution to the hungry. Partner agencies are typically very local in nature, with services tailored to the “unique needs of the community” (Alameda County Community Food Bank 2022).

Food banks play a major role in reducing hunger in the United States. According to a 2012-13 survey of 200 food banks, 46.5 million people in 15.5 million households are reached each year – 1 in 7 people in the United States (Weinfield et al. 2014). In comparison, 47.6 million people participated in SNAP during the year. In FY 2018, SNAP reached 20.1 million households.

Food bank resources are more quickly and easily accessed in comparison to federal nutrition programs. Participation in SNAP requires application and eligibility confirmation, with later re-certification requirements required to retain benefits. In contrast, food bank partner agencies typically operate akin to a local public good, with resources provided to all who choose to visit during operating hours. Food banks often publish lists of partner agencies’ food programs and operating hours to their websites. More recently, several larger food banks (including the Alameda County Community Food Bank) have introduced emergency food web tools or helplines, designed to pair callers with groceries or hot meals on the day the call is made. While SNAP requires advanced planning and confirmation of eligibility to enroll in the program, food bank resources are readily accessible by members of the community when needed.

Although food banks aim to serve as complements to federal nutrition programs, some households utilize their resources as substitutes to such programs. Households living in areas without food banks compare the perceived benefits of SNAP participation to the utility costs of applying for and utilizing the benefits when considering program take-up. In areas with food banks, the choice set expands. Households now have the outside option of receiving food from food bank resources, an action potentially accompanied by its own utility costs. For households on the margin of participating in SNAP, the presence of a local food bank may reduce the incentive to take up the federal program. Food banks actively work to negate

any such complementary effect by reducing the utility costs of SNAP take-up.

Food banks desire – and in many cases actively work – to increase utilization of federal nutrition programs within their community. Many food banks engage in SNAP outreach, providing information on the program’s benefits and helping individuals navigate the application process. For example, the Alameda County Community Food Bank offers phone and web services for checking eligibility and completing applications, hosts community events to teach about the program, and even mails pre-filled applications to clients who are known to be eligible.

SNAP Benefit Cycle

The “SNAP benefit cycle” refers to the front-loaded benefit utilization patterns displayed by a non-negligible portion of program households. SNAP benefits are distributed on a monthly basis, with participants receiving physical stamps prior to the 1990’s and more recently transfers via electronic benefit card balances to be redeemed at approved retailers throughout the benefit month.² Maximum allotments range from \$192 per month for a one member household up to \$1,155 for a family of 8 (roughly \$144 per person), with the full amount given on the assigned transfer day. Roughly one-third of program participants receive the maximum benefit allotment for their household size, an amount “intended to be sufficient to cover their food purchases for the month” (Castner and Henke 2011). While many program participants draw down their balance uniformly over the benefit month, a considerable share exists that instead rapidly exhausts balances each month.

First detailed anecdotally, the extent and severity of the benefit cycle has more recently been confirmed through empirical studies. Wilde and Ranney (2000) provided the first credible evidence of national-level patterns. Using Consumer Expenditure Diary Survey (CEX) and Continuing Survey of Food Intake by Individuals (CSFII) 1989-91 data, the authors show a sizable spike in mean daily food-at-home expenditures per person in the first three days after benefit distribution, with no observable change in food away from home (Wilde and Ranney 2000). The food-at-home effect is concentrated among households who grocery shop infrequently (once per month), with smoother consumption observed for households who shop more frequently. Models using Nationwide Food Consumption Survey (NFCS) data find a 0.73% per day reduction in market value of food consumed over the last week, indicating a shift toward lower cost food over the benefit month (Shapiro 2005). Later work using two years’ worth of scanner data from three Nevada supermarkets found that reductions in food quantity purchased can explain the full decline in expenditures over the benefit cycle (Hastings and Washington 2010).³

²As the exact date of transfer varies across and even within states, the “benefit month” refers to the month-long period between benefit transfer. In California, the benefit month can begin between the first and the tenth of each month.

³Since SNAP benefits in Nevada are all distributed on the first of every month, the authors are able to know the elapsed time since benefit issuance without merging in administrative program data. However, none of the amount of monthly benefits, carryover balance, nor the amount of benefits spent at other stores

In 2011, the United States Department of Agriculture (USDA) provided internal confirmation of the cycle. The “Benefit Redemption Patterns in the Supplemental Nutrition Assistance Program” aim was twofold: first to analyze the impact of the 2009 SNAP expansions under the American Recovery and Reinvestment Act of 2009 (ARRA), before describing general differences in redemption patterns across space and time (Castner and Henke 2011). The authors find that households made more frequent (if slightly smaller) purchases following the 13.6% ARRA benefit increase, with the share of benefits redeemed at supermarkets and supercenters roughly constant. More importantly for understanding of the benefit cycle, the study compares redemption patterns by day and week since benefit issuance in fiscal year 2003 to those in 2009 (split into pre-ARRA months of October through March and post-ARRA months of April through September). A random sample of 10,000 households per state and month reveals that 62 percent of households redeemed between 51 and 100 percent of benefits in the first week after issuance. An average household redeemed 18-22% of monthly benefits on the first day, 56-62% within the first week, and 77-82% by the end of week two (Castner and Henke 2011). While these patterns are quite constant across households with different benefit amounts, rapid benefit depletion early in the month was more likely for households without children or elderly, with disabled family members, without external earning sources, and in counties with persistent poverty. Further, households in non-metropolitan areas and in the south or mid-Atlantic regions redeemed benefits more rapidly.

While valuable due to providing official evidence of the benefit cycle’s existence, the 2011 USDA study’s conclusions are narrowed due to limitations of the data and lack of controls for confounding factors. The Anti-Fraud Locator for EBT Redemption Transaction System (ALERT) files “do not contain a record of the amount or date of a household’s benefit issuance” which results in benefit amounts being imputed from sequential transactions. The sample is limited to 10,000 randomly selected households per month for each state, purely to limit computational intensity. When comparing redemption patterns across regions, states are differentiated based on 1) the size of their caseloads and 2) the number of stores per square mile. One can reasonably expect both these factors to be related to general economic conditions of a particular state. Further, the reporting of average values for entire states fails to account for the considerable heterogeneity in local conditions within states.

Since 2005 many papers have verified the existence of the SNAP benefit cycle, all the while adding to the list of affected dimensions. Table 2.1 summarizes a subset of these findings. Building on the already-discussed findings, affected households are found to make extensive margin adjustments in addition to modifying caloric intake. Using data from a nationwide survey, (Todd 2015) documents a decline of nearly half an eating occasion per day in the last week of the benefit cycle prior to the ARRA expansion. Confirming evidence is found using the Bureau of Labor Statistic’s Time Use Survey 2006-2008, where the probability of an entire day without food, perhaps the most extreme extensive margin response, increases

are known. Further, the lack of variation in SNAP distribution prevents cleanly identifying a SNAP timing effect separate from a general “first of the month” effect.

nearly 1% over the benefit month (Hamrick and Andrews 2016). Additionally, the timing of benefit distribution within the week can impact the caloric and nutrient makeup of monthly consumption baskets as well (Castellari et al. 2017). These intra-month cycles and their effects are seen across the country, despite substantial variation in disbursement schedules from state to state (Goldin, Homonoff, and Meckel 2022).

Evidence on cyclicalities in food consumption is consistent with patterns in food expenditures. Early work using the 1989-91 CSFII estimated a 0.45% decline in caloric intake per day since benefit receipt (Shapiro 2005). Declines in energy consumption over the benefit month are larger for working-age adults and those with lower real benefits (Todd and Gregory 2018). SNAP participants are also more likely than non-participants to forego all eating occurrences in a day, with the probability of skipping all daily meals increasing over the benefit month (Hamrick and Andrews 2016). Studies focused on women have found substantial declines over the benefit month in fruit, vegetable, and dairy consumption (Sanjeevi, Freeland-Graves, and George 2018), energy intake and consumption of milk products (Tarasuk, McIntyre, and Li 2007), and healthy eating scores among African American women (Kharmats et al. 2014).

Damon, King, and Leibtag (2013) examine the impact of distribution timing on food purchasing behavior. Combining Nielsen Homescan and CES data from 2003, low and high income households in areas with early SNAP disbursement are compared to those living in areas where benefit distribution is staggered. As the data do not identify wage income or participation in social safety net programs, the authors rely on variation generated by comparing groups split on two dimensions: high versus low income (top 25% of income distribution versus bottom 25%), and early versus staggered SNAP disbursement areas (first 10 days versus first 15 days). The authors find limited evidence that low-income households decrease their food expenditures in week 4 relative to the high-income households, primarily through a reduction in grocery store expenditures (Damon, King, and Leibtag 2013). While these results are consistent with those of Wilde and Ranney (2000) and Shapiro (2005), the lack of program participation data limits the authors' ability to attribute the pattern specifically to SNAP rather than other programs or more general economic conditions.

FoodAPS

Historically, the dearth of micro-level SNAP data has limited researchers' abilities to rigorously explore the purchasing and consumption behavior of SNAP participants at an individual or household level. In an attempt to fill this gap, the United States Department of Agriculture's Economic Research Service (ERS) conducted the National Household Food Acquisition and Purchase Survey (FoodAPS) from 2012 to 2013. The nationally-representative survey collected a one-week snapshot from households of "food purchased or otherwise acquired for consumption at home and away from home, including foods acquired through food and nutrition assistance programs" (USDA ERS 2022). Importantly, the survey gathered information on SNAP participation and benefit amounts, as well as which food purchases were made using these benefits.

Many recent studies have taken advantage of the panel structure of FoodAPS to confirm the existence of the SNAP benefit cycle and delve further into its extents and consequences; Table 2.2 summarizes these papers. Seemingly regardless of the chosen modeling techniques, researchers have confirmed cycles in food expenditures (Damon, King, and Leibtag 2013; Dorfman et al. 2018; Smith et al. 2016). Decreases in the purchasing of healthy and perishable foods accompany the expenditure declines throughout the benefit cycle, with a greater share of food obtained for free from schools (Whiteman, Chrisinger, and Hillier 2018; Smith et al. 2016). This is primarily a scale change, as the number of free food acquisition events stays largely constant throughout the month (Tiehen, Newman, and Kirilin 2017). SNAP participants appear more aware of their hardship shortly before and shortly after benefit disbursement: reporting of food hardship increases by 25 percentage points during the “salience” window of 3 days prior to and 4 days after benefit distribution (Gregory and Smith 2019).

One recent work focuses on possible mechanisms behind the benefit cycle. Smith et al. (2016) tests two main hypotheses: short-run impatience and income fungibility. If SNAP participants have strong preferences for consumption in the current period relative to later weeks, then spending will be higher when benefits are received. Resources are then likely to be low toward the end of the benefit month, leaving the individual at greater risk of negative income shocks (Smith et al. 2016). Alternatively, if SNAP participants treat SNAP income differently than cash income and respond to SNAP increases differently than paycheck increases, the higher marginal propensity to spend out of SNAP would exacerbate any perceived impatience. Analysis of the FoodAPS data suggest that the former is primarily at play, with households spending 96 cents of every food dollar on food at home the day that benefits are issued (regardless of the fund source); SNAP benefits exhibit a higher marginal propensity to spend on food-at-home, suggesting a lack of fungibility (Smith et al. 2016).

Coping Strategies During the Benefit Cycle

Families employ a range of tactics throughout the month to cope with the benefit cycle. Switching to lower cost and lower quality items, borrowing funds from family members, skipping or cutting meals, and foregoing other home or medical expenses are all strategies frequently employed by those troubled by food insecurity (Castner and Henke 2011). The use of food banks stands as a primary coping mechanism. In areas with well-developed food banks, households that have exhausted their monthly SNAP benefits can obtain a bundle of groceries to boost available calories and provide a source of energy to sustain themselves. Additionally, households who are well-informed of the benefits of local food banks know that an option exists for additional food resources, and can alter their SNAP benefit spending patterns in a way that accounts for the availability of food bank benefits.

Although the existence of the SNAP cycle and its persistence over time and space has been well documented, few studies have examined the extent to which household coping strategies can mitigate the impacts of premature SNAP benefit exhaustion. In particular, no studies have rigorously evaluated the interplay between local public goods, such as food banks, and

federal poverty programs. A recent working paper by Laurito and Schwartz (2019) focuses on the interaction *between* federal nutritional assistance programs, investigating how the national school lunch and breakfast programs attenuate food insecurity resulting from the SNAP benefit cycle. Through event study and panel fixed effects models on the USDA's FoodAPS dataset, the authors find evidence of a compensatory response within school meal take-up. School breakfast and lunch participation is found to be 22-23% higher among 11-18 year olds in the final days of the SNAP benefit month in comparison to the second week, and lowest immediately after benefit disbursement (Laurito and Schwartz 2019). School meal programs exhibit a social pressure dimension akin to food banks, with students reporting that avoiding teasing or bullying is a main contributor to lack of take-up (Woodward et al. 2015). Additionally, students are more likely to eat school meals if they observe their friends participating (Marples and Spillman 1995). Results for children ages 5-10 find no similar decrease in school meal participation and own meal provision following SNAP benefit receipt, potentially suggesting that needy households are employing other coping strategies to provide their youngest children with sufficient food.

2.3 The USDA Food Acquisition and Purchasing Survey (FoodAPS)

The USDA FoodAPS dataset has increased understanding of the purchasing behavior, nutritional intake, and food-insecurity of SNAP households; can it provide insights into the role that food banks play in the benefit cycle? Following discussion of the FoodAPS instrument, I summarize the characteristics of SNAP households by food bank utilization and discuss differences in average acquisition patterns.

Instrument Design

FoodAPS targeted all food acquisition and purchasing activity among four key groups: households receiving SNAP benefits, non-SNAP households with income below the federal poverty guideline, non-SNAP households with income between 100 and 185 percent of the poverty guideline, and non-SNAP households above 185% of the poverty guideline. For each household sampled, all household members were asked to complete journals on all acquisition events for both food-at-home and food-away-from-home. For each acquisition event, respondents noted the location of the event, whether the food was obtained for free, and all items obtained. In the case of food-at-home, barcodes for all food items brought into the home were scanned using a provided scanner, providing detailed UPC-level information on food obtained for consumption at home.

USDA's ERS designed the instrument to obtain a complete picture of food in the home; measuring free food was thus a necessity as "free food may make up a significant share of some households' weekly food consumption." Indeed a substantial share of sampled households obtained food for free during the observation week: in total, 22% of all reported food-at-

home acquisitions were free. The free food share was even higher among SNAP households, who reported 30% of their food acquisitions coming from free sources. Survey workbooks reminded respondents to indicate whenever free food was brought into the home and take detailed records of those items. Households scanned barcodes for all free packaged goods, while items without barcodes were described on blue workbook pages. Figure 2.3 provides an example of a blue food-at-home (FAH) workbook page. Figure 2.4 shows how food banks and pantries were included as one of ten main food acquisition locations. Presence of a unique food bank barcode suggests that the USDA anticipated notable use of emergency food networks by respondents when designing the survey.

In addition to documenting food-at-home acquisitions, FoodAPS obtained both self-reported and administrative measures of contemporaneous SNAP participation. During the initial interview, each household's primary respondent was asked if they or any members of the household were currently enrolled in SNAP, or if any household members had participated in the last 12 months. If a respondent answered affirmatively to either question, they were asked for the date of last benefit receipt. Households were also asked for consent to link self-reported SNAP participation to state-level administrative files.⁴ Consenting households were linked to two types of SNAP administrative data: caseload files and the Anti-fraud Locator using EBT Retailer Transactions (ALERT) system. Caseload data containing SNAP case identification numbers, names, addresses, benefit allotments, and issuance dates for the period March through November 2012 were used for matching. The ALERT data contains records for each use of an EBT card, and included location, card number, purchase amounts, and remaining balances.⁵ In total, current SNAP participation was confirmed through administrative records for 1,308 households.⁶

To provide variation in timing since receipt of SNAP benefits, participating households were surveyed at various points throughout the month. While each household reported food acquisition events for one contiguous week, the time both during the week and in the sampling period were selected at random. Figure 2.5 illustrates the variation in sampling period and the source of variation in the days since SNAP benefit receipt. For sampled SNAP households, the selection of the survey week was independent of when SNAP benefits are received. As a result, composing a measure of days since disbursement of SNAP benefits – or “Days Since SNAP” – results in two dimensions of variation. First, two SNAP participants who receive their benefits on the same day could have been sampled in different weeks, resulting in different “Days Since SNAP” for each of their observation days (the yellow line represents one possible observation day for the calculation). Second, SNAP participants sampled on the same week of the month may have different “Days Since SNAP” due to differences in the date of benefit disbursement. In states where all benefits are distributed on the same day, this variation comes from comparison to residents of other states with

⁴122 households who reported SNAP participation did not consent to administrative matching.

⁵While SNAP issuance dates are not directly observed in the ALERT data, it can be approximated well using increases in remaining balances between transactions, given the frequency in transactions made by households in the sample.

⁶Current participation was defined as receipt of benefits within 36 days of the survey week's end.

different disbursement schedules. In states like California, where benefits are distributed randomly across a period of the month, this dimension of variation does not depend on spatial variation and can arise between two SNAP participants in the same census block.

Characteristics of Food Purchasing and Acquisition

To analyze the effect of food bank utilization among FoodAPS households, I combine several of the instrument's datasets. I begin with the Individual file (FI), which contains detailed demographic information for all family members. These records report general information on age, employment, household role, health, as well as children's school attendance and participation in federal nutrition programs. Individual-level records are matched to the household characteristics using the Household file (FH), which reports family size, household-aggregated income measures, and describes assets held by the family. The Household file also reports current SNAP participation along with the date and amount of last benefits, allowing construction of a measure for the number of days since SNAP benefit receipt.

The analysis data are completed by incorporating purchase and acquisition information from the food-at-home Event file (FAHF). The event file contains one record for every food-at-home acquisition event during the survey and combines information obtained from survey workbooks, handheld scanners, and purchase receipts. Each record reports the event date, the type of location (i.e. supermarket, big box store), distance and travel time to that location, and payment method. Paid events include the total expenditure for the event as well as dollar amount of payments made out of SNAP, WIC, or other federal nutrition assistance programs. Free events were determined when respondents checked the "free" workbook page box or reported a \$0 event total. Use of food banks or pantries was determined through use of the corresponding place barcode and based on the place name on the workbook page. To conduct my analysis, I aggregate the event file to the survey-day level and merge it with the FI and FH files. After removing records for erroneously-reported food-away-from-home events and limiting the sample to confirmed SNAP participants, I am left with records for 4,799 events from 1,216 households.⁷

Table 2.3 summarizes characteristics of surveyed SNAP-enrolled households, splitting the sample according to whether families used food banks. On average, households contained slightly more than three family members. Households that did not use food banks were split fairly evenly between adults (1.7) and children (1.4), while food bank households exhibited a higher adult-to-child ratio (1.9 to 0.9). Household incomes averaged just over \$2,000 per month – or just below 120% of the federal poverty guideline – for both groups. Estimated monthly SNAP benefits average \$278 for the food bank sample and \$343 for the non-food bank households, likely reflecting the larger mean family size and greater number of children.⁸

⁷SNAP households were one of four targeted groups and represented 32.76% of all sampled households. Event records correspond to 3,462 unique household-days.

⁸Although SNAP participation is confirmed using state administrative data, benefit amounts are not directly observed. As a result, the USDA ERS uses four different specifications of the MATH SIPP+ Microsimulation Model of SNAP eligibility to estimate monthly benefits. The first two runs assumed that

Primary respondents are overwhelmingly female and unmarried, with mixed education levels concentrated below a four-year college degree. Over half of these respondents identify as white, followed by 20% black, 14-20% other race, and only 2-4% asian.

While households using food banks appear comparable to other sampled households based on gender, race, and education, differences do arise when considering household assets and food security measures. Although similar shares of respondents in each group own their own home, live in public housing, or are self-employed, a substantially smaller share of food bank households own a vehicle (48%) in comparison to those who never use food banks (72%). Additionally, households that visit food banks are very food-insecure at nearly twice the overall rate, with 36% of households reporting very low food security and 72% agreeing that they ran out of food in the last month and were unable to afford a balanced diet (compared to half the non-food bank sample). While these disparities are suggestive of households turning to food banks when food-insecurity becomes too high, it does not account for patterns in income flows and time preferences throughout the month.

Turning to the nature and timing of food acquisition events reveals several patterns. Table 2.4 reports characteristics of every food acquisition event among surveyed SNAP households. Acquisitions are spaced roughly evenly throughout the survey week, with 20% of events taking place on the first weekday. While the share of acquisitions from dollar stores, convenience stores, and superstores is roughly even across groups, food bank households do a disproportionate share of their shopping from grocery stores or supermarkets that accept SNAP. However, nearly one-third of food acquisition events for food bank households are without charge, with 19.4% of their events involving accessing food banks. Given the frequency of free food acquisitions, it is not surprising that food bank households on average spend half as much per acquisition as non-food bank households (\$16.27 versus \$30.96) and use SNAP half as often (23.1% of the time versus 49.6%).

Aggregating event information to the household-day level shows further differences in how and when households acquire food. Table 2.5 reports daily aggregates for event volumes, expenditures, and timing relative to SNAP benefit distribution. Aggregating purchases and acquisitions to the daily level accounts for the fact that households frequently make more than one acquisition per day and choose stores in a complementary manner. On average, households obtain food-at-home 1.4 to 1.5 times per day. Food bank households spent just over \$23 per day on food-at-home, with \$10 coming out of SNAP benefits and \$13 out of pocket. Non-food bank households spent considerably more on average, with total daily expenditures of \$42 with \$24 out of SNAP. Households who used food banks were also considerably more likely to obtain food for free, whether it be from food banks or other sources. 41% of days with food acquisitions involved at least one free event for food bank

all household members belong to the same SNAP unit, while runs 3 and 4 behave more realistically and allow for multiple SNAP units within a household. I choose to report results from model 3, which uses reported income as the basis for calculating benefits – model 4 instead multiplies reported net earnings by a factor of 1.4 to approximate gross earnings (USDA ERS). I consider this to be the more conservative estimate of benefits, and rely on the assumption that there are no systematic differences in reporting error between households who used food banks and those that did not.

users, while other households obtained free food on only 5% of days. Trips on weekdays and holidays are similar between groups (0.72-0.75 events for weekdays, 0.02-0.03 on holidays). Most informatively, food bank users made substantially fewer acquisitions early in the SNAP benefit month. Only 12% of their acquisition days took place in the first week after SNAP benefit disbursement, in contrast to 28% for non-food bank users.

A consistent theme seen across these tables summarizing the FoodAPS data is a dearth of food bank use. In total, 4,826 households were sampled. SNAP participants make up only a subsample of respondents: program participation could only be confirmed for 1,216 households. Although this group represents the most food-insecure subsample, few use food banks during the survey week. Only 25 unique households access food banks during the seven observed days, resulting in 139 acquisition events over 93 household-days by households that used food banks. While the timing of SNAP benefit distribution is still plausibly exogenous for these families, there is a substantial lack of identifying variation from which effects of food bank access can be estimated. Therefore all empirical findings that follow should be considered representative of only the surveyed households and should not be generalized to the broader SNAP population.

2.4 Food Bank Utilization and Food-At-Home Expenditures, Visualized

Tables 2.3-2.5 provided preliminary evidence on potential differences between the characteristics and shopping patterns of food bank-using and non-food bank households. In this section, I present graphical analyses that explore the dynamic behavior of surveyed households throughout the SNAP benefit month. These figures illustrate that food bank use occurred within the FoodAPS household by both SNAP and non-SNAP households, with variation both over the month and between the groups. Further, once I concentrate on the sample of SNAP households, I observe clear patterns of immediate benefit draw-down. When splitting the SNAP sample according to food bank use, this effect largely disappears for households who took advantage of these local nutrition assistance resources.

Figure 2.6 plots the average number of food bank visits per week for SNAP and non-SNAP households who ever use food banks in the sample by week in the benefit month. Note that, since we are including non-SNAP households, we are limited to just one source of variation in this plot (as only households on SNAP have a benefit month that may differ from the calendar month). While food bank utilization is on average low, those who visited a food bank display considerably higher overall utilization at the start and end of a month. On average, households sampled in the first week of the month visited food banks 0.9 to 1.05 times per week for SNAP and non-SNAP households, respectively. These numbers fall for both groups in the second week (roughly 0.75 weekly visits for both groups) before rising in both the subsequent weeks for non-SNAP households. In contrast, SNAP households' usage of food banks continue to fall through week 3 (0.71 visits) before rising again in the fourth

week (0.84 visits).

Next, I focus on the SNAP portion of the sample and expenditures over the benefit month to explore patterns in daily expenditures by payment source. Figure 2.7 plots the variation in daily food-at-home expenditures made using SNAP benefits and non-SNAP payment sources throughout the benefit month. The green line plots expenditures out of SNAP benefits, while the red series shows out of pocket expenditures. Both series are measured relative to the date of benefit disbursement: SNAP benefits are received on day 0, while 14 indicates that two weeks have elapsed since the household received their benefits. On the day of benefit issuance, SNAP households on average spend over \$125 out of their monthly benefits – roughly 20% of maximum benefits for a family of four. Expenditures out of SNAP fall dramatically on day two to roughly \$50, and continue on a gradual decline until the final days of the benefit month. This striking early behavior is consistent with hyperbolic discounting, while increased expenditures in the final week of the benefit month may represent exhaustion of a buffer stock left for emergency purposes (Shapiro 2005). out of pocket expenditures appear much more consistent throughout the month and display a mild compensatory response to SNAP benefit exhaustion. Daily expenditures on food-at-home from non-SNAP sources begin at roughly \$22 per day on the date that SNAP benefits are distributed, display a slight positive trend with more variation in the middle of the month, and finish around \$25 per day at the end of the benefit month. Here we see evidence that spending out of pocket to compensate for exhausted SNAP benefits is not a primary strategy employed by the average SNAP respondent.

Splitting the SNAP sample according to food bank use suggests a stark difference in front-loaded food-at-home expenditures. Figure 2.8 provides a first look into whether food banks may stand as a possible coping mechanism to smooth consumption throughout the month and avoid rapid exhaustion of SNAP benefits. Here I focus on total food-at-home expenditures over the course of the benefit month, now splitting our respondents according to whether they accessed a food bank at any point during the observation week. The green series plots average daily food expenditures by the number of days since SNAP benefits were received for households that did not utilize food banks, while the red series reports equivalent daily expenditures for food bank users. Given previous studies of food bank users have found substantial stigma to first utilization of food banks that diminishes soon after use, the red series represents at least a subset of households who have experienced what food bank benefits looked liked, improved their information set, and made the decision to continue utilizing those resources. Non-food bank households display a nearly identical pattern as before, spending over \$150 on food-at-home on the day they received SNAP benefits. These households dramatically reduce daily expenditures beginning the next day, with expenditures remaining roughly constant through the remainder of the benefit month. In contrast, households who used food banks only spent approximately \$68 *combined* out of SNAP and other sources on the first day of the benefit month. Expenditures by food bank households peak on the fourth day of the benefit month and fall to below \$10 per day at the end of the first week. By the second week of the benefit month, food bank users demonstrate a near linear trend, spending roughly \$20 on food-at-home from all sources each day.

2.5 Empirical Strategy

I next use information on daily expenditures and free food acquisition events to determine whether counter-cycles exist in expenditure patterns for households using food banks. The following analysis takes advantage of plausibly exogenous variation in benefit issuance and survey timing to test whether food bank resources mitigate the SNAP benefit cycle effects on food expenditures and caloric intake. Results suggest that SNAP households who utilized food banks during the sample week did not exhibit the same level of benefit exhaustion on the day of benefit issuance. To harness the variation in SNAP benefit receipt and FoodAPS data collection timing, I estimate variants of the following empirical model of daily food expenditures:

$$Exp_{idw} = \alpha_i + \mathbf{Days\ Since\ Snap}'_{idw} \boldsymbol{\theta} + \boldsymbol{\tau}_{dw} + \mathbf{X}'\boldsymbol{\gamma}_i + \mathbf{R}_i + \epsilon_{idw} \quad (2.1)$$

The dependent variable Exp_{idw} measures the daily food expenditures – either total expenditures, expenditures out of SNAP benefits, or out of pocket expenditures – by household i on day d of week w . The vector $\mathbf{Days\ Since\ Snap}_{idw}$ consists of a set of indicator variables for purchases that occur a certain number of days since SNAP benefits were received by the household. $\boldsymbol{\tau}$ is a vector of time controls, including day-of-week and week-of-year fixed effects as well as an indicator for whether the surveyed day is a federal holiday. To counteract potential survey fatigue, I also include controls for each day of the survey week. \mathbf{X} controls for a vector of household characteristics: age, marital status, gender, education level, and race of the primary respondent, along with combined household income and counts for the number of adults, kids, and seniors in the home. \mathbf{R} includes census region fixed effects to account for unobserved differences in expenditure patterns that vary across regions of the country. Alternate specifications rely on the panel nature of the survey to soak up unobserved heterogeneity through household fixed effects, \mathbf{HH}_i , in place of the region and household controls. ϵ_{idw} is an idiosyncratic shock comprised of all remaining sources of variation in daily food-away-from-home expenditures.

In this specification, $\boldsymbol{\theta}$ are the parameters of interest and measure the average changes in daily expenditures throughout the SNAP benefit month. To start, I replicate preferred semi-parametric bin specifications common to the FoodAPS literature. I also test the hypothesis of a “Day Zero Break” through a more parsimonious specification. Across all models I omit the fourth week of the benefit month, resulting in $\hat{\theta}$ coefficients being interpreted relative to the final week of the benefit month. Early benefit draw-down will be reflected in large, positive coefficient estimates for the day of benefit receipt (day zero), and potentially positive coefficients for the remaining days of the first week of the benefit month (days 1-6). Perfectly smooth consumption patterns would appear as zero coefficient estimates for all periods relative to week 4.

Identification of $\boldsymbol{\theta}$ relies on the randomization of survey timing and utilized control structure. As households surveyed by FoodAPS were monitored for one randomly-selected week

during the survey period, SNAP-participating households had an equal probability of being observed at the start of the benefit month as at later points. This randomization ensures that observable and unobservable household characteristics are uncorrelated with a household's location in the benefit month at the time of survey administration. Additionally, I leverage the plausible exogeneity in the timing of SNAP benefit disbursement within and across states to provide variation in households' positions within the benefit month even if located in the same census region. Stated another way, after accounting for time controls, census region, and household characteristics (or household fixed effects), I require that the idiosyncratic error term ϵ_{idw} is not systematically correlated with the number of days that have elapsed since a household received their SNAP benefits.

To estimate the potential consumption-smoothing effect of food banks, I modify Equation 2.1 by allowing the expenditure effect of time elapsed since SNAP benefit receipt to differ for households that visit food banks:

$$Exp_{idw} = \alpha_i + \mathbf{Days\ Since\ SNAP}'_{idw} \boldsymbol{\theta} + \mathbf{Days\ Since\ x\ Food\ Bank}'_{idw} \boldsymbol{\beta} + \boldsymbol{\tau}_{dw} + \mathbf{X}'_i \boldsymbol{\gamma}_i + \mathbf{R}_i + \epsilon_{idw} \quad (2.2)$$

The vector **Days Since x Food Bank** interacts each of the **Days Since SNAP** variables with an indicator for household food bank utilization. In Section 2.6 I present results using both a measure of immediate utilization during the survey week and recent utilization within the previous 30 days. The coefficients $\boldsymbol{\beta}$ measure the difference in average daily expenditure X days since benefit receipt for households that utilize food bank resources relative to those that do not.

To start, I replicate preferred semi-parametric bin specifications common to the FoodAPS literature. I also test the hypothesis of a “Day Zero Break” through a more parsimonious specification. Across all models I omit the fourth week of the benefit month, resulting in $\hat{\beta}$ coefficients being interpreted relative to the final week of the benefit month. To identify the $\hat{\beta}$ coefficients, I rely on the plausible exogeneity in the timing of SNAP benefit disbursement day assignments and benefit from the random choice of sampling weeks within the FoodAPS instrument. That is, after accounting for time controls, census region, and household characteristics (or household fixed effects), the idiosyncratic error ϵ_{idw} is not systematically correlated with the number of days that have elapsed since a household received their SNAP benefits.

2.6 Results

I next present results characterizing the consumption patterns observed throughout the SNAP benefit cycle and the impacts of food bank use on these patterns. I first replicate prior studies by examining the cyclical patterns present in food-at-home expenditures by SNAP households. Next, I present results obtained from estimating variations of Eq. 2.2 that

explore the extent to which food bank utilization – either during the survey week or within the previous 30 days – modifies the cyclical patterns observed in food-at-home expenditures.

Food-at-Home Expenditures and the SNAP Benefit Cycle

Turning first to overall patterns in food-at-home expenditures throughout the SNAP benefit cycle, I find evidence of greatly-increased expenditures immediately upon benefit receipt that matches prior studies. Table 2.6 reports results obtained from estimating Eq. 2.1 for the sample of currently-enrolled SNAP households surveyed in the FoodAPS instrument. Column (1) presents coefficient estimates and standard errors from the regression of daily food-at-home expenditures in U.S. dollars on four measures of the time elapsed since SNAP benefit receipt. $0\ DSS$ and $1-6\ DSS$ are dummy variables equal to one when households are on the day of or 1-6 days since benefit receipt, respectively. $Week\ 2$ ($Week\ 3$) indicates that a household is in the second (third) week of the benefit month. I cluster standard errors at the household level to account for potential differences in error variances arising within the household. Column (2) introduces the vector of household characteristics \mathbf{X} , while Column (3) further adds time controls τ_{dw} . Column (4) extends Column (3) with the addition of census region fixed effects, while Column (5) utilizes household fixed effects. Columns (6) and (7) repeat the specification of Column (5), but focus on specific payment sources. Column (6) estimates Eq. 2.1 using expenditures out of SNAP benefits as the dependent variable, while Column (7) uses all out of pocket (OOP) expenditures.

All specifications provide evidence of extensive shopping on the day that benefits are received, followed by immediate drops in spending that persist throughout the benefit month. Results from Column (1) indicate that, when a household received their SNAP benefits, they spent an average of nearly \$92 more than an average day in the final week of the SNAP benefit month. The point estimate for this day-zero spike stays stable between \$90.42-\$92.26 in Columns (3)-(5) as controls are added and maintain high statistical significance, confirming the extensive expenditures immediately upon receipt of SNAP benefits. Results in Columns (6) and (7) show that this spending is entirely driven by purchasing out of SNAP. On the day benefits are received, SNAP expenditure increase by the entire \$91 while out of pocket expenditures are indistinguishable from typical patterns at the end of the benefit week. After the date of benefit receipt, expenditures immediately fall; average food-at-home expenditures on days 1-6 of the benefit month drop by roughly \$75, to between \$13.78 and \$16.68 higher than in the fourth week and are again entirely driven by expenditures out of SNAP benefits. All coefficient estimates within the first week of receipt in Columns (1)-(6) are highly statistically significant.

I observe weaker evidence that household food-at-home expenditures fall during the middle half of the benefit month. Point estimates indicate week 2 daily expenditures different than those in Week 4 by $-\$5.33$ to $\$0.98$, with the largest decrease observed in Column (5) when household fixed effects are included (also the only point estimate statistically significant at the 10% level). More consistent estimates and significance are obtained for the Week 3 effect, with daily purchases estimated to be between $\$3.60$ and $\$5.32$ lower than in Week

4, with 5/7 coefficients statistically significant at the 10% level. Once again no evidence of changes in expenditure patterns are observed for out of pocket expenditures.

These findings parallel prior analyses of the SNAP benefit cycle using the FoodAPS instrument. Econometric analyses have documented that the propensity to spend out of SNAP falls sharply after the day of benefit receipt (Smith et al. 2016), with a large share of SNAP participants conducting the bulk of their monthly food acquisition events on this day (Whiteman, Chrisinger, and Hillier 2018; Dorfman et al. 2018). Evidence that take-up of other nutritional assistance programs – including both free school breakfasts and lunches – increases sharply among school-age children in SNAP households during the final half of the benefit month (Laurito and Schwartz 2019) motivates the following inquiry into whether utilization of local nutrition assistance resources modifies behavior throughout the benefit month.

Food Bank Use and the SNAP Benefit Cycle

Results in Table 2.7 explore the extent to which households utilizing food banks display different temporal patterns in food-at-home expenditures. Columns (1) and (2) present coefficient estimates obtained from estimating Eq. 2.2 where “Days Since SNAP” variables are interacted with an indicator for whether a household obtained food from a food bank during the survey week (Current FB Use). Columns (3)-(4) present comparable estimates for households that used food banks within the previous thirty days (Previous FB Use). Columns (1) and (3) match the specification of Column (4) in Table 2.6 and includes all household controls, time controls, and census region fixed effects. Columns (2) and (4) parallel Column (5) of Table 2.6, replacing region fixed effects with household-specific ones.

Estimates for current food bank utilization suggest that, while overall daily expenditures are lower for food bank households, a greater share of food-at-home purchasing occurs later in the first benefit week – with weaker evidence of smoother consumption further into the benefit month. The estimates of Column (1) match the spending patterns from Column (4) of Table 2.6 for non-food bank households, with extensive acquisition occurring on the date that households receive their SNAP benefits (\$90.87-\$94.90 additional daily expenditure, highly statistically significant across models). Expenditures fall sharply for non-food bank households across the remaining days in the first week of the benefit month but remain \$12.79-\$13.67 higher than in the fourth week. Smaller magnitude decreases are estimated in Weeks 2 and 3 for these families, with point estimates between -\$2.33 and -\$5.08, with only 2/8 coefficients statistically significant at the 10% level.

In contrast, current food bank-utilizing households display a smoother expenditure pattern throughout the first week of the benefit month. Households that utilized food banks during the survey week in Columns (1)-(2) spent between \$49.58 and \$53.51 less on the day of benefit receipt. This effect is highly statistically significant and more than halves the day zero spike observed by other SNAP participants. The pattern reverses for the remaining days in the first week of the benefit month, with current food bank households spending between \$28.15-\$41.44 more per day than other SNAP households 1-6 days since benefit receipt. I

observe no evidence of additional expenditure smoothing in either the second or third benefit week for these households. Further, no evidence of similar expenditure smoothing is observed in Columns (3) and (4) for households who visited food banks in the previous thirty days rather than during the survey week. Across both models, none of the benefit timing interaction terms are statistically distinguishable from zero, suggesting that food bank acquisitions likely serve to offset purchases in the short-run rather than provide a long-run backstop of non-perishable food supplies.

Table 2.8 explores how the evidence on food bank's role in the SNAP benefit cycle differs with alternate timing specifications. Table 2.8 presents estimates of Eq. 2.2 using two alternate "Days Since SNAP" specifications. Columns (1) and (2) specify weekly bins for being one, two, and three weeks since SNAP benefit disbursement, with interactions for current and previous food bank visitation, respectively. Column (3) tests a "Day Zero Break" approach, specifying a linear time trend since benefits were received and a dummy variable for the day of receipt. The extent to which the $Days\ Since \times FB$ and $0\ DSS \times Prev\ FB$ coefficients differ from zero indicate the differential linear trend or alternate behavior on the day of benefit receipt for food bank households. All columns include the full set of time controls τ_{dw} and household fixed effects.

Estimates for the weekly specifications from Columns (1) and (2) provide further evidence that the expenditure-smoothing behavior of food bank households primarily occurs within the first week rather than later in the benefit month. Across columns (1) and (2), coefficients on un-interacted weekly indicators show high initial spending in week one with expenditures monotonically decreasing throughout the benefit month. Relative to benefit week 4, daily expenditures are estimated to be \$48.36-\$49.80 higher in week 1, \$31.24-\$31.92 higher in week 2, and \$16.86-\$18.01 higher in week 3 for non-food bank households (all coefficients statistically significant beyond the 1% level). In contrast, all interaction terms for food bank households other than week 3 in Column (1) are statistically indistinguishable from zero with negative point estimate reflective of the lower overall spending by food bank households.

The *Day Zero Break* specification in Columns (3) and (4) provides additional evidence of the differential behavior of food bank households on the day of benefit receipt. For households who do not use food banks, total expenditures on food-at-home in Columns (3) and (4) are predicted to be \$73.56 to \$76.90 higher than any other day in the benefit month's deviation from the linear trend of -\$0.81 per day. In contrast, in Column (3) the full "Day Zero" expenditure shift for households who used food banks during the survey week is negative, with food bank households spending below the implied linear trend on the day of benefit receipt, with no observed difference in the linear trend. I measure no differential patterns among households that utilized food banks within the previous thirty days, suggesting once again that the expenditure-smoothing impact of food banks is relatively short-run.

Table 2.9 focuses on patterns present across payment methods. All columns report coefficient estimates and standard errors from estimation of Eq. 2.2, where total expenditures are replaced as the dependent variable with expenditures out of SNAP benefits in Columns (1)-(2) and out of pocket expenditures in Columns (3)-(4). All estimates include interactions for "Current" use of food banks during the survey week and include the full set of time con-

trols. Prior analysis in Table 2.6 showed that spending out of SNAP drove the expenditure patterns throughout the benefit month for all participating households surveyed. Models in Table 2.9 extend this evidence to allow for differences in patterns by food bank use.

Estimates presented in Table 2.9 again confirm that the early spikes in daily expenditures are driven solely by spending out of SNAP benefits. Coefficient estimates for non-food bank households in Columns (1) and (2) match those observed in Table 2.6 and once again show extensive draw-down immediately upon benefit receipt, with minimal cyclicity in out of pocket expenditures observed. Households that utilized food banks during the survey week show decreased spending out of SNAP of \$34.60-\$42.14 on the day of benefit issuance. I find no statistically significant differences in weeks 2 and 3. Weaker evidence of patterns is observed for out of pocket spending, with only the $1-6 \text{ DSS} \times \text{Used FB}$ coefficient of \$17.14-\$18.94 significant to at least the 10% level in both Columns (3) and (4). Point estimate magnitudes are comparable to the negative day zero differential, matching earlier findings that suggest smoothing primarily within the first week of the benefit cycle.

2.7 Conclusion

Food banks are a powerful, local source of nutritional assistance, yet relatively little is understood about the role their resources play in the SNAP benefit cycle. In this paper I provide evidence suggesting differential behavior patterns immediately upon receipt of benefits for SNAP households that utilize food banks. To obtain rich information on food purchasing and acquisition behavior by SNAP-enrolled households across the United States, I employ data from the USDA's FoodAPS instrument. After summarizing prior studies documenting the presence of the SNAP benefit cycle and the expenditure patterns within the FoodAPS sample, I discuss the observed differences in household characteristics and behavior by food bank use. Next, I present visual evidence of variation in food bank utilization over time among surveyed households and stark differences in day-zero SNAP benefit utilization for households that visited food banks during the survey period.

Following the presentation of visual evidence, I leverage variation in timing since SNAP benefit receipt and food bank use among surveyed households to estimate empirical models of daily food-at-home expenditures. I find that, while typical SNAP households spend \$91-92 more on the day of benefit receipt than in the final week of the benefit month, households that use food banks spend \$50-54 less that day. This decrease in day-zero expenditure by food bank clients is offset by higher average spending on the next six days, indicating that the expenditure smoothing exhibited by these households primarily occurs in the first week of the benefit month. Subsequent analyses exploring alternate timing specifications and payment methods confirm that the observed expenditure patterns are driven by payments out of SNAP benefits, and that the expenditure smoothing observed by food bank households is largely isolated to the first week following benefit issuance.

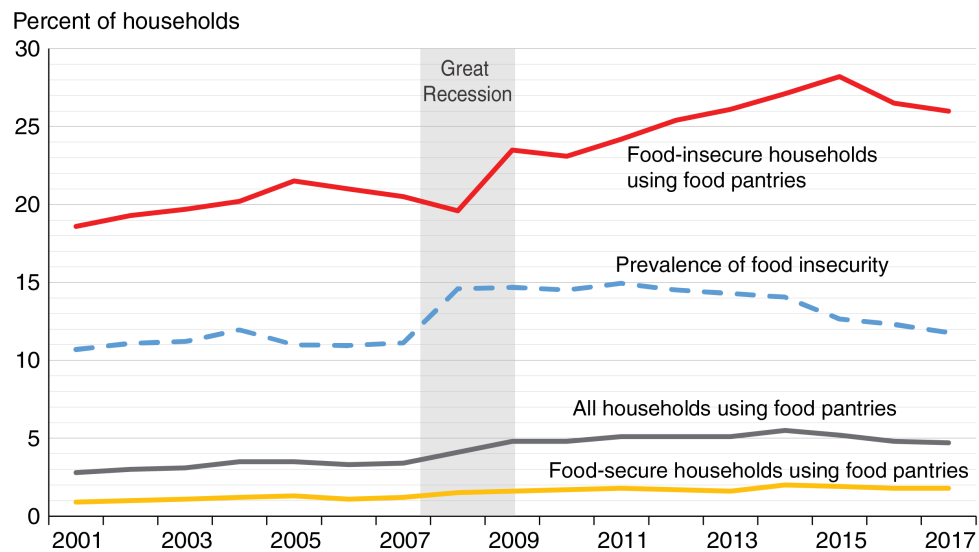
While the presented evidence is suggestive of an expenditure-smoothing role of food bank use for SNAP households, the generalization of these findings is constrained by coverage

limitations of the FoodAPS instrument. FoodAPS was highly successful in its ability to capture a complete picture of food purchasing and acquisition behavior by households. In particular, its focus on SNAP participants as one of its four key groups and free food as an important type of acquisition event resulted in the best-to-date representation of how food-needy households substitute between paid and free events throughout the benefit month. While the focus on free food events provides a representative sample of households that used food banks during the sample period, food bank users were not a key group of interest and represent only a small portion of the total sample. Further, the factors determining households' decisions to visit food pantries are largely unobserved, and are not brought about by exposure to a policy or experiment. The lack of a source of quasi-random variation in food bank use places strong identifying assumptions on the estimated models for identification of the causal impacts of food bank use on food-at-home expenditures. Instead, the estimates presented herein should be taken primarily as suggestive evidence and motivation for the construction of new datasets targeted at characterizing food bank use and the sources of variation in food bank take-up.

Future research can continue to expand our knowledge of the role food banks play in the domestic nutrition assistance landscape. Studies that can improve our understanding of the behavioral barriers to take-up of either local food assistance resources or federal aid programs will help target outreach and provide important insights for tackling perceptions of social stigma. Direct partnership with food banks to study patterns in demand for local assistance can expand our knowledge of the role these resources play in combating short-term or chronic hunger within their communities. In addition, future studies that link food bank utilization with SNAP enrollment and follow households throughout the benefit month can provide insight into the role of food banks as substitute for or complement to federal assistance.

Figure 2.1: Food Pantry Use for Food-insecure Households in 2017

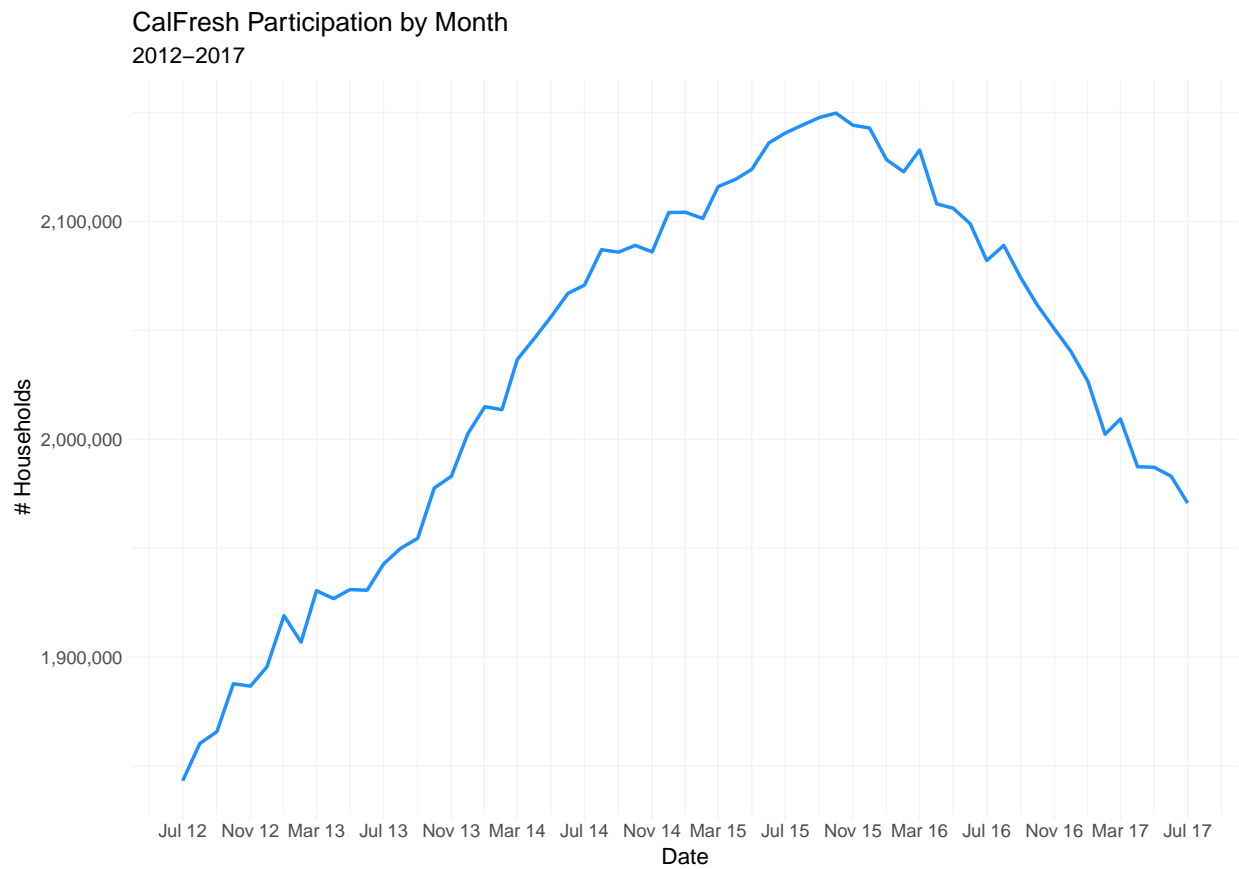
Food pantry use for food-insecure households in 2017 was about five times the rate for all U.S. households



Source: USDA, Economic Research Service using data from U.S. Department of Commerce, U.S. Census Bureau, 2001-2017 Current Population Survey Food Security Supplements.

Source: USDA Economic Research Service.

Figure 2.2: CalFresh and CFAP Participation



This figure plots the number of households participating in CalFresh, the California implementation of SNAP, by month from July 2012 through July 2017. Data Source: California Department of Social Services.

Figure 2.3: FoodAPS Primary Respondent Workbook, Food-at-Home Entry

Foods and Drinks Brought into the Home

Complete one **BLUE** page for each **PLACE** where you got food that you brought home

(√) DAY you brought food home	<input type="checkbox"/> Mon	<input type="checkbox"/> Tue	<input type="checkbox"/> Wed	<input type="checkbox"/> Thu	<input type="checkbox"/> Fri	<input type="checkbox"/> Sat	<input type="checkbox"/> Sun
Name of PLACE where you got food:							
Location:							
Name of PERSON who got the food:							
(√) Did you . . .							
Use store or manufacturer's coupons?	<input type="checkbox"/> yes		<input type="checkbox"/> no				
Use a store loyalty card or a frequent shopper card?	<input type="checkbox"/> yes		<input type="checkbox"/> no				
Save your receipt?	<input type="checkbox"/> yes		<input type="checkbox"/> no				
Enter total paid including tax and tip	\$ _____						
(√) How did you pay? Check ALL that apply							
<input type="checkbox"/> Cash	<input type="checkbox"/> WIC amount: \$ _____		<input type="checkbox"/> Debit card	<input type="checkbox"/> Credit card			
<input type="checkbox"/> Check	<input type="checkbox"/> SNAP EBT amount: \$ _____		<input type="checkbox"/> TANF EBT	<input type="checkbox"/> Free			
(√) Did you SCAN the food and drinks?							
<input type="checkbox"/> ALL		<input type="checkbox"/> None			<input type="checkbox"/> Some		
List ALL foods and drinks you COULD NOT SCAN							
Description (Please be as specific as possible)	Write size or amount if known (Ounces, grams, lbs, etc.)			How many?			

TAPE RECEIPT
HERE

*If you need more lines, use another Blue page.






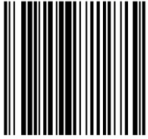









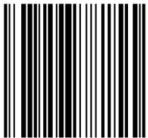

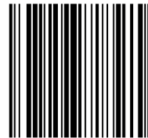





QUESTIONS? Call 1-866-275-8659

Office Use

Source: USDA Economic Research Service.

Figure 2.4: FoodAPS Primary Respondent Workbook, Food-at-Home Source Barcodes

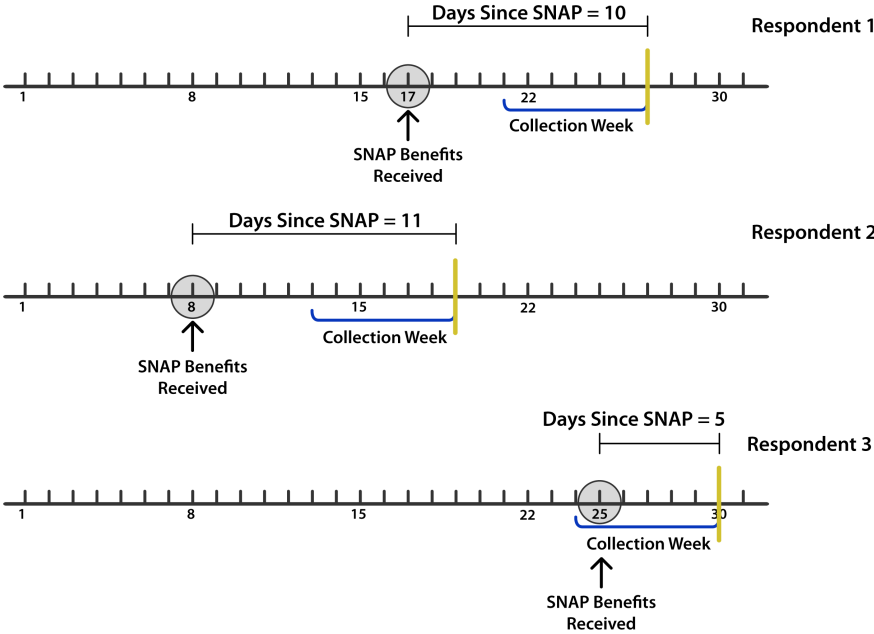
PLACES – Scan a place before scanning food from that place

<p>Convenience store, corner store, bodega</p>    <p>P-1001</p>	<p>Grocery store or supermarket</p>   <p>P-1006</p>
<p>Big box or discount store</p>    <p>P-1002</p>	<p>Liquor or package store</p>   <p>P-1007</p>
<p>Drug store or pharmacy</p>   <p>P-1003</p>	<p>Wholesale club</p>    <p>P-1008</p>
<p>Farmers' market / farm stand</p>   <p>P-1004</p>	<p>Garden, hunting, fishing</p>    <p>P-1009</p>
<p>Food bank or pantry</p>   <p>P-1005</p>	<p>Other place</p>  <p>Bakery, delicatessen (deli), fish or meat market, or any other store not listed.</p> <p>P-1010</p>

Questions? Call us toll-free at 1-866-275-8659


Source: USDA Economic Research Service.

Figure 2.5: Variation in FoodAPS Instrument




This figure walks through the variation in timing since SNAP benefit receipt present in the FoodAPS dataset. The grey circle indicates the day that a respondent received SNAP benefits. “Collection Week” shows the week the household was surveyed, and “Days Since SNAP” counts the number of days elapsed between the receipt of benefits and the end of the collection week.

Figure 2.6: Food Bank Utilization by SNAP Participation



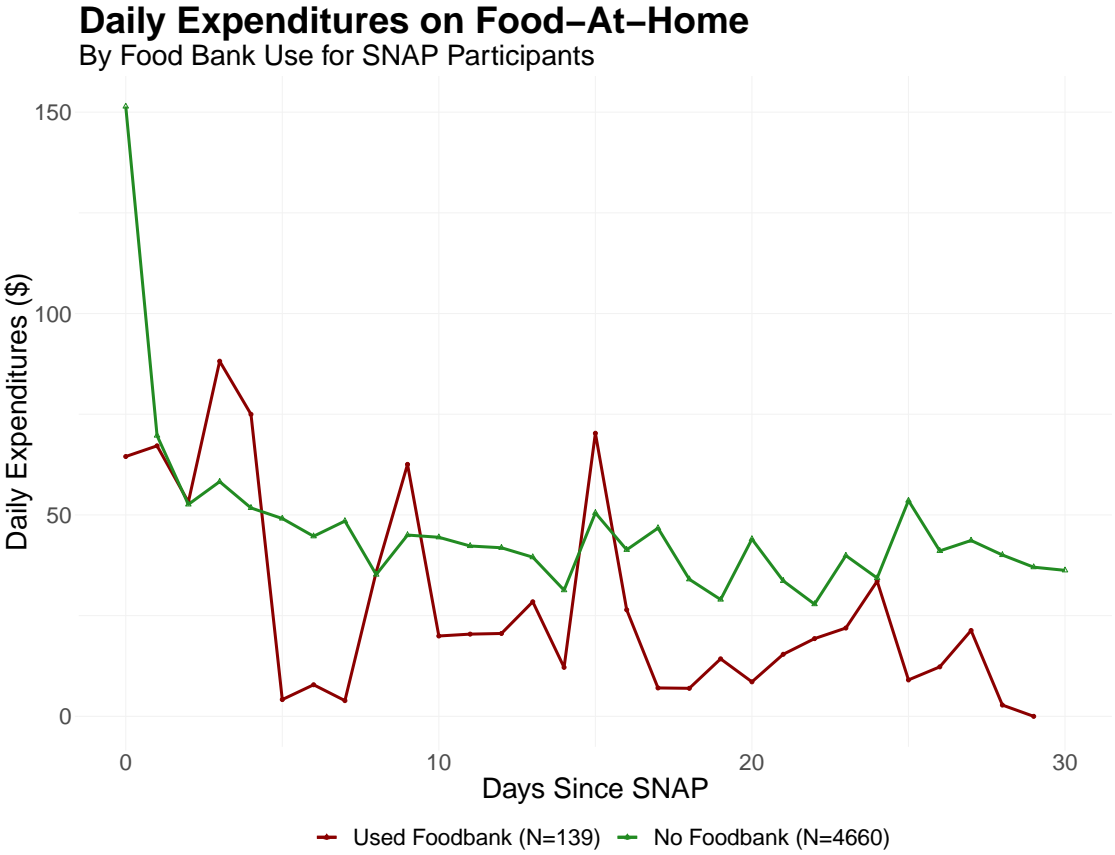
This figure plots the average number of food bank visits per week for SNAP households (blue line) and non-SNAP households (orange line) among FoodAPS households that ever utilized food banks.

Figure 2.7: Daily Food-at-Home Expenditures



This figure plots the average daily food-at-home expenditures in US dollars by payment source for SNAP-participating households. The green series plots expenditures made using SNAP benefits, while the red series reports expenditures made using a non-SNAP payment source.

Figure 2.8: Daily Food-at-Home Expenditures by Food Bank Use



This figure plots the average daily food-at-home expenditures in US dollars by whether a household utilized food banks during the FoodAPS survey week. The green series plots daily expenditures made by households that did not utilize food banks during the collection week, while the red series reports expenditures made by households that visited a food bank.

Table 2.1: Evidence of the SNAP Benefit Cycle

Paper	Data Source	Model	Findings
Wilde and Ranney (2000)	CEX, CSFII	Engel Curve, Endogenous Switching	Mean food-at-home expenditures rise in the 3 days following benefit disbursement. Smooth consumption for high-frequency shoppers.
Shapiro (2005)	CSFII, NFCS	Quasi-hyperbolic discounting	Caloric intake declines 10-15% over the benefit month, explained by time preferences.
Hastings and Washington (2010)	NV grocery store scanner data	Household fixed effects	Total purchases fall across weeks in benefit month, steepest declines for storable and perishable goods.
Castner and Henke (2011)	USDA Program Data	Data Visualization, Conditional Means	Majority of households redeem over half of benefits within the first week, 80% of benefits by end of week two.
Damon, King, and Leibtag (2013)	Neilsen Homescan and CES (11 MSAs)	Household by week fixed effects	Low income households in early disbursement states purchase less food in week 4 of every month, esp. at grocery stores.
Todd (2015)	2007-10 National Health and Nutrition Examination Survey	Linear model, days since benefit receipt; ARRA	38% decline in caloric intake in final two days and eating occasions per day fall by 0.46 in last week (prior to ARRA).
Goldin, Homonoff, and Meckel (2016)	Kilts-Nielsen Retail Scanner	Household fixed effects	Cycles in utilization parallel issuance dates, no effect on retailer pricing.
Hamrick and Andrews (2016)	BLS American Time Use Survey	Logit	Probability of a day without food increases .8% over benefit month.
Castellari et al. (2017)	Neilsen Homescan in Nevada	Household fixed effects	More purchases on date of disbursement, 5% increase in beer purchases when benefit distribution falls on a weekend.

This table summarizes findings from prior studies focused on identifying the existence and consequences of the SNAP Benefit Cycle.

Table 2.2: Evidence of the SNAP Benefit Cycle in FoodAPS

Paper	Model	Findings
Smith et al. (2016)	Maximum Likelihood, marginal propensity to spend on food at home	Evidence consistent with hyperbolic discounting, against fungibility of SNAP
Tiehen, Newman, and Kirlin (2017)	Data Summary	Food spending declines sharply over the benefit month, but no increase in the # times free food is acquired.
Dorfman et al. (2018)	Finite Mixture Model	Benefit cycle driven by a minority of participants spending 2/3 of benefits in first 4 days.
Whiteman, Chrisinger, and Hillier (2018)	Univariate regression on “days since SNAP,” limited controls	Bulk of food acquisition occurs directly after benefit receipt, expenditure and caloric intake decline over the benefit month; mean healthy eating scores and vegetable consumption falls with days since benefit receipt.
Beatty et al. (2019)	Household and state fixed effects	Presence of other income streams has limited effect on benefit cycle.
Gregory and Smith (2019)	Individual Logit and random effects, in “salience window” (first 4 days, last 3 days of benefit month)	25% increase in reporting food hardship in salience window
Laurito and Schwartz (2019)	Day and month fixed effects	School breakfast and lunch take-up higher among schoolchildren in SNAP households during last 2 weeks of benefit month.

This table summarizes findings from prior studies focused on identifying the existence and consequences of the SNAP Benefit Cycle.

Table 2.3: Household Characteristics by Food Bank Use

	Used Food Banks		No Food Bank Use	
	Mean	Std. Dev.	Mean	Std. Dev.
Household Size	3.1200	2.1276	3.4635	1.9264
# Kids	0.8800	0.9274	1.4106	1.4599
# Seniors	0.2400	0.4359	0.1503	0.4101
# Non-senior Adults	1.9200	1.0376	1.6558	0.9757
Household Monthly Income	\$2,058.36	\$2,005.14	\$2,072.53	\$1,951.59
% Poverty Guideline	116.8323	80.6973	119.2202	103.1001
Estimated Monthly SNAP Benefits	\$277.80	280.43	\$343.38	293.86
Female PR	0.7600	0.4359	0.8010	0.3994
PR Married	0.1200	0.3317	0.2846	0.4514
PR has <HS Ed.	0.2000	0.4082	0.2746	0.4465
PR has HS Ed.	0.3200	0.4761	0.3275	0.4695
PR has Some College	0.4800	0.5099	0.3123	0.4636
PR has Bachelor's	0.0000	0.0000	0.0705	0.2561
PR has > Bachelor's	0.0000	0.0000	0.0143	0.1187
PR White	0.5600	0.5066	0.6549	0.4756
PR Black	0.2000	0.4082	0.1948	0.3962
PR Asian	0.0400	0.2000	0.0151	0.1221
PR Other Race	0.2000	0.4082	0.1352	0.3421
Owns home	0.4000	0.5000	0.3090	0.4623
Public Housing	0.1600	0.3742	0.1360	0.3430
Has Vehicle	0.4800	0.5099	0.7244	0.4470
Self-Employed	0.1200	0.3317	0.0907	0.2873
Food Security: Low	0.2800	0.4583	0.2561	0.4367
Food Security: Very Low	0.3600	0.4899	0.1788	0.3834
At Times not Enough to Eat	0.2000	0.4082	0.1553	0.3624
Food Ran Out in Last Month	0.7200	0.4583	0.4836	0.4999
Couldn't Afford Balanced Diet	0.7200	0.4583	0.5147	0.5000
Skipped/Cut Size of Meals	0.8000	0.4082	0.6927	0.4616
# Households	25		1191	

Note: This table reports summary statistics for household characteristics from the FoodAPS instrument for households that are currently enrolled in SNAP, split by whether households used food banks during the survey period. Bold mean values for food bank households correspond to mean differences with two-sided p-values under 0.05.

Table 2.4: Characteristics of Food-at-Home Acquisition Events

	Used Food Banks		No Food Bank Use	
	Mean	Std. Dev.	Mean	Std. Dev
Survey Day 1	0.2014	0.4025	0.1942	0.3956
Survey Day 2	0.1295	0.3370	0.1730	0.3783
Survey Day 3	0.1727	0.3793	0.1511	0.3582
Survey Day 4	0.1295	0.3370	0.1326	0.3392
Survey Day 5	0.0791	0.2709	0.1189	0.3237
Survey Day 6	0.1583	0.3663	0.1084	0.3109
Survey Day 7	0.1295	0.3370	0.1219	0.3272
Supermarket or Grocery Store	0.2446	0.4314	0.3541	0.4783
Dollar Store	0.0576	0.2337	0.0745	0.2626
Superstore	0.2950	0.4577	0.3479	0.4763
Convenience Store	0.0576	0.2337	0.0659	0.2481
Food Bank	0.1942	0.3970	0.0000	0.0000
Accepts SNAP	0.6475	0.4795	0.9017	0.2977
Acquired for Free	0.2932	0.4570	0.0423	0.2013
Paid with Cash	0.4396	0.4991	0.3736	0.4838
Paid with Credit or Debit Card	0.1942	0.3970	0.1620	0.3685
Used SNAP as Payment Source	0.2308	0.4237	0.4964	0.5000
Total Paid, \$	16.2663	38.2407	30.9594	46.1773
Amount Paid from SNAP, \$	49.1411	80.7859	39.8033	53.8609
Amount Paid from TANF, \$	26.7175	9.4376	25.8606	28.8050
Amount Paid from WIC, \$	10.8220	6.8064	21.5842	25.2136
Observations	139		4660	

Note: This table reports summary statistics for food-at-home acquisition events from the FoodAPS instrument for households that are currently enrolled in SNAP, split by whether households used food banks during the survey period. Bold mean values for food bank households correspond to mean differences with two-sided p-values under 0.05.

Table 2.5: Characteristics by Survey-Day

	Used Food Bank		No Food Bank Use	
	Mean	Std. Dev.	Mean	Std. Dev.
# Acquisition Events	1.4946	0.8025	1.3832	0.7163
Daily Expenditures, \$	23.0876	46.7565	42.3911	61.6384
Daily Exp. out of SNAP, \$	10.0396	41.4708	24.5152	53.4476
Daily Exp. out of pocket, \$	13.0481	21.8808	17.8759	34.4988
Obtained Food for Free	0.4066	0.4939	0.0539	0.2259
SNAP Week 1	0.1183	0.3247	0.2847	0.4513
SNAP Week 2	0.2903	0.4564	0.2686	0.4433
SNAP Week 3	0.2581	0.4399	0.2015	0.4012
SNAP Week 4	0.3333	0.4740	0.2452	0.4303
Weekday	0.7527	0.4338	0.7186	0.4497
Holiday	0.0215	0.1458	0.0294	0.1689
Observations	93		3369	

Note: This table reports summary statistics for daily food-at-home acquisitions by SNAP households surveyed in the FoodAPS instrument, split by whether households used food banks during the survey period.

Table 2.6: Food-at-Home Expenditures and the SNAP Cycle

	Total Exp.					SNAP	OOP
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
0 DSS	91.70*** (9.206)	92.26*** (9.011)	91.47*** (9.102)	91.59*** (9.094)	90.42*** (10.35)	91.09*** (9.397)	0.7038 (4.069)
1-6 DSS	16.51*** (3.913)	16.68*** (3.877)	14.27*** (4.001)	14.22*** (3.985)	13.78*** (5.098)	16.66*** (4.371)	-1.111 (2.681)
Week 2	0.7455 (2.089)	0.9815 (2.078)	-2.056 (2.272)	-2.040 (2.274)	-5.331* (2.995)	-3.964 (2.419)	-1.853 (1.948)
Week 3	-3.595 (2.353)	-3.951* (2.249)	-3.999* (2.349)	-4.023* (2.347)	-5.322* (2.942)	-4.732** (2.104)	-0.2175 (2.072)
HH Cntrl.	No	Yes	Yes	Yes	No	No	No
Time Cntrl.	No	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	No	No	Yes	No	No	No
HH FE	No	No	No	No	Yes	Yes	Yes
N	3,446	3,446	3,446	3,446	3,446	3,446	3,446
Adj. R ²	0.13	0.15	0.16	0.16	0.27	0.32	0.2027

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered by household. These models use Eq. 2.1 to estimate the change in daily food-at-home expenditures relative to the time since SNAP benefits were received. The dependent variable $Exp_{i,d,w}$ measures the daily expenditures in US dollars by household i on day d of benefit week w on food for consumption at home. $0\ DSS$ and $1 - 6\ DSS$ are dummy variables equal to one when households are on the day of or 1-6 days since benefit receipt, respectively. $Week\ 2$ ($Week\ 3$) indicates that a household is in the second (third) week of the benefit month. Coefficients on these “Days Since SNAP” variables therefore measure the difference in average daily expenditures relative to that in the final week of the benefit month. $HH\ Cntrl.$ include average household income, the number of children, seniors, and adults in the house, as well as controls for the primary respondent’s gender, ethnicity, and education level. $Time\ Cntrl.$ include a dummy variable for federal holidays as well as fixed effects for the survey day, day of the week, and week of the year. $Region\ FE$ and $HH\ FE$ indicate the presence of census region or household fixed effects, respectively.

Table 2.7: Food-at-Home Expenditures, the SNAP Cycle, and Food Bank Use

	Current FB Use		Previous FB Use	
	(1)	(2)	(3)	(4)
Food Bank HH	-18.68*** (4.94)		-5.36 (3.96)	
0 DSS	91.74*** (9.19)	90.87*** (10.47)	94.90*** (9.96)	94.43*** (11.35)
0 DSS × FB	-53.51*** (20.50)	-49.58*** (17.92)	-26.85 (22.88)	-28.82 (26.51)
1-6 DSS	13.67*** (4.04)	13.23** (5.17)	13.14*** (4.41)	12.79** (5.56)
1-6 DSS × FB	28.15** (10.94)	41.44*** (10.95)	6.97 (10.04)	8.44 (13.46)
Week 2	-2.37 (2.29)	-5.08* (3.04)	-2.33 (2.47)	-5.05 (3.35)
Week 2 × FB	14.94 (13.30)	-0.104 (11.84)	2.00 (5.59)	-2.19 (6.62)
Week 3	-4.09* (2.40)	-4.83 (3.05)	-3.83 (2.59)	-3.69 (3.27)
Week 3 × FB	4.52 (6.44)	-8.15 (10.02)	-1.96 (5.00)	-10.88 (6.61)
HH Cntrl.	No	Yes	No	Yes
Time Cntrl.	Yes	Yes	Yes	Yes
Region FE	Yes	No	Yes	No
HH FE	No	Yes	No	Yes
N	3,446	3,446	3,446	3,446
Adj. R ²	0.16	0.27	0.16	0.27

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered by household. These models use Eq. 2.2 to estimate the change in daily food-at-home expenditures relative to the time since SNAP benefits were received by Food Bank (FB) use. The dependent variable Exp_{idw} measures the daily expenditures in US dollars by household i on day d of benefit week w on food for consumption at home. 0 DSS and 1–6 DSS are dummy variables equal to one when households are on the day of or 1-6 days since benefit receipt, respectively. Week 2 (Week 3) indicates that a household is in the second (third) week of the benefit month. Coefficients for the interaction terms including $\times FB$ denote the estimated difference in expenditures for households utilizing food banks, either during the survey week (Current FB Use) or within the previous thirty days (Previous FB Use). See Table 2.6 notes for description of control specifications.

Table 2.8: The SNAP Cycle and Food Banks, Alternate Timing

	(1)	(2)	(3)	(4)
<i>Weekly</i>				
Week 1	48.36*** (7.10)	49.80*** (7.51)		
Week 1 × FB	-2.64 (20.37)	-12.07 (19.26)		
Week 2	31.24*** (6.05)	31.92*** (6.44)		
Week 2 × FB	-21.86 (19.89)	-10.90 (15.44)		
Week 3	16.86*** (4.29)	18.01*** (4.59)		
Week 3 × FB	-22.25** (9.42)	-14.28 (10.60)		
<i>Day Zero Break</i>				
Days Since SNAP			-0.81*** (0.27)	0.09 (0.80)
Days Since × FB			-2.57 (1.80)	0.09 (0.80)
0 DSS			73.56*** (10.61)	76.90*** (11.38)
0 DSS × Prev FB			-94.72*** (36.27)	-30.96 (27.46)
FB Measure	Current	Previous	Current	Previous
Time Controls	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
Observations	3,446	3,446	3,446	3,446
Adj. R ²	0.21	0.24	0.27	0.27

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered by household. These models use Eq. 2.2 to estimate the change in daily food-at-home expenditures relative to the time since SNAP benefits were received. The dependent variable Exp_{idw} measures the daily expenditures in US dollars by household i on day d of benefit week w on food for consumption at home. The *Weekly* control structure includes dummy variables for being in the first, second, and third weeks of the benefit month and interactions with whether a household utilized a food bank during the survey week (Current) or in the prior thirty days (Previous). The *Day Zero Break* model in Column (3) specifies a linear trend and a day zero dummy for the day of benefit receipt, and interacts both with the measure of previous food bank use. All models include the full set of time controls and specify household fixed effects.

Table 2.9: The SNAP Cycle and Food Banks, Payment Types

	SNAP Exp.		Out of Pocket Exp.	
	(1)	(2)	(3)	(4)
Food Bank HH	-12.75*** (3.151)		-5.977 (3.779)	
0 DSS	93.74*** (8.871)	91.47*** (9.506)	-1.812 (3.205)	0.8063 (4.124)
0 DSS × Used FB	-42.14** (16.49)	-34.60** (16.74)	-10.09 (6.515)	-16.19* (8.714)
1-6 DSS	19.58*** (3.563)	16.47*** (4.435)	-4.952*** (1.889)	-1.427 (2.707)
1-6 DSS × Used FB	12.66 (14.58)	19.77 (13.24)	17.14* (9.405)	18.94* (10.63)
SNAP Week 2	1.141 (1.844)	-4.098* (2.463)	-2.812* (1.513)	-1.496 (1.974)
SNAP Week 2 × Used FB	10.77 (12.43)	9.734 (9.624)	3.931 (5.526)	-9.303 (6.345)
SNAP Week 3	-3.391* (1.787)	-4.283* (2.192)	-0.4776 (1.838)	-0.1346 (2.145)
SNAP Week 3 × Used FB	5.704 (3.512)	-6.934 (8.244)	-2.264 (4.532)	-2.195 (6.123)
FB Measure	Current	Current	Current	Current
Region FE	Yes	No	Yes	No
HH FE	No	Yes	No	Yes
Time Controls	Yes	Yes	Yes	Yes
N	3,446	3,446	3,446	3,446
Adj. R ²	0.20	0.32	0.04	0.20

* p < 0.10, ** p < 0.05, *** p < 0.01. Standard errors are clustered by household. These models use Eq. 2.2 to estimate the change in daily food-at-home expenditures relative to the time since SNAP benefits were received. The dependent variable Exp_{idw} measures the daily expenditures in US dollars by household i on day d of benefit week w on food for consumption at home. All models include the full set of time controls and specify either census region fixed effects and household characteristics or household fixed effects.

Chapter 3

Are we #Stayinghome to Flatten the Curve?

3.1 Introduction

Beginning December 2019, the novel coronavirus SARS-CoV-2 (COVID-19) spread rapidly around the world and in the U.S., prompting dramatic policy responses. Local, state, and national governments around the world faced an extensive set of policy instruments with which to fight the pandemic and limit the virus' impact on their constituents. Because many regions exhibited exponential growth in coronavirus cases, policymakers increasingly implemented aggressive stay-at-home mandates that sought to reduce transmission through human interaction and “flatten the curve” (Mervosh, Lu, and Swales 2020).

By March 31, the U.S. reported the highest number of confirmed COVID-19 cases globally, with over 67% more cases than the next closest country (*Coronavirus disease (COVID-19) dashboard* 2020). 30 states had implemented stay-at-home policies by this date, with 42 states and Washington D.C. ultimately adopting a stay-at-home order. Understanding how effective these policies were at modifying residents' behavior (and ultimately reducing the propensity for COVID-19 transmission) is of key importance for the maintenance of current mandates and the design of future pandemic control policies.

The relevant benefits of non-pharmaceutical interventions (NPI) – such as quarantining infected households, closing schools, and banning social events or large gatherings – have largely been informed by existing mathematical models (Hatchett, Mecher, and Lipsitch 2007). In addition, some anecdotal and historical evidence supports their efficacy. In California's San Francisco Bay Area, the first area of the country to implement stay-at-home mandates, doctors reported “fewer cases than expected” after two weeks of social distancing (Kahn and Marinucci 2020). Analysis of internet-connected thermometers suggests that new fever rates on March 23 were below those at the start of the month, while state hospitalization rates showed a commensurate decline in growth rates. Washington state officials reported similar reductions in COVID-19 transmission as a result of the state's containment

strategies (Baker 2020). Exploration of death rates and NPI roll-out in 17 U.S. cities during the 1918 influenza pandemic support these claims, finding that implementation of multiple social distancing practices intended to reduce infectious contacts early in the outbreak led to 50% lower peak death rates and flatter epidemic curves relative to cities that did not implement such policies (Hatchett, Mecher, and Lipsitch 2007).

Recent simulations provide further insight into the benefits of social distancing. While epidemiological models of the U.K. and U.S. suggest that techniques for mitigating exposure of those most at risk may drastically reduce peak load on the healthcare system and cut COVID-19 deaths by half, such techniques on their own might not be enough to prevent the healthcare system from being overwhelmed. Some argue that, in this case, a combination of social distancing, self-quarantine of infected people, and suspension of schools would need to be maintained until a vaccine is available to prevent a rebound (Ferguson et al. 2020). Other experts call for widespread testing coupled with digital contact tracing as a means to reduce viral spread while minimizing harmful social and economic side-effects. Simulations based on a moderate mitigation policy (comprising 7-day isolation following any symptoms, a 14-day quarantine for the household, and social distancing for all citizens over age 70) find that, had it been implemented in late March, it would have reduced potential U.S. deaths by 1.76 million (Greenstone and Nigam 2020). Given that this simulated policy is less stringent and maintained for a shorter duration than many of the policies that states actually implemented, the benefits from existing stay-at-home mandates (either directly from reduced COVID-19 deaths or indirectly due to decreased transmission of other illnesses) could be substantially larger.

Careful empirical study of these stay-at-home mandates is critical for understanding the benefits to the policies, given the visibility and extent of their costs. Even before mandates limited economic activity, GDP forecasts suggested an economic contraction in the U.S. of 24% (McCabe 2020). Concerns over these costs prompted comments from the executive branch regarding relaxation of restrictions and allowing non-essential businesses to reopen, which quickly prompted opposition from public health experts (Finucane and Andersen 2020) and many economists (IGM 2020).

This paper contributes to the existing literature first by documenting widespread changes in U.S. mobility patterns surrounding the COVID-19 pandemic. Using cellular location data, we examine the ways in which travel behavior changed across the United States in response to the virus' spread (Unacast 2020). We show that tremendous nationwide reductions in travel activity levels and human encounter rates occurred prior to statewide mandates, suggesting residents were already responding to perceived risks and more local policies. Prior to any state implementing a statewide mandate, average travel distances had already fallen by 16 percentage points, the human encounter rate by 63 percentage points, and non-essential visits by 39 percentage points relative to pre-COVID-19 levels, providing evidence that extensive social distancing preceded many of the policies designed to induce such behavior.

Next, we provide empirical evidence that statewide stay-at-home mandates induced further reductions in travel behavior and greater increases in social distancing. States' policies combined closures of non-essential businesses with instructions for all residents to remain at

home except for the purchase of necessities (i.e. groceries or medicine), with the goal of limiting “unnecessary person-to-person contact” (New Jersey 2020) and to “mitigate the impact of COVID-19” (California Executive Order N-33-20 2020). We first estimate a differences-in-differences model that isolates the effect of statewide mandates by comparing differences before and after mandate implementation within states and between early-adopting, later-adopting, and control states. Using this framework, we test whether states’ stay-at-home policies induced further changes in mobility and daily human encounters. Decomposition of the treatment effect estimates reveals that a majority of estimate weights fall on comparisons between mandate and non-mandate states, with effect heterogeneity primarily found in the size of travel reductions.

After discussing results from the difference-in-differences model, we present estimates from event study models that directly examine the dynamic effects of stay-at-home mandates. In addition to typical unweighted event studies, we implement weighted event studies that directly balance states on differences in pre-adoption mobility trends (Ben-Michael, Feller, and Rothstein 2019). These methods show that while visits to non-essential businesses returned to previous levels within ten days of mandate implementation, reductions in average travel distance and human encounter rates persisted throughout the observed portions of mandates. Both event study methods yield similar overall treatment effect estimates as in the difference-in-differences model, supporting the finding that statewide mandates induced further reductions in travel activity even after considerable pre-mandate reductions.

Across all methods, we find evidence that residents reduced daily activity, even before mandates, and that patterns differed by state. Moreover, we estimate significant additional reductions in travel and increased social distancing in response to stay-at-home mandates. We estimate a 7.0 percentage point reduction in average distance traveled, a 2.1 percentage point decline in non-essential visits, and a 3.5 percentage point reduction in the daily rate of human encounters after the average stay-at-home mandate was implemented. These represent meaningful changes, as non-essential visits and human encounter rates fell a further 27–31% relative to the reductions observed before states implemented mandates, with average travel distances changing by an additional 140%.

Taken together, our results provide evidence that the implementation of non-pharmaceutical interventions in the form of statewide stay-at-home mandates encouraged additional social distancing and reduced the opportunities for person-to-person transmission. Our estimates of mandate-induced behavior changes provide important insight into the effectiveness of recent policies designed to “flatten the curve” and stop the spread of COVID-19. We provide evidence that these policies further modified travel behavior and reduced the frequency of in-person contact. Further, our findings of substantial reductions in mobility prior to state-level policies convey important implications for the relaxation and eventual removal of these pandemic policies.

3.2 Data

Mobility Data

We obtain travel activity and social distancing data from the analytics company Unacast. To understand how well different communities are social distancing, Unacast uses cellular location data for 15-17 million identifiers per day to construct three measures of behavior in response to COVID-19 policies (Unacast 2020). Each measure is aggregated to the state-by-day level and is defined as the daily percentage point change relative to that day of week's average for the pre-COVID-19 period of February 10 through March 8 (henceforth referred to as the baseline period). While all data are published publicly to their Social Distancing Dashboard in the form of figures and maps (Unacast 2020), we obtained the balanced panel of state-by-day observations for the period of February 24 through April 29, 2020 directly from Unacast.

Unacast receives location data from millions of mobile devices. Location information is received through authorized applications, Wi-Fi or Bluetooth connections, and A-GPS positions. Obtained information includes the location of the device at a given point in time (latitude, longitude, and elevation) along with the mobile device make, model, and operating system, the corresponding application gathering the data, GPS accuracy value, and the direction and rate of travel. Each state-day observation we use is calculated using position information. Taken together, these three measures paint a comprehensive picture of behavior changes in response to states' stay-at-home mandates.

The first metric we use, the change in average distances traveled (\dot{ADT}), provides a measure of overall changes in travel activity during the COVID-19 period. To create this measure, the average distance traveled across all devices assigned to a state on a given day is compared to the state's average for that day of the week during the pre-COVID-19 baseline period. A value of $\dot{ADT}_{it} = 0$ indicates that the average distance traveled for individuals assigned to state i on date t was identical to the baseline distance for that day of the week, while a value of -7 conveys that, on average, devices assigned to the state traveled an average distance 7 percentage points less than typical levels. As a result, \dot{ADT} measures the percentage point change in average distance traveled relative to more typical behavior prior to the pandemic, accounting for pre-existing differences in states' average mobility patterns or residents' propensity to travel throughout the week.

Changes to average distances traveled give a sense of broad transformations to travel behavior. Declines in relative activity will yield negative values of \dot{ADT} , which can reflect both intensive margin (shorter distances traveled for the same frequency of trips) and extensive margin (some trips foregone entirely) adjustments. Reductions in \dot{ADT} following mandate implementation reflect compliance on average with states' guidances to work from and stay at home except for essential activities, while positive or nonexistent changes would suggest a disconnect between private behavior and public policy.

The second metric we use is the change in visits to non-essential businesses, defined as Non-Essential Visits (\dot{NEV}). To the extent that non-essential businesses closed following

stay-at-home mandates, we expect to see reductions in the number of trips residents take to these types of retail or service businesses. Our utilized measure of the change in visits to non-essential businesses ($N\dot{E}V$) offers a similar comparison to $A\dot{D}T$ that is targeted at travel to the types of businesses most heavily impacted by stay-at-home mandates. Businesses likely to be deemed “non-essential” include department stores, spas and salons, fitness facilities, event spaces, and many others. To improve accuracy, non-essential businesses are defined according to group definitions in both the Unacast SDK and the OpenStreetMaps POI’s. $N\dot{E}V$ is calculated by dividing a state’s average number of visits to non-essential businesses by its day-of-week baseline level. A value of $N\dot{E}V_{it} = 2$ indicates a two percentage point increase in visitations to non-essential businesses relative to baseline norms for that weekday in a given state.

Finally, we use changes in the rate of unique human encounters ($E\dot{N}C$) as a measure of social distancing. While $A\dot{D}T$ and $N\dot{E}V$ provide information on two potential margins for adjusting travel behavior, neither directly captures changes in human-to-human interaction. As COVID-19 transmission primarily occurs through “close contact from person-to-person,” having a measure of potential human encounters allows us to further understand whether reductions in travel distance and business visitations translate into fewer opportunities for viral transmission (CDC 2020).

$E\dot{N}C$ measures the change in the rate of unique human encounters per square kilometer relative to the state’s baseline levels. Following (Pepe et al. 2020), one unique encounter is produced every time two devices assigned to a given state are observed within a 50 meter radius circle of each other for no more than 60 minutes.¹ Dividing the state-level sum of encounters for the day by the state’s square kilometer of land area provides the state’s daily rate of unique human encounters. Finally, this encounter rate is normalized by the state’s average encounter rate for that day of the week during the baseline period.² As a result, an encounter rate equal to the state’s baseline rate for that day of the week results in a value of $E\dot{N}C_{it} = 0$, while a value of $E\dot{N}C_{it} = -12$ indicates a 12 percentage point reduction in the encounter rate for state i on date t relative to the state’s pre-COVID-19 level.

Table 3.1 provides summary statistics for the three mobility measures across the sample period. The table is divided into four panels. The first panel corresponds to the majority of the pre-COVID-19 baseline period of February 24 to March 8, 2020.³ The second panel summarizes behavioral changes for the rest of March, during which the majority of states enacted their stay-at-home policies. The third panel covers the period from April 1-29, 2020,

¹For data quality reasons, we drop observations for Washington D.C. from our analysis of human encounter rates. The point estimates reported in this paper are unchanged when D.C. is included, albeit with larger standard errors.

²In contrast to $A\dot{D}T$ and $N\dot{E}V$, which are normalized by the state’s day-of-week average from the entire pre-COVID-19 baseline period of February 10 through March 8, due to data limitations $E\dot{N}C$ is only normalized using the state’s day of week average for February 24 through March 8.

³While the full baseline period used for normalization extends back to February 10, our state-by-day mobility panel only begins on February 24 and does not include daily observations for February 10-23. However, Figure 3.1 shows that travel patterns were largely indistinguishable from baseline levels prior to early March, providing evidence that behavior did not change substantively during the month of February.

which includes adoption of the final ten mandates and the observed post-adoption period. The final panel provides the average, median, and standard deviation for the total sample. The three columns report summary statistics for the changes in average distance traveled ($\dot{A}DT$), non-essential visits ($\dot{N}EV$), and the unique human encounter rate ($\dot{E}NC$).

In the top panel, covering February 24 through March 8, 2020, we see that the daily average distance traveled was larger than pre-COVID-19 baseline levels by 0.79 percentage points. Non-essential visits are 1.78 percentage points lower than the pre-COVID-19 baseline, with the rate of human encounters on average at baseline in the third column. Looking beyond the mean or median provides evidence of substantially heterogeneity across states, with the first quartile for all measures reflecting reduced activity while the third quartile reveals higher travel.

For the final three weeks of March 2020, all three mobility measures experience large decreases relative to pre-COVID-19 levels, attesting to average reductions in both travel and social interactions. During this period we observe 75th percentile changes of 12.93 percentage point reductions in average distances traveled, 21.06 percentage point reductions for non-essential visits, and encounter rate declines of 48.66 percentage points below baseline levels. Travel reductions become even more dramatic in the month of April; average reductions in all three travel activity measures during April exceed in magnitude the equivalent March declines.

Our utilized travel measures display very high correlations with travel data produced by other sources. Comparing our utilized travel change variables with the measure of state-level changes in retail and recreation travel from Google’s COVID-19 Community Mobility Reports, we find average correlations of 0.95 with $\dot{A}DT$ and 0.98 with $\dot{N}EV$ (Google 2020). Changes in non-essential visits are nearly perfectly correlated with the Google measure, with no state exhibiting a correlation below 0.96. Correlations for $\dot{A}DT$ remain high but exhibit greater variation across states. 31 states have correlations above 0.95, including California and New York at 0.97 and 0.98, respectively. Wyoming displays the lowest correlation at 0.75. These strong relationships across data providers suggest that our results are indicative of widespread changes to general mobility patterns and not spurious results arising from anomalies in our chosen data source.

Stay-at-Home Mandate Data

To denote periods before or after a state implemented a “stay-at-home order,” we obtain the date each statewide policy was issued for all fifty states and the District of Columbia (Mervosh, Lu, and Swales 2020). We define our early adopters as the first four states to pass a stay-at-home mandate: California, Illinois, New Jersey, and New York. The second group comprises the 38 late adoption states and D.C.. The last group comprises the eight remaining states that never implemented statewide mandates: Arkansas, Iowa, Nebraska, North Dakota, Oklahoma, South Dakota, Utah, and Wyoming. The observed stay-at-home mandates all consisted of a mix of specific non-pharmaceutical interventions; each observed policy closed or placed considerable limits on non-essential businesses and required residents

to stay at home except for essential activities. Essential services include grocery stores, gas stations, pharmacies, banks, laundry services, and business essential to government functions (covid19.ca.gov 2020). Throughout this paper, we refer to all mandates that implement this combination of policies as a “stay-at-home mandate.”

While we focus our attention on statewide stay-at-home policies, many county and local policies had already been implemented and were already affecting individual-level mobility. Six San Francisco Bay Area counties required residents to stay-at-home beginning March 17, two days prior to the statewide mandate. By mid-March, schools of all levels had begun closing their doors and transitioning to online instruction. On March 9, Stanford University moved classes online “to the extent possible,” with Harvard and many other institutions swiftly following suit (Kadvany 2020). Further, business leaders including Google, Microsoft, Twitter, Facebook, and Amazon transitioned some or all of their employees to working remotely well before statewide mandates entered into effect (Aten 2020). As a result, the behavioral changes following statewide stay-at-home mandate adoption represent only a partial response to the suite of actions and policies undertaken to combat the spread of COVID-19. Our estimated “mandate effects” that follow therefore capture the behavioral responses specific to statewide stay-at-home mandates and underestimate the effect of all combined policies. If local, county, and business policies had already incentivized residents to stay at home, then we would expect a reduced response to later statewide mandates (which would be reflected in small magnitude estimates in our models). Any estimated mandate effects that follow reflect mobility responses in addition to those already realized by pre-existing policies.

3.3 Empirical Strategy

Difference-in-Differences under Staggered Adoption

To determine the effect of statewide stay-at-home mandates on the outcome of interest, we begin by estimating the following difference-in-differences model under staggered adoption:

$$Y_{sd} = \alpha + \beta SAH_{sd} + \eta_s + \delta_d + \varepsilon_{sd} \quad (3.1)$$

The outcome Y_{sd} denotes the change in a given measure of travel activity (ADT , NEV , or ENC) for state s on date d relative to the state’s baseline level for that day of the week. Each outcome is expressed as a function of a constant α , whether a state has a statewide mandate in effect, and both time and unit fixed effects. SAH is an indicator equal to one if state s has a stay-at-home mandate in place on date d and zero otherwise. In the sections that follow, we consider the two cases where SAH includes variation for just the first four states to adopt statewide mandates and for all states that ever adopted a statewide mandate. The vector of state fixed effects η_s controls for time-invariant characteristics of states that affect the outcome, while date fixed effects δ_d control for factors affecting the outcome on a given date common to all states (i.e. executive branch press conferences or daily

changes in worldwide COVID-19 deaths/hospitalizations). The term ε is an idiosyncratic error comprised of unobserved determinants of changes in the outcome that are not controlled for by the variables specified in the linear Eq 3.1. Rather than include state-specific time trends, our preferred specifications use outcomes residualized of timing cohort pre-trends (Goodman-Bacon 2018).

The coefficient β measures the difference in the change in average outcome for states that implemented a stay-at-home mandate relative to the change in activity in states that had yet to implement or never implemented such policies, controlling for state and time-varying factors that also correlate with the outcome of interest. In this way $\hat{\beta}$ provides an estimate of the average treatment effect for treated states (ATT). Models are estimated using data for the entire sample period of February 24 through April 29, 2020, covering the observed pre-mandate, adoption, and post-adoption periods.

This empirical approach allows us to identify the relationship between stay-at-home mandates and daily changes in each of the outcomes of interest while explicitly controlling for other confounding factors that are specific to each state or date. Employment rates prior to COVID-19 or the shares of a given local population previously working from home are controlled for with η , while day-to-day changes in factors common to all states – motivated by new information on the virus’ spread and nationwide media coverage or federal appeals for social distancing – are absorbed by δ . The mandate effect β is identified under the assumption that, after controlling for cohort-specific pre-trends, common day-to-day trends, and time-invariant state characteristics, stay-at-home mandates are as good as random. Equivalently, the day-to-day outcome changes from typical pre-COVID-19 levels in states that had yet to adopt or never adopted a mandate are what the change in the outcome would have been for treated states absent a stay-at-home mandate. Given the time-varying nature of adoption, we can express this underlying assumption as the weighted average of parallel trends for each simple two-by-two difference-in-differences estimators (Goodman-Bacon 2018). Our approach is identified using changes in the outcome that differ from typical pre-COVID-19 levels for a state and the state’s average change during the COVID-19 time. A remaining source of bias would be if the early mandate states were trending differently than the control states before March 18 in ways that differed from trends after March 18, or if similar differences in trends exist across all mandate states. Standard errors of the estimated parameters are clustered by state to account for variation in state policies potentially affecting the variance of the residual ε .

Unweighted and Weighted Event Studies

To directly model the dynamic nature of mobility responses to statewide stay-at-home mandates and relax assumptions of the difference-in-differences approach, we employ two event study methods. First, we estimate traditional event studies equivalent to the preferred difference-in-differences specifications:

$$Y_{sd} = \alpha + \sum_{k=\underline{k}}^{\bar{k}} \beta_k \cdot \text{Days Since}_{sd}^k + \eta_s + \delta_d + \varepsilon_{sd} \quad (3.2)$$

where

$$\text{Days Since}_{sd}^k = \begin{cases} \mathbf{1}\{d \leq SAH_s + k\} & \text{if } k = \underline{k} \\ \mathbf{1}\{d = SAH_s + k\} & \text{if } \underline{k} < k < \bar{k} \\ \mathbf{1}\{d \geq SAH_s + k\} & \text{if } k = \bar{k} \end{cases} \quad (3.3)$$

In this fashion the difference-in-differences estimator is decomposed into $\bar{k} - \underline{k} + 1$ individual coefficients relating the impact of being k days relative to when a state adopted its statewide mandate on date SAH_s . Event-time effects are identified under identical conditions as the staggered difference in differences model of Eq. 3.1 while relaxing the assumption of a time-invariant treatment effect. To ensure that dynamic treatment effects are identified separately from time trends even in models that omit never-treated states, the endpoints \underline{k} and \bar{k} are binned to include all dates that fall either before \underline{k} or after \bar{k} (Schmidheiny and Siegloch 2019). Following convention, we drop the $k = -1$ bin and normalize all event-time effects relative to the day prior to mandate adoption.

Second, we estimate weighted event studies to further account for differences in pre-adoption mobility between states (Ben-Michael, Feller, and Rothstein 2019). Weighted event studies extend the synthetic control method (Abadie, Diamond, and Hainmueller 2010) to the staggered adoption event study framework and cleanly nest within the fixed effects approaches of Eq. 3.1 and 3.2. To correct for imperfect pre-treatment balance, weighted event study augments a “partially pooled” SCM estimator with a fixed effects outcome model. Synthetic controls are constructed based on the balance of residualized pre-treatment outcomes; in this way, the approach softens the difference-in-differences identifying assumptions to balance of treated units against a weighted combination of still-untreated states and builds upon recent research on doubly-robust estimators with an extension to the staggered adoption setting (Sun and Abraham, in press; Arkhangelsky and Imbens 2019; Ben-Michael, Feller, and Rothstein 2019; Chernozhukov, Wuthrich, and Zhu 2020).⁴

We report results for both traditional unweighted event studies and weighted event studies in the form of event study graphs. Each figure presents the estimates for the event-time coefficients, along with 95% standard errors clustered at the state level for unweighted event studies and jackknife 95% standard errors for weighted event studies (Arkhangelsky and Imbens 2019).

⁴For each treated unit k , \widehat{ATT}_{jk} can be thought of as a doubly-weighted difference in differences estimator, wherein the change in the treatment unit j is obtained as the difference between the treatment unit’s outcome k periods post-adoption and its pre-period average, and the change in the control group is the average for equivalent changes for all donor units, weighted by partially pooled synthetic control weights. Averaging \widehat{ATT}_{jk} across all treated units at a given point in event time yields a period-specific treatment effect \widehat{ATT}_k that can be thought of as equivalent to the dynamic ATT obtained from an unweighted event study design. Standard errors are obtained using a jackknife approach (Arkhangelsky and Imbens 2019).

3.4 Results

Overall Changes to Mobility Patterns

Across the United States, COVID-19 upended daily routines. As a result of revised work-from-home guidelines, school closures, family needs, layoffs, and state policies, travel behavior changed dramatically in the U.S. from February through April 2020. Figure 3.1 plots over time the changes in average distance traveled (\dot{ADT}), visits to non-essential businesses (\dot{NEV}), and the unique human encounter rate (\dot{ENC}) per day for all U.S. states, measured as the percentage point change relative to typical pre-COVID-19 baseline levels. The solid line plots the average for the first four states to implement mandatory stay-at-home policies: California (implemented March 19), Illinois (March 21), New Jersey (March 21), and New York (March 22). The dotted line plots the average for the 39 states that adopted stay-at-home mandates later in the sample period, while the dashed line plots the daily average for the remaining 8 states that had yet to adopt a stay-at-home mandate by April 29th.⁵

Travel behavior and social interactions in late February and through the first week of March look largely typical nationwide. Distance traveled and visits to non-essential businesses exhibit only small fluctuations relative to baseline activity levels for all states and groups. The change in the human encounter rate exhibits much greater variation throughout the week, increasing over the course of the work week before falling considerably over the weekend. Despite this greater within-week variation, the average human encounter rate for all states finishes the work week of March 2-6 above baseline levels.

Beginning the week of March 9, residents across the country began deviating from typical travel patterns. By Wednesday March 11, residents of all states had begun reducing their distances traveled, trips to non-essential businesses, and encounters with others relative to pre-COVID-19 norms. Initially, changes to mobility patterns in early-adoption states are largely indistinguishable from those for other states; by March 15, residents across all three groups had reduced travel distance by 8 to 13 percentage points, unique human encounters by 28 to 29 percentage points, and visits to non-essential businesses by 12 to 17 percentage points.

By March 18, before the first statewide mandate went into effect, the rate of travel activity decline had grown considerably in. The change in travel distances fell further, to between -12 and -23 percentage points, with changes between -34 to -49 percentage points for non-essential visits. Unique human encounters had already fallen between -61 and -71 percentage points relative to pre-COVID-19 baseline levels, a dramatic indicator of extensive social distancing occurring even before statewide orders required such behavior. By the start of April, travel behavior and social interactions had largely stabilized at levels 35 to 65 percentage points below previous norms, with within-week patterns and gaps between early adopter, later adopter, and never-adopter states remaining stable as well.

⁵Massachusetts adopted a stay-at-home “advisory” that recommended but did not require that residents stay home. Our reported results include Massachusetts as a stay-at-home state, and the findings are robust to the exclusion of Massachusetts’ policy.

Figure 3.1 provides initial evidence that changes in travel behavior are correlated with the decisions of whether and when to adopt stay-at-home mandates. Following the start of statewide mandate adoption on March 19, residents of early adopter states exhibit larger magnitude reductions every single day through April 29 across all three measures. Each week during this period, mean encounter rates in early adoption states are consistently 10 to 16 percentage points lower than in states that never adopted mandates, with a larger weekly gap in travel distance (between 13 and 19 percentage points) and a similar 12 to 22 percentage point gap for non-essential visits. Trends for later-adopting states fall between early and never-adopting states, with late adopters displaying declines between 5 and 7 percentage points larger in magnitude than never-adopters for human encounter rates, between 7 and 10 percentage points for travel distance, and between 6 and 10 percentage points for non-essential visits.

Effect of Stay-at-Home Mandates on Daily Mobility and Social Distancing

Figure 3.1 provides preliminary visual evidence that residents across the country drastically reduced travel activity and engaged in extensive social distancing prior to the adoption of statewide stay-at-home mandates. We next present estimates of empirical models designed to identify any additional changes in mobility and social distancing patterns attributable to states' stay-at-home mandates. We begin by presenting results of the staggered difference-in-differences treatment effect estimates in Table 3.2 before discussing the unweighted and weighted event study results in Figures 3.2 to 3.4.

Table 3.2 presents difference-in-differences estimates following Eq. 3.1 for the effect of stay-at-home mandates on travel activity across all three mobility measures. Columns (1) - (4) report coefficient estimates for the treatment effect restricted to only the first four adopters' mandates (CA, IL, NJ, and NY compared to all other states) while columns (5) - (8) report estimates using variation from all 43 adopting areas to identify the treatment effect. Columns (1) and (5) report treatment effects with state and date fixed effects, while columns (2) - (4) and (6) - (8) utilize dependent variables residualized of timing cohort pre-trends (Goodman-Bacon 2018).

Comparing column (1) to column (2) illustrates the bias present when models fail to control for cohort-specific trends in the presence of staggered adoption timing and dynamic treatment effects. In column (1) we estimate a -4.5 percentage point change in average distance traveled due to the first four states' early mandates. Adding controls for adopters' pre-trends across timing cohorts in column (2), the treatment effect estimate on SAH_{it} changes sign and loses all statistical significance.⁶ Once we account for these early differences in COVID-19 outbreak trajectories between early and later-adopting cohorts, we fail to identify any differential effect of early adopters' mandates relative to later policies across all measures in columns (2) - (4).

⁶The same patterns hold true for both $N\dot{E}V$ and $E\dot{N}C$ when cohort pre-trends are included.

Columns (5) to (8) report equivalent estimates using adoption of all statewide mandates to identify the treatment effect estimate \widehat{ATT}^{SAH} . Here, with much greater variation in adoption timing, ATT estimates are identified both through comparisons of changes in treated units to changes in states that never adopted a mandate and through comparisons between various pairs of states treated at different times. Comparing column (5) with column (6), we once again see the change in point estimates when accounting for cohort pre-trends. In this case, however, controlling for differences in pre-trends does not eliminate statistical significance and yields an ATT estimate statistically indistinguishable from that in column (5) ($p = 0.15$).

In our preferred specifications for all statewide stay-at-home mandates in columns (6) to (8), we estimate large magnitude treatment effects relative to pre-mandate behavior changes. Looking at estimates for changes in average distances traveled, we observe a treatment effect estimate of -6.99 percentage points in column (6), statistically significant beyond the 1% level. That is, once a statewide mandate is implemented, we estimate a 6.99 percentage point reduction in the change in average distance traveled relative to control states. This represents an additional 140% decline relative to the pre-mandate change of -4.98, and an additional 24% reduction relative to average changes over the entire sample period (-29.29). In column (7) we estimate a -2.15 percentage point change in visits to non-essential businesses per day relative to control states (corresponding to 31% and 6% additional reductions relative to pre-mandate and full sample average changes, respectively). Turning next to changes in human encounter rates in column (8), we obtain an ATT estimate of a -3.506 percentage point decline per day after a mandate is implemented. Once again the treatment effects are statistically significant beyond the 1% level. This mandate effect corresponds to a 27% additional reduction relative to pre-mandate average changes of -14.85, and a 7% reduction relative to the -53.57 average change observed over the full sample period.

Decomposing the Difference-in-Differences Treatment Effect

A potential concern of the difference-in-differences estimator relates to the weighting of individual periods. Under staggered adoption, the estimated treatment effect can be expressed as a weighted average of all unique two-period by two-group difference-in-difference estimators (Goodman-Bacon 2018). Weights are implicitly assigned to each timing cohort and unit, proportional to the variance of the treatment indicator in each period and the size of each cross-sectional group. A key implication of these weights is a favoring of units treated near the middle of the sample period, with non-convexity indicating a potential for negative weights (Borusyak and Jaravel 2017; de Chaisemartin and D’Haultfœuille 2020; Sun and Abraham, in press). A resulting consequence is that negative (positive) treatment effects could also be obtained even when the effects of stay-at-home mandates for all adopting states are positive (negative) (Callaway and Sant’Anna 2020).

To shed light on the implicit weighting of the difference-in-differences ATT estimates presented in Table 3.2, we decompose the treatment effect estimates from columns (6) to (8) of Table 3.2 into their component two-by-two comparisons (Goodman-Bacon 2018). We find

that over half the overall difference-in-differences estimate’s weight is placed on comparisons of mandate states to never-adopter states. With adoption timing spanning March 19 to April 8, two-by-two comparisons can be made across many more cohorts and donor pools; in total, 18 comparisons are made between treatment cohorts and never-treated states, with 306 different comparisons between early and later adopters. More than half the overall ATT weight is given to comparisons of treatment cohorts versus pure control units, comprising 56-57% of the estimate across activity measures. The remaining weight is split evenly between comparisons of timing cohorts, with 21-22% of ATT weight given to comparisons of early treated units against later treated units still in the donor pool, and to later treated units post-treatment relative to previously-treated states. Except for early versus later treated units for average distances traveled, we observe consistently negative average ATT estimates across all three mobility measures, showing that treatment effect heterogeneity primarily concerns the size of reductions in travel activity rather than the sign of activity changes.

Unweighted and Weighted Event Studies of Mobility and Social Distancing

To investigate the dynamic nature of mobility and social distancing responses to stay-at-home mandates and to address concerns regarding imbalances in changes to mobility patterns during the pre-mandate period, we present results of unweighted and weighted event studies. First, unweighted event studies avoid the implicit weighting concerns of the difference-in-differences estimator and allow an understanding of how treatment effects evolve and persist over time. Next, weighted event studies extend these benefits and allow for the comparison of stay-at-home states to a synthetic control state balanced on pre-mandate mobility change trajectories.

Figures 3.2, 3.3, and 3.4 plot unweighted and weighted event study graphs for average distances traveled, visits to non-essential businesses, and the rate of human encounters. The left panel in each figure reports coefficient estimates and 95% confidence intervals from an unweighted event study following Eq. 3.2 using outcomes residualized of cohort pre-trends. The right panel reports results from an equivalent weighted event study with the weight between separate and pooled synthetic control weights set equal to $\nu = \sqrt{q^{pool}}/\sqrt{q^{sep}}$, the ratio of the square roots of pooled to separate SCM imbalance (Ben-Michael, Feller, and Rothstein 2019).⁷ The x-axis of each plot reports event time, indicating the number of days elapsed since a state’s stay-at-home mandate entered into effect. An event time of zero indicates the first full day a state’s mandate was in effect. The unweighted event study relies on the parallel trends assumption required for the overall difference-in-differences approach, while the doubly-robust approach of the weighted event study necessarily imposes balance on changes in pre-treatment outcomes.

⁷While an interior ν of 0.01 – 0.99 offers substantial imbalance reductions relative to the pooled or separate SCM cases, the optimal choice of ν is not immediately obvious. Estimating weighted event studies over the range of ν allows us to better understand how sensitive the overall ATT estimate is to the shift in weight from separate SCM for each state to a purely pooled SCM approach.

The unweighted event study (left panel) plots day-to-day ATT estimates averaged across all adopting states. These estimates are obtained from a two-way fixed effects control approach akin to columns (6), (7), and (8) in Table 3.2, with a vector of dummy variables for being each of $k \in \{-24, 21\}$ days relative to mandate adoption. The day prior to adoption ($k = -1$) is normalized to zero, such that all point estimates are interpreted as a differential change in a given travel outcome on the k^{th} day since mandate adoption relative to the day immediately preceding adoption. 95% confidence intervals clustered at the state level are reported in the gray band. Estimates statistically distinguishable from zero in the post-period measure the daily treatment effect of stay-at-home mandates on mobility patterns. Non-zero estimates in the pre-mandate period ($k < 0$) are evidence that the difference-in-differences parallel trends assumption is likely violated and that residents of adopting states were already differentially modifying their travel behavior relative to residents of control states prior to stay-at-home mandates requiring such modification.

The unweighted event study for average distance traveled (left panel of Figure 3.2) reveals the extent of differences in pre-trends for mandate versus control states. Despite controlling for cohort pre-trends, the unweighted event study (left panel) still displays large pre-treatment differences between mandate and control states, reflecting the patterns seen in Figure 3.1. In all periods 24 to 10 days prior to mandate adoption, adopting states display markedly greater average travel distances between 4 and 14 percentage points higher relative to control states. Pre-period point estimates remain positive but lose statistical significance for periods -9 to -2. Once a mandate is adopted, we observe an immediate reduction of 4 percentage points. This effect remains stable for 9 days before gradually attenuating over the remaining post periods and becoming indistinguishable from zero after 18 days.

The weighted event study results for average distance traveled (right panel of Figure 3.2) demonstrate the improved comparison achieved by internalizing pre-trends and its impact on treatment effect estimates. Using control units matched on prior mobility changes with $\nu = 0.6$ (60% of the weight given to the pooled synthetic controls and 40% to the individual controls) produces a greatly-improved overall pre-treatment match, with balance achieved in ten additional periods relative to the unweighted event study. Immediately after mandate adoption, average travel distances fall discontinuously by 10 percentage points in mandate states. In contrast to the unweighted event study, this estimated effect persists across all post-mandate periods. Averaging event day-specific ATT estimates across the entire mandate period yields an overall weighted event study ATT estimate of -8.97. Comparing to the previous estimates from Table 3.2, we find that eliminating the implicit weighting of difference-in-differences estimators and internalizing pre-trends yields an ATT estimates roughly 25-60% greater in magnitude.

Turning to estimates for changes in non-essential visits in Figure 3.3 and human encounter rates in Figure 3.4, we observe greater similarity between unweighted and weighted event study estimates. Pre-treatment point estimates are statistically insignificant in all periods for non-essential visits and in 17 of 23 periods for human encounter rates. The unweighted event study for non-essential visits in the left panel of Figure 3.3 shows a 6 percentage point reduction on the first day of mandate adoption that rapidly attenuates and becomes sta-

tistically insignificant after 8 days. The weighted event study in the right panel displays greater balance with narrow confidence intervals in all pre-treatment periods, with smaller magnitude point estimates that are lower variance but statistically indistinguishable from those in the unweighted event study. The unweighted event study for human encounter rate changes in Figure 3.4 presents some evidence of differential trends, followed by an immediate and persistent treatment effect between -5.5 to -4.5 percentage points per day throughout the post-adoption period. Pre-treatment balance once again improves in the weighted event study, with no distinguishable difference between mandate states and their synthetic controls in any pre-treatment period. The weighted event study yields treatment effect point estimates that begin slightly larger in magnitude than but statistically indistinguishable from those in the unweighted event study. The weighted model estimates statistically significant declines in human encounter rates between -3 and -6 percentage points for the first ten days following mandate adoption. Although point estimate magnitudes never rise above -2.5 for days 11 to 21, increased noise relative to the unweighted event study leads to statistically insignificant estimates in all but two of these eleven periods.

Once again overall weighted event study ATT estimates are in line with our findings from difference-in-differences models. The weighted event study for non-essential visits yields an overall mandate effect of -0.84 (compared to -2.15 for column (7) of Table 3.2), with an overall ATT of -4.14 for encounter rates (compared to -3.51 for column (8) of Table 3.2). These estimates are highly robust to the specific choice of ν ; overall ATT estimates fall between -9.94 and -8.96 for average distance traveled, between -1.08 and -0.81 for non-essential visits, and between -4.17 and -3.51 for human encounter rates across the space of $\nu \in [0, 1]$.

The consistency of both weighted and unweighted event study estimates with results from difference-in-differences approaches provides confirming evidence that statewide stay-at-home mandates elicited further reductions in travel activity by affected residents. Event studies yield additional detail as to how these responses evolved, showing that reductions occurred immediately upon policy implementation and largely persisted even as residents were subject to the policies for three full weeks. This pattern is especially true for average distance traveled and human encounters, suggesting that residents of mandate states continued to socially distance, a key avenue for reducing the potential transmission of COVID-19. The persistence of mandate effects under the weighted event study approach is confirmation that our findings actual, behavioral responses and not merely the result of pre-trend imbalance.

3.5 Conclusion

Temporarily closing non-essential businesses and mandating that residents stay at home except for essential activity is the prime policy instrument currently employed by states to promote social distancing and slow the spread of COVID-19. If effective, these policies will have reduced strain on the medical system and provided much-needed time for the development of pharmaceutical treatments that can reduce transmission rates and end the pandemic. If unsuccessful, states will have incurred large economic costs with few lives

saved. Whether these mandates cause people to stay at home and engage in social distancing is a key requirement of a successful policy. Knowing whether such policies will have their intended effect is of increasing policy relevance, as all but eight states eventually adopted such policies. Understanding whether and how individuals reduce travel activity and maintain social distance in response to stay-at-home mandates is the primary empirical question we tackle in this paper. We establish two empirical findings.

First, we find that, by the time the average adopter had implemented its statewide mandate, residents had already reduced travel by considerable amounts relative to pre-COVID-19 levels. Average travel distances had already fallen by 16 percentage points, human encounter rates by 63 percentage points, and non-essential visits by 39 percentage points before the first statewide mandate came into effect, providing evidence of extensive social distancing occurring even before such behavior was required by statewide orders.

Second, we find evidence that adoption of state-level stay-at-home mandates induced further reductions across three travel activity measures. The staggered difference-in-differences models estimate a reduction in average distance traveled of 5.51 percentage points, a decline in visits to non-essential businesses of 2.15 percentage points, and a decrease in the rate of unique human encounters of 3.51 percentage points relative to pre-COVID-19 baselines. Estimated magnitudes remain highly comparable when directly accounting for differences in pre-mandate behavior for treatment and control states. Through the weighted event studies that construct control units to balance pre-treatment travel behavior net of state fixed effects, we find large, statistically significant drops immediately following mandate implementation across all measures that persist for the duration of the sample period for distances traveled and human encounters. Resulting estimates of the overall mandate effects mirror those obtained from the difference-in-differences models, and are similarly sized for any mix of pooled and separate synthetic control weights.

As travel activity is a main source of social interaction beyond one's immediate family (Silvis, Niemeier, and D'Souza 2006) and travel to non-work locations increases the probability of co-location with others, these reductions in distances traveled likely reflect commensurate decreases in physical interactions with those outside of one's immediate family. Our estimates for changes in unique human encounters support this notion, providing evidence of further social distancing once states adopted a stay-at-home mandate. Further, these findings are not limited to the Unacast mobility measures; use of Google's COVID-19 Community Mobility Reports estimates similarly large and statistically significant effects of statewide mandates. All this provides consistent evidence that, during their first three weeks, stay-at-home mandates have had the intended effect of inducing greater social distancing than would have occurred otherwise, helping to reduce the opportunities for communication of COVID-19 within adopting states.

Our work complements that of Gupta et al. (2021), who similarly explore changes in mobility patterns in response to COVID-19 policies. The authors find comparably large changes to a broad set of mobility measures in March 2020 prior to the adoption of county and state-level stay-at-home mandates (Gupta et al. 2021). Some state and many county-level event studies that they estimate reveal parallel trend violations; our work demonstrates


techniques to correct for differences in pre-period mobility trends. Our weighted event study results show the sensitivity of treatment effect estimates to potential imbalances in comparison groups in traditional difference-in-difference and event study methods.

Overall, our estimates suggest that residents subject to stay-at-home mandates on average responded as desired to curb the spread of COVID-19. Our empirical approaches isolate the mandate effect from other drivers of daily changes in travel activity levels and from pre-existing trends, and control for a host of potential confounding factors that differ between states that adopted policies relative to other states and those yet to adopt policies. Under these rigorous control approaches, we find persistent evidence of state mandates inducing further reductions in travel activity even after considerable earlier declines around the country. Further, our estimates are average treatment effects in response to statewide mandates only. Given the extent of prior school closures, new work from home abilities, and county-level stay-at-home policies, our findings represent only a portion of the way individuals responded to COVID-19 policies. As a result, our estimates represent a considerable lower bound on individuals' comprehensive responses to all COVID-19 policies.

These findings have important policy implications for the fight against COVID-19. First, individuals on average responded as intended to statewide mandates. Despite considerable prior reductions, residents heeded their states' directives and reduced travel activity. Second, the declines in economic activity directly attributable to statewide mandates may be much smaller than previously thought. Because individuals around the country had already more than halved the quantity of trips taken to non-essential retail and service businesses, much of the lost business and resulting unemployment would have likely still occurred even if states had not adopted their stay-at-home policies. Further, as the mandate-induced reductions in visits to non-essential businesses amount to only one-tenth of the overall reductions since COVID-19 arose, it is likely that loosening or removing statewide policies may not be sufficient to induce mobility patterns to quickly return to pre-COVID-19 levels. Further policies will be needed to ensure that individuals can safely resume activity and return to local businesses.


Future work can establish whether these changes in mobility have significant health effects, taking into account the benefits from avoided hospitalizations and other indirect health benefits from reduced travel activity and social distancing. Because reductions in travel distance and increased social distance likely decrease exposure to other potentially deadly illnesses, estimates of the health benefits due to stay-at-home policies likely underestimate their overall impact. We support continued efforts to obtain accurate counts of the mortality and morbidity consequences from COVID-19 to help ensure future research can provide sufficient policy guidance in the case of future pandemics.

Figure 3.1: Changes in Travel Activity and Social Distancing




Each series represents the change in each day’s mobility measure relative to pre-COVID-19 levels for the given group of states. The solid line corresponds to the average change for the four states that implemented stay-at-home mandates by end-of-day March 22 (California, Illinois, New Jersey, and New York). The dotted line plots the average for the 39 states that adopted statewide mandates at later points, while the dashed line represents the average for the eight states that never adopted a statewide mandate. The first panel plots changes in average distance traveled, the second changes in unique human encounters per square kilometer, and the third changes in visits to non-essential businesses. The gray bars designate weekend days. The vertical lines indicate the dates of the first and last statewide stay-at-home mandates (March 19 and April 8).

Figure 3.2: Unweighted and Weighted Event Studies for Changes in Average Distances Traveled




The unweighted event study (left panel) plots regression coefficients for dummy variables equal to one for being k days away from the first effective date of each statewide stay-at-home mandate, with 95% confidence intervals represented in the gray band. A point estimate of -10 indicates a 10 percentage point greater decline in the average distance traveled per day for a state k days since mandate adoption relative to the day prior to mandate adoption ($k = -1$). The right panel plots equivalent point estimates and jackknife 95% confidence intervals from a weighted event study, with partially pooled synthetic controls constructed to match treated units on residualized pre-treatment outcomes with 60% of synthetic control weights obtained from pooled versus individual synthetic control weights ($\nu = 0.6$).

Figure 3.3: Unweighted and Weighted Event Studies for Changes in Visits to Non-Essential Businesses



The unweighted event study (left panel) plots regression coefficients for dummy variables equal to one for being k days away from the first effective date of each statewide stay-at-home mandate, with 95% confidence intervals represented in the gray band. A point estimate of -10 indicates a 10 percentage point greater decline in the average distance traveled per day for a state k days since mandate adoption relative to the day prior to mandate adoption ($k = -1$). The right panel plots equivalent point estimates and jackknife 95% confidence intervals from a weighted event study, with partially pooled synthetic controls constructed to match treated units on residualized pre-treatment outcomes with 21% of synthetic control weights obtained from pooled versus individual synthetic control weights ($\nu = 0.21$).

Figure 3.4: Unweighted and Weighted Event Studies for Changes in the Unique Human Encounter Rate



The unweighted event study (left panel) plots regression coefficients for dummy variables equal to one for being k days away from the first effective date of each statewide stay-at-home mandate, with 95% confidence intervals represented in the gray band. A point estimate of -10 indicates a 10 percentage point greater decline in the average distance traveled per day for a state k days since mandate adoption relative to the day prior to mandate adoption ($k = -1$). The right panel plots equivalent point estimates and jackknife 95% confidence intervals from a weighted event study, with partially pooled synthetic controls constructed to match treated units on residualized pre-treatment outcomes with 18% of synthetic control weights obtained from pooled versus individual synthetic control weights ($\nu = 0.18$).

Table 3.1: Summary Statistics on Travel Behavior and Social Distancing

	$\dot{A}DT$ Distance Traveled	$\dot{N}EV$ Non-Essential Visits	$\dot{E}NC$ Human Encounters
<u>Before Mandates, February 24th to March 8th</u>			
Average	0.79	-1.78	0.00
25th percentile	-1.59	-5.09	-10.79
Median	0.77	-1.93	-2.35
75th percentile	3.54	1.63	9.92
<u>March 8 to March 31st</u>			
Average	-25.43	-41.26	-60.14
25th percentile	-37.51	-59.51	-77.40
Median	-25.38	-48.21	-67.85
75th percentile	-12.93	-21.06	-48.66
<u>April 1 to April 29th</u>			
Average	-40.92	-59.02	-78.85
25th percentile	-47.87	-65.44	-85.25
Median	-39.98	-58.92	-80.53
75th percentile	-33.33	-53.13	-73.58
<u>Total Sample</u>			
Average	-26.59	-40.43	-55.22
Median	-30.93	-51.75	-71.59
Standard Deviation	19.98	26.23	34.87
N	3366	3366	3300

Source: Unacast. This table reports summary statistics for the changes in the average distance traveled $\dot{A}DT$ (Column 1), non-essential visits $\dot{N}EV$ (Column 2) and rate of unique human encounters $\dot{E}NC$ (Column 3). Data cover the period from February 24th to April 29th. Each observation is measured at the state-by-day level and represents an aggregate of mobile device-level travel and social distancing behavior on a given day.

Table 3.2: Statewide Stay-at-Home Mandates, Travel Activity, and Social Distancing

	Early SAH States				All SAH States			
	$\dot{A}DT$ (1)	$\dot{A}DT$ (2)	$\dot{N}EV$ (3)	$\dot{E}NC$ (4)	$\dot{A}DT$ (5)	$\dot{A}DT$ (6)	$\dot{N}EV$ (7)	$\dot{E}NC$ (8)
SAH_{it}	-4.45** (2.07)	2.50 (2.11)	2.63 (1.93)	3.05 (4.60)	-5.51*** (1.04)	-.99*** (1.45)	-2.15** (0.88)	-3.51*** (0.98)
\bar{Y}	-26.59	-29.29	-36.23	-53.57	-26.59	-29.29	-36.23	-53.57
\bar{Y}^{pre}	-3.98	-4.98	-7.02	-14.85	-3.98	-4.98	-7.02	-14.85
N	3,366	3,366	3,366	3,300	3,366	3,366	3,366	3,300
Adj. R ²	0.93	0.90	0.96	0.95	0.93	0.90	0.96	0.95

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the state level. These models estimate the effect of statewide stay-at-home mandates on travel activity and social distancing. The dependent variables measure the percentage point changes in average distances traveled, visits to non-essential businesses, and unique human encounters for the same day of the week relative to the pre-COVID-19 baseline level (average of Feb 10 - Mar 8). A coefficient of one indicates a marginal effect of a 1 percentage point increase in travel relative to pre-COVID-19 levels, controlling for time and the average COVID-19 mobility change in the state during the sample period. All models include state and date fixed effects and residualize outcomes of timing cohort pre-trends (Goodman-Bacon 2018).

Bibliography

- (BAWSCA), Bay Area Water Supply Conservation Agency. 2021. Lawn be gone! <https://bawasca.org/conserve/rebates/lawn>.
- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. Synthetic control methods for comparative case studies: estimating the effect of california's tobacco control program. *Journal of the American Statistical Association* 105 (490): 493–505. <https://doi.org/10.1198/jasa.2009.ap08746>.
- Alameda County Community Food Bank. 2022. About us: food distribution. Accessed April 15, 2022. <https://www.accfb.org/about-us/food-distribution/>.
- Allcott, Hunt. 2011. Social norms and energy conservation. *Journal of Public Economics* 95 (9): 1082–1095.
- Andreoni, James. 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. *The Economic Journal* 100:464–477.
- Arkhangelsky, Dmitry, and Guido W. Imbens. 2019. *Double-robust identification for causal panel data models*. arXiv: 1909.09412 [econ.EM].
- Asci, Serhat, and Tatiana Borisova. 2014. The effect of price and non-price conservation programs on residential water demand. Selected paper, 2014 AAEA Annual Meeting, Minneapolis, MN.
- Aten, Jason. 2020. Microsoft, google, and twitter are telling employees to work from home because of coronavirus. should you? *INC* (March 6, 2020). <https://www.inc.com/jason-aten/microsoft-google-twitter-are-telling-employees-to-work-from-home-because-of-coronavirus-should-you.html>.
- Baker, Mike. 2020. Coronavirus slowdown in seattle suggests restrictions are working. *New York Times* (March 29, 2020). Accessed April 1, 2020. <https://www.nytimes.com/2020/03/29/us/seattle-washington-state-coronavirus-transmission-rate.html>.
- Ben-Michael, Eli, Avi Feller, and Jesse Rothstein. 2019. *Synthetic controls and weighted event studies with staggered adoption*. arXiv: 1912.03290 [stat.ME].

- Blackman, Allen, Shakeb Afsah, and Damayanti Ratunanda. 2004. How do public disclosure pollution control programs work? evidence from indonesia. *Human Ecology Review* 11 (3): 235–246.
- Blomquist, Sören, and Håkan Selin. 2010. Hourly wage rate and taxable labor income responsiveness to changes in marginal tax rates. *Journal of Public Economics* 94 (11-12): 878–889.
- Blundell, Richard, Alan Duncan, and Costas Meghir. 1998. Estimating labor supply responses using tax reforms. *Econometrica* 66 (4): 827–861.
- Borusyak, Kirill, and Xavier Jaravel. 2017. *Revisiting event study designs*. Working Paper. <https://doi.org/10.2139/ssrn.2826228>.
- Brent, Daniel A., Joseph H. Cook, and Skylar Olsen. 2015. “do the right thing:” the effects of moral suasion on cooperation. *Journal of the Association of Environmental and Resource Economists* 2 (4): 597–627.
- Briefel, Ronette, Jonathan Jacobson, Nancy Clusen, Teresa Zavitsky, Miki Satake, Brittany Dawson, and Rhoda Cohen. 2003. The emergency food assistance system – findings from the client survey. *USDA-ERS Electronic Publications from the Food Assistance and Nutrition Research Program*, no. E-FAN-03-007.
- Buck, Steven, Maximilian Auffhammer, Hilary Soldati, and David Sunding. 2020. Forecasting residential water consumption in california: rethinking model selection. *Water Resources Research* 56.
- California Budget and Policy Center. 2014. Use of Food Assistance Varies Widely Across California Counties. https://calbudgetcenter.org/app/uploads/140731_CalFresh_TableAndMap.pdf.
- California Department of Social Services. 2022. Eligibility and issuance requirements. Accessed April 15, 2022. <https://www.cdss.ca.gov/inforesources/cdss-programs/calfresh/eligibility-and-issuance-requirements>.
- California Department of Water Resources. 2015. Executive update: hydrologic conditions in california (04/01/2015), accessed June 6, 2020. <http://cdec.water.ca.gov/cgi-progs/reports/EXECSUM>.
- California Executive Order B-29-15. 2015. https://www.waterboards.ca.gov/waterrights/water_issues/programs/drought/docs/040115_executive_order.pdf.
- California Executive Order N-33-20. 2020. <https://covid19.ca.gov/img/Executive-Order-N-33-20.pdf>.
- Callaway, Brantly, and Pedro H. C. Sant’Anna. 2020. *Difference-in-differences with multiple time periods*. arXiv: 1803.09015 [econ.EM].

- Castellari, Elena, Chad Cotti, John Gordanier, and Orgul Ozturk. 2017. Does the timing of food stamp distribution matter? a panel-data analysis of monthly purchasing patterns of us households. *Health Economics* 26 (11): 1380–1393.
- Castner, Laura, and Juliette Henke. 2011. Benefit redemption patterns in the supplemental nutrition assistance program. U.S. Department of Agriculture, Food and Nutrition Service, Office of Research and Analysis. <https://fns-prod.azureedge.us/sites/default/files/ARRASpendingPatterns.pdf>.
- CDC. 2020. *CDC updates COVID-19 transmission webpage to clarify information about types of spread*. Centers for Disease Control and Prevention. <https://www.cdc.gov/media/releases/2020/s0522-cdc-updates-covid-transmission.html>.
- Chernozhukov, Victor, Kaspar Wuthrich, and Yinchu Zhu. 2020. *Practical and robust t-test based inference for synthetic control and related methods*. arXiv: 1812.10820 [econ.EM].
- City of Sacramento. 2021. Landscape rebates. <https://www.valleywater.org/saving-water/rebates-surveys/landscape-rebates>.
- Coleman-Jensen, Alisha, Matthew P. Rabbitt, Christian A. Gregory, and Anita Singh. 2021. Household food security in the united states in 2020. *USDA-ERS Economic Research Report*, no. 298.
- Colin, Cameron A., and Douglas L. Miller. 2015. A practitioner’s guide to cluster-robust inference. *The Journal of Human Resources*, 50 (2): 317–372.
- covid19.ca.gov. 2020. *Coronavirus (covid-19) in california stay home order faq*. Accessed April 5, 2020. <https://covid19.ca.gov/stay-home-except-for-essential-needs/#top>.
- Cuff, Denis, and Paul Rogers. 2015. Big difference in water use between wealthy california communities and everyone else. *East Bay Times* (June 20, 2015). Accessed June 6, 2020.
- Cunyngham, Karen. 2020. Reaching those in need: estimates of state supplemental nutrition assistance program participation rates in 2017. *USDA Report*.
- Dal Bó, Ernesto, and Pedro Dal Bó. 2014. “do the right thing:” the effects of moral suasion on cooperation. *Journal of Public Economics* 117:28–38.
- Damon, Amy L., Robert P. King, and Ephraim Leibtag. 2013. First of the month effect: does it apply across food retail channels? *Food Policy* 41:18–27.
- Coronavirus disease (COVID-19) dashboard*. 2020. <https://covid19.who.int/>. World Health Organization. Accessed March 21, 2020.
- de Chaisemartin, Clément, and Xavier D’Haultfoeuille. 2020. *Two-way fixed effects estimators with heterogeneous treatment effects*. Technical report. arXiv: 1803.08807 [econ.EM].

- Dorfman, Jeffrey H., Christian Gregory, Zhongyuan Liu, and Ran Huo. 2018. Re-examining the snap benefit cycle allowing for heterogeneity. *Applied Economics Perspectives and Policy* 41 (3): 404–433.
- Dwenger, Nadja, and Lukas Treber. 2018. Hohenheim Discussion Papers in Business, Economics and Social Sciences 21-2018.
- EBMUD. 2021. Save like a pro. Accessed September 27, 2021. <https://www.ebmud.com/water/conservation-and-rebates/residential/save-pro/>.
- EPA, U.S. 2021. Watersense products: showerheads. Accessed September 26, 2021. <https://www.epa.gov/watersense/showerheads>.
- Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, et al. 2020. Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. *Imperial College COVID-19 Response Team*, <https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf>.
- Ferraro, Paul J., Juan Jose Miranda, and Michael K. Price. 2011. The persistence of treatment effects with norm-based policy instruments: evidence from a randomized environmental policy experiment. *American Economic Review: Papers Proceedings* 101 (3): 318–322.
- Finucane, Martin, and Travis Andersen. 2020. Experts say relaxing social distancing controls could be ‘catastrophic’. *The Boston Globe* (March 24, 2020). <https://www.bostonglobe.com/2020/03/24/metro/experts-say-its-too-soon-relax-social-distancing/>.
- Flamer, Keith. 2021. Best eco-friendly washing machines from consumer reports’ tests. *Consumer Reports* (May 26, 2021). Accessed September 26, 2021. <https://www.consumerreports.org/washing-machines/best-eco-friendly-washing-machines-consumer-reports-tests-a5430738463/>.
- Goldin, Jacob, Tatiana Homonoff, and Katherine Meckel. 2022. Issuance and incidence: snap benefit cycles and grocery prices. *American Economic Journal: Economic Policy* 14 (1): 152–178.
- Goodman-Bacon, Andrew. 2018. *Difference-in-differences with variation in treatment timing*. Working Paper, Working Paper Series 25018. National Bureau of Economic Research. <https://doi.org/10.3386/w25018>. <http://www.nber.org/papers/w25018>.
- . in press. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.
- Google. 2020. *Google COVID-19 community mobility reports*. <https://www.google.com/covid19/mobility/>.

- Greenstone, Michael, and Vishan Nigam. 2020. Does social distancing matter? *Becker Friedman Institute Working Papers* 2020 (26). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3561244.
- Gregory, Christian A., and Travis A. Smith. 2019. Saliency, food security, and snap receipt. *Journal of Policy Analysis and Management* 38 (1): 124–154.
- Gupta, Sumedha, Thuy Nguyen, Shyam Raman, Byungkyu Lee, Felipe Lozano-Rojas, Ana Bento, Kosali Simon, and Coady Wing. 2021. Tracking public and private responses to the covid-19 epidemic: evidence from state and local government actions. *American Journal of Health Economics* 7 (4). <https://doi.org/10.1086/716197>.
- Hamrick, Karen S., and Margaret Andrews. 2016. Snap participants' eating patterns over the benefit month: a time use perspective. *PLOS ONE* 11 (7): e0158422.
- Hastings, Justine, and Ebonya Washington. 2010. The first of the month effect: consumer behavior and store responses. *American Economic Journal: Economic Policy* 2 (2): 146–162.
- Hatchett, Richard J., Carter E. Mecher, and Marc Lipsitch. 2007. Public health interventions and epidemic intensity during the 1918 influenza pandemic. *Proceedings of the National Academy of Sciences* 104 (18): 7582–7587. <https://doi.org/10.1073/pnas.0610941104>.
- He, Chunyang, Zhifeng Liu, Jianguo Wu, Xinhao Pan, Zihang Fang, Jingwei Li, and Brett A. Bryan. 2021. Future global urban water scarcity and potential solutions. *Nature Communications* 12.
- Hoynes, Hilary, Erin Bronchetti, and Garret Christensen. 2017. The real value of SNAP benefits and health outcomes. University of Kentucky Center for Poverty Research Discussion Paper Series 104. <https://fns-prod.azureedge.us/sites/default/files/ARRASpendingPatterns.pdf>.
- Hoynes, Hilary, and Diane Whitmore Schanzenbach. 2015. NBER Working Paper Series 21057. <https://doi.org/10.3386/w21057>.
- IGM. 2020. *IGM economic experts panel, policy for the covid-19 crisis*, March 27, 2020. <http://www.igmchicago.org/igm-economic-experts-panel/>.
- IPCC. 2021. “Summary for Policymakers.” In: *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)].
- Ito, Koichiro. 2014. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review* 104 (2): 537–563.

- Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. Moral suasion and economic incentives: field experimental evidence from energy demand. *American Economic Journal: Economic Policy* 10 (1): 240–267.
- Johnson, Matthew S. 2020. Regulation by shaming: deterrence effects of publicizing violations of workplace safety and health laws. *American Economic Review* 110 (6): 1866–1904.
- Kadvany, Elena. 2020. Stanford tells 7,000 undergrads to leave campus; class will be online only next quarter. *Palo Alto Online* (March 6, 2020). Accessed March 6, 2020. <https://paloaltoonline.com/news/2020/03/06/stanford-cancels-in-person-classes-two-students-possibly-exposed-to-coronavirus-in-self-isolation>.
- Kahan, Dan M., and Eric A. Posner. 1999. Shaming white-collar criminals: a proposal for reform of the federal sentencing guidelines. *The Journal of Law & Economics* 42 (S1): 365–392.
- Kahn, Debra, and Carla Marinucci. 2020. Bend it like the bay area: doctors see flatter curve after 2 weeks of social isolation. *Politico* (March 30, 2020). Accessed April 1, 2020. <https://www.politico.com/states/california/story/2020/03/30/bend-it-like-the-bay-area-doctors-see-flatter-curve-after-2-weeks-of-social-isolation-1269663>.
- Kharmats, Anna Y, Jessica C Jones-Smith, Yun Sang Cheah, Nadine Budd, Laura Flamm, Alison Cuccia, Yeeli Mui, Angela Trude, and Joel Gittelsohn. 2014. Relation between the supplemental nutritional assistance program cycle and dietary quality in low-income african americans in baltimore, maryland. *The American Journal of Clinical Nutrition* 99 (5): 1006–1014.
- Klaiber, H. Allen, V. Kerry Smith, Michael Kaminsky, and Aaron Strong. 2014. Measuring price elasticities for residential water demand with limited information. *Land Economics* 90 (1): 100–113.
- Laurito, Agustina, and Amy Ellen Schwartz. 2019. Does school lunch fill the “snap gap” at the end of the month? *Southern Economic Journal* 86 (1): 49–822.
- Mann, Michael E., and Peter H. Gleick. 2015. Climate change and california drought in the 21st century. *Proceedings of the National Academy of Sciences* 112 (13): 3858–3859.
- Marples, Carol Ann, and Diana-Marie Spillman. 1995. Factors affecting students’ participation in the cincinnati public schools lunch program. *Adolescence* 30 (119): 745–754.
- McCabe, Caitlin. 2020. Analysts slash gdp estimates as coronavirus ripples through economy. *The Wall Street Journal* (March 20, 2020). Accessed March 21, 2020. <https://www.wsj.com/articles/analysts-slash-gdp-estimates-as-coronavirus-ripples-through-economy-11584735139>.

- Mervosh, Sarah, Denise Lu, and Vanessa Swales. 2020. See which states and cities have told residents to stay at home. *New York Times* (March 23, 2020). Accessed March 26, 2020. <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html>.
- Nataraj, Shanthi, and W. Michael Hanemann. 2011. Does marginal price matter? a regression discontinuity approach to estimating water demand. *Journal of Environmental Economics and Management* 61:198–212.
- Nemati, Mehdi, Steven Buck, and Hilary Soldati. 2017. *The Effect of Social and Consumption Analytics on Residential Water Demand*. 2017 Annual Meeting, February 4-7, 2017, Mobile, Alabama 252738. Southern Agricultural Economics Association.
- New Jersey. 2020. *Executive order no. 107*, March 21, 2020. <https://nj.gov/infobank/eo/056murphy/pdf/EO-107.pdf>.
- Olmstead, Sheila M. 2009. Reduced-form versus structural models of water demand under nonlinear prices. *Journal of Business & Economic Statistics* 27 (1): 84–94.
- Pepe, Emanuele, Paolo Bajardi, Laetitia Gauvin, Filippo Privitera, Brennan Lake, Ciro Cattuto, and Michele Tizzoni. 2020. Covid-19 outbreak response: a first assessment of mobility changes in italy following national lockdown.
- Perez-Truglia, Ricardo, and Ugo Troiano. 2014. NBER Working Paper 21264. <http://www.nber.org/papers/w21264>.
- Pratt, Bryan. in press. A fine is more than a price: evidence from drought restrictions. *Journal of Environmental Economics and Management*.
- PRISM Climate Group, Oregon State University. Created 4 Feb 2004. <http://prism.oregonstate.edu>.
- Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz. 2012. The elasticity of taxable income with respect to marginal tax rates: a critical review. *Journal of Economic Literature* 50 (1): 3–50.
- Sanjeevi, Namrata, Jeanne Freeland-Graves, and Goldy Chacko George. 2018. Relative validity and reliability of a 1-week, semiquantitative food frequency questionnaire for women participating in the supplemental nutrition assistance program. *Journal of the Academy of Nutrition and Dietetics* 117 (12): 1972–1982.
- Schlenker, Wolfram, and Jason Scorse. 2011. Does being a “top 10” worst polluter affect environmental releases? evidence from the u.s. toxic release inventory. Presented at the Stanford Institute for Theoretical Economics.
- Schmidheiny, Kurt, and Sebastian Siegloch. 2019. *On event study designs and distributed-lag models: equivalence, generalization and practical implications*. CESifo Working Paper Series 7481. CESifo. https://ideas.repec.org/p/ces/ceswps/_7481.html.

- Shapiro, Jesse M. 2005. Is there a daily discount rate? evidence from the food stamp nutrition cycle. *Journal of Public Economics* 89 (2–3): 303–325.
- Silvis, Julia, Deb Niemeier, and Raissa D’Souza. 2006. *Social networks and travel behavior: report from an integrated travel diary*. 11th International Conference on Travel Behavior Research: Tokyo.
- Smith, Travis A., Joshua P. Berning, Xiaosi Yang, Gregory Colson, and Jeffrey H. Dorfman. 2016. The effects of benefit timing and income fungibility on food purchasing decisions among supplemental nutrition assistance program households. *American Journal of Agricultural Economics* 98 (2): 564–580.
- SoCal WaterSmart. 2021. Turf replacement program. <https://socalwatersmart.com/en/residential/rebates/available-rebates/turf-replacement-program/>.
- Sun, Liyang, and Sarah Abraham. in press. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.
- Tarasuk, Valerie, Lynn McIntyre, and Jinguang Li. 2007. Low-income women’s dietary intakes are sensitive to the depletion of household resources in one month. *The Journal of Nutrition* 137 (8): 1980–1987.
- Tiehen, Laura, Constance Newman, and John A. Kirlin. 2017. The food-spending patterns of households participating in the supplemental nutrition assistance program: findings from usda’s foodaps. *USDA-ERS Economic Information Bulletin*, no. 176.
- Todd, Jessica E. 2015. Revisiting the supplemental nutrition assistance program cycle of food intake: investigating heterogeneity, diet quality, and a large boost in benefit amounts. *Applied Economic Perspectives and Policy* 37 (3): 437–458.
- Todd, Jessica E., and Christian Gregory. 2018. Changes in supplemental nutrition assistance program real benefits and daily caloric intake among adults. *Food Policy* 79:111–120.
- Tsikas, Stefanos A. 2021. Bringing tax avoiders to light: moral framing and shaming in a public goods experiment. *Behavioural Public Policy*, 1–24.
- U.S. Bureau of Economic Analysis. 2021. *Gross domestic product: implicit price deflator [gdpdef]*. <https://fred.stlouisfed.org/series/GDPDEF>.
- Unacast. 2020. *Covid-19 toolkit, social distancing dashboard*. <https://www.unacast.com/covid19/social-distancing-scoreboard>.
- USDA ERS. 2022. Foodaps national household food acquisition and purchase survey. Accessed April 19, 2022. <https://www.ers.usda.gov/data-products/foodaps-national-household-food-acquisition-and-purchase-survey/>.

- Valley Water. 2021. River-friendly landscape (turf replacement, smart controller, irrigation upgrade) rebates. <https://www.valleywater.org/saving-water/rebates-surveys/landscape-rebates>.
- Weinfield, Nancy S., Gregory Mills, Christine Borger, Maeve Gearing, Theodore Macaluso, Jill Montaquila, and Sheila Zedlewsk. 2014. Hunger in America 2014: National Report Prepared for Feeding America. <http://help.feedingamerica.org/HungerInAmerica/hunger-in-america-2014-full-report.pdf>.
- Whiteman, Eliza D., Benjamin W. Chrisinger, and Amy Hillier. 2018. Diet quality over the monthly supplemental nutrition assistance program cycle. *American Journal of Preventive Medicine* 55 (2): 205–212.
- Wichman, Casey J. 2014. Perceived price in residential water demand: evidence from a natural experiment. *Journal of Economic Behavior & Organization* 107:308–323.
- Wichman, Casey J., Laura O. Taylor, and Roger H. von Haefen. 2016. Conservation policies: who responds to price and who responds to prescription? *Journal of Environmental Economics and Management* 79:114–134.
- Wilde, Parke E., and Christine K. Ranney. 2000. The monthly food stamp cycle: shopping frequency and food intake decisions in an endogenous switching regression framework. *American Journal of Agricultural Economics* 82 (1): 200–213.
- Woodward, Jenny, Pinki Sahota, Jo Pike, and Rosie Molinari. 2015. Interventions to increase free school meal take-up. *Health Education* 115 (2): 197–213.
- Yoo, James, Silvio Simonit, Ann P. Kinzig, and Charles Perring. 2014. Estimating the price elasticity of residential water demand: the case of phoenix, arizona. *Applied Economic Perspectives and Policy* 36 (2): 333–350.

Appendix A

Appendix to Chapter 1

Figure A.1: East Bay Municipal Utility District (EBMUD) Warning Letter



September 3, 2015

IMPORTANT INFORMATION - EXCESSIVE WATER USE ORDINANCE

Dear Customer:

RE: Account No. [REDACTED]

As you are aware, California is in a critical drought. EBMUD and the State of California have declared a critical drought emergency and adopted mandatory restrictions on outdoor water use. EBMUD also has adopted an Excessive Water Use Ordinance that prohibits single family residences from using more than 80 units per two month billing period (approximately 1,000 gallons per day). You are receiving this letter because EBMUD is reaching out to all residential water users who our records show have consistently used a high volume of water (more than 750 gallons per day) in recent months, which means the household is exceeding or could exceed the allowable limit. We want to work with you to help you conserve water. Sometimes high use is the result of unnoticed leaks and we urge you to check your water system for leaks and let us know if you find and repair a leak.


We realize some high-volume users have made substantial investments in water conserving fixtures and landscapes and if you have done that, thank you. However, water use cutbacks are required to ensure that adequate water resources remain available for all customers during this time of severe drought that potentially will continue into 2016. You may have received a letter in June advising you of EBMUD's Excessive Water Use Penalty Ordinance (no. 364-15) and encouraging you to take additional measures to conserve. Today's letter reminds you, as a customer whose recent water use has been high, that **water use above 80 units (approximately 1,000 gallons per day) during a two-month billing cycle is prohibited and has consequences.**

Financial Penalty. Excessive use will result in a financial penalty of \$2.00 for each unit in excess of the 80 unit maximum. If you receive a water bill with any applicable penalty, that bill serves as the notice of violation of EBMUD's rules and regulations.

Public Disclosure Requirements. The California Public Records Act (Government Code sections 6250-6276.48) requires EBMUD to release the name, home address and water use data of customers using water in violation of EBMUD's regulations and ordinance if that data is requested pursuant to a Public Records Act request. In the past EBMUD has received public

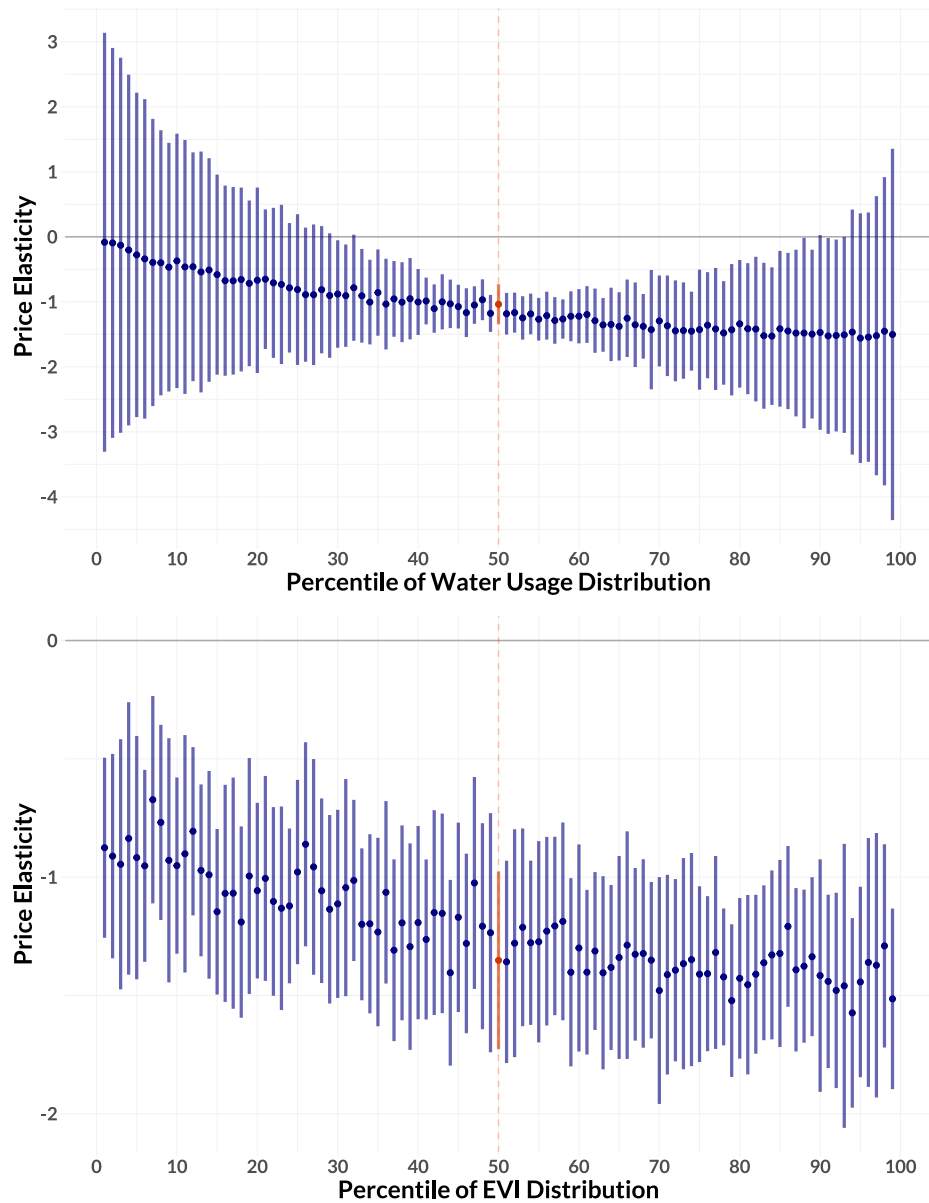
Source: EBMUD. This figure shows a copy of the warning letter sent to high-usage households during the Excessive Use Penalty Ordinance (EUP), informing them of the consequences of violating the policy and need to conserve water.

Figure A.2: Average and Marginal Water Prices over Time.



This figure plots the mean average (top panel) and marginal (bottom panel) prices per one hundred cubic foot unit (CCF) of water for detached single family homes in the EBMUD service area from April 2012 through May 2019. The dotted line reports the average (marginal) prices for customers that violated the EUP while the policy was in effect (indicated by the orange band) and were exposed to fees, moral suasion, and potentially public shame (“EUs”). The solid line reports average (marginal) prices for all other single family customers that were not directly exposed to EUP policy instruments (“Control Households”).

Figure A.3: Price Elasticity Heterogeneity by Household Water Usage



This figure plots the coefficients and 95% confidence intervals for the overall price elasticity at each percentile of average household water consumption in the top panel and average parcel vegetation (EVI) in the bottom panel. Estimates are obtained using models that follow Eq. 1.1 where the price instrument is interacted with each percentile of a given variable's distribution. Each point estimate is calculated as the sum of the percentile's interaction term and the coefficient on the baseline term (50th specified as the omitted group) with the standard error of the sum calculated from the relevant components of the variance-covariance matrix.

Table A.1: Residential Water Price Schedules over Time

Fiscal Year	Service Charge	Marginal Prices			Elevation Surcharges	
		Tier 1	Tier 2	Tier 3	Band 2	Band 3
2012	\$25.24	\$2.28	\$2.83	\$3.47	\$0.43	\$0.88
2013	\$26.74	\$2.42	\$3.00	\$3.68	\$0.46	\$0.93
2014	\$29.34	\$2.66	\$3.29	\$4.04	\$0.50	\$1.02
2015	\$32.12	\$2.91	\$3.60	\$4.42	\$0.55	\$1.12
2016	\$38.68	\$2.95	\$4.06	\$5.36	\$0.60	\$1.24
2017	\$41.38	\$3.16	\$4.34	\$5.74	\$0.64	\$1.33
2018	\$45.20	\$3.45	\$4.74	\$6.27	\$0.70	\$1.45
2019	\$49.26	\$3.76	\$5.17	\$6.83	\$0.76	\$1.56

This table reports marginal prices and elevation surcharges within the EBMUD service area for Fiscal Years 2012-2019. “Marginal Prices” represent the cost of consuming a volume of water equal to one hundred cubic feet of water (or CCF, equal to 748 gallons) for units consumed within consumption tier 1 (0-14 CCF), tier 2 (15-32 CCF), and tier 3 (33+ CCF). “Elevation Surcharges” report the surcharge paid for every CCF used by homes in elevation band 2 (approximately 200-600 feet) and band 3 (\approx 600+ feet).

Table A.2: Drought Surcharges, Fiscal Year 2016

Drought Stage	Consumption Tier		
	Tier 1	Tier 2	Tier 3
0 (Normal)	\$0	\$0	\$0
1 (Moderate)	\$0	\$0	\$0
2 (Significant)	\$0.23	\$0.31	\$0.40
3 (Severe)	\$0.59	\$0.79	\$1.03
4 (Critical)	\$0.73	\$0.99	\$1.3

This table reports EBMUD drought surcharges for fiscal year 2016. While the drought surcharges were in effect, marginal prices for units consumed in a given tier increased by the listed amount.

Table A.3: EBMUD Elevation Bands

Band	Description	Surcharge	% Sample
1	Homes in pressure zones that do not require pumping (approximately 0-200 feet above sea level)	None	56.24%
2	Homes in pressure zones requiring pumping (approximately 200-600 feet above sea level)	12-20% MP	35.54%
3	Homes in pressure zones requiring “considerable” pumping (approximately 600+ feet above sea level)	23-42% MP	9.21%

This table reports information and approximate cutoff elevations for the EBMUD service area elevation bands. Elevation thresholds are approximate as elevation band assignment is ultimately determined by pressure zone membership, which is influenced by elevation as well as water pressure and pumping requirements. Surcharges are reported as the percent of marginal prices faced across all consumption tiers.

Table A.4: Comparison Across Elevation Band Cutoffs, Alameda County

	Band 2 (200-600ft)			Band 3 (> 600ft)		
	Below (190-199ft)	Above (201-210ft)	P-Val	Below (590-599ft)	Above (601-610ft)	P-Val
Mean Temp (°C)	14.900 (2.976)	14.909 (2.952)	0.03	14.556 (3.279)	14.531 (3.260)	0.48
Total Precip. (mm)	95.593 (116.816)	96.254 (117.418)	0.16	104.306 (126.020)	103.635 (125.283)	0.02
Home Size (sqft)	1,569 (560)	1,617 (580)	0.00	2,437 (892)	2,374 (817)	0.00
Lot Size (sqft)	5,643 (2,604)	6,393 (48,267)	0.00	11,523 (13,050)	14,030 (34,873)	0.00
# Bedrooms	2.962 (0.955)	2.987 (0.924)	0.00	3.423 (1.149)	3.420 (1.089)	0.91
# Bathrooms	1.725 (0.755)	1.794 (0.810)	0.00	2.503 (0.946)	2.491 (0.913)	0.26
Owner-Occupied	0.632 (0.227)	0.659 (0.225)	0.00	0.891 (0.097)	0.899 (0.103)	0.00
HH Income >\$200K	0.125 (0.136)	0.129 (0.132)	0.00	0.302 (0.142)	0.341 (0.189)	0.50
Observations	153,973	180,685		73,037	63,456	

This table reports summary statistics for homes in Alameda County just on either side of the EBMUD service area elevation band cutoffs. “Below” (“Above”) columns report the means for homes within 10 feet of elevation below (above) the approximate cutoffs for elevation bands 2 and 3. “P-Val” reports the p-value for the test of equality of means above and below each boundary.