Title

Essays on the Economics of Environmental Policy

Author

Somers, Jackson Chandler

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on the Economics of Environmental Policy

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

in

Economics

by

Jackson Chandler Somers

Committee in charge:

Professor Mark Jacobsen, Chair
Professor Judson Boomhower
Professor Jennifer Burney
Professor Richard Carson
Professor Joshua Graff-Zivin

2022

The Dissertation of Jackson Chandler Somers is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

DEDICATION

I dedicate this to my parents, who without their support, even when they didn’t fully understand what I was doing, believed it was the best thing for me. Without their unique support and understanding of me, I would not have made it this far.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to extend my deepest gratitude to Professor Mark Jacobsen for his support and guidance. His advice has proven invaluable and his guidance, while never forceful, was consistently correct and what was best for me. There are insufficient words to describe my thanks. Thank you for being the advisor I needed and for providing so much assistance.

I am also deeply grateful to Professor Judson Boomhower. Without his environmental course at the end of my second year, I would have never chosen to pursue the field of environmental economics. His guidance and encouragement were incredibly beneficial to me, especially as I began to learn the field.

From both of you, I learned a lot about being a good academic, a helpful critic, how to shift perspectives on problems, and much about myself and who I'd like to be.

I would also like to extend my gratitude to Professor Joshua Graff-Zivin for his assertive but generous input and his constant support. A good balance between being straightforward, kind, and helpful is challenging to implement, but he does so very well. I really appreciate his concern with my work and my career.

I extend my gratitude to Professor Richard Carson for his expansive knowledge and numerous suggestions. His clear interest in all topics is inspirational and greatly benefitted my work. His input is always welcome, even when there is an overwhelming amount.

Finally, I would like to extend my gratitude to Professor Jennifer Burney for her helpful and kind input on projects, especially on particulars that were outside of economics. Her feedback always benefitted my work and her care is very much appreciated.

All of you made me a better person and academic. I thank you.

I cannot begin to express my thanks to my wife, Alexis, who supported me the entire time. In the numerous lows and few highs, she was always there to reassure me. There were many late nights, ones where I spent little time with her. I deeply appreciate that sacrifice. I thank you for bearing with the many drafts that were read over, in particular for dealing with my consternation at poorly explained portions. I am truly lucky to have someone as supportive as she is.

I’m extremely grateful to my mother and father, Jocelyn and Jamie. Without either of them, I would not have gotten this far. My mother for her unrelenting belief in me and insistence in my abilities, in all situations. My father for his calm and consistent support throughout this, being the rock that he is.

I gratefully acknowledge my friends throughout my experience in the department: Erica Chuang, Zhiyun Jiang, Camila Navajas-Ahumada, Greg Raiffa, and Mauricio Romero for their friendship, feedback, and support throughout the program.

Thanks also to my friend Brian van den Broek, who having already acquired a Ph.D. understood the many personalities and experiences of the process. He often reassured me or commiserated with my many experiences and all the while was a great supporter.

Thanks to my lifelong friend Benjamin Umholtz, who while outside of academia, was an avid supporter and always believed in me. His constant presence in my life has been a blessing.

Finally, thank you to my mother-in-law and father-in-law, Lisa and Dan, who allowed me to join their family and supported me as one of their own children. I am deeply grateful for the support and consistent inquiries into my life and how it was going.

Chapter 1, in part is currently being prepared for submission for publication of the material. The dissertation author was the sole author of this chapter.

Chapter 2, in part is currently being prepared for submission for publication of the material. The dissertation author was the sole author of this chapter.

Chapter 3, in part is currently being prepared for submission for publication of the material. The dissertation author was the sole author of this chapter.

VITA

2012	Bachelor of Arts in Economics, University of California, Berkeley
2018	Master of Arts in Economics, University of California San Diego
2022	Doctor of Philosophy in Economics, University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays on the Economics of Environmental Policy

by

Jackson Chandler Somers

Doctor of Philosophy in Economics

University of California San Diego, 2022

Professor Mark Jacobsen, Chair

This dissertation consists of three chapters, linked by the theme of analyzing household behavioral responses to environmental policies. Chapter 1 investigates the behavioral response of households receiving a composting program expansion, in particular when they can throw food scraps into their compost, and analyzes the effect on methane emissions reductions. Chapter 2 investigates the effect of carbon cap and trade permit and low-carbon fuel standard credit prices on consumer gasoline purchases, and compares the effect size to gasoline taxes and gasoline price increases unrelated to either policy. Chapter 3 investigates the heterogeneity in responses to water quality violations by examining how different demographic groups respond to these violations.

1. Do Composting Programs Work and Are They Worth It?

1.1 Introduction

Mounting scientific evidence indicates that urgent action to reduce greenhouse gas emissions is required in order to mitigate large, widespread climate damages. Enacting the necessary reductions is a burdensome policy challenge that requires implementing effective abatement investments. In particular, methane emissions reduction has recently garnered a substantial amount of attention in international climate change agreements. One additional ton of methane released into the atmosphere does approximately 80 times as much damage over a 20 year period as one additional ton of CO2 (Pachauri et al. (2014)), with even higher estimates such as from Errickson et al. (2021), making it a valuable target to rapidly lower emissions over the short-term. Over 100 countries, including the United States, have pledged to cut methane emissions by 30% by the year 2030.[1] It is important to choose policies that will lead to positive net benefits, i.e. to have the policy benefits outweigh the costs. This paper evaluates the efficacy of one proposed policy to reduce methane emissions: diverting household organic waste (compost) away from landfills.

Landfills in the United States produce 17.4% (4.56 MMT of methane) of all US anthropogenic methane emissions (EPA (2021b)). Landfill methane emissions are generated by the decomposition of organic materials. A common policy solution to mitigate landfill methane

[1] Lisa Friedman, "More Than 30 Countries Join US Pledge to Cut Methane Emissions" *The New York Times*, October, 11, 2021, https://www.nytimes.com/2021/10/11/climate/methane-global-climate.html

emissions is to divert organic materials generated by households and businesses – materials such as yard waste, food waste, and soiled paper. However, the number of households that participate in organic materials diversion and the extent to which they participate are both unknown quantities. Further, these programs are costly for jurisdictions to implement, resulting in a tension between the benefits of avoided methane and the cost of the program. In this paper, I investigate household behavior in response to an organics program expansion and conduct an analysis to estimate the cost of these programs per ton of carbon dioxide equivalent emissions (CO2e) avoided for counties across the United States.

In this article, I examine the rollout of a residential curbside food scraps and soiled paper products (FS&SP) program in Austin, Texas. All households serviced by the city of Austin waste department, Austin Resource Recovery (ARR), were provided with weekly curbside yard trimmings services prior to the implementation of the FS&SP program.[2] Once the program began, households were additionally allowed to dispose of FS&SP into their organics bins. All households were provided with a new organics bin and charged for the service, but participation in setting out and using the bin was not mandatory. Households were also provided with a pamphlet explaining the program and detailing the new materials allowed in the bin.

To identify the causal effects of adding FS&SP to organics disposal, I obtained administrative data to estimate a difference-in-differences model that uses the staggered, route-level rollout of the FS&SP program in Austin to compare organics disposal quantities before and after the program began. In addition to graphical evidence showing the validity of the identification strategy, the city of Austin chose pilot sites to be representative of the city's demographics and expanded further program phases around the pilot routes. This alleviates concerns for a bias in site selection.

The addition of FS&SP to the organics bin increased weekly organics diversion by 2.3 pounds per household per week, or about a 45% increase in organics diversion. The set-out rate

[2] It is important to understand the distinction between yard trimmings and organics collection in this setting. Yard trimmings means small branches and leaves from a yard, while organics is yard trimmings plus food waste and soiled paper. An organics expansion means including food scraps and/or soiled paper into yard trimmings.

for this program is approximately 30%, suggesting most residents don't set out their organics bin in any given week. Conditional on participation, households dispose of about 7.67 pounds of FS&SP per week. This corresponds to nearly 100% of FS&SP that the average household in Austin sent to a landfill *prior to* the FS&SP program.

A household diverting organic waste will necessarily stop putting organic matter into the garbage. Yet, it is not obvious how household recycling behavior will change. Specifically, soiled paper and cardboard are *not* recyclable[3] but they *are* compostable, so recycling totals may decline if households fix this common error.[4] Alternatively, households may recycle more than before (if that is feasible) due to pro-environmental behavior spillovers, as argued in Ek and Miliute-Plepiene (2018) and Alacevich et al. (2021). I find that households see a decline in both garbage and recycling disposal. This suggests that households are not only diverting their organic matter from the garbage, but they are also fixing their errors in recycling soiled paper and cardboard by diverting them into the organics bin.

Next, I formulate the social cost-benefit of a FS&SP program, which allows me to estimate the cost per ton of CO2e avoided. Recent literature within economics has focused on evaluating programs and policies with emissions reduction as a primary goal, typically finding programs to be costlier than anticipated. Examples of program evaluations include the Weatherization Assistance Program (Fowlie et al. (2018)), "Cash For Coolers" (Davis et al. (2014)), CAFE standards (Jacobsen (2013)), "Cash for Clunkers" (Busse et al. (2012), Jacobsen and Van Benthem (2015)), and promoting energy efficient technology (Allcott and Taubinsky (2015)).

To calculate the net benefit across the contiguous United States, I take the household-level estimate and match route-level Austin covariates to county covariates across the United States to estimate the implied diversion behavior of that county. As methane is emitted by landfills, the

[3]Soiled paper and cardboard, such as a greasy pizza box, have an altered fiber structure from being soiled and are of sufficiently low quality that they cannot be recycled.

[4]Fixing this disposal error has a welfare benefit in that fewer loads of recyclable material are sent to a landfill after processing at a materials recovery facility. Recyclable material that contains too much contaminated material is not recyclable.

emissions estimate is affected by a landfill having a landfill gas (LFG) capture system in place; these systems reduce LFG emissions by 50-75% on average. LFG capture systems are legally required by EPA to be installed in large landfills.[5] This results in locales that have high methane emissions potential (i.e. large cities) also disposing of their garbage into landfills with LFG systems. Welfare considerations are necessarily altered by LFG systems, since the emissions damages from the largest potential emitters are significantly reduced.

My baseline results suggest that with the city of Austin's program cost per household and estimated diversion behavior, the cost per ton of CO2e avoided for the City is \$541. Across the United States, the cost per ton of CO2e ranges from \$117 to \$3410 per ton of CO2e. The conventionally accepted social cost of carbon (SCC) is \$51 per ton of CO2e (IWG (2021)). While these estimates cast doubt upon cost viability of FS&SP programs, a discussion is necessary. This analysis suggests that FS&SP programs result in a substantially higher cost per ton of CO2e avoided than conventional SCC values. One ton of food scraps releases approximately 0.045 - 0.055 tons of methane in a landfill, suggesting around \$183 to \$224 in damages per ton of food waste, but this is greatly outweighed by the cost of the program.

Greenhouse gas damages per ton are subject to a great amount of uncertainty and it is possible organics diversion programs will be necessary to reduce the chance of a potential climate catastrophe. The social cost of carbon, the global warming potentials (GWP) horizon, the costs incurred by the jurisdiction, and household diversion behavior are all key variables in the net benefit equation. A cost of \$0.37 per household per month for the service or a diversion amount of approximately 27 pounds per household per week (which is substantially over 100% of all household FS&SP) would result in a net benefit of zero (that is, the costs are equal to the benefits). Household diversion cannot exceed 100% of total household FS&SP – which is typically about 7 to 10 pounds per household per week. This means that the estimate of necessary household diversion is unattainable. On the other hand, given the uncertainty of the SCC, it could be substantially too low and if so, the value of organics expansion programs are dramatically

[5] Installation is required for landfills that produce 25,000 metric tons of CO2e per year.

understated. This implies that the programs could be worthwhile if the SCC is underestimated.

My results suggest the national median estimated FS&SP program cost is $536 per ton of CO2e avoided. To place my results in context, I compare this to other *existing* policies whose aim is to offset carbon emissions. For a broad overview of program costs per ton of CO2e, see Gillingham and Stock (2018). In his analysis of CAFE standards, Jacobsen (2013) estimates that the cost of the CAFE standards are $222 per ton of CO2 avoided. Fowlie et al. (2018) examine the Weatherization Assistance Program in the US and find that the cost of residential energy efficiency standards are approximately $200 per ton of CO2 avoided. In studying the Energy Star program, Datta and Gulati (2014) find a cost range of $140 to $352 per tonne of carbon saved. The most recent Regional Greenhouse Gas Initiative auction had a clearing price of $9.30 per cap-and-trade allowance, which implies a price of $9.30 per ton of CO2e (RGGI (2021)). Meanwhile, the California-Quebec cap-and-trade program had an implied price of $23.30 per ton of CO2e (CARB (2021)). This suggests that a FS&SP program is more expensive per ton of CO2e than other existing policy solutions.

I contribute to our understanding of organics programs, which are becoming increasingly popular.[6] To the best of my knowledge, this is the first paper to provide a causal estimate of the impact of an organics program in any context. Pertaining to waste, my paper is most similar to Fullerton and Kinnaman (1996). They study the effect of the implementation of a pay-as-you-throw program, finding a modest decrease in total garbage disposed. Additionally, Ek and Miliute-Plepiene (2018) and Alacevich et al. (2021) study a composting program rollout, but are concerned with behavioral spillovers onto other disposal behavior. My paper directly estimates household diversion behavior. An estimate of this nature is missing in the literature and is important for assisting cities in assessing the value of implementing a FS&SP program. While there will be implementation differences across cities (such as whether or not to fine residents

[6] For example, some large cities that have already implemented these programs are Seattle, WA, Portland, OR, San Francisco, CA, and Toronto, Canada. These programs are mandatory, while New York City has a voluntary FS&SP program. Many smaller cities in the United States have organics programs in place, see Steeter and Platt (2017). California, Vermont, and Massachusetts have mandatory organics diversion laws in place as well.

for non-compliance or to ban organics from the garbage), understanding approximately how households respond to receiving organics services is important. This allows us to evaluate the efficacy of the programs, assess whether a city should implement one, and understand which cities should mandate organics laws.

My paper contributes to understanding household demand for pro-environmental services. By providing households with the ability to divert organics, households take up this voluntary behavior due to warm-glow (Andreoni (1990)) *and* by revealed preference are increasing their own welfare through diversion. The benefits to a household from organics diversion are almost exclusively due to less waste being sent to landfills. Therefore, warm-glow is likely the force behind households diverting their FS&SP upon receiving the program expansion. In particular, organics diversion can be seen to have a "positive frame": to divert organics waste benefits *all* households, while not doing so maintains the status quo of a household's waste being disposed of into the garbage (Andreoni (1995)). The "positive frame" amplifies household organics diversion behavior. Additionally, households voluntarily correct their recycling mistakes, further suggesting that household participation is due to warm-glow.

My paper also augments a large environmental and engineering literature on landfill methane emissions. I calculate the emissions reduction from FS&SP diversion and the cost per ton of CO2e avoided using an *empirical* estimate of diversion, which is the first of its kind. As a comparison, most literature estimating methane emissions reductions use hypothetical diversion estimates for their analyses, for example EPA (2010), Cruz and Barlaz (2010), Shahid and Hittinger (2021), Pai et al. (2019). Despite a large literature on methane emissions and the significant impact of methane as a greenhouse gas, there is little on household organics disposal behavior and the impact it has on the viability of organics programs. Further, instead of assuming a representative landfill to calculate the emissions changes, I use landfill-level data to estimate the emissions avoided from implementing a FS&SP program. Given that landfill methane emissions are a large source of total US methane emissions and the current international scrutiny of methane emissions, it is important and pertinent to understand if diverting household organic material from landfills is a cost-effective solution.

The paper now proceeds. Section 2 discusses the background on the Austin curbside organics program. Section 3 features a model to motivate organics diversion behavior and which types of households might engage in diversion. Section 4 provides background on the data used and the analysis. Section 5 presents the results of the regression and additional analyses. Section 6 provides background on methane emissions, the cost-benefit analysis, and the results. Section 7 discusses possible explanations for the findings and how the policy might be adjusted. Finally, Section 8 concludes.

1.2 Background on the Austin Curbside Organics Program

Austin is a large city home to nearly one million people in southeastern Texas. About half of the residents of Austin are served by the local waste department, Austin Resource Recovery (ARR). ARR services all non-private community residences with fewer than 5 units. If a residence has 5 or more units on the property, it must seek out private collection services.[7]

From here on, "residences" will refer to those that ARR services, i.e. it will not include properties with 5 or more units. Prior to the implementation of the expanded organics program in Austin, residences received curbside garbage, curbside single-stream recycling, and curbside yard trimmings services, along with other services such as twice-yearly brush pickup. In 2011, the city of Austin proposed a zero-waste "Master Plan" with the goal of Austin becoming a zero waste city by the year 2040 (The City of Austin (2011)). The report briefly details the rollout of their organics expansion, which was to expand the program across different portions of the city over time.

The organics program expansion built upon the weekly yard trimmings service by additionally allowing food waste, plant matter, and soiled paper into the organics bin.[8] Prior

[7]This seems to be fairly common across cities. Many have the cutoff of 5 units, while a few others (such as Denver, CO) have a 7 unit cutoff.

[8]As the city of Austin puts it, "If it grows, it goes!"

to the expansion, only grass clippings, leaves, and small branches (i.e. yard trimmings) were allowed in the organics bin. All residents who received the organics expansion were provided with a 32-gallon organics bin.[9] Organics collection, before and after the expansion, was weekly and year-round.

The City of Austin contracts with the composting business Organics By Gosh (OBG) to accept the organic matter diverted by the residents. The organic matter is processed by OBG and the compostable soil created at the end of this process is sold to residences and businesses. The City pays for OBG to accept the organic matter, however the tipping fee is much lower than that of the landfill. Finally, it is important to note that the composting process generates almost no pollution.

The expansion consisted of two pilots and four phases. For the pilots and the first three phases, residents were supplied with a new organics bin one week prior to the beginning of the FS&SP addition. When a household received the new bin, on top was an information packet describing the program expansion and what belongs in the organics bin. Beginning in the fourth phase, residents had the information packets mailed to them.

The pilot routes chosen across the city were intended to be representative of the Austin populace. Expansions for further phases were outgrowth from the initial routes and subsequent phases. In all, the City had two pilots and four "phases", resulting in a total of six different (actual) phases.

The fact that this program is a curbside program is non-trivial. Among cities with organics programs, 133 out of 270 are curbside only, while 101 are drop-off only.[10] Curbside programs provide collection to households at the curb of their residence, while drop-off programs requires residents to deliver their organics to a centralized drop-off location somewhere in the city. Typically there are several drop-off centers throughout a city. Comparing these two diversion

[9]Prior to phase 1, residents were provided with 96-gallon containers. ARR determined these were too large for residents and began providing smaller 32-gallon bins.

[10]Nora Goldstein, "Residential Food Scraps Collection Access In The U.S.", *Biocycle*, October, 22, 2021, https://www.biocycle.net/residential-food-scraps-collection-access-in-the-u-s/

policies, curbside programs require substantially less effort from households to divert organic materials. A curbside program provides the best opportunity to engender household participation in the program.

Table 1.1 shows the dates the phases began, the number of households in each phase, and the number of routes after each phase was rolled out. The two pilots only included 5 new routes and constituted about 7700 people each. Subsequently, the phase rollouts included significantly more routes and households. The original begin date for phase 4 – the final phase – was delayed from 9/30/2020 to 2/8/2021 because of COVID-19.

1.3 Model

This section lays out the household's problem, the firm's problem, the government's problem, and the environmental damages. All of this determines the social planner's problem for the full cost-benefit at the end.

1.3.1 Households

Formalizing a household's decision can help guide how households determine their organics diversion decision. A household must decide between disposal into the garbage or organics diversion. The household will execute this decision by considering the costs and benefits of organics diversion and garbage disposal. Specifically, even in a very simple setting where more organics diversion results in less garbage, both of these margins matter. For example, a household can derive benefits from either less waste going to the landfill or from diverting more material to the compost. Therefore, considering both quantities is essential to examine the drivers of organics diversion behavior. Further, we should expect heterogeneity in household disposal behavior due to differences in effort costs, awareness of proper disposal/diversion behavior, and different pro-environmental preferences. In the econometric model, this means that covariates are important factors to consider in extrapolating the results to other counties in the United States.

However, this model omits certain considerations. For example, I assume there is

Table 1.1. Organics Phase Rollout Table

	Pilot 1	Pilot 2	Phase 1	Phase 2	Phase 3	Phase 4
Start Date	2012-12-19	2014-02-18	2017-10-02	2018-06-25	2019-10-09	2021-02-08
HH Count, total	7,781	15,422	53,777	93,075	152,000	207,000
Number of organics routes	5	10	39	67	106	133

Note: This table shows the pilot and phase rollout dates, the total number of households included in the program, and the number of organics routes after each rollout.

no feedback effect, that is participating households do not consume even more products resulting in organic waste.[11] I omit signaling concerns, such as a household wanting to look pro-environmental or experiencing social pressure and so putting out their organics bin or participating as a consequence of these effects. I omit alternative margins of disposal as well, in particular using the garbage disposal for food waste or an at-home composter.[12]

Prior to receiving the organics expansion, a household could either use an at-home composter for their FS&SP or they could dispose of it into the garbage bin. There is the additional option that soiled paper could be disposed into the recycling bin, leading to only food scraps being placed into an at-home composter or the garbage.

Let G be all waste disposed of by a household in pounds per week, with y being the pounds per week of FS&SP disposed. Prior to the expansion, the lack of FS&SP disposal can be viewed as a quota set equal to zero. This means that no FS&SP is diverted. Let $b(y, G-y)$ be the benefits of disposal to a household. The first slot describes the benefit from disposing of waste in something other than the organics bin and the second slot represents the benefit of disposing of FS&SP into the organics bin.[13] Let $c(y, G-y)$ encompass the effort costs of disposal. This model assumes that the level of G remains fixed throughout. Note that with the implicit quota of no FS&SP diversion, we have $y = 0$. Omitting a numeraire consumption good, prior to the expansion we have

$$u_{\text{prior}} = b(0, G) - c(0, G)$$

With expansion, now we have $0 \leq y \leq y_{\max} < G$, where $y_{\max}$ represents the total amount of divertable FS&SP by a household. Using the above utility function and proper assumptions,

[11] However, this is an intriguing avenue of future research.

[12] In Austin, a household could receive up to a \$75 rebate from the purchase of an at-home composter or chicken coop. In total, approximately 2% of households received rebates/vouchers for at-home composters or chicken coops over an eight year period, from 2009 to 2017.

[13] It should be noted that y is FS&SP and *not* the larger classification of organics. This is because the problem of concern is on FS&SP diversion and not all organics. To include yard trimmings is trivial, however the model becomes less clear in its description.

there exists a y^* which maximizes utility resulting in:

$$u_{\text{post}} = b(y^*, G - y^*) - c(y^*, G - y^*)$$

Note that $y^* > 0 \iff u_{\text{post}} \geq u_{\text{prior}}$. This suggests a household will only undertake FS&SP diversion if it results in higher utility. However, this model of behavior can have households that do not participate ($y = 0$) or fully participate ($y = y^*$) and these represent important and desirable edge cases to fully capture household behavior.

Observe that with G fixed, we can parameterize the equation to $b_G(y) - c_G(y)$. To achieve a unique maximum, we need $b_G(y) - c_G(y)$ to be continuous and strictly quasiconcave on the interval $[0, \bar{y}]$. This guarantees the existence of a maximum on the *closed* interval.

The implementation of the organics expansion, not considering pecuniary costs, results in a Pareto improvement. Households that wish to participate in FS&SP diversion are now able to, while those with their highest utility at $y^* = 0$ can refrain from participating.

Examining the benefits and costs of diversion is important. The costs of disposal can incorporate many things, such as inertia, effort costs, disutility of organic matter decaying, and so forth. It seems reasonable to expect there to be inertia present in disposing of FS&SP to something other than the organics bin. Further, the typical process of FS&SP diversion is to have a small container for the FS&SP until it gets full, at which point it is taken to the outside organics bin. The small container can generate smells and pests. This contributes to the costs of FS&SP disposal. This makes the case that one should expect households to find it requires more effort and is therefore costlier to engage in FS&SP diversion behavior than to simply dispose of the material into the garbage.

Consider a household that is observed with $y^* > 0$. The above paragraph argues that it is unlikely for the case to be that $c(0, G) > c(y^*, G - y^*)$, implying that the benefits of diversion behavior to a household drive the diversion decision. There are no particular private tangible benefits to organics diversion. The clearest benefits to a household to begin FS&SP diversion comes primarily through warm-glow, pro-environmental preferences, or social pressure. In the first case, marginally less garbage results in higher benefits. That is, $\frac{\partial b(\cdot, G-y)}{\partial (G-y)}$ is negative, implying that more garbage results in lower benefits. This can be interpreted solely as a warm-glow from lowering disposal to a landfill. In the second case, having more organics diversion results in higher benefits, that is $\frac{\partial b(y,\cdot)}{\partial y} > 0$. This can be interpreted as either a warm-glow from more compostable matter to be converted into soil or as social pressure, specifically a household wishes to signal that they are diverters of FS&SP.[14] The drivers of a positive y^* largely derive from warm-glow or pro-environmental preferences.[15]

Next, consider the pecuniary costs of disposal. In addition to relying on a household's goodwill, an additional measure a jurisdiction can implement is to offer smaller garbage bins for a lower price.[16] A household may be motivated to divert FS&SP of an amount y that results in a lower garbage bin price tier, incentivizing households to participate when they may otherwise lack intrinsic motivation. However, many households will not be on the margin to benefit from a smaller garbage bin. It might simply be that the household produces too much waste, they do not care about composting or the smaller bin, there is too much variance in their weekly waste loads, or their value on time is too high to separate the waste.

In the case of constant prices, we can have households where $y^* = 0$, whereas a step-function pricing scheme we can have a positive y^* that results in $b(y^*, G-y^*) - c(y^*, G-y^*) + (p_G^i - p_G^{i-1}) > u(0, G)$, where i is the tier of the garbage bin. In this case, the household undertakes organics diversion because of the pecuniary benefit received, namely a lower monthly garbage bin bill. This is facilitated because diverting organics necessarily reduces the total

[14]One possibility is that a household, in an effort to signal their moral virtue but to avoid the displeasure of diversion, simply puts their organics bin out as a signal to their neighbors but never diverts any waste. This is entirely possible, however we began with the presumption that $y^* > 0$, so these cases are excluded.

[15]A possible tangible, non-pecuniary motivation is free compost, which some jurisdictions offer. However, this is not the case in the City of Austin. Further, unless there are strict protocols, free-riding is an easily imaginable option for a household that desires compost soil but does not want to divert their organic material.

[16]As of 2021, in the City of Austin, there is a \$18.80 user fee. A 24-gallon bin additionally costs \$3.85 per month, 32-gallon costs \$5.10 per month, 64-gallon costs \$10.25 per month, and 96-gallon costs \$30.70 per month. In the City of Seattle, which does not mention a fixed user fee, a 19-gallon bin is \$31.50, a 32-gallon bin is \$40.95, a 64-gallon bin is \$81.55.

amount of waste disposed of in other ways.

It is important to note that it is nearly universal to make recycling and organics bin costs a requirement, that is households do not receive an opt-in/opt-out choice, they are required to have the bins.[17] The mechanism to increase diversion behavior arises from providing a choice in garbage bin size, which incentivizes more diversion from a household in order to get a smaller garbage bin. Once a household receives the organics program expansion and the new bin, they do not factor the program cost into their decision; however, a social planner does consider this additional cost.

Another pecuniary cost – one which is not implemented by the City of Austin – is fines. These are in place in San Francisco, Seattle, and Portland. Fines are typically paired with bans on organic material in the garbage in order to incentivize proper organics diversion. These programs also have tagging programs which inform residents they have made an error in their disposal behavior. In particular, these programs that require organics diversion inspect garbage bins to verify households are not simply putting out their organics bins in order to appear to be in compliance.

1.3.2 Firm

There are two firms in this model. The first is the landfill and the second is the composter. Both industries are assumed to operate in perfect competition.

The landfill has the profit function:

$$\pi^L(y,g) = p_t \cdot (y+g) + p_E \cdot f(y,g) - c(y+g)$$

Where g is the amount of non-compostable material taken in by the landfill, y is the amount of compostable taken in by the landfill, p_t is the tipping fee charged per ton of waste, p_E is the price

[17]However, some jurisdictions provide a choice in the size of organics bin as well. The City of Austin does not. Further, cities that do provide a choice in organics bin size typically have the incremental costs as substantially less than that of garbage bin sizes. Therefore, all else equal, a household would rather divert more and get a larger organics bin while getting a smaller garbage bin than remain at the larger garbage bin.

of energy, $f(y,g)$ is a function that determines the amount of electricity/natural gas generated by y and g. These are not summed because the rate of methane generation differs between the two waste categories. The function $c(y+g)$ is the landfill's cost of disposal.

For a change in food waste, the marginal condition is:

$$p_t + p_E \cdot \frac{\partial f}{\partial y} = \frac{\partial c}{\partial y}$$

The landfill will maximize profits, in particular before the program it will choose (y^*_{NP}, g^*_{NP}) and after the program it will choose (y^*_P, g^*_P). "NP" stands for "no program" and "P" stands for "with a program". In both cases, $\pi(y^*_{NP}, g^*_{NP}) = \pi(y^*_P, g^*_P) = 0$, by perfect competition.

Composters typically make many different compost soils, each using different mixtures of organics material. I simplify this case to be a compostable soil that uses diverted household organics material and one that does not.[18] In the case of the composter, the firm will maximize:

$$\pi^C(y_h, y_o) = p_C \cdot f_h(y_h, y_o) + p_o f_o(y_o) + p_s \cdot y_h - c(y_h, y_o)$$

The variable y_h represents the organics diverted from households to the composter, while y_o represents the other (typically organics) materials used in producing compost. The price p_C is the price that compost made from household organics sells for, while p_o is the price of compost that uses entirely non-household organics. The function f_h uses household organics and non-household organics to produce compostable soil, while f_O uses non-household organics to produce compostable soil. These soils are not interchangeable. Due to the low profitability of household organics, the composter chargers a tip fee p_s for each ton of y_h received. Finally, the function $c(y_h, y_o)$ is the cost function.

[18] Processing household organics is non-trivial because the ratio of yard trimmings to organic matter has to be within a narrow range. Further, there are typically additional materials added to aid the decomposition process.

Profit maximization has the condition:

$$p_C \cdot \frac{\partial f_h}{\partial y_h} + p_s = \frac{\partial c}{\partial y_h}$$

In this case, the price p_s allows the composter to achieve zero profit. This cost is charged to the consumer through the local government. Again, by the perfect competition, profits before and after the program are zero.

1.3.3 Government and Environment

Each week, in total households divert:

$$W_f = \frac{\int_0^{\bar{y}} y \cdot N(y) dy}{2000}$$

$N(y)$ is a function outputting number of households that dispose of y pounds of waste per week. It is on the interval $[0, \bar{y}]$, with $\bar{y}$ equal the maximum amount of divertable FS&SP. The integral is divided by 2000 to have W_f in short tons of FS&SP.

The governing jurisdiction's equation is:

$$W_f \cdot (p_s - p_t) + p_j \cdot N - c(W_f, N) = 0$$

The difference $W_f \cdot (p_s - p_t)$ represents the savings from diverting FS&SP waste to a composter. The amount p_j is the price to each household for the program. The function $c_j(W, N)$ represents the cost of the program given the weight W and the number of households in the jurisdiction N. The jurisdiction chooses p_j so that the jurisdiction's problem is equal to zero.

The environment avoids damages equal to:

$$\text{Damages Avoided} = \text{CH4}(W_f) \cdot \text{GWP} \cdot \text{SCC}$$

Where CH4 is a function that outputs the estimated tons of methane emissions given the tons of FS&SP diverted. The variable GWP is global warming potentials, which in this case is 80 tons of CO2e per ton of methane. This describes the equivalent amount of CO2e that would need to be emitted into the atmosphere to do as much damage as one ton of methane. The SCC is $51 per ton of CO2e, which describes the damages of one additional ton of CO2 into the atmosphere. All of this results in Damages Avoided being in units of dollars per ton of CO2e.

Assume that the household utility is normalized to $b(0,G)-c(0,G)=0$. Let $y=y_h$. Time notation will be suppressed for clarity of the equations. This results in the social planner solving:

$$\begin{aligned}\max_{y,y_o,g}\int_0^{\overline{y}}[b(y,g-y)&-c(y,g-y)-p_j]N(y)dy\\&+\pi^L(y,g)+\pi^C(y,y_o)\\&+W_f\cdot(p_s-p_t)\\&+[\mathrm{CH4}(W_f)\cdot\mathrm{GWP}\cdot\mathrm{SCC}]\end{aligned}$$

Then the cost-benefit equation simplifies to:

$$\begin{aligned}\text{Net Benefit}=\int_0^{\overline{y}}[b(y,g-y)&-c(y,g-y)-p_j]N(y)dy\\&+W_f\cdot(p_s-p_t)\\&+[\mathrm{CH4}(W_f)\cdot\mathrm{GWP}\cdot\mathrm{SCC}]\end{aligned}$$

Where the firm profits are equal to zero because of perfect competition. This net benefit will be used to estimate the cost per ton of CO2e avoided.

The econometric model will estimate the average household diversion of w_f per week. With this, $w_f\cdot N=W_f$, where N is the total number of households, provides the total amount of

FS&SP diverted.

From the household model arises two hypotheses of interest.

Hypothesis 1: A lower price of garbage disposal for smaller bins will incentivize households to downsize their garbage bins, resulting in a leftward shift in the distribution of garbage bins.

Hypothesis 2: A household engaging in a positive amount of FS&SP diversion behavior will decrease the amount of improper disposal to the recycling bin. Namely, if we interpret the variable G as recyclable material and y as the potential FS&SP that is compostable and potentially recyclable, we will observe a decrease in improperly recycled material. This hypothesis will be tested with the available data.

1.4 Data and Analysis

1.4.1 Data

This paper uses route-level waste data provided by ARR from January 2010 to January 2020. This data includes all routes ARR services for each waste type collected, including garbage, recycling, yard trimmings, organics, brush, street sweeping, and more, at the route-day level. Routes are assigned pickup days that remain constant so long as the route is not phased out and barring any holidays. Nearly all routes with the same pickup day are adjacent. Routes for different waste types (e.g. garbage and recycling) do not have the same name nor do they overlap perfectly.

January 2020 was chosen as the stopping point for the analysis due to COVID-19. During this time, many businesses – in particular restaurants – closed down. These closures caused more households to eat in, thereby generating more waste within households, but less waste in the commercial sector. This represents a non-permanent change in household behavior that can vary across households, and as such I stop short of the “beginning” of COVID-19 in Austin. Further, COVID-19 caused numerous employment shortfalls in the waste industry across the nation due to employees with safety concerns, resulting in waste pickup becoming less reliable. Because these issues affect the reliability of the data, I do not conduct the analysis past January of 2020.

Most of the route names, for all types of waste, changed over time. Many route shapes also changed over time. These changes were implemented in order to alter the names for clarity, balance out the routes as the population changed over time, or balance out responsibility for supervisors over the routes. For the organics service maps, which include organics and yard trimmings routes, I use the phase 3 map as the master template. This map was the most recent map prior to COVID-19. I overlay the phase 3 map with the previous organics/yard trimmings maps and assign area shares to the routes for any previous iterations of the organics/yard trimmings maps. For routes from the phase 3 map that are derived from a larger yard trimmings route, I assign average disposal rates using the fraction of the phase 3 area the newer, smaller route occupied across the older routes.

Further, I overlap the phase 3 map with the available garbage and recycling maps during the time periods that those maps were in effect. For each week, I first aggregate the particular waste type to the week-level for each route because, for example, a Monday organics route may in part receive Tuesday recycling service. I multiply these area shares by the household service count in each area and the amount of waste disposed and assign this number to the appropriate phase 3 organics route. Since recycling services are bi-weekly, I divide the total amount over two weeks in half and assign the resulting number to each of the two weeks.

Finally, the Census covariates are at the Census tract level. I overlap the 2010 Census tract map with the phase 3 route maps, constructing area shares for each Census tract overlapping with a phase 3 route. For each covariate and a given route, I weight the value of the covariate by area share of an overlapping Census tract to provide an area share weighted-average for the covariates of a given route.

Prior to the program implementation, all organics routes were yard trimmings routes (i.e. only yard trimmings were allowed as organics disposal). As the program expanded, the existing yard trimmings routes were broken up into multiple "organics routes", which allowed disposal of both yard trimmings and FS&SP. The phase 3 organics and yard trimmings map can be seen in Figure 1.1. This map includes routes for both yard trimmings only and organics (which includes yard waste service). The routes beginning with "Y" are yard trimmings routes, while the routes beginning with "O" are organics routes. On the map, the north-most routes are all Wednesday pickups (hence the "W" as the second letter), while the routes just south are Thursday pickups, then Friday below that, and so forth. In this map, the routes with names in black are routes that existed prior to phase 3, while the routes with red names are those that were added for the phase 3 rollout. The numbers below the route names state the number of households served within that route.

Occasionally, the data has a one week gap. These, according to ARR, are due to data logging errors and not due to missed pickup collection. Further, these errors are infrequent and random, therefore they are not a concern for biased data. The route "OF25" was dropped from the data. It did not overlap with routes initially serviced by ARR, so there is no prior data available.

For this particular map, there 8 new Monday routes servicing 11,233 households, 8 new Tuesday routes servicing 11,306 households, 8 new Wednesday routes servicing 10,848 households, 8 new Thursday routes servicing 10,069 households, and 7 new Friday routes servicing 9,612 households.

The dataset consists of 64,927 total route-week level observations. This consists of 131 routes, of which 25 are yard trimmings routes and 106 are organics routes. Routes are of varying size by area and household count. The smallest organics route by household count is 929 households (OH2) and the largest is 2125 households (OH12). Yard trimmings route are larger than organics routes in both area and population served. This is because they pick up less material per household.

Table 1.2 presents summary statistics of the pilot and phases of the organics expansion. The statistics suggest that phases are well-balanced across covariates, with very little variance in average values. This is a positive sign and evidence against any biased choice in routes. Within

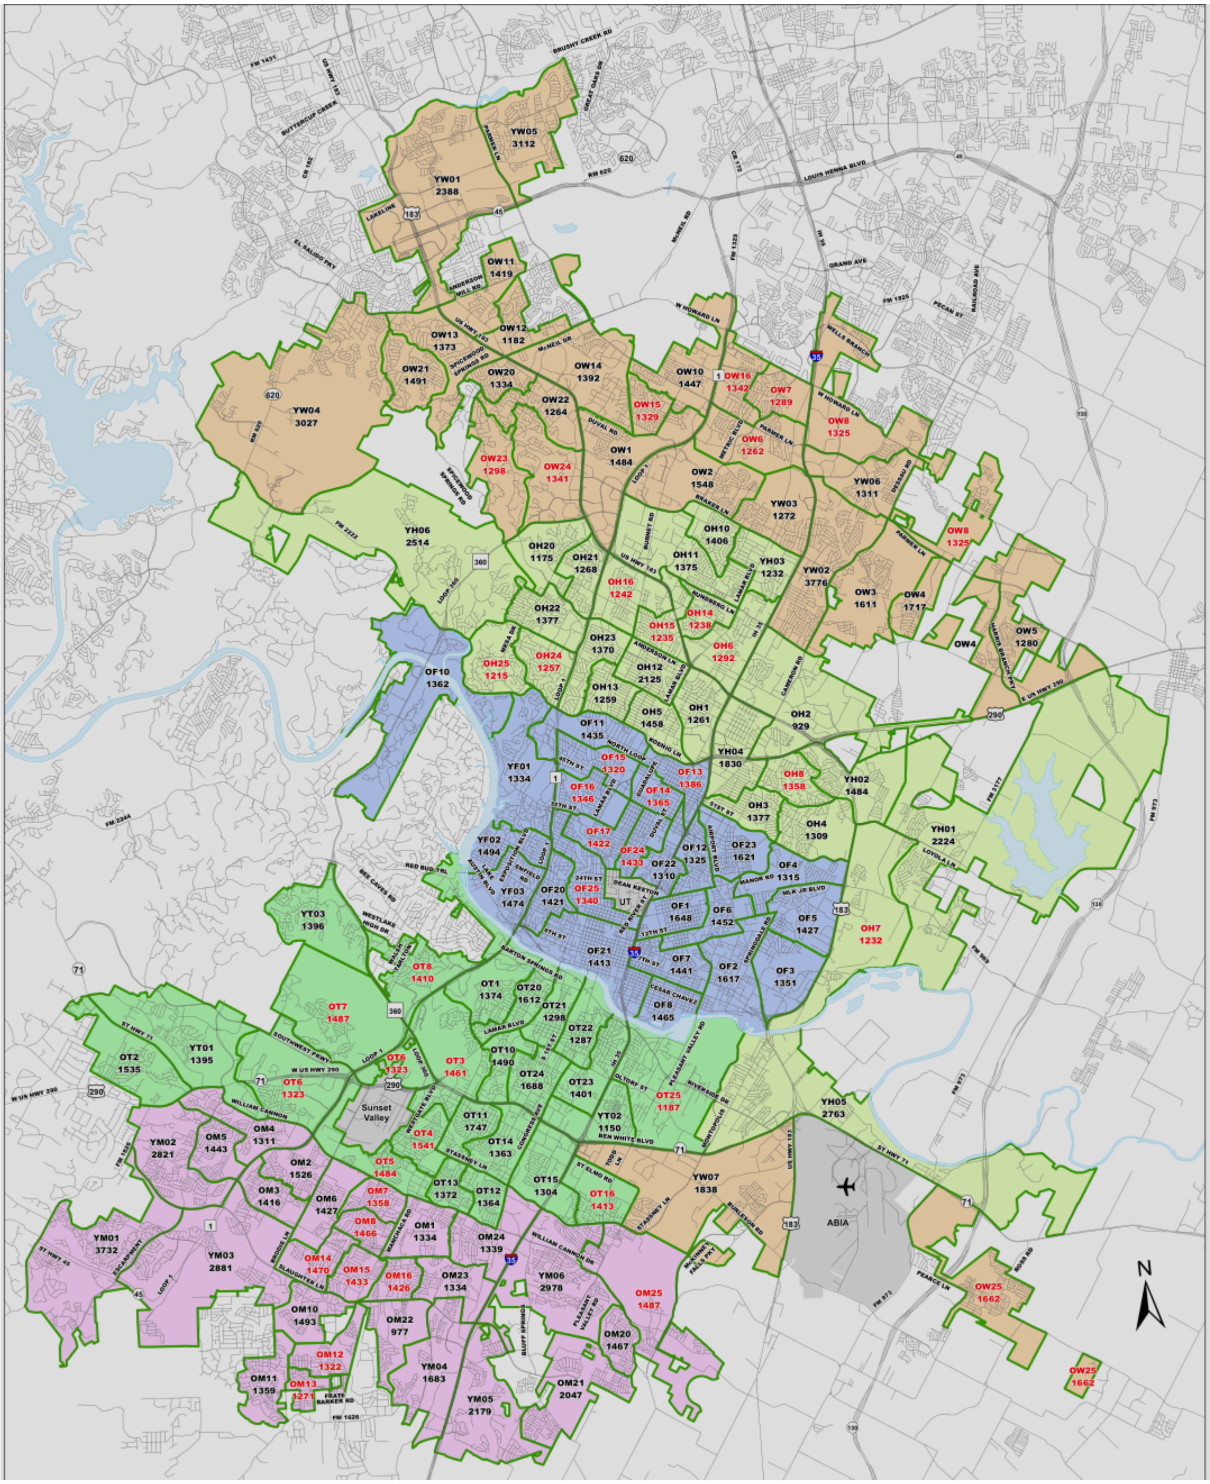

Figure 1.1. Phase 3 Organics Routes

Note: This is the map of the phase 3 rollout. The routes beginning with the letter "O" are organics routes, while those beginning with "Y" are yard waste routes (no organics). The routes in black were already in place prior to the rollout of phase 3 and the routes in red are the new routes added from the phase 3 rollout. The numbers below the route names are the number of households served by ARR within the route.

Table 1.2. Summary Statistics

Phase Year	Pilot 2012 & 2014	Phase 1 2017	Phase 2 2018	Phase 3 2019
Avg Route Pop	4547	4783	5732	5773
SD	(2250.93)	(2287.02)	(2804.22)	(4766.07)
Pct Bachelor Deg	0.474	0.513	0.501	0.505
SD	(0.153)	(0.155)	(0.144)	(0.195978)
Median income	65760	66416	70908	66207
SD	(12604)	(14904)	(18664)	(20143)
Avg HH size	2.47759	2.42984	2.3948	2.45199
SD	(0.42)	(0.339)	(0.3)	(0.431)
Med. home value	301170	317354	322243	315448
SD	(93067.4)	(103957)	(149806)	(152237)
Frac. White only	0.4968	0.5347	0.5205	0.5207
SD	(0.172)	(0.173)	(0.159)	(0.186)
Frac. Latino	0.3797	0.3043	0.3173	0.324
SD	(0.169)	(0.149)	(0.132)	(0.186)
Frac. Black only	0.0657	0.076	0.0742	0.061
SD	(0.0481)	(0.075)	(0.0543)	(0.044)
Frac. Asian only	0.0334	0.057	0.0617	0.0669
SD	(0.027)	(0.041)	(0.041)	(0.043)

Note: This table presents summary statistics of the pilot and phases of the organics program expansion in Austin, Texas.

a phase, routes have substantial heterogeneity as the table suggests, this is because rollouts occurred simultaneously across different sectors of the city. As Table 1.1 displays, as phases continued to be added the expansions included larger numbers of households, but even still Table 1.2 shows the phase covariates remained relatively balanced over time.

1.4.2 Estimating Equation

The analysis uses the staggered rollout to identify the effect of the curbside organics program. The estimate seeks to identify the estimated increase in organics composting once a route is a recipient of the program. The regression run is:

$$y_{rt} = \beta \cdot \mathbf{1}\{\text{Received Organics Expansion}_{rt}\} + \Theta X_{rt} + \alpha_r + \delta_t + \varepsilon_{rt} \tag{1.1}$$

The outcome variable y_{rt} is the pounds of organic waste disposed of per household in the route r in time period t.[19] The variable y_{rt} represents organics diverted. When the variable $\mathbf{1}\{\text{Received Organics Expansion}_{rt}\}$ is equal to 0, y_{rt} consists solely of yard trimmings weight. When the curbside organics program indicator switches to a value of 1, the variable y_{rt} becomes the sum of yard trimmings and the newly allowed FS&SP. The coefficient β is the additional amount of organic material diverted once the program is expanded to include the route. Route fixed effects are found in α_r while a variety of time fixed effects are in δ_t. The δ_t is a stand-in for different time fixed effects used, such as month, week, and year or week-by-year. The matrix X_{rt} is the set of demographic covariates described above, including garbage and recycling per household per week. Standard errors for this regressions are clustered at the route level.

The coefficient β can be interpreted as the additional pounds of organics diverted per week per household when the route receives the organics program expansion. That is, it reflects the additional material, beyond yard trimmings, diverted to the organics bin.

There are several potential confounders to properly estimating β. The staggered rollout

[19]For complete clarity, the subscript t is week-year, it begins on January 1st, 2010 and ends on January 31st, 2020.

allows me to characterize the counterfactual diversion behavior for routes that receive the organics expansion. The routes that receive the expansion are the treated group, while those that have not yet received the expansion are the control group. The control routes allow me to establish how much the treated groups would have disposed of had they not received the organics expansion. Without the control group, it would be significantly more challenging to disentangle the effect of receiving the organics expansion from any possible shocks or underlying trends.

Across routes, we should expect differences in household behavior. While some of this behavior will be captured by the covariates, there may be unobservable differences across routes. For instance, one route could contain more households with pro-environmental preferences than another route. Another example is that one route might have a higher diversion rate prior to the program than another route. The route fixed effects alleviate these issues by controlling for differences across routes.

Food waste and soiled paper are not generated at a constant rate throughout the year. For example, major holidays/events such as Independence Day, Thanksgiving, or the Super Bowl all generate more food scraps than other times of the year. Week fixed effects control for this seasonality. Further, we might expect that a supply shock, such as a poor harvest increasing food prices, would lower household consumption of food, thereby reducing food scraps. Shocks such as this are controlled by the addition of year fixed effects.

The results from Equation (1) will be used for the methane emissions avoided estimate. A slightly modified formulation of Equation (1) will also be used for covariate analysis and the impact on garbage and recycling behavior.

1.5 Results

This section presents the results of the econometric estimate. Additional results analyzing spillovers onto recycling and the impact of specific covariates are also discussed. Appendix 4 contains the results by estimating the regression using the estimation strategy from Callaway and Sant'Anna (2021), to avoid the pitfalls of TWFE. However, the results do not change in any substantive manner.

1.5.1 Main Results

Figure 1.2 shows the effect of the organics program expansion rollout. Prior to the program, only yard waste was allowed in organics disposal. Once the program rolled out, residents were then allowed to dispose of food waste and soiled paper in addition to yard waste in the organics bin. This plot shows the estimates for each week, from 51 weeks before to 51 weeks after, with the 52nd weeks before and after being pooled. The values shown in the plot are normalized to the week before an expansion began in a given route. Upon the program taking effect, the plot displays a sudden increase that sustained throughout the program. Also, the plot shows that the program did not affect organics disposal behavior prior to the program expansion.

It is possible that households may require some time to learn what materials can be added to the organics bin or to develop a system/"rhythm" to begin diverting FS&SP. Alternatively, households may initially participate and taper off. For example, they may find it onerous to maintain the altered disposal behavior or find the smell problematic. Figure 1.2 answers both of these questions and suggests there is minimal learning and there is potentially a small amount of tapering; however if there is any it is not large. Further, running a regression with four 13-week bins (i.e. one bin for weeks 0-12, one for weeks 13-25, one for weeks 26-38, and one for weeks 39-51 suggests that the only tapering would be from weeks 39 to 51. However, an indicator for all weeks after the first year has a higher estimate than weeks 39 to 51, so this could simply be a result of noise.[20]

The figure shows that participating households begin diverting immediately and then maintain that level of engagement thereafter. A striking observation that the plot shows is that households appear to have readily understood the program and proper diversion behavior. This plot and this data display an average household within a route, so certain changes such

[20]These results can be provided upon request.

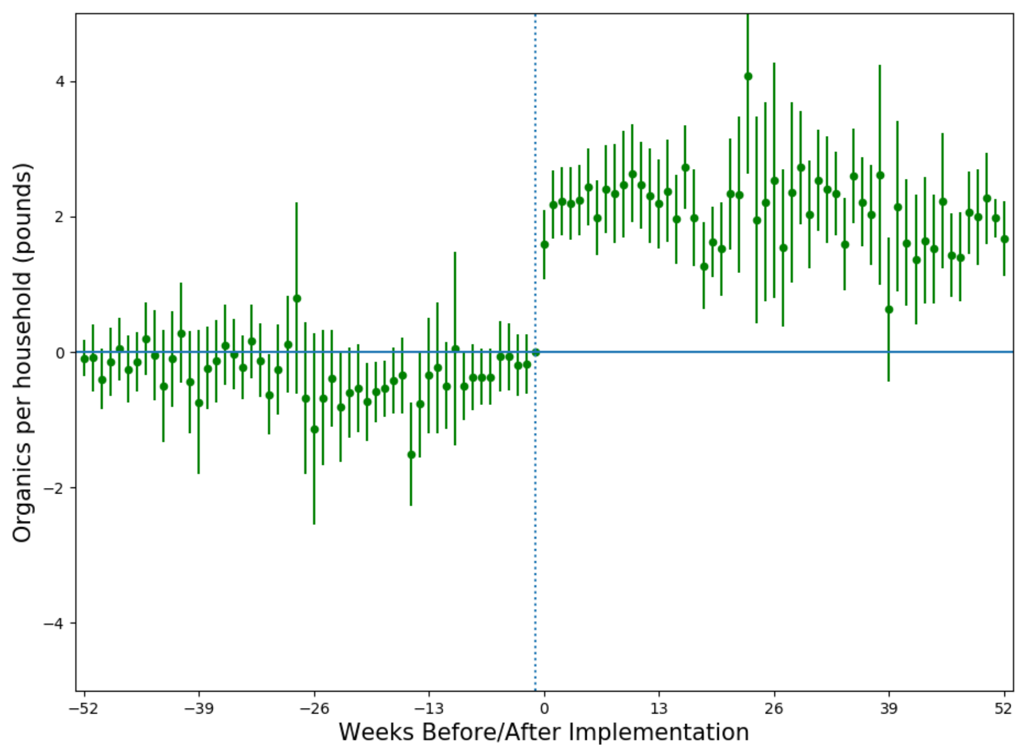

Figure 1.2. Effect of Organics Expansion: Before/After Expansion

Note: This plot shows the before and after of a phase rollout on the increase in the pounds of organics disposed of per household. All values are relative to the week prior to the rollout. Weeks before 52 weeks prior and 52 weeks after are pooled.

as a household engaging less in a given week while another engages more will average out and are not observable. However, the graph becomes exhibits more variance past the 15 week post-implementation mark. This is likely due to the fact that Phase 3 data does not extend much beyond 15 weeks because it is very late in the timeframe of the dataset.

This plot also provides evidence that households did not change their behavior prior to the program expanding to their residence. A household, learning of this program, could in theory increase their organics disposal behavior prior to the rollout. The plot suggests that households do not anticipate the program rollout and do not change their disposal behavior beforehand.

The results of the regression can be found in Table 1.3. The first three columns of the regression are without covariates, while the second three include covariates. The first and fourth columns have fixed effects for the route, the particular week of collection, year of collection, and the day of the week of collection. The second and fifth columns have the same time fixed effects just mentioned, but include a route-week fixed effect. Finally, the third and sixth columns have route, week-by-year, and day-of-the-week fixed effects. The second set of three columns have covariates of pounds of garbage per week per household, pounds of recycling per week per household, average family size, average household size, median income, median age, median home value, percentage of people with at least a high school degree, percentage of people with at least bachelor's degree, percent of the population over age 65, percent male, percent white, percent black, and percent population over the age of 25.

The table describes the results for the coefficient β as presented in Equation 1. The table shows that receiving the organics program increases weight to the organics bin by between 2.25 and 2.43 pounds per week per household. The estimates are very stable across specifications. Prior to the organics program, yard trimmings weight was about 5 pounds per week per household, suggesting that the total weight of the organics bin increased by about 45%.

The stability of the estimate is striking. The variety of fixed effects used and the inclusion of covariates have minimal impact on the estimate. In examining individual routes for heterogeneity, Figure 1.3 shows the results of running the regression interacting each organics route

Table 1.3. Organics Program Implementation

	Dependent variable:					
	Pounds of Organics Disposed Per Household					
	No Covariates			Covariates		
	(1)	(2)	(3)	(4)	(5)	(6)
INDC = 1	2.26616***	2.24574***	2.24614***	2.42917***	2.41060***	2.39701***
	(0.05340)	(0.05341)	(0.05060)	(0.05770)	(0.05737)	(0.05381)
FEs						
Week	X			X		
Year	X	X		X	X	
Day	X	X		X	X	
Route	X	X	X	X	X	X
Week x Year			X			X
Route x Week		X			X	

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

Note: Data series from January 2010 to January 2020. The dependent variable is average pounds of organics per household per week. The primary variable of interest is a binary variable indicating if a route has received the organics expansion yet or not. The first three columns do not include covariates; the final three columns do include covariates. Standard errors are clustered at the route level.

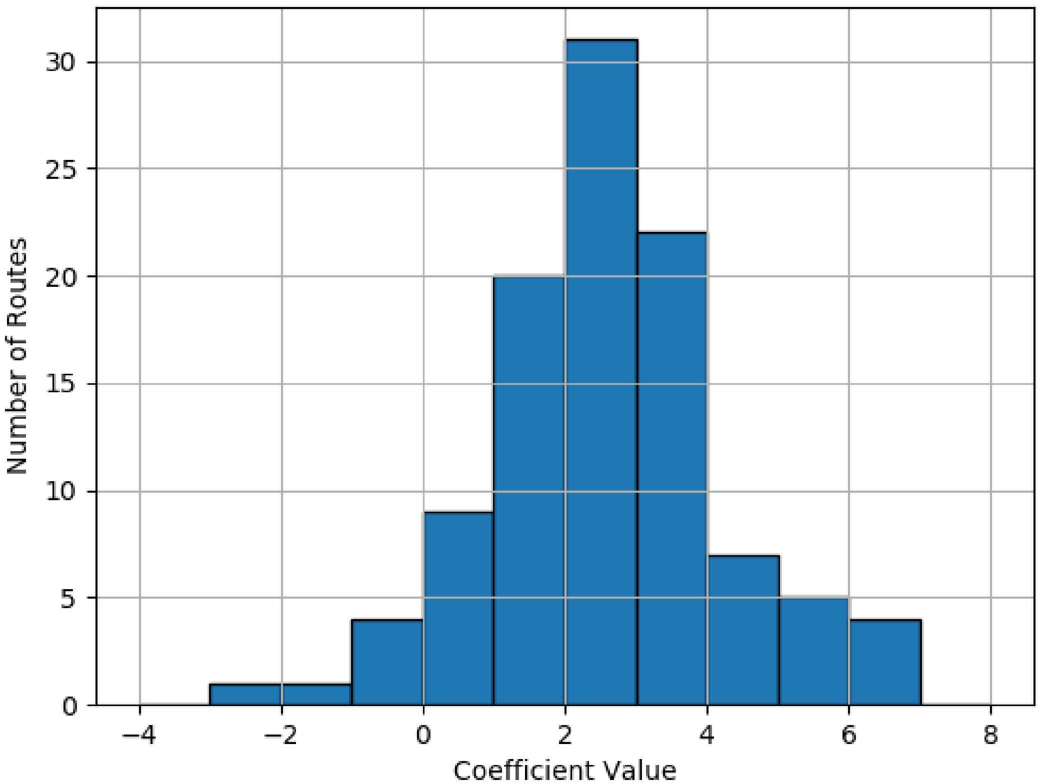

Figure 1.3. Histogram of Coefficients

Note: This plot shows the results of interacting the program expansion indicator with every organics route.

with the indicator for receiving the organics expansion. The figure shows that the coefficients are clustered between 2 and 3, which is where the main estimate β resides. The distribution is not skewed and looks roughly normal. In addition, a reasonable number of negative values occur in line with statistical chance. This plot shows variation across routes, with some routes participating more than other routes. This plot displays the variation across routes, which will be exploited for the county covariate matching below.

Analyses of the Austin waste stream in 2012 estimate about 7.2 pounds of organic matter in the waste stream *after* accounting for the diversion of yard trimmings. The city of Austin estimates approximately a 29% participation rate (i.e. 29% of people actually set out their bins) in the organics program, which, using the estimates found in Table 1.3, suggests an estimate of

$$\frac{2.3}{0.29} = 7.59 \text{ additional pounds per week per household that uses the organics bin.}$$

This coincides closely with the estimated maximum amount of disposable household organic matter. It suggests that conditional on a household participating, they are disposing of (nearly) all of their organic material into the organics bin.

Most cities in the United States only collect city-level waste data. The results from Austin can be compared to overall averages of other large cities that have added FS&SP to their organics. The city of Seattle, WA disposed of approximately 4.81 pounds of FS&SP per household per week in 2016, while the city of Portland, OR disposed of approximately 2.2 pounds of food scraps (not including soiled paper) per household per week in 2016. Other program factors matter here, such as laws banning organics from the garbage or the maturity of the program. Further, other cities have organics programs, such as New York City, but they require a neighborhood to achieve a certain number of sign-ups before it is rolled out. Their estimates are around 10 pounds of FS&SP per household per week, but the selection effect biases collection estimates relative to city-wide programs.

1.5.2 Additional Results

The Austin residential waste stream, as of the most recent 2019 materials audit, was 77.4% yard trimmings, 15.2% food scraps, 4.9% soiled paper, and 2.5% non-compostable material.[21] My estimate does not account for non-compostable material nor is it known what the percentage of the organics stream was non-compostable prior to adding FS&SP. Assuming there was no new contamination upon adding FS&SP, this implies that 15.2/20.1 = 75.6% (or about 1.74 lbs per household per week) of newly diverted waste is food scraps, while 4.9/20.1 = 24.4% (or about 0.56 lbs per household per week) of newly diverted waste is soiled paper and soiled cardboard.

Adding FS&SP to organics diversion also impacts household garbage and recycling disposal behavior. Because households are diverting what was once garbage, households should

[21] These audits typically sample several hundred households over the course of many weeks. The auditors separate the waste into categories, weigh them, and report. The margin of errors for the percentages are not listed.

Table 1.4. Organics Program Implementation: Effect on Garbage and Recycling

	Dependent variable:					
	Pounds of Waste Disposed Per Household					
		Garbage			Recycle	
	(1)	(2)	(3)	(4)	(5)	(6)
INDC = 1	-0.74237^{***}	-0.76093^{***}	-0.71313^{***}	-0.61905^{***}	-0.61621^{***}	-0.89319^{***}
	(0.05421)	(0.05892)	(0.05735)	(0.05667)	(0.05828)	(0.05649)
FEs						
Week	X			X		
Year	X	X		X	X	
Day	X	X		X	X	
Route	X	X	X	X	X	X
Week x Year			X			X
Route x Week		X			X	

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

Note: Data series from January 2010 to January 2020. The dependent variable for the first three columns is the average pounds of garbage disposed per household per week. The second three columns feature average pounds of recyclable material per household per week. The primary variable of interest is a binary variable indicating if a route has received the organics expansion yet or not. The coefficients represent the increase/decrease in disposal behavior after receiving the organics expansion. All regressions feature the same covariates as in Table 1.3. Standard errors are clustered at the route level.

experience a decline in the weight of their garbage disposal (assuming there is not a strong rebound effect). Meanwhile, it is ambiguous how households' recycling behavior will respond. There could be no change in household recycling behavior since material that is recyclable *is not* compostable and vice versa. By adding FS&SP to organics diversion, this could generate spillovers causing households to recycle more. Finally, households could recycle less because they were previously recycling certain material improperly (in this case it would likely be soiled paper and soiled cardboard, which is not recyclable[22]). Table 1.4 tries to resolve this question. Here I use the pounds of garbage or recyclable material disposed of per household per week as the dependent variable to estimate the impact on garbage and recycling behavior when a household receives FS&SP services. Otherwise, the regression is the same as in Equation 1.

As anticipated, the total amount of garbage disposed declines by about 0.75 pounds per household per week. Note that this is not a pound-for-pound decrease compared with the increase in organics disposal. Interestingly, households also recycle less material after the FS&SP program begins. The estimated portion of soiled paper and soiled cardboard in the organics waste stream is about 0.56 pounds per household per week, while the estimate in Table 1.4 showing the change in recyclable material with the implementation of the program is nearly identical in weight at about 0.6 pounds per household per week. This suggests that at least some of the decrease observed in household recycling behavior comes from diverting improperly placed soiled paper and soiled cardboard from recycling into the organics bin. It is not clear if this effect arises from the information pamphlet provided with the new bin or from the expansion simply including soiled paper and allowing households to dispose of the material properly. Either way, this is a positive incidental benefit to expanding organics programs.

Household size is arguably the only demographic characteristic where it is reasonably clear to conjecture the correlation between it and organics diversion. In general, larger households have more waste (because they have more people) and so it is reasonable to assume that larger

[22]The particulars of this are simple: the fibers of soiled paper are weaker than unsoiled paper. The lesser integrity of soiled paper makes it unusable.

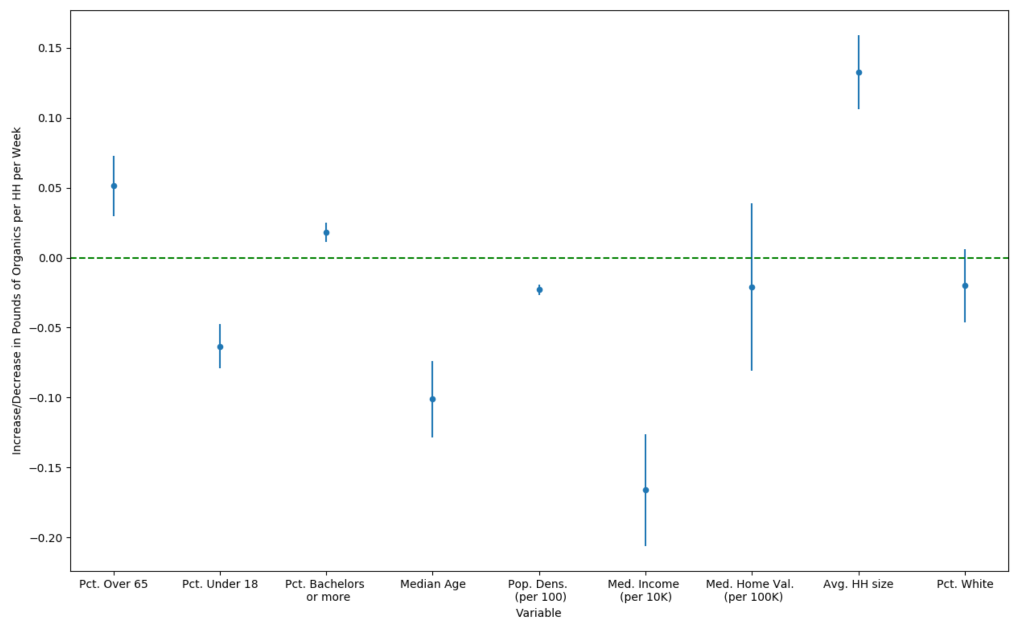

Figure 1.4. Covariate Marginal Effects

Plotted are the marginal effects of each of the covariates after receiving the organics program expansion.

households will also divert more waste (all else equal). Otherwise, correlations are not so clear. For example, richer people may have a higher valuation of time and so do not find it worthwhile to participate in an organics program; alternatively, richer people might be *more* environmentally aware and so participate more than poorer people. Households with a higher income (typically) have a higher time value, which suggests that it is expected for them to engage in less organics disposal behavior. Alternatively, they could eat out more and so do not have as much FS&SP waste in the first place. People with bachelor's degrees might be more aware of the benefits of organics diversion and so participate more; or, they might not believe it is worthwhile and so do not participate.

The plot in Figure 1.4 shows the marginal effect of each of nine demographics. This estimate is obtained by running Equation (1), but including interaction terms of each of the demographics with the indicator for receiving the organics program expansion. Two of these route-level covariates are insignificant and close to zero: median home value and the percent of the population that identifies as white. Unsurprisingly, routes with larger household sizes have more organics disposal. Adding one additional individual to a household increases disposal by 1.32 pounds per week (however, the graph shows a 1/10th increase in household size for graphical clarity). Routes with older populaces tend to experience more organics disposal – about 0.05 pounds per week more for a 1% increase; on the other hand, a one year increase in the median age of a route causes disposal to decrease by 0.10 pounds per week. This is interesting and perhaps seems contradictory. However, generations can differ and in particular, a higher median age does not necessarily imply more people over age 65, only that there are more older people. For example, if a route has more people further along in their middle age versus younger people, this will increase the median age but not affect the percent over age 65. Meanwhile, routes where there are more children present have lower organics disposal, about 0.06 pounds less per week. This seems plausible if families are simply too busy to have the time to separate their waste (or cannot generate family-wide compliance). In a similar contradiction, routes with a higher percentage of bachelor's degrees have slightly more organics disposal – a one percent increase is associated with an increase of 0.018 pounds per week – while routes with higher median incomes disposed of less, approximately 0.16 pounds less for a $10,000 increase in median income.

1.6 Methane Emissions Impact

The remainder of the paper now focuses on using the empirical estimate from above to estimate the cost of an organics program expansion per ton of CO2e across the United States.

1.6.1 Landfill and Methane Emissions Background

A landfill is a firm that effectively sells space to place waste. The waste is systematically buried. The important features are that waste is buried in the landfill after which it is covered with a clay, dirt, or synthetic material (usually at the conclusion of the day). If a landfill has a capture system, there are pipes throughout the landfill that capture the methane. Landfills are complex and their construction affects several aspects of decomposition. Appendix 2 provides additional details on landfill construction.

Almost all material in a landfill decomposes to some degree. Both the rate of decay and the total percent of decomposition within a landfill varies by material type. Organic matter, such as food scraps, paper, leaves, and branches decompose rapidly in a landfill. For example, it takes a banana peel about one month to begin to decompose, while an aluminum can takes a century or two, and styrofoam containers begin decomposition anywhere between 500 and 1 million years (but are non-biodegradable). Not all of a given material decomposes in a landfill. For example, about 50% of food waste decomposes in a landfill, while 50% of food waste will never decompose and remains in the landfill. Due to their rapid decomposition, diversion of organic material from a landfill, relative to all other materials, will rapidly affect landfill methane emissions.

Some landfills have landfill gas capture systems in place. EPA mandates that a landfill producing more than 25,000 metric tons of greenhouse gases (in CO2e) must install a landfill gas capture system onsite. There were 559 operational LFG collection projects and 357 shutdown projects in 2020. These systems have substantial fixed costs to install. The methane can either be put to use or flared. In the case of flaring, the methane produces no additional value as it is simply burned off. Otherwise, a landfill can convert the methane to a usable form, such as natural gasoline or electricity. Electricity can be sold to the grid, while natural gas (depending on the type) can be sold or used onsite as fuel for waste trucks. In this way, a landfill can generate an additional revenue stream by converting methane into a usable resource.

Generally, landfill gas capture systems do not capture methane emissions from materials that decompose quickly as effectively as from materials that decompose slowly (Barlaz et al. (2009) and Cruz and Barlaz (2010)). This raises an issue with food scraps because food scraps decompose quickly (relative to other organic materials, excluding grass and leaves – which also decompose rapidly). Therefore, using the average capture rate of a landfill for FS&SP methane emission estimates will overstate the amount of methane captured, implying that more methane will be saved from the diversion of food scraps compared to diversion of an average waste composition. Assuming average landfill moisture conditions, "Bulk MSW", a catch-all category for the average US waste composition, has an estimated typical collection efficiency of 85%, while food scraps and mixed organics have an estimated collection efficiency of 66% and 65%, respectively.

Landfill emissions are measured in three ways. The first way is with engineering equations provided by EPA to estimate the methane emissions from the landfill, while the second and third are empirical estimates from airplanes/drones and satellites (Duren et al. (2019) and Jacob et al. (2016)). The IPCC and EPA have standardized emissions equations that represent "good practice", along with empirical estimates (or ranges) of the parameters used in the equations. Many model and parameter details can be found within the IPCC Solid Waste Disposal chapter (IPCC (2006)). However, the accuracy of this bottom-up estimation method has faced substantial scrutiny (of Sciences (2018)) and is considered "low confidence" in its accuracy. The report suggests alternatives, primarily calling for a mixture of the latter two empirical measurements to augment engineering data estimates. However, this data is not readily available for most landfills.

The standard first-order decay model for landfill methane emissions is additively separable in the waste types. This means that methane generation from one waste type is independent of any other waste. Emissions are only dependent on the weight of the material, the time the waste is in place, and the specific waste type decay parameters. Background information on the equation used to estimate emissions can be found in Appendix 1, along with explanations of the relevant parameters. The derivation of the model assumes that diversion of one type of material does not affect diversion of any others. In particular, a household diverting food waste does not impact how said household disposes of any other types of waste. This allows for the household waste stream excluding food scraps to be identical over time. Therefore, the only impact on landfill methane emissions as a result of diverting food waste results from food waste. These assumptions have been scrutinized, as mentioned above. However, it is the current standard approach and there is no data available to be able to construct a better model or estimates for the United States.

Methane emissions are considered non-biogenic. Note that a landfill emits many different gases, but primarily CO2 and methane. CO2 is not counted in emissions inventories because it is of biogenic origin and is accounted for in “Agriculture, Forestry, and Other Land Use”. Further, because the material does not decay under natural conditions (it instead is buried in a landfill).

The equation to estimate methane emissions in year t is:

$$CH4_t = \sum_{i=0}^{t} \sum_{s<t} W_s \cdot DOC \cdot DOCf \cdot MCF \cdot F \cdot (16/12) \cdot (e^{-k(t-s-1)} - e^{-k(t-s)})$$

The input W_s represents tons of waste disposed of in year s. The remaining inputs are described in Appendix 2; they are standard parameters used to estimate landfill methane emissions in the United States. Material in a landfill, as described above, does not decay immediately; methane is slowly released over time. This equation captures that a decreasing proportion of the waste decays every year. Over time, if waste is disposed of in the landfill, then waste at different times will decay simultaneously. For example, if in the first year 10,000 tons of food waste is disposed of, then in the second year, approximately 11% of the *decomposable* food waste will have decayed.[23] In the second year, if another 10,000 tons of food waste are disposed of, then in the third year 11% of the decomposable tons disposed of in the second year will decompose *and* approximately 11% of the remaining decomposable tons from the first year will decompose as well. Thorough explanations of the equation and its parameters can be found in IPCC (2006) and EPA (2020), while a more basic explanation of the variables in the equation can be found in Appendix 1.

As methane passes through the intermediate cover, some of it is oxidized. This oxidation reduces the amount of methane emissions. Further, by including the empirical capture rate estimated by landfills from GHGRP, this further reduces methane emissions, resulting in end result emissions of:

$$CH4_t^{\text{emitted}} = CH4_t \cdot (1 - \text{capture rate}) \cdot OX$$

1.6.2 Cost-Benefit Equation

The implementation of the food scraps program takes food waste from the landfill and provides it instead to the composter. However, this comes at a cost to the jurisdiction providing the new food scraps service. Unfortunately, estimating the internal benefits and costs households experience is not feasible in this setting, so it is omitted from the cost-benefit equation.

The cost-benefit equation, in "per ton of organics disposed", is:

$$\begin{aligned} \text{Net Benefit} &= \Delta\text{CH4 emissions damages} \\ &+ \text{City tip fee reductions} \\ &- \text{Cost of the program} \end{aligned}$$

[23] This is assuming use of an average landfill food waste decay value of k = 0.12.

The variable **ΔCH4 emissions** represents the change in CH4 emissions when the program begins. All else equal, this will be a benefit, as CH4 emissions will be reduced because there is less decomposable material in the landfill. This benefit is calculated as

$$\text{CH4 emissions value} = \sum_{t=1}^{T=50} \text{CH4}_t \cdot GWP \cdot SCC$$

Where the GWP value used is 80 tons of CO2e per ton of methane, as consistent with the most recent IPCC estimate. However, recent literature, such as in Errickson et al. (2021), have estimated US methane costs at $8,290 per tonne of methane (using the GWP figure of 80, this is approximately $103 per ton of CO2e, implying about double the cost I use in this estimation). Given the equation, the SCC can be backed out to estimate the cost per ton of CO2e for the program. The variable **City tip fee reductions** represents the per ton difference in tipping fees to the city by diverting waste to a composter from a landfill. In general, this will be a benefit as composting tipping fees are almost always cheaper than landfill tipping fees. The negative value for **Cost of the program** is straightforward: it costs the city money to run the program. This cost includes additional trucks, labor, overhead, and purchasing new bins.

Certain costs are excluded from this analysis. An important cost that is excluded is transportation costs. Typically, large cities must transport waste a great distance to large landfills. Diverting waste might reduce the distance traveled to deliver the organics. It is assumed the distance between the two is approximately the same, which anecdotally is true, except in extreme cases. Pollution from electricity utilities is also not included because landfills are not marginal to electricity production.

The **Cost of the program** can vary for jurisdictions. For example, Austin's cost of $3.75/HH/month might be high because of their waste contracts or the department's cost structure. As a comparison, the city of Seattle in April of 2009 began to allow residents to divert FS&SP into the organics bin. Simultaneously, Seattle's waste department (Seattle Public Utilities) increased pickup frequency to weekly from bi-weekly (every other week). As a consequence of these two changes, organics bin rates increased. If it is assumed that the cost of a 32-gallon bin doubles because collection frequency doubled, then the cost increase from adding FS&SP to the yard trimmings bin cost households an additional \$3.42/HH/month.[24]

In the case of the city of Austin, Texas, these values can readily be filled in to evaluate the program. The component **ΔCH4 emissions** is calculated using the empirical household organics diversion estimate. The landfill that services Austin has a landfill gas capture system in place, but the methane is flared, thereby resulting in no additional use for the captured methane. The **City tip fee reductions** component was acquired through discussions with employees at ARR, where the landfill tip fee (for the year 2019) was \$23 per ton of waste and the composting tip fee (for the year 2019) was \$16 per ton of waste, resulting in a savings to the city of \$7 per ton of waste diverted to the composter. While other cities may have differences in costs associated with transportation (such as the landfill used being significantly farther from the city than the composter), this is not the case with Austin. The landfill and composter are on opposite sides of the city, but approximately equidistant in this regard, resulting in no additional transport costs. Finally, the **Cost of the program** is estimated at about \$9 million in total for FS&SP. There is a component for the yard trimmings program, but that program was already in place prior to the addition of food scraps to the organics bin (and prior to the start date of this analysis). The cost of the program per ton of CO2e avoided is \$546, while the cost of the program per ton of FS&SP is \$791.

In order to match the county covariates to Austin covariates, I run a regression in the form of Equation 1, except I interact the indicator of receiving an organics program with each of the following covariates: percent of households with at least a bachelor's degree, median income, average household size, average family size, percent under age 18, percent age 65 and over, and percent of the population identifying as white. I adjust median income by the cost of living index in order to more reliably represent income of a county. Then, I sum up the resulting estimates

[24]Prior to the addition of FS&SP, compost cost \$5.35/HH/month for up to 128-gallons of waste. If the cost is assumed to be linear, this is implies a 32-gallon bin would have cost \$1.34/HH/month. Doubling this, we get \$2.68/HH/month. After the change, a 32-gallon bin cost \$6.10/HH/month, so 6.10 - 2.68 = \$3.42.

in order to obtain the implied estimate for pounds of organic waste diverted per household per week. I input this estimate into the cost-benefit equation to estimate yearly tons of food waste and get the estimate for county methane emissions.

For counties I use only the contiguous United States. I omit Alaska because of the existence of numerous, very small landfills scattered across the state. Accurately matching with the demographic composition of these areas is not feasible, given their unique characteristics. In particular, these landfills typically serve fewer than 500 people and are not subject to the same federal regulations as landfills that service a larger population. For similar demographically unique reasons to Alaska, I do not include Hawaii either.

There are a total of 3,108 counties in the contiguous United States. This data comes from the Census, along with the same covariates from the Census as used for the City of Austin. My data has a total of 1,716 open landfills. This data comes from the Waste Business Journal, a trade journal that collects landfill data (at least) every year. This data contains information such as landfill name, location coordinates, tipping fee by year, disposal tonnage by year, and ownership. In addition, I download data from the Greenhouse Gas Reporting Program (GHGRP) hosted by EPA. This data consists of similar data as the Waste Business Journal, however it is voluntary, so many landfills are missing from the data. While the GHGRP data lacks the more granular business details such as tipping fees, it does contain emissions and capture system data, such as if a capture system is in place, what kind of capture system is in place, collection efficiency, and emissions estimates for the landfill. I merge the Waste Business Journal data with the GHGRP data to construct a landfill database. For landfills not in the GHGRP but in the Waste Business Journal database, I assign them as having no capture system in place.

For the largest 30 counties by population, I match their true landfills used if that information is available. Very large areas produce a lot of waste and this waste either must go to a "mega landfill" or to a few different large landfills. If we take New York City, for example, all waste is exported out of the state to Pennsylvania, Virginia, and Michigan. Connecting this is important because these landfills might have different collection systems in place than nearer, smaller landfills, which would generate inaccurate estimates of methane emissions. The remaining counties are matched to the nearest 3 landfills because they typically deliver to their nearest landfills. I assume that counties send 1/3rd of their total waste to each landfill.

1.6.3 Emissions Costs

I present my results in a sequence. The first map only includes methane emissions. All maps thereafter include the full cost-benefit analysis results. The cost-benefit analysis uses a SCC of $51 per ton of CO2e and a GWP20 of 80 where relevant. These parameters can be adjusted and do not necessarily represent the true damages. In particular, as the results will show, these values will result in a consistently negative net benefit across the counties of the contiguous United States. Given the GWP used is an optimistic choice, using a lower GWP will result in smaller benefits from avoided methane emissions.

For all of these maps, I assume a time frame of 50 years. For estimating methane emissions, it is common practice to use at least a 50 year window in order for the model to converge on a near-constant rate of methane emissions. All amounts presented should be interpreted as being over a 50 year time frame. In particular, the methane emissions are monotonically increasing up to the converging amount. The plot shown in Figure 1.5 describes methane emissions from a landfill. In this plot, I assume a rate of one ton of food scraps disposal per year for 50 years, at which point disposal ceases. It can be seen that the plot begins to converge around 30 years and the slope is nearly flat around the 50th year of disposal. Methane emissions immediately begin to drop off once disposal ends.

In all of the plots, there will be two kinds of dots overlain on the maps. The light blue dots represent landfills with some kind of landfill gas capture system in place, while the dark blue dots represent landfills with no LFG capture system in place.

The largest counties by population are the greatest emitters of methane pollution simply because of the large number of people within the county, all of whom dispose of waste. An initial guess might suggest that large metropolitan areas should implement organics programs,

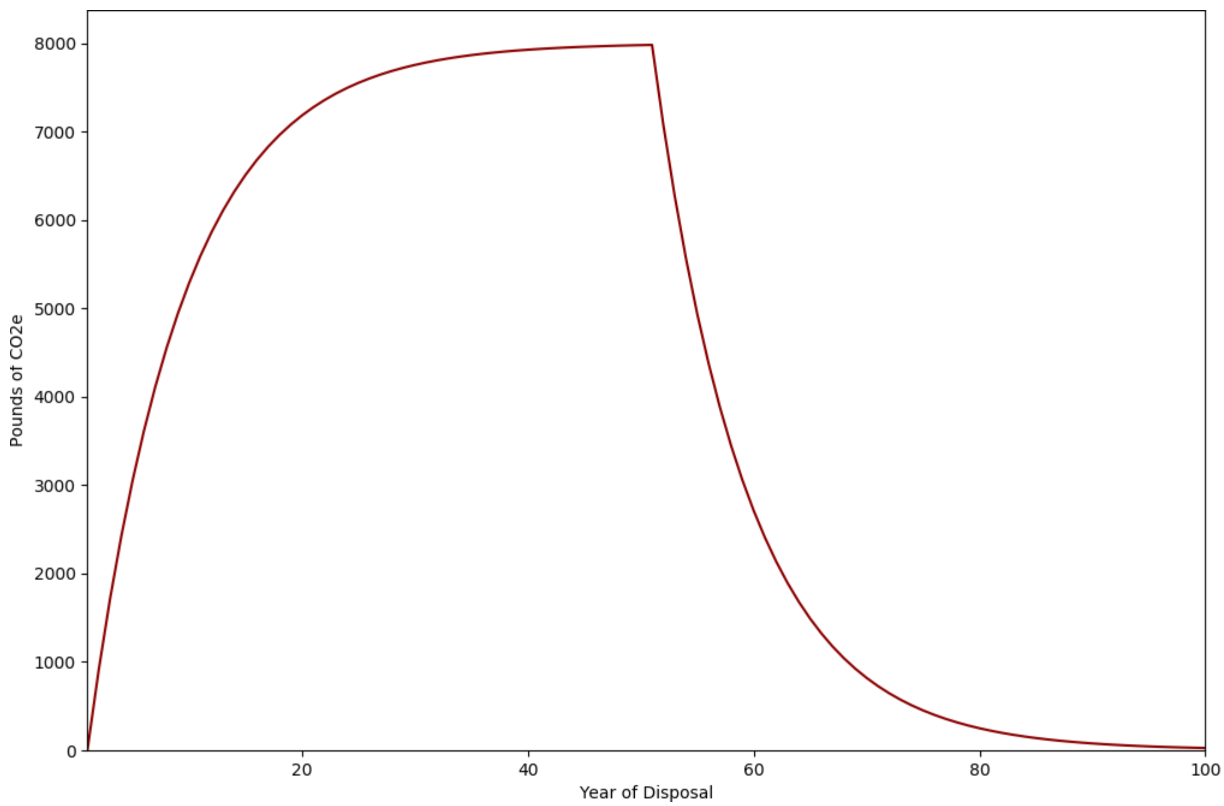

Figure 1.5. Example Methane Emissions from Landfill

This graph shows the estimate methane emissions from a landfill that receives one ton of food scraps per year for 50 years. The convergence of methane emissions begins to become clear around the 30th year. I use conventional values here, in particular I assume a decay rate (k) of 0.12.

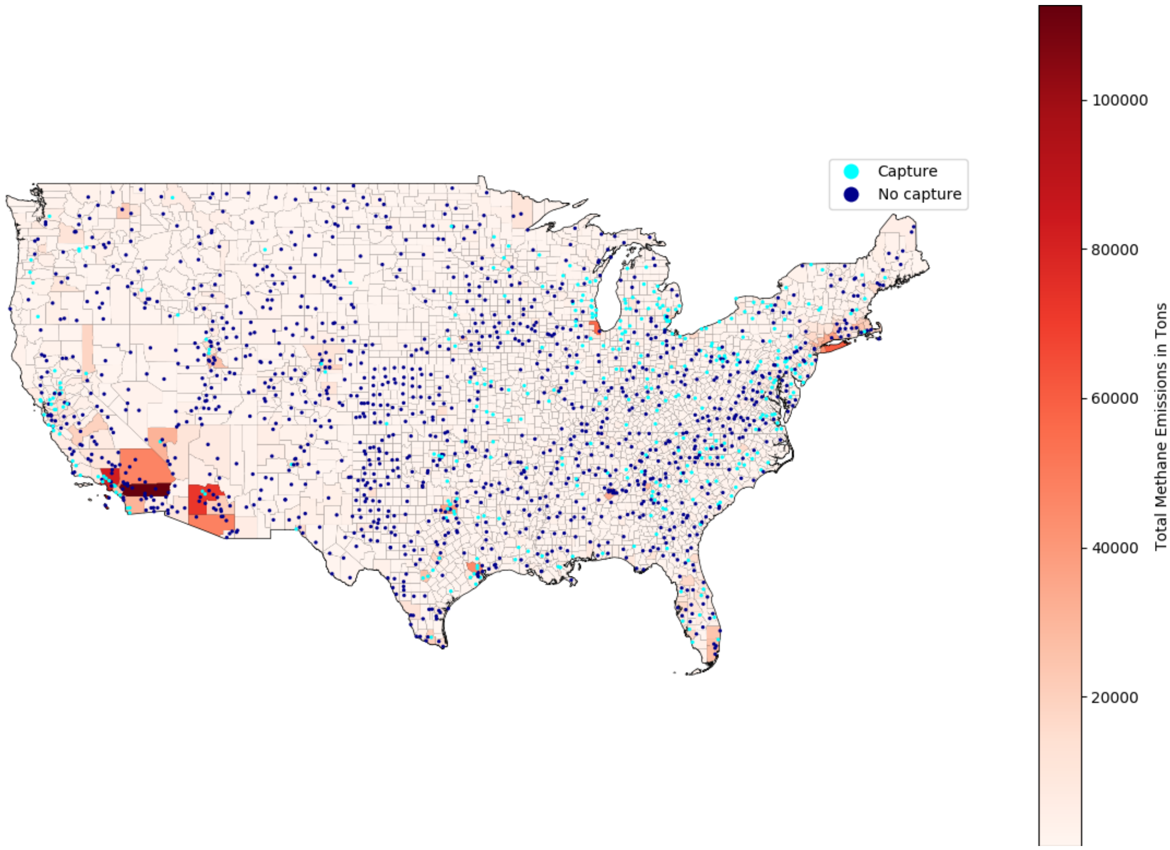

Figure 1.6. Estimated Methane Emissions from FS&SP Diversion

This graph shows the estimated methane emission tonnages by counties across the United States accounting for capture rates.

while smaller counties should not. However, the requirement that sufficiently large landfills must install an LFG capture system will stymie this guess because large population centers effectively have to use landfills with capture systems in place.

Figure 1.6 displays estimated emissions while accounting for landfill gas capture systems in place. In the case with no LFG capture system, the maximum methane emissions declines from approximately 454,000 tons of methane emitted, which amounts to emission damages of $1,816,000,000, to a maximum when including LFG capture systems of 126,000 tons of methane emitted, which equals $504,000,000 of damages.

Plots now proceed with full net benefit calculations. Figure 1.7 shows the cost per ton of CO2e avoided. The values in the plot are much higher than the SCC of $51 per ton of CO2e.

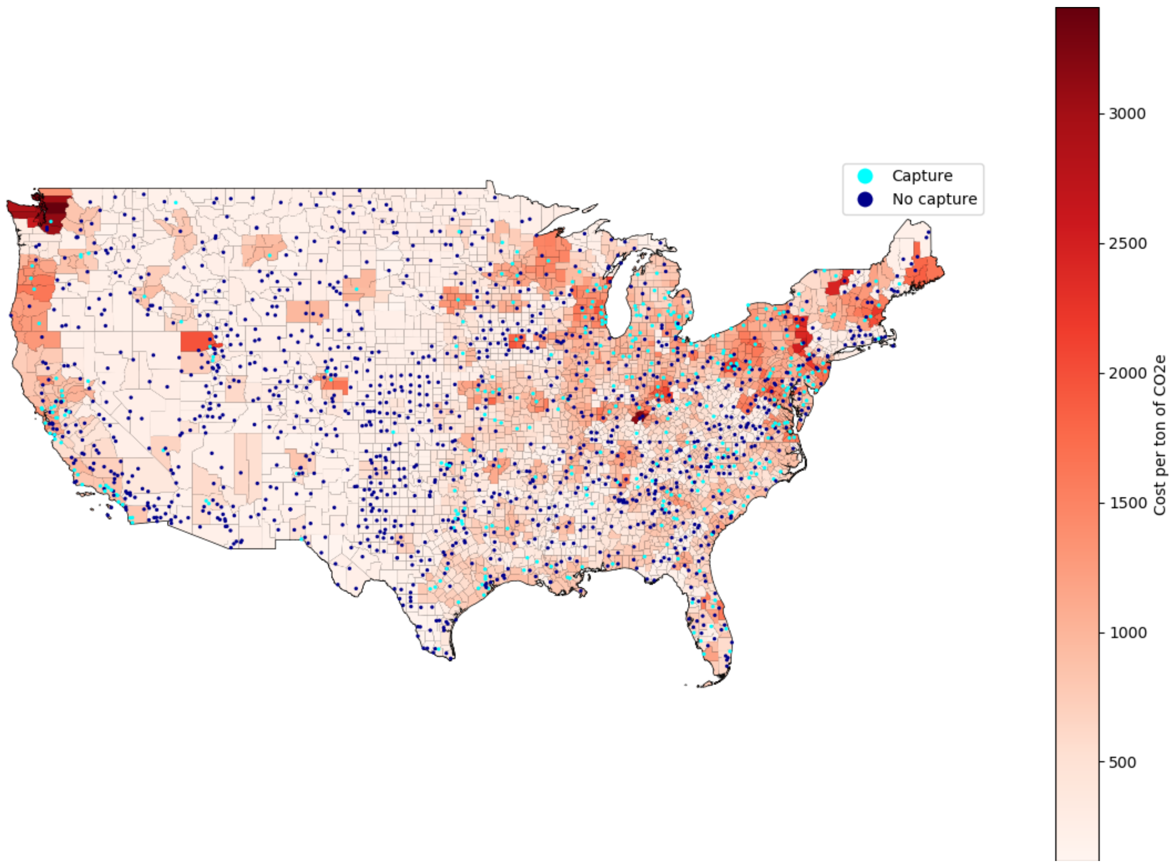

Figure 1.7. Cost per ton of CO2e Avoided

This graph shows the estimated cost per ton of CO2e avoided for counties across the United States. The median and mean values are $536 and $627 per ton of CO2e avoided, respectively. The monthly household cost for the FS&SP program is the same as in Austin, Texas of $3.75/HH/month.

There are several reasons for this. The first is that, as Figure 1.6 displays, while large population centers emit a lot of methane, they also tend to have LFG capture systems in place. This reduces the social value of an organics program because diverting waste will release substantially less methane into the atmosphere than without the capture system. The second reason is that Austin does not have perfect compliance with their organics program, which this calculation extends to other regions. The third reason and arguably the most important is simply the cost of the program. While it is not a particularly large amount per resident, at a cost of $3.75 per household per month, on average a household must dispose of $3.75 worth of FS&SP equivalent emissions per month to break-even. This amounts to approximately 9.7 pounds of FS&SP per week *assuming* no capture system is installed. This amount is near or exceeds the average total FS&SP disposable within households.

The values shown can be compared to the SCC is an estimate that is subject to uncertainty and, in particular, because the tail risks have substantial consequences for the planet, understanding what SCC values can justify an organics program expansion is important. For example, to limit warming to 1.5 degrees Celsius, de Coninck et al. (2018) suggests a carbon price range in 2030 should be between $135-$5500 per ton of CO2e. The estimates from this map are near the lower end of this estimate.

Figure 1.8 shows the necessary weekly disposal amount per household required to achieve a net benefit of zero. While many areas across the United States *might* produce the necessary amount of FS&SP if disposing at a 100% rate, much of the United States exceeds 100% of their weekly FS&SP in total. The median weekly diversion rate per household is estimated at 27 pounds. This suggests that even with a 100% FS&SP diversion rate (which again, is 7-10 pounds per week per household at most), these programs do not pass a cost-benefit test at their current costs or at the chosen SCC. The next two graphs explore the net benefit of an organics program expansion and monthly household costs to break-even.

Figure 1.9 displays the net benefit per ton of food waste disposed. As the plot shows, in all counties the net benefit is negative. The median net benefit is $528 per ton of FS&SP

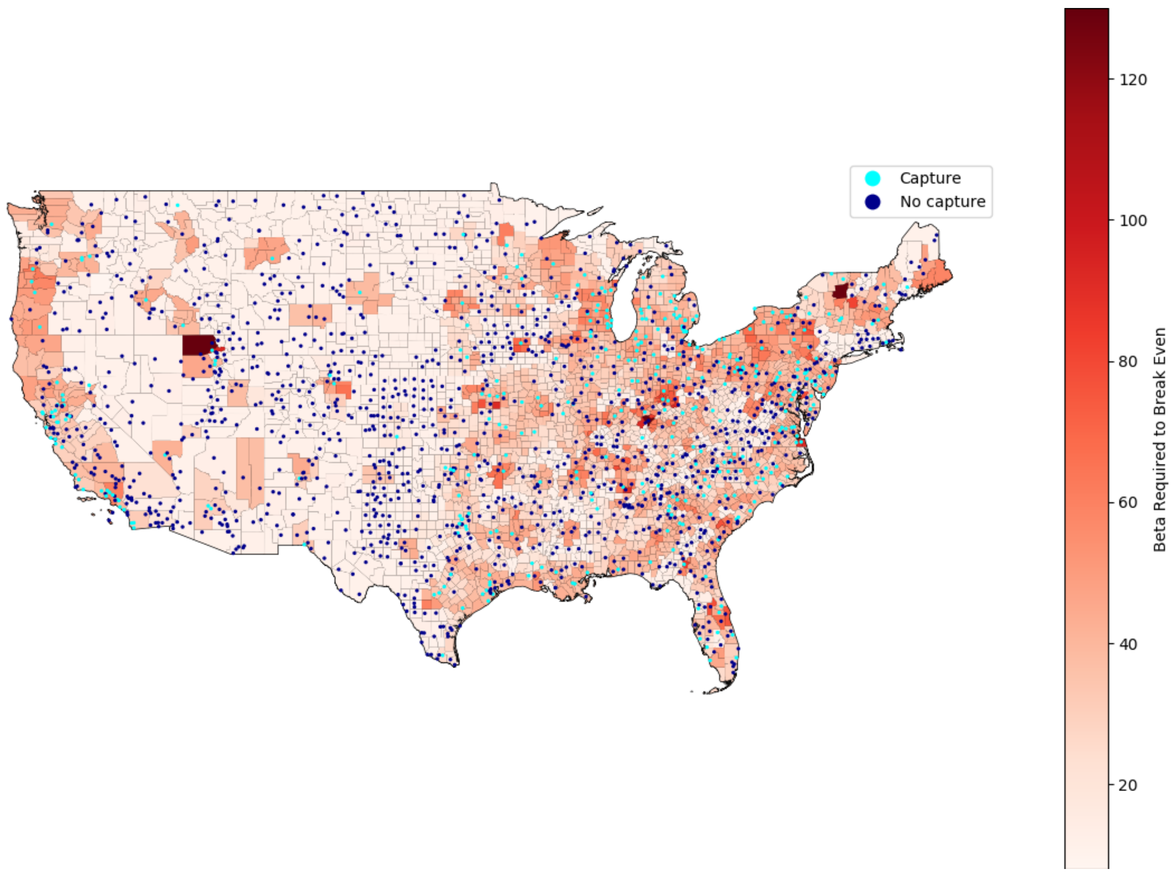

Figure 1.8. Weekly Household FS&SP Disposal Rate Needed

This graph shows the estimated required weekly FS&SP disposal rate of households per month needed for counties across the United States to have a net benefit of zero, using a monthly household cost of $3.75

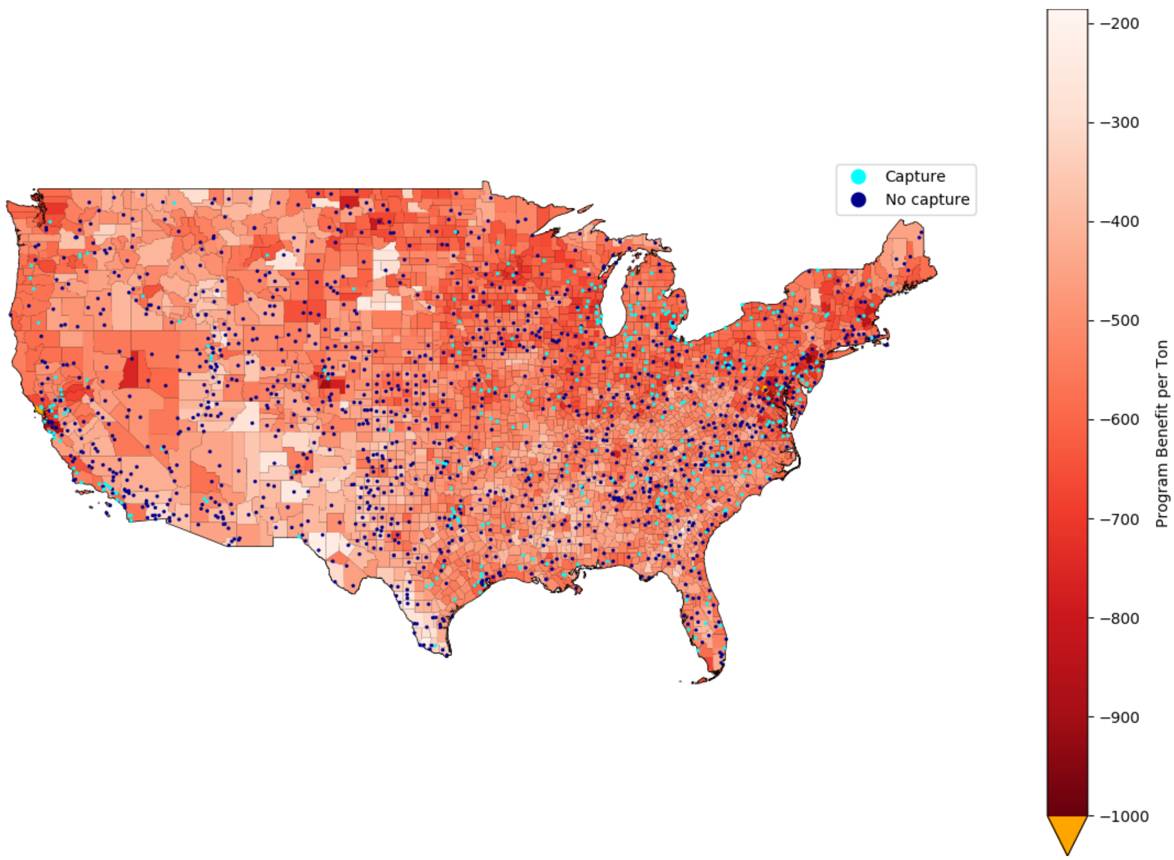

Figure 1.9. Net Benefit by County: Cost of \$3.75/HH/mo.

This graph shows the estimated net benefit for counties across the United States, using the same monthly household cost for the FS&SP program as Austin, Texas.

disposed. Many of the smaller counties in the United States have a lower SCC required; this is primarily due to the fact that their landfills typically don't have a LFG capture system installed. The counties with the highest cost per ton of CO2e are those with large populaces that are served by landfills with LFG capture systems.

These plots all show not only the variation of uptake for an organics program, but also the importance of the landfill where the waste is sent. If a landfill has a LFG capture system then an organics program has a harder time passing a cost-benefit test, even more so if the system is more efficient than average. However, while this conclusion might be valid, this does *not* examine the costs of a landfill itself, which might be substantially more in whole than diverting waste.

Now, I turn my attention to the issue of landfill methane emissions underestimation. The underestimation of landfill emissions in total is important. With methane emissions comprising a large part of the social benefit, getting this estimate right is essential. While the cost-benefit largely remains negative, it presents a less stark result.

In Figure 1.10, I double the total emissions tonnages. This is in line with upper ranges from Table 1.1 in of Sciences (2018), where underestimates are within this range. The graph still depicts a large cost per ton of CO2e avoided, however it is near the upper end of estimates from Nordhaus (2017) for more stringent temperature increase goals at 2.5 degrees Celsius. while it is on the lower end of estimates provided in de Coninck et al. (2018).

Table 1.5 summarizes the map information, with the columns displaying the 1st, 25th, 50th, 75th, and 99th percentile of the cost of a FS&SP program per ton of CO2e avoided. Each block is a specific estimate. The first block of the table shows the "Base Estimate", where only standard values are used for the estimation. In the second portion, the emissions from landfills are doubled. As we would expect, when emissions are doubled, then the value of FS&SP doubles, resulting in the cost per ton of CO2e avoided to halve. This also shows that most of the value from FS&SP diversion is a result of avoided methane emissions. The third block shows what happens if methane capture is lowered by one-third because FS&SP decomposes more quickly

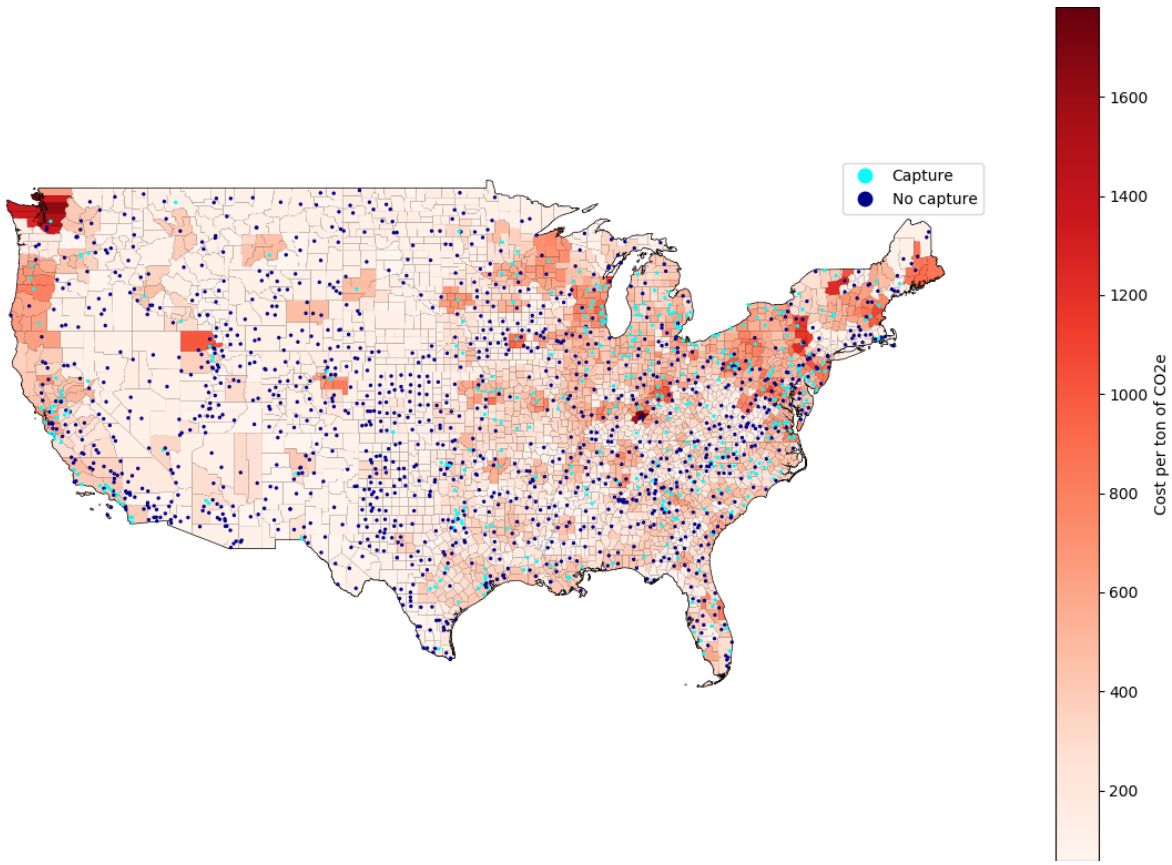

Figure 1.10. Cost per ton of CO2e Avoided: Doubled Emissions

This graph shows the estimated cost per ton of CO2e avoided for counties across the United States assuming the landfill emissions rate is doubled from the estimate a standard FOD equation produces, using a monthly household cost of $3.75

Table 1.5. Results of Cost-Benefit

Estimate	*Cost($) / tCO2e Avoided*				
			Percentiles		
	1st	25th	50th	75th	99th
Base Estimate	161	240	536	879	2120
Double Emissions Estimate	81	120	268	439	1061
Lower Capture Estimate	161	238	348	430	621
Constant Capture Rate (70%)	503	665	718	777	986
No Capture Estimate	152	201	217	235	298
Using GWP100 (= 32)	403	600	1348	2167	5275
Max Diversion Rate (10 lbs/HH/week)	55	62	151	236	463
No Tip Fee Difference	164	245	547	895	2171

Note: This table displays the 1st, 25th, 50th, 75th, and 99th percentiles of the estimates. The first row displays the base estimate. The remaining rows make certain assumptions within the cost-benefit model to explore the effect of different margins. Descriptions of each row are provided in the text.

than most waste, resulting in lower landfill gas capture. The lower percentiles are very close to the base estimate because these are landfills that are already not capturing their emissions or have low capture rates. Higher percentiles have a much lower cost per ton of CO2e avoided because their capture rates are lowered.

The remaining rows make certain assumptions about the cost-benefit equation to display how particular parameters matter. In the fourth row, the capture rate is set to 70%.[25] This results in a rightward compression of the distribution of cost per ton of CO2e avoided. The lower cost counties have a dramatic increase (because they have no capture) while the higher cost counties have a decrease (because they have higher empirical capture rates). In the fifth row, I assume that landfills have no capture. This results in a leftward compression of the distribution, reducing the cost per ton of CO2e avoided.

The sixth row uses the GWP100 figure of 32 tons of CO2e per ton of methane instead of 80. This dramatically increases the cost per ton of CO2e avoided. This should be anticipated, since the emissions damages are now much lower. In the seventh column, I assume a "maximum" diversion rate of 10 pounds per household per week. This greatly reduces the cost per ton of CO2e avoided, suggesting that in many regions of the United States, if maximal FS&SP can be achieved, then the costs are sufficiently low to justify the programs. I believe this estimate is particularly noteworthy; while there will be variation across counties in the maximum FS&SP diversion rate, if a county can achieve 100% participation and diversion (or close to it), the programs can be readily justified. In the eighth row set the difference between the tip fee at the landfill and the composter to zero. The tip fee at the composter is always lower than the landfill, however it can be seen this is a rather small portion of the cost.

This table shows that the distribution of methane polluting behavior has a long right tail. This is due to the very high populace counties, which dispose of high quantities of FS&SP. As can be seen, if the evidence that emissions from landfills are undercounted is correct, then the

[25]This is because in EPA (2020), the emissions for a landfill with no capture is calculated to be 1.39 tons of CO2e per short ton of food scraps, while capture with electricity use is 0.42 tons of CO2e per short ton of food scraps. The implied capture rate from the difference of these values is approximately 70%.

cost per ton of CO2e avoided is substantially lower than the baseline estimate.

1.7 Discussion

The purpose of this section is to discuss what alternatives are feasible given the results of my cost-benefit analysis.

The City of Austin's organics expansion is younger than other cities' programs. For example, the City of Seattle implemented their food scraps program in April of 2009 (and changed collection from bi-weekly to weekly). On January 1st, 2015, Seattle then banned all organics from the garbage, making organics disposal mandatory. Seattle's organics program already had a reasonably high participation rate, but the requirement of mandatory organics disposal increased the number of households subscribed and slightly increased the disposal rate.

It is entirely possible that household curbside organics is *not* worthwhile, yet commercial organics programs are worthwhile. While this is a topic for a different paper, supermarkets produce a large amount of organics waste. Implementing organics collection or diversion for this has several benefits. First, it is generally cheaper to dispose of organics in bulk. Second, the organics stream from a supermarket is largely food waste, which can be used uniquely as a specific input for certain composting/reuse methods. Finally, the food waste doesn't have to be disposed of; in particular, some commercial composting policies encourage food deemed safe to be provided to homeless shelters and food banks.

In some cases, the incentive for a jurisdiction to expand an organics program is ambiguous. The social benefits received from a reduction in methane emissions is global, not local, so this is typically not factored into local benefits. In certain parts of the United States, such as the northeastern United States, landfill tipping fees are very high (typically around $100 per ton of waste or higher). In many cases, the jurisdictions must pay this fee. Diverting organics to local farms or larger-scale composting/anaerobic digesters is typically much cheaper. In the case of California, jurisdictions are implementing composting programs because the laws mandate them.

In the case of Austin, it is not as clear. Anecdotally, the government incentive to implement the program largely seems to come from a warm-glow or (relative) over-valuation of environmental benefits. Certain cities, such as Austin or San Francisco, are also subject to certain perceptions nationally which they may be inclined to fulfill.

In some cases, such as Massachusetts' law to have a state-wide commercial composting program, jurisdictions are persuaded by (claims of) increases in employment and revenues. In addition, they can cite to business owners their costs will decrease because organics hauling is cheaper than garbage hauling.[26] This is pertinent for commercial organics, but it is not true for residential programs.

An alternative to composting not discussed in this article is anaerobic digestion (AD). With AD, organic material – be it of a singular kind (e.g. food waste) or mixed (called co-digestion) – is placed into an enclosed reactor, where microbes break down the matter within. This process generates *both* biogas and "digestate"[27], which can be converted into fertilizer or other products. The production of biogas in an AD is typically higher than in a landfill, between 50-75%. The resultant digestate has both solid and liquid forms and if treated properly can be converted into a fertilizer, compost, or an addition to soil. Most AD systems are located on farms, with it being fairly rare for an AD facility to be stand-alone. These are a viable alternative to composting, with the additional by-product of biogas. In addition, industrial-scale composting centers typically take up many acres of space in order for the organic material to aerobically decompose into compost, which is a similar constraint faced by landfills seeking to expand.

1.8 Conclusion

Methane emissions are substantially more costly than CO2 emissions, especially over a short horizon. This paper explores one methane emission mitigation strategy: household

[26] Per ton, this is true. However, hauling is typically priced by bin size, not weight, so it is possible for businesses to be spending more if the cost of the two bins required (even if the business acquires a smaller garbage bin) is more than the single, larger garbage bin.

[27] This is the by-product of anaerobic digestion after the organics have been processed.

diversion of organic matter – such as food scraps, leaves, and grass – from landfills. However, jurisdictions must pay for these programs, which can be costly. I first provide novel findings regarding household responses to receiving an expanded organics program. Households do participate in this program (though not all participate); on average, households that participate divert nearly all of their organics. Further, households fix mistakes they make regarding recycling once they are provided with the proper method of disposing of their seemingly-but-not-actually recyclable material.

With this estimate, I wrote a cost-benefit equation to explore where in the United States an organics program (expansion) would be viable. In addition to the tension between the social benefit of avoided methane emissions and cost of the program, the largest landfills emitting methane typically have landfill gas capture systems in place. All of this results in a needed cost per ton of CO2e avoided of approximately \$536 per ton of CO2e or an additional diversion rate of approximately 27 pounds per week per household, which is significantly greater than 100% of organic matter that could be disposed of under the program expansion to achieve a break-even net benefit.

While an expanded organics program does appear to fail a cost-benefit test, this does not mean it is not viable. The estimates here are founded on specific assumptions that may not prove to be true. In particular, if there is a true social cost of carbon and it is much higher, then the estimated damages from methane emissions will increase substantially. Also, if a jurisdiction can implement the program rather cheaply (as a monthly cost to a household), then the programs can also be worthwhile.

The conclusions of this article should concern proponents of household organics diversion, but the proper conclusion is not dismal. Alternative methane-reducing options, such as reducing enteric fermentation or oil and gas well and pipe leaks may be more viable option (Hausman and Muehlenbachs (2019)). However, the implementation of organics programs requires careful investigation of how to lower the cost of the programs and how to generate high level of participation across a jurisdiction's populace. Both of these are essential to the viability of organics programs.

1.9 Appendix 1: First-Order Decay Equation

The first order decay equation used is considered good practice by the IPCC (IPCC (2006)). The equation embodies a linear relationship between waste in place and landfill methane emissions. While this estimation method has faced much scrutiny, it is standard and more robust estimation methodologies do not have data available across the United States.

The equation in year t is:

$$CH4_t = \sum_{i=0}^{t}\sum_{s<t} W_s \cdot DOC \cdot DOCf \cdot MCF \cdot F \cdot (16/12) \cdot (e^{-k(t-s-1)} - e^{-k(t-s)})$$

Each of the acronyms must be described and more details can be found in IPCC (2006). First, pick a particular waste type, such as food waste, wood, grass, and so forth. The weight of this is represented by W. A certain percent of each waste type is degradable organics carbon, whereas the remainder is not; this is what DOC represents. For food waste, the number is 15%, while for soiled paper it is 40%. A certain percentage of this DOC is also decomposable; the standard value is 50% and is represented by $DOCf$. The value MCF is the “methane correction factor”, a value between 0 and 1 that is 1 for all landfills in the United States, but in lower- and middle-income countries this value will typically be less than one. The factor itself is a representation of how effective the anaerobic decomposition in the landfill is; the worse a landfill is managed, the less anaerobic decay that likely takes place. The term F is the percentage of landfill emissions that is methane – the standard percentage is 50% and this is typically not a debated value. Between 50-60% of landfill emissions are methane while 40-50% are CO2, with a middling amount of NO2. As discussed in Section 5.1, a landfill emits both CO2 and methane but only the methane is considered in landfill methane inventory calculations. The $\frac{16}{12}$ number is the weight ratio of methane to CO2. Finally, the $e^{-k(t-s-1)} - e^{-k(t-s)}$ is used to calculate the amount of waste in place for a specific year. In this case, since $s < t$, it is the amount of waste that decomposes in the year s. The term k represents the decay rate; a larger k provides a faster decay rate and a smaller k is a slower decay rate. If waste is deposited in a given year, say year 1, then each year thereafter it decomposes. However, every year less and less material from year 1 decomposes, by weight. In order to determine the half-life, $t_{1/2}$, of the waste type in place it is calculated as $t_{1/2} = ln(2)/k$.

This equation does not include an index for waste material type. The equation above is the same for all materials, only the parameter values change. That is, for each waste type, calculate CH4_t and sum all of the resulting emissions values to obtain the total emitted.

The above equation describes the potential CH4 emissions from a landfill. However, several factors can cause landfill methane emissions to be lower or higher. In this case, both the landfill gas capture rate and the oxidation factor impact methane emissions, both of which lower emissions. In the first case, a landfill must first have a capture system in place. These capture systems can vary in their effectiveness. As methane travels upward to escape the landfill, it passes through the intermediate cover layer. This oxidizes some of the methane.

$$CH4_t^{\text{emitted}} = CH4_t \cdot (1 - \text{capture rate}) \cdot OX$$

Where OX is the oxidation factor, typically assumed to be 0.9. The capture rate used for a landfill is calculated as the average capture rate from 2017 to 2019 based on the GHGRP data. The mean collection efficiency for all landfills is 45%, however for landfills with LFG capture systems the collection rate is 67%.

1.10 Appendix 2: Additional Methane Emissions Details

Leaving aside legal permitting, a modern "sanitary" landfill begins by digging a very large, deep hole. The bottom and sides of the cell are lined with a low-permeability layer, such as densely-packed clay. Atop this is a (typically plastic) liner. Both of these layers are intended to prevent leachate[28] from leaking out of the landfill into nearby groundwater. Above the liner

[28] Leachate is created by water percolating through the waste in a landfill and "leaching" out some of the components of the waste passed through. Less technically, it can be thought of as "garbage juice".

is typically a geotextile mat to prevent the gravel placed on top of this mat from damaging the plastic liner. In between this is a series of leachate collection pipes, which collect leachate and (usually) pipe it out of the landfill. On top of the gravel layer is a soil layer that ranges in thickness of one to two feet. Finally, waste is placed on top of all of this. Waste is placed every day, where it is compacted (typically by a bulldozer driving over the waste pile) and dirt is placed on top at the end of every day, called a cover. This represents the closure of the cell.

Figure 1.11 shows the four phases of pollution from decomposition in a landfill. There is a fifth phase for decomposition, but it is not displayed in this graph as it is the end phase where the landfill has stabilized and has a limited number of chemical reactions.

The first phase is aerobic. Oxygen-consuming bacteria break down the waste in a landfill, which largely results in carbon dioxide. Over time, the oxygen in the landfill is exhausted and phase two begins.

Phase two, as with all following phases, is anaerobic. This phase is a transition phase from an aerobic to anaerobic environment. The bacteria in the waste pile produce alcohols (e.g. methanol and ethanol) and acids (e.g. acetic, lactic, and formic acids). This conversion produces carbon dioxide and hydrogen. The pH level of the landfill lowers, while the acids interact with certain available nutrients and produce nitrogen and phosphorus.

Phase three exhibits bacteria consuming the available acids produced in phase two and producing acetate. The consumption of acids raises the pH level in the landfill. These bacteria produce compounds for methanogenic bacteria to consume. The acetate-producing bacteria produce nutrients for the methanogenic bacteria, while the methanogenic bacteria consume the carbon dioxide and acetate, which allows the acetate-producing bacteria to thrive.

Phase four features the same processes as phase three, but this phase is marked by the relative stabilization of the composition and production of methane and carbon dioxide. The precise composition split for methane and carbon dioxide is not known, but it is typically around 50%, with a spread of about 5-10%.

The length of these phases depends on the waste composition within the landfill. However,

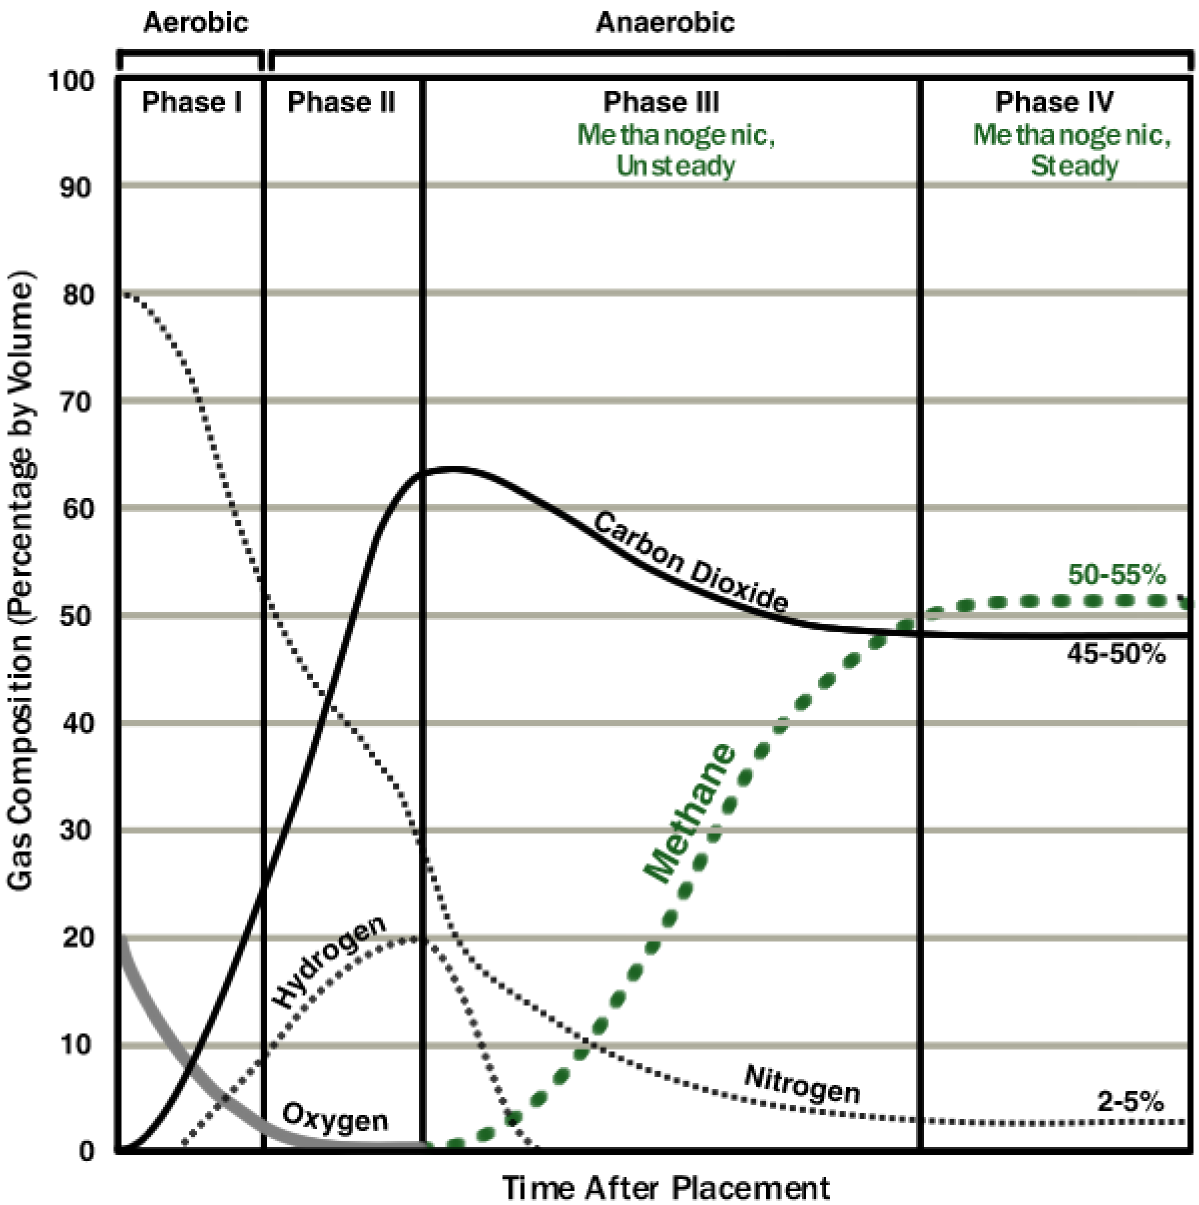

Figure 1.11. Phases of Landfill Decomposition

This graph shows the four phases of decomposition in a landfill. The phases are described in the text.

gas production typically lasts for at least 20 years while emissions from the landfill can last for much longer. If a landfill is composed of a large number of organics, gas production will typically be much longer. Different portions of a landfill will be in different phases of this decomposition process.

The stage of the waste cell impacts methane collection efficiency. Emissions collection in a cell does not start immediately, however I am abstracting from this problem for these purposes. The thickness of the intermediate cover used affects methane oxidation as does particular features of the soil and climate, such as dry soil or a dry climate. The thickness of intermediate cover is typically not tracked by government entities nor are the intermediate cover soil characteristics. Precisely the extent to which these variables affect landfill methane emissions is not currently known.

1.11 Appendix 3: Additional Maps

The plot seen in Figure 1.12 shows the necessary household monthly organics service fee required for the net benefit to be equal to zero. There is a strong correlation between a higher household cost (which means that it would be “cheaper” to implement the program) and whether the county is serviced by a landfill without a LFG capture system. For example, large-populace counties such as those containing Seattle, Los Angeles, Austin, or Philadelphia require much lower household costs to be viable (implying it is more expensive for them to implement) because they are serviced by large landfills with LFG capture systems in place.

1.12 Appendix 4: Alternative Specifications

Recent econometric literature has called into question the validity of two-way fixed effects (TWFE) models in the presence of treatment effects varying over time (Goodman-Bacon (2021), Sun and Abraham (2021), Callaway and Sant’Anna (2021)). The recent literature has described a significant problem with current TWFE models: the treatment effects may be biased

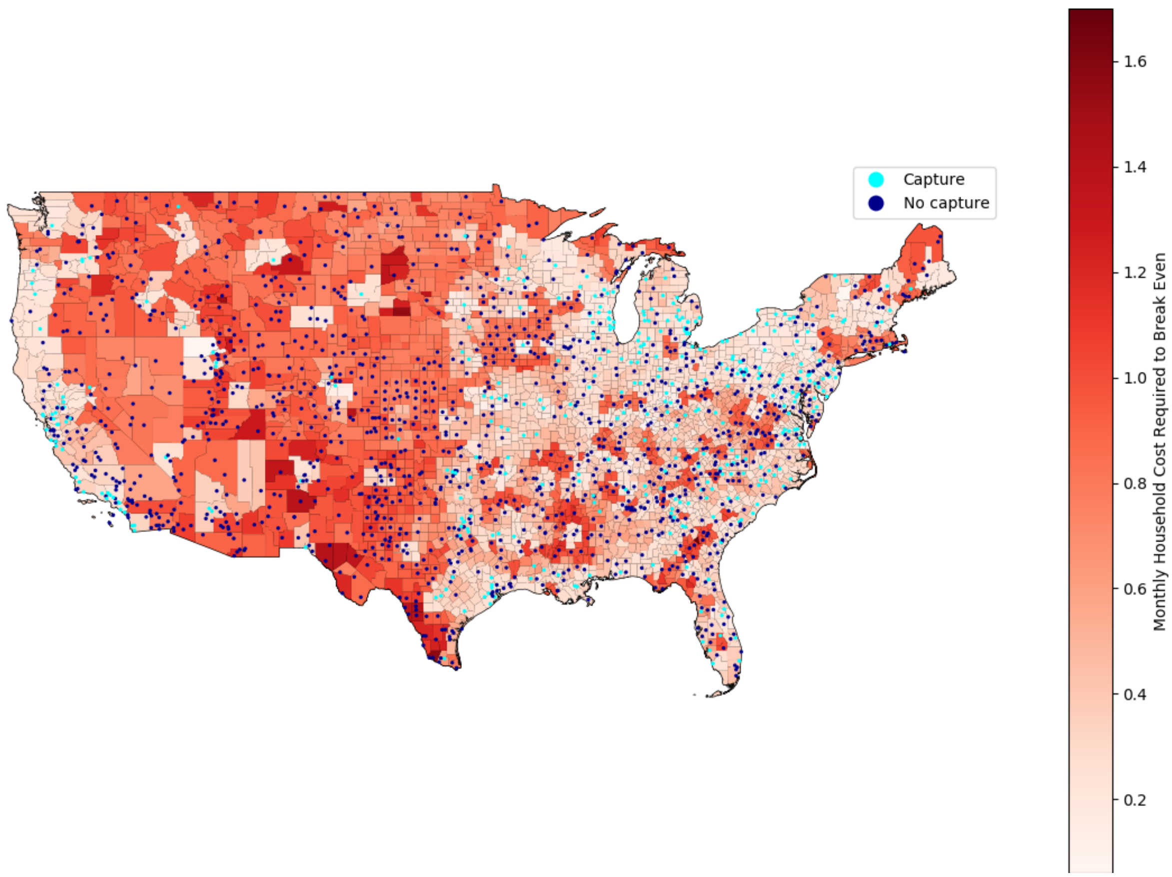

Figure 1.12. This graph shows the estimated household cost per month needed for counties across the United States to have a net benefit of zero.

and even more worryingly, the treatment effects may have units assigned negative weights. This problem does not arise in standard two-by-two difference-in-differences models, only in the generalized version of difference-in-differences, which is when TWFE models are typically used.

As a robustness check to these recent concerns, I estimate a regression using the assumptions and methodology of Callaway and Sant'Anna (2021). The authors estimate a treatment effect, $Y_t(g)$ that depends on the group g and the time period t. The group g is defined by the time at which treatment begins for a group of treated units. The authors avoid the issues of TWFE models by estimating an average treatment on the treated for each group-time pair, $ATT(g,t)$. In the simplest version, these $ATT(g,t)$'s are actually two-by-two difference-in-differences estimates comparing the treated g to some control group, comparing some post-treatment period t to the pre-treatment period $g-1$ (the g denotes both the group and the treatment timing).

Following the paper's assumptions, I assume there is no treatment anticipation (Assumption 3 in the paper) because households were not permitted to begin diverting their FS&SP into the organics bin beforehand. In addition, I present results using both "Not-Yet-Treated" groups (Assumption 5) and "Never-Treated" groups (Assumption 4). Although all groups are eventually treated, they are not treated in my sample horizon. Further, the results are nearly identical. To avoid any issues, I will use the results from the "Not-Yet-Treated" regression. I use the "doubly robust" estimator provided in the paper, which exploits both outcome regression and inverse probability weighting components of the $ATT(g,t)$ estimator.

The authors present several schemes for aggregating the $ATT(g,t)$'s to provide a swathe of estimates of the effect of the policy under consideration. For estimating group level effects, they propose: The authors propose the measure,

$$\theta_W^O = \frac{1}{\kappa} \sum_{g \in \mathscr{G}} \sum_{t=2}^{\mathscr{T}} 1\{t \geq g\} ATT(g,t) P(G = g | G \leq \mathscr{T})$$

where $\kappa = \sum_{g \in \mathscr{G}} \sum_{t=2}^{\mathscr{T}} \mathbf{1}\{t \geq g\} P(G = g | G \leq \mathscr{T})$, to ensure the weights in the second term sum to one. This is the "simple" ATT term in the tables below. This measure represents the weighted average of the $ATT(g,t)$'s. This estimate avoids the potential issus with a TWFE estimator, however it does not account for treatment timing. That is, some groups may be exposed to treatment longer than others.

To accommodate this, the authors first propose the estimate:

$$\theta_{sel}(\tilde{g}) = \frac{1}{\mathscr{T} - \tilde{g} + 1} \sum_{t=\tilde{g}}^{\mathscr{T}} ATT(\tilde{g},t)$$

where $\mathscr{T}$ is the final time period in the sample. The $\theta_{sel}(\tilde{g})$ represents the average treatment effect among the units in group $\tilde{g}$ over all periods from the beginning of their treatment to the end of the sample period. This measure can then be aggregated to provide an overall treatment effect.

$$\theta_{sel}^{O} = \sum_{g \in \mathscr{G}} \theta_{sel}(g) P(G = g | G \leq \mathscr{T})$$

where each $\theta_{sel}(g)$ is defined in the equation above. This term better accounts for groups that are in the treatment for longer to estimate the overall ATT. This estimate is called the "Group" ATT in the tables.

There are three results in the table presented. The first result will show the "simple" aggregation, where the average of the group-time treatment effects are calculated, weighted by the number of units in each treatment group. The third result will show the "group" aggregation, which takes the average treatment effect across each group. The second result aggregates to an overall effect by averaging the group effects while weighting them by the length of time treated.

Table 1.6 contains results using both "not yet treated" units and "never treated" units. The preferred specification is for "not yet treated" units as all the units are eventually treated in the rollout, even though the entirety is not in the purview of this sample. Regardless, the estimates for both groups are nearly identical. The format of the two tables is identical, with the "simple" aggregate first, then the "group" aggregate, and finally the estimate for each treatment group.

Table 1.6. Callaway & Sant'Anna Estimation: Main Estimate

	Dependent variable:	
	Pounds of Organics Disposed Per Household	
	Not Yet Treated	Never Treated
ATT "Simple"	2.6959***	2.4965***
	(0.5264)	(0.5915)
ATT "Group"	2.7502***	2.6233***
	(0.3094)	(0.399)
Pilot 1	4.8851***	4.8090***
	(1.1293)	(1.3095)
Pilot 2	1.6611	1.1927
	(3.2542)	(3.4216)
Phase 1	2.3223***	2.1230***
	(0.4289)	(0.5213)
Phase 2 (1st week)	2.5467***	2.4451***
	(0.8204)	(0.8971)
Phase 2 (2nd week)	2.7150***	2.5193***
	(0.4871)	(0.5507)
Phase 3 (1st week)	3.2540***	3.2517***
	(0.5584)	(0.5777)
Phase 3 (2nd week)	2.9898***	2.9898***
	(0.5655)	(0.5671)

This table displays the results of running the doubly-robust estimation procedure as described in Callaway and Sant'Anna (2021). This set of results shows the estimation by using the data according to the true beginning treatment time (i.e. because some phases happened in consecutive weeks). The "Not Yet Treated" column shows the results by using the control group for units that are not yet treated and the "Never Treated" column shows the results by using the control group for units that are never treated.

Table 1.6 shows the results of the estimation strategy. Note that because the unit of analysis is at the route-week level and that the rollout for both Phase 2 and Phase 3 were done in adjacent weeks, both of these phases have a "1st week" and a "2nd week" estimate. Additional results are provided in Table 1.7 by shifting all of the rollouts to the earlier week or the later week. The results from this table do not change conclusions in a substantive manner. The primary difference is that the treatment timing is slightly different, so the $ATT(g,t)$'s are either a little larger or a little smaller than the results from Table 1.6.

The "simple" ATT is 2.696 additional pounds of diverted material for the "not yet treated" specification and 2.497 additional pounds for the "never treated" group. The parameter for ATT "Group" additionally weights by length of time in the treatment; the results are 2.75 pounds of additional diverted material for the "not yet treated" group and 2.62 pounds for the "never treated" group. While both of these are larger estimates than the TWFE estimate, they are within one standard error and are very similar estimates. The similarity between the estimates suggests that proceeding with the TWFE estimate is not unreasonable.

The breakdown of the group ATTs reveal more about the rollouts. The first pilot has the highest group estimate of all at 4.89, however the second pilot has the lowest at 1.66. Given the small number of routes in each of these pilots, this is not terribly surprising. These two average out to approximately 3.27 pounds of additional organic material diverted, suggesting that they are still above average, but within the margin of error. The phases initially suggest that as phases were rolled out, each phase diverted more additional material on average. However, these are again all easily within the margin of error, making conclusions difficult to draw. The aggregate ATTs are both reasonably close to the TWFE estimates, suggesting that the potential bias in the TWFE estimator is not too large.

Due to the rollouts sometimes occurring in different weeks, even if they are in the same phase, and the structure of the data, some estimates are split into two different weeks. I again present three different sets of results. This table takes the consecutive weeks for Phases 2 and 3 and moves them "down" or "up" so they all occur in the same week.[29] This does not change results in any substantial manner. The reason for doing shifting the start date "up" or "down" is because there is only one "Phase 2", but the rollout happened in different (but consecutive) weeks. Obtaining an estimate for the entirety of Phase 2 and Phase 3 using this estimator is to provide additional information on the treatment group dynamics.

Table 1.7 has a similar format to Table 1.6, however it is split in two. The first set of results show moving the dates for Phase 2 and 3 "down", while the second set of results show moving the dates "up". The first estimate in the two sections shows the result of the "simple" aggregation. In the "down" results, the "not yet treated" estimate is 2.57, while the "never treated" estimate is 2.4. These are even closer to the TWFE estimates than those in Table 1.6. The ATT "group" estimates are very close to the "simple" estimates, with the "not yet treated" group estimate of 2.575 and the "never treated" group estimate of 2.471. The group dynamics are similar to those in Table 1.6 as well, with Pilot 1 being the highest estimate, with 4.88 and 4.81 for "not yet treated" and "never treated", respectively. Meanwhile, Pilot 2 has the lowest estimates for "not yet treated" and "never treated" at 1.66 and 1.193, respectively. The phase results suggest the Phase 1 and Phase 2 estimates are both very close to the main TWFE estimates. The "not yet treated" group has an estimate for Phase 1 of 2.32 and for Phase 2 of 2.27. The "never treated" group has estimates for Phase 1 and Phase 2 of 2.123 and 2.17, respectively. For Phase 3, both specifications have the same estimate of 2.915 pounds of additional organic material diverted. One interesting result from these specifications is that the estimate for Phase 3 is much larger than Phase 1 and Phase 2, but still within the margin of error. The particular reason for this is unclear. One particular proposal is that because Phase 3 was the most recent phase rolled out in the data, there is the least amount of data available.

For the "up" portion of Table 1.7, the ATT "simple" estimate for the "not yet treated" specification is 2.68 and the "never treated" specification is 2.48. The ATT "group" results are

[29] E.g. Phase 2 occurred in weeks 446 and 447 of the data. Moving it "down" moves the rollout in week 447 to 446. Moving it "up" moves the week 446 rollout routes to week 447. Phase 3 is the same for weeks 512 and 513 in the sample.

Table 1.7. Callaway & Sant'Anna Estimation: Simultaneous Rollouts

	Dependent variable:	
	Pounds of Organics Disposed Per Household	
	Not Yet Treated	Never Treated
DOWN		
ATT "Simple"	2.5743***	2.3987***
	(0.5283)	(0.5616)
ATT "Group"	2.5748***	2.4709***
	(0.3603)	(0.3732)
Pilot 1	4.8817***	4.8090***
	(1.1218)	(1.3844)
Pilot 2	1.658	1.1927
	(3.2671)	(3.3327)
Phase 1	2.319***	2.1230***
	(0.4140)	(0.5061)
Phase 2	2.2683***	2.1698***
	(0.5395)	(0.6535)
Phase 3	2.9149***	2.9149***
	(0.581)	(0.5360)
UP		
ATT "Simple"	2.6796***	2.4759***
	(0.5214)	(0.5599)
ATT "Group"	2.7247***	2.5934***
	(0.3294)	(0.3707)
Pilot 1	4.8834***	4.8090***
	(1.1486)	(1.2871)
Pilot 2	1.6606	1.1927
	(3.2336)	(3.1619)
Phase 1	2.3223***	2.1230***
	(0.3945)	(0.5210)
Phase 2	2.6313***	2.4345***
	(0.4)	(0.5657)
Phase 3	3.0547***	3.0547***
	(0.5543)	(0.5456)

This table displays the results of running the doubly-robust estimation procedure as described in Callaway and Sant'Anna (2021). The first set of results shows the result from shifting the start week of the Phase 2 and Phase 3 rollouts to the first week (of two) of the rollout (446 and 512, respectively). The second set of results show the results by shifting the start week for all Phase 2 and Phase 3 to the second of the rollout (447 and 513, respectively). The "Not Yet Treated" column shows the results by using the control group for units that are not yet treated and the "Never Treated" column shows the results by using the control group for units that are never treated.

2.73 and 2.6, a larger estimate than the "simple" estimate, but not much larger. The estimate for Pilot 1, Pilot 2, and Phase 1 are very similar to the "down" portion of the table. In fact the "never treated" estimates are identical, but the SEs are different. Given how group-time estimates are calculated, the "not yet treated" group estimates being different makes sense.[30] The estimate for Pilot 1 for the "not yet treated" specification is 4.88, practically the same as the "down" portion. Pilot 2 has "not yet treated" at 1.661, compared to the "down" portion of 1.66. Phase 1 for the "not yet treated" group is 2.322, while the "down" specification has 2.319. Now, for Phase 2 and 3 both "not yet treated" and "never treated" have different estimates than the "down" set of results. In both cases, the "up" estimates are larger. For Phase 2, the "not yet treated" estimate is 2.63, while the same estimate for the "down" group is 2.27. The "never treated" group estimate is 2.435, while the "down" group is 2.17. Finally, for Phase 3, both "not yet treated" and "never treated" have an estimate of 3.055, while in the "down" portion of the table the estimate is 2.915.

Overall, these results are simply an upper and a lower bound on the baseline estimates. The disparity between the two sets of results arises because the $ATT(g,t)$'s effectively estimate a canonical two-by-two difference-in-differences. Given the chosen week changes, we should expect variation in the baseline control group and observe that while this happens, these estimates are both readily within the error bars of each other. Therefore, these differences are not a cause for concern.

Chapter 1, in part is currently being prepared for submission for publication of the material. The dissertation author was the sole author of this chapter.

[30]The "never treated" group does not change, so the estimates should be the same. The "not yet treated" group also does not change, however the timing changes. Each ATT for a given group-time is therefore estimated with a slightly different group of control estimates.

2. Heterogeneity in Consumer Response to Water Quality Violations[31]

2.1 Introduction

Regulated community water systems (CWS) in the United States provide drinking water to approximately 94% of Americans. In 2017, 22 million people – 7.2% of all CWS customers – experienced a violation by their CWS provider in the United States. These violations can cause negative health effects, such as bacterial infections, lead poisoning, or "blue baby syndrome" from nitrates. Eight percent of CWSes serve populations greater than 10,000, yet provide water to 82% of the population served, while 27% of all CWSes serve populations ranging from 501 to 3,300 serve, yet provide water to about 7% of the population.[32] The dispersion in the number of customers results in heterogeneity in the populations served, both within a CWS and across CWSes. A water quality violation can affect a large number of households, requiring many households to cease consuming their tap water (if unfiltered) and finding alternative means to access clean drinking water. As such, understanding which groups respond to these violations is important in order to identify inequities regarding who responds and to consider which groups may not be responding to these violations and why. This paper investigates how different demographic groups respond to water quality violations.

[31]Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[32]Further information can be found in Tiemann (2017).

CWSes incur violations by not testing for contaminants or exceeding a certain predetermined threshold for a contaminant in their water supply.[33] CWSes are legally required to report any violations to their regulatory entity (the state or federal government). Water quality violations are placed in one of three tiers. Tier 3 is the lowest (and least severe) tier; the CWS must inform consumers of the violation within one year of its occurrence. As CWSes must send out a yearly Consumer Confidence Report (CCR) to each household served, these violations are typically contained therein.[34] Tier 2 is the middle tier, where consumers must be informed within one month of the violation. These are typically included in a household's monthly water bill. Finally, Tier 1 is the highest (and most severe) tier.[35] Consumers must be informed within 24 hours of the violation. The reports informing consumers must be conveyed in multiple ways, including mail, TV news, paper news, and if necessary, phone calls. In addition, if the CWS customers are linguistically diverse, they must be provided with resources to acquire the information in their preferred language. Any of these violations must be remediated according to the rules set forth by the regulatory body. The immediacy of reporting Tier 1 violations results in customers of the violating CWS being able to respond rapidly to their circumstances. However, this assumes that all customers see the violation notice and respond accordingly.

This paper examines the heterogeneity in responses at the household level to water quality violations. Our data gives us a unique perspective by focusing on household-level responses to water quality violations. We explore the demographic heterogeneity in response to these contaminants, by considering important factors in responses such as race, education, and income. Depending on the severity of the contamination, consumers may respond to these violations by substituting from tap water to other sources (such as bottled water) or by investing in a cleaning technology. Different demographic groups may respond differently depending on certain factors.

[33]The threshold depends on the contaminant. The thresholds are measured in either maximum contaminant level (MCL), maximum residual disinfectant level (MRDL), or treatment technique (TT).

[34]The CCR must be submitted every year by July 1.

[35]Examples of Tier 1 violations include MCL exceedance of E. coli, MCL nitrate violations, and chlorine dioxide MRDL exceedance, among others. Most Tier 2 violations are failure to report tests for "serious" contaminants or exceedances not considered Tier 1 violations.

For example, higher-income households can more readily afford to install a filtration system inside their home to avoid any concerns about violations, whereas lower-income households do not necessarily have the capital or the ability to install these systems. They must purchase water on a frequent basis to avoid drinking and using contaminated water. Similarly, there is evidence that certain groups – such as black households or lower-income households (Pierce and Gonzalez (2017)) – exhibit more mistrust in their tap water than other racial groups or higher-income households.

To identify the effect of a water quality violation on household purchases, we exploit the timing of violations occurring in CWSes by comparing households that experienced a violation to those that did not experience a violation. Our identifying assumption is that, conditional on covariates, households do not systematically have different water bottle purchasing behavior. Furthermore, households are not able to anticipate water quality violations in their systems. In cases where the violations are more common with a CWS, fixed effects across time and units alleviate these concerns.

Households react most strongly to *Bacteria* and *Lead & Copper Rule* violations. Our findings show that white and black households both respond to *Bacteria* violations, consuming approximately 7% and 8% more ounces of bottled water when a violation occurs, respectively. When a *Lead & Copper Rule* violation occurs, black households consume approximately 23% more ounces of bottled water. Meanwhile, when a household with income below the median faces a *Lead & Copper Rule* violation, they consume approximately 13% more ounces of bottled water for the duration of the violation. We find results that augment and differ from previous work, such as Shimshack et al. (2007), finding educated households (with at least some college education) respond to *Bacteria* violations, while we also find that less educated households respond to *Other* violations – which includes categories such as turbidity[36] – by increasing their consumption by 13% for the duration of a violation. Our results overall show an important

[36]This is a measure of how cloudy the water is. The concern arises from the correlation between higher turbidity water and higher levels of pathogens.

variety of responses, however we show that only certain groups respond. These results are a cause for concern for both households that are responding and those that are not. The households that are responding are purchasing a luxury good, yet we do not observe only higher-income households purchasing bottled water. As a consequence, these violations may put additional financial strain on households. The groups that do not respond are also a cause for concern and we are not able to identify why there is a lack of a response. If the households are simply immune to the violations – perhaps because of filtration – there is no cause for concern. However, if households are simply ignoring the violations or are unable to respond due to a lack of means, this is an important concern.

Consumers can respond to water quality violations in several ways. For the moment, assume that individuals are consuming tap water. They can ignore the violation and continue to drink tap water. This comes with the health risks associated with consuming that particular contaminant. Alternatively, they can substitute to drinking bottled water or other liquids, or subscribe to a water delivery service. This comes at a pecuniary cost. Lastly, they can install a water filtration system. This typically requires a higher fixed cost than bottled water and requires replacement of the filters, along with other possible inconveniences (such as being slow to refill), while also requiring them to have the financial means to do so. All of the potential responses for consumers are pecuniary, allowing us to directly see the costs incurred. Our data allows us to observe all of these margins to some degree.

The financial circumstances of a household may restrain their ability to adapt in the most effective manner possible. Lower-income households will not necessarily be able to make large purchases for products that will filter their water long-term. These households may not own their homes and so simply are not permitted to install filtration systems. In this regard, we may expect poorer households to purchase more water bottles to avoid contaminated water relative to higher-income households. However, higher-income households can more easily afford to purchase additional bottles of water, so it is possible that higher-income households have a stronger behavioral response than lower-income households.

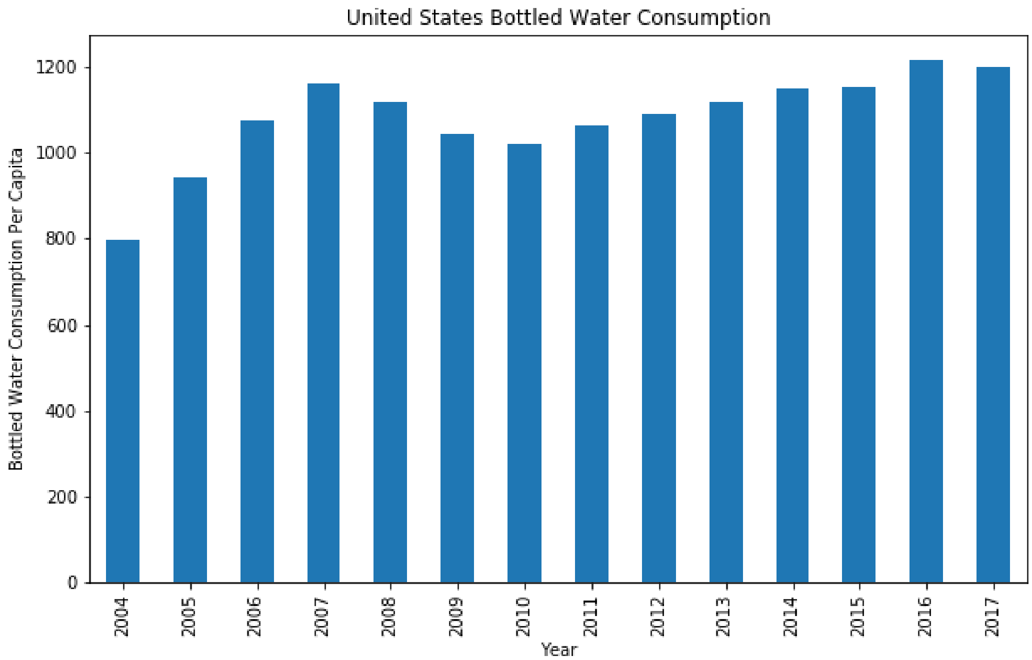

Figure 2.1. US Bottled Water Consumption

Note: This is plot shows ounces of bottled water consumption per capita per year in the USA.

An additional issue when observing household responses to these violations is that the household must see the violation notification sent by the CWS. If a household does not see the violation or ignores it, they may not respond as a result of not seeing the violation rather than a lack of means to respond. It is important to be able to distinguish the two, however with our data we are not able to do so. Disentangling the two is very important for improving the violation notification system and our results suggest a good starting place for exploring the system.

The effect of a water quality violation may have a lasting impact. Once a violation is remedied, a household may not revert to the same behavior as before the violation occurred. For a variety of reasons – such as trust in the tap water/water company or a newfound preference for bottled water – a household may consume a different amount of bottled water after the violation. As a result, some households may have a permanent response to a violation, causing us to observe different consumption levels across demographic groups.

Over time, bottled water consumption per capita has increased in the United States, as seen in Figure 2.1. This trend is present in most states as well. Due to the lack of data, mapping CWS service areas to zip codes is a challenge and so we only examine certain states where this data is available.[37] The states in our dataset represent 33% of the United States population and about 20% of all PWSes in the United States.

Figure 2.2 shows the number of violations per year by category in our dataset. Overall, there are between 10,000 and 20,000 violations per year, with a majority of them in Tier 3. Tier 1 violations are the least common violations. This suggests that most violations are not harmful to consumers and are primarily failures to report. However, even a few Tier 2 and Tier 1 violations can affect a large number of households for an extended period of time, causing acute negative health responses.

To the best of our knowledge, this is the first paper to use household-level data to estimate the response to water quality violations on a large-scale. While our analysis is similar to that of Allaire et al. (2018), Allaire et al. (2019), and Graff-Zivin et al. (2011), we are able to use household-level data to examine heterogeneity in the response to these violations. Our analysis at this level allows us to investigate precisely the kind of households that respond to water quality violations. With these heterogeneous responses, we are able to consider how households respond and relate this to tap water mistrust (e.g. Pierce and Gonzalez (2017)). Education, income, and race are important factors associated with mistrust in tap water and our results suggest that it is precisely lower-income households and black households that respond strongly to water quality violations. These are among the groups that have deeper mistrust of tap water. Relatedly, our analysis is useful for concerns about the reach of Consumer Confidence Reports and who responds (Bennear and Olmstead (2008)). Our analysis points to groups that might not be responding to these violations and it is important to investigate precisely why this is the case.

Our paper contributes to the increasingly important academic and social discussion on inequality. In particular, we focus on consumption inequality (Attanasio and Pistaferri (2016)).

[37]The eight states we have in our dataset are Arizona, Arkansas, California, Kansas, New Jersey, Oklahoma, Pennsylvania, and Texas.

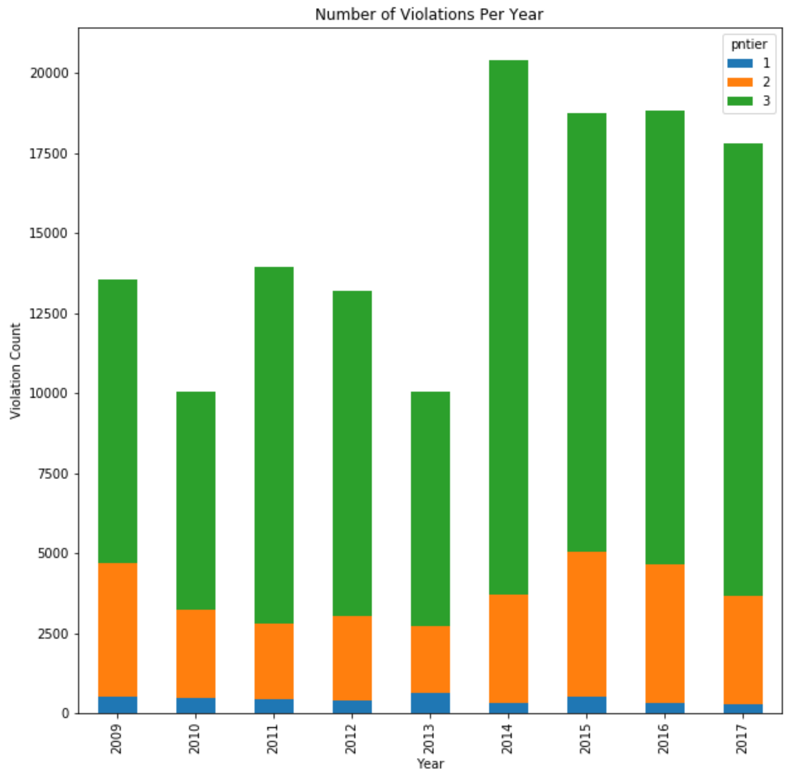

Figure 2.2. Water Quality Violations by Tier

Note: This figure shows the number of water quality violations each year by tier.

Our research suggests that there is a surprising heterogeneity in responses to water quality violations that are not necessarily in step with income, which one might expect given water bottles are a luxury good. On the other hand, for example, we find that it is primarily lower-income households who respond to Lead & Copper Rule violations. Additionally, we find that black households are the primary respondents to Lead & Copper Rule violations, while both black and white households respond to violations pertaining to excess bacteria levels. Our article also contributes to the literature on mortality inequality (Currie and Schwandt (2016)). Water quality violations affect health outcomes and our results suggest that some household demographic groups are aware of these adverse effects. These violations will not necessarily have acute health effects, yet the exposure necessarily imposes an additional cost to health, even if it is not guaranteed.

The paper now proceeds. Section 2 briefly discusses background information on the Safe Drinking Water Act. Section 3 describes the datasets used. Section 4 presents the analysis used. Section 5 provides the results of both analyses, a hurdle model and a regression. Section 6 discusses possible explanations for the findings and important questions for policy. Finally, Section 7 concludes.

2.2 Background

The United States federal government enacted the Safe Drinking Water Act (SDWA) on December 16, 1974. This required the EPA to set standards for drinking water quality for all public water systems (PWS)[38] and oversee implementation of the law across states and territories.[39][40] Over time, the SDWA has been amended, typically increasing its oversight and tightening restrictions. All states except Wyoming and "Indian Country" have primacy, meaning they are the enforcing bodies. States can and typically do have more stringent requirements beyond federal requirements. This means that many exceedances or reporting violations will be in higher tiers than the federal law stipulates, require testing of additional contaminants, or more frequent testing.

When a CWS conducts the required test on their water supply and the test exceeds the legal amount, they are obligated to report this to the state.[41] Once the state receives the notification of an exceedance, different actions are taken depending on the contaminant and the severity. In general, the state will provide a timeline and instructions for remediation to the CWS along with the required date for a public notification to be sent. Sometimes a fine is administered,

[38] A public water system has three subcategories: CWSes, non-transient non-community CWSes (e.g. schools), and transient non-community CWSes (e.g. campgrounds).

[39] This does not cover private wells.

[40] Bottled water is not covered by this act; that is under the purview of the Food and Drug Administration.

[41] It should be noted that precisely who conducts the test and who notifies the state of the violation is state dependent. For example, in some states a third-party (usually the laboratory conducting the tests) might retrieve the sample from the CWS, while in other states the CWS will send in the sample to the laboratory.

though this is infrequent and usually only if it's a repeat warning or violation.

An important review of the history of the SDWA is contained in Tiemann (2017). This paper summarizes the legal aspects and requirements behind the laws. The Governmental Accountability Organization in GAO (2011), conducted a review of the data provided by states to the Safe Drinking Water Information System (SDWIS), finding a large amount of underreporting, typically due to state department resource constraints, the lack of a streamlined data entry process, or due to inadequate guidance, but also occasionally due to "sympathy" for PWSes (i.e. if a small CWS incurs a violation not seen as a health hazard, let it violation do not report it to lessen the legal repercussions, in particular if the CWS is "trying"). A majority (91%) of these reporting errors are due to "compliance determination" errors, meaning a violation occurred but the primacy agency did not issue a violation to the CWS in violation and did not report this violation to the Federal EPA SDWIS.

2.3 Data

Our data on consumers comes from the Nielsen Homescan Consumer Panel dataset, covering the years 2009 to 2017. This dataset is gathered by The Nielsen Company, consisting of a "stratified, proportionate" sample of about 60,000 United States households. Each household reports all of their household purchases to Nielsen via a scanner. Various quality assurance measures are taken to motivate households to submit their purchases and to verify the accuracy of their purchases and the prices. Households purchases are characterized by "trips", in which there are typically multiple purchases at a store or stores. Each item purchased corresponds to a unique UPC code and relevant product information, along with the price and quantity purchased. UPC codes allow for detailed item-level identification of the product. In addition, the Consumer Panel dataset tracks household characteristics yearly. Most of the household characteristics are binned; for example income is in bins of at least five-thousand dollars and race is coarsely listed in five groups.

The Nielsen dataset does not allow us to fully observe the universe of purchases by consumers. If a consumer does not purchase from stores within the Nielsen data scope, then it is not contained in the dataset. Specifically, if a consumer procures water from a delivery service or from their refrigerator/freezer door, this is not observed. As mentioned, we observe bottled water purchases from a large number of stores, water filter purchases, and water filtration unit purchases. In this paper, we only present the results on bottled water purchases.

Our data source on CWS violations comes from the EPA SDWIS violations database. This database stores all violations by all PWSes in the United States. Quarterly, each state submits their SDWIS violations to the SDWIS database. EPA retains this data for at least five years. The data consists of unique PWS identifiers, the violation start date, population served, the violation end date, the contaminant code and type, and the violation code and type.

The SDWIS violations database is not without issues. Chief among them is underreporting. "For example, we estimate that the 14 states audited in 2009 did not report or inaccurately reported about 54,600 – or 84 percent – of the monitoring violations (typically Tier 3 violations) committed by community water systems to SDWIS/Fed" GAO (2011). Fortunately, more severe infractions experienced less underreporting. Meanwhile, the estimated percentage of unreported or inaccurately reported health-based violations was between 12% and 40%. Overall, this underreporting implies that some households are unknowingly exposed to (possibly harmful) contaminants. As reporting is biased one way, the total cost of water quality violations is going to be an underestimate. We drop observations with known data entry issues (e.g. very small CWSes) to try to minimize this bias.

Zip code level water company service areas are not readily available. For states that provide geographic information system (GIS) matched PWS zip code data, we use this to determine service areas. We match the polygons representing service areas to the Census-provided polygons representing ZIP Code Tabulation Areas (ZCTA).[42] In addition to the state-

[42] In general, a ZCTA is exactly the zip code. Most of the time the difference arises in non-residential areas, for example where a business/businesses has/have a very large factory or warehouse.

provided GIS data, we gathered data from the past two Unregulated Contaminant Monitoring Rule (UCMR) data collections.[43] EPA chooses 25 to 30 unregulated contaminants to be tested by PWSes in the United States. Testing occurs over a two year period. PWSes that serve more than 10,000 people must test for the contaminants, while PWSes below the cutoff are randomly selected as testers. This data contains all zip codes served by the PWSes participating in the particular round of the UCMR program.

2.4 Model

Our model estimates the treatment effect of a household experiencing a water quality violation. In particular, due to zeroes, we are not able to apply the more often used log-linear regression. Hence, our model takes the form:

$$y_{ht} = \alpha + \sum_{c \in C} \beta_c \times \sum_{w \in W_z} [\mathbf{1}\{c_{tw}\} \frac{p_{wz}}{p_z}] + \Theta X_{ht} + \Lambda P_{zt} + \delta_t + \phi_z + \varepsilon_{ht}$$

Here, y_{ht} is the outcome of interest – in most cases ounces of bottled water purchased – for household h in time period t. The parameters β_c measure the effects of having a particular contaminant group – outlined below – occurring within the region z (that household h lives in) at time t. Here, z will either be county or ZCTA. The indicator for a violation is multiplied by the ratio of p_{wz}, the estimated population served in zip code z by water company w, over p_z, the population of the zip code. This ratio represents the fraction of the population served by the violating water company in a zipcode. This provides us with the correct coefficient β because not all households in a zip code are served by the same water company. The matrix X_{ht} is a vector of covariates, such as income or household size, while P_{zt} are weather variables, including average temperature and rainfall. The time fixed-effects δ_t are year-by-month fixed-effects and the area fixed-effects ϕ_z.

The primary contaminant categories are *Chemicals and Elements*, *Nitrates*, *Bacteria*,

[43]"UCMR 4" took place from 2018-2020 and "UCMR 3" took place from 2013-2015.

LCR, and *Other*. We constructed these categories manually. Each contaminant is assigned a particular four-digit number. We created the groups based their first digits. For example, all bacterial codes begin with a "3", while all chemicals begin with a "2". In cases where there are too few observations, such as with radioactive contaminants, we placed them into the nearest category according to their characteristics.

The *Other* category consists of turbidity and violations of the Surface Water Treatment Rule. *LCR* is the Lead and Copper Rule,[44] which mandates an "action level" (rather than the standard maximum contaminant level (MCL) or maximum contaminant level goal (MCLG)) where a CWS must take action to remediate the high levels of lead or copper. The category of *Bacteria* consists of bacteria that cause illnesses, such as E Coli (one of the most common bacteria violations).

One issue is that for a large fraction of trips, households do not purchase any bottled water. Further, as a violation can last a long period of time, households make their purchases at different times. Both of these imply that a large fraction of the data consists of zeroes. As such, our analysis also runs a hurdle model for the first stage. The hurdle model works by assuming that the "zeroes" (those who do not purchase bottled water) and the non-zero values (those who purchase bottle water) are the result of different data-generating processes. The first-stage estimates these DGPs, then separately estimates a standard regression in the second stage. With this, directly interpreting coefficient estimates is incorrect and requires the computation of marginal effects. For more information, see Cameron and Trivedi (2005). Our primary results are for the hurdle model, yet we also include the results from the regression without the hurdle model as well.

We windsorize our sample by excluding the top 0.05% individual shopping trips including bottled water. This is to avoid outliers artificially increasing coefficient estimates. Further, we remove CWSes with a population served under 500. Smaller CWSes are subject to less stringent reporting requirements and have substantially more variability in reporting accuracy. This helps

[44]In December 2016, after the Flint, Michigan water crisis, this rule was updated in the Water Infrastructure Improvements for the Nation Act (WIIN). This added a public notification requirement for LCR action level violations, typically within 24 hours of the discovery of the violation.

alleviate the underreporting concerns mentioned earlier.

Further, the data is not always complete. In some instances, a violation start date is provided, but no violation end date is provided. In other cases, no contaminant type is provided, so we cannot identify the tier for the violation. In either instance where the missing data cannot be readily inferred, it is dropped. Finally, if a violation start date is more than four years from the compliance start date, it is dropped. This is because SDWIS systems can become backlogged, so the occurrence and the submission time can have large gaps.

2.5 Results

The results section is presented in two subsections. The first subsection shows the specification when the regression is run with a hurdle model, while the second subsection shows the results without the hurdle model. The results between the two methods are largely the same and many conclusions and implications follow.

Hurdle

The results presented are interpreted as the marginal effect $\frac{dy}{dx}$, the effect of a particular violation on weekly ounces of bottled water purchased. The tables presented include two different columns: one where the violation is not included in the first stage and one where it is included. The second stages for both are the same. These are included because the interpretation of the result is slightly different, however both interpretations are important. Both estimations in the first stage include covariates of race, an indicator for a child under age 6, education, household income, minimum and maximum temperature in the location, and precipitation in the location. In addition, we include county fixed-effects and year-month fixed-effects. However, only the second column includes the violation times the population ratio affected in the first stage. The second stage includes the violation treatment times the population ratio, education, income, and household size, along with year-month fixed-effects and county fixed-effects. We then estimate the marginal effects from this estimation.

Table 2.1. Hurdle with Violation Only

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Vio.	No Vio. in First Stage	With Vio. in First Stage
	(1)	(2)
Bacteria	1.905**	4.645**
	(0.807)	(1.882)
Chem/Ele	1.430*	0.796
	(0.865)	(1.423)
LCR	2.564	3.712
	(1.730)	(3.181)
Nitrate	-1.017	-3.41
	(2.006)	(2.938)
Other	1.083	1.666
	(0.976)	(1.446)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group. The first column does not include the violation in the first stage of the hurdle model; the second column does include the violation in the first stage of the hurdle model. Standard errors are clustered at the zip code level.

The interpretation of the two columns differ slightly and either can be appropriate depending on the circumstance. The first column – which does not include the violations in the first stage – can be interpreted as the effect of a violation on households that are predicted to be purchasers of bottled water. The second column, on the other hand, restricts the interpretation by additionally requiring that these households are responding to the water quality violations by purchasing bottled water (i.e. in the former, we are only requiring the hurdle to consider “not purchaser” versus ”purchaser”. In the latter, “purchaser” now becomes akin to “purchaser and responder to violations”). As a result, in many instances the magnitude of the second column results will be larger than the first column. For this section, the first column will be emphasized more than the second.

Table 2.1 presents the results of the hurdle model with no interactions on the treatment variable for a violation. The estimate for *Bacteria* is smaller than those from the regression at 1.9 ounces of bottled water per week. The estimates for *Chemicals/Elements* are 1.43

additional weekly bottles of water consumed. On average consumers increase their bottled water consumption by about 65 ounces over the average duration of a *Bacteria* violation and by approximately 56 ounces of bottled water over the length of a *Chemicals/Elements* violation. The second column shows that households that respond to these violations have a stronger response to Bacteria violations – 4.65 ounces of bottled water per week – or approximately 181 extra ounces of bottled water purchased per week.

The coefficient estimate on *Nitrate* is negative in both specifications. Nitrates are usually found in water contaminated by fertilizer runoff and the acute toxicity from nitrates is low. Young infants are most affected by this risk, but other populations are not particularly at risk. Thus, these coefficients seem reasonable, especially given they are not significant.

Table 2.2 shows the results from the hurdle model by interacting a violation category with a race category. In this, for ease of exposition, we have three groups: white households, black households, and households that identify as any other race. The reason for this choice is that including more race categories does not aid with the exposition of the results. The table shows that two groups react to violations: white households and black households; however the two groups respond to different violations. In the case of white households, they react to *Bacteria* violations in both specifications. We estimate in the first specification an additional 2.3 ounces of bottled water purchased and in the second specification an additional 4.93 ounces. The mean duration for a bacteria violation is approximately 239 days or about 34 weeks, suggesting that the average white household purchases about 90 extra ounces of bottled water per violation – approximately a 7% increase in consumption. At approximately \$0.08 per fluid ounce of water, this suggests white households spend an additional \$7.18 per bacteria violation.

Black households respond to *LCR* violations in both specifications. Additionally, in the first specification they respond to *Bacteria* and *Chem/Ele* violations. In response to *LCR* violations, which have a mean duration of 39 weeks, we find black households purchase 8.73 additional ounces of bottled water in the first specification and 9.19 additional ounces of bottled water in the second specification. This amounts to a 23% increase in consumption for the duration

Table 2.2. Hurdle with Violation and Race Interaction

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Race, Vio.	No Vio. in First Stage	With Vio. in First Stage
	(1)	(2)
White, Bacteria	2.299**	4.934**
	(0.929)	(1.996)
Black, Bacteria	3.274*	2.857
	(1.720)	(3.007)
Other, Bacteria	-3.687	3.91
	(3.755)	(8.158)
White, Chem/Ele	0.797	0.659
	(0.892)	(1.404)
Black, Chem/Ele	7.453***	5.41
	(1.870)	(4.454)
Other, Chem/Ele	0.874	-3.595
	(2.304)	(4.071)
White, LCR	1.422	3.461
	(2.521)	(4.351)
Black, LCR	8.728***	9.186**
	(1.760)	(4.492)
Other, LCR	-0.0694	-4.734
	(2.261)	(5.015)
White, Nitrate	-1.534	-3.882
	(2.185)	(3.567)
Black, Nitrate	3.961	5.634
	(5.538)	(17.560)
Other, Nitrate	1.385	-2.212
	(3.675)	(5.396)
White, Other	1.099	2.021
	(1.091)	(1.688)
Black, Other	1.321	1.173
	(1.554)	(3.081)
Other, Other	1.514	0.133
	(1.336)	(3.919)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group, interacted with a category for race. The first column does not include the violation in the first stage of the hurdle model; the second column does include the violation in the first stage of the hurdle model. Standard errors are clustered at the zip code level.

of the violation or approximately \$27.24 extra spent on bottled water. In the first specification, there is a somewhat less precise estimate that black households respond to *Bacteria* violations, with an estimate of 3.27 additional ounces of bottles of water purchased per week. This is approximately an 8% increase in bottled water consumption and a \$10.20 increase in bottled water purchases for the duration of the violation. Further, they respond to *Chem/Ele* violations with a point estimate of 7.45, with a length of 39 weeks, which suggests a 20% increase in bottled water consumption from these violations.

The disparity of these results is interesting and suggests that different households respond to different violations – a theme for the remainder of this section. Notably, if most households undertake permanent avoidance behaviors such as water filtration, then we should not observe different groups responding to only some violations but not all violations. Namely, the water filtration or permanent avoidance behavior should mean we do not observe any reaction. This table suggests that white and black households respond to different violations, but also do not respond to certain violations and that many have not undertaken permanent avoidance behavior.

In particular, it is revealing that these households respond to different violations and that most households do not respond at all. This table suggests that black households respond to most violations and do so rather strongly, which is an additional burden of the violations occurring. This relates to the general mistrust in tap water among the black community – who already have the highest baseline level of bottled water consumption of any race. However, white households appear to be sensitive to bacteria violations, but not other violations. This table also disentangles Table 2.1 by breaking out which groups are responding. It suggests that some groups respond more strongly than others. Considering these results in addition to the income results below and within the categories is a cause for concern.

Table 2.3 shows the results from interacting "high income" households – the upper 50% of the income distribution – with the contaminant groups and does the same for low income households. The responses to both *Bacteria* and *LCR* are both solely from households below the median income threshold. The first column of the table shows that the reaction by "low

Table 2.3. Hurdle with Violation and Income Interaction

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Income, Vio.	No Vio. in First Stage	With Vio. in First Stage
	(1)	(2)
High, Bacteria	1.72	3.775
	(1.317)	(2.433)
Low, Bacteria	2.161*	5.657**
	(1.202)	(2.230)
High, Chem/Ele	1.346	2.125
	(1.084)	(1.823)
Low, Chem/Ele	1.454	-0.344
	(0.970)	(1.868)
High, LCR	0.692	1.276
	(1.612)	(3.376)
Low, LCR	4.464**	6.224
	(2.113)	(3.858)
High, Nitrate	-1.537	-1.212
	(1.542)	(2.532)
Low, Nitrate	-0.517	-5.108
	(3.471)	(4.507)
High, Other	0.471	-0.58
	(1.152)	(2.273)
Low, Other	1.566	3.472
	(1.323)	(2.124)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group, interacted with a binary variable for being above or below the median income in our sample. The first column does not include the violation in the first stage of the hurdle model; the second column does include the violation in the first stage of the hurdle model. Standard errors are clustered at the zip code level.

income" households for *Bacteria* is 2.16 additional ounces of bottled water. In the case when the violation is included in the first stage, this estimate more than doubles to 5.57. Households below the median react to *LCR* violations by purchasing 4.46 additional ounces of bottled water per week. For both *Bacteria* and *LCR* violations, households below the median income respond more strongly than households above the median income. In particular, households below the median income increase their bottled water purchases by 13% or approximately $13.91 for the duration of an *LCR* violation. Meanwhile, accordingly, these households purchase 6% more ounces of bottled water per week according to the first specification and 16% more according to the second specification.

This table suggests that the household response depends on income and in particular, it is predominantly households below the median income that are responding to these violations. Given bottled water is a luxury good and these households typically have lower disposable incomes, this indicates that these violations put an additional financial strain on households who are not as well equipped to afford the avoidance behaviors that need to be undertaken. This is concerning for equity as it suggests the brunt of the violations falls on lower-income households, while higher-income households are not responding presumably because many have undertaken avoidance behaviors already.

Table 2.4 shows the results of the hurdle model with contaminants interacted with education. We have two categories for education: up to a high school degree, and some college or more. For the education variable, Nielsen includes both female and male heads' of household education. We constructed a variable that is the maximum of these two variables. The results show that *Bacteria* and *Other* violations generate avoidance behavior. In particular, households with education of at least some college respond to *Bacteria* violations at 1.91 additional ounces of bottled water purchased per week. When we include the violation variable in the hurdle model, at least some college has a stronger response, as expected, of 3.97 ounces. Meanwhile, households with up to a high school degree have a less precise response of 7.44 additional ounces. Respectively, these are a 13% and 20% increase in bottled water consumption. Further,

Table 2.4. Hurdle with Violation and Education Interaction

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Education, Vio.	No Vio. in First Stage	With Vio. in First Stage
	(1)	(2)
Up to HS Deg., Bacteria	1.906	7.436*
	(1.609)	(3.859)
Some College and Up, Bacteria	1.908**	3.965**
	(0.929)	(1.835)
Up to HS Deg., Chem/Ele	1.266	-2.373
	(1.163)	(2.851)
Some College and Up, Chem/Ele	1.46	1.574
	(0.999)	(1.663)
Up to HS Deg., LCR	1.635	-1.359
	(2.669)	(7.472)
Some College and Up, LCR	2.783	4.811
	(1.727)	(2.977)
Up to HS Deg., Nitrate	1.149	-7.186
	(5.287)	(8.649)
Some College and Up, Nitrate	-1.415	-2.674
	(1.947)	(2.674)
Up to HS Deg., Other	2.803**	5.236*
	(1.417)	(3.085)
Some College and Up, Other	0.592	0.749
	(1.090)	(1.710)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group, interacted with a binary variable for education status. The first column does not include the violation in the first stage of the hurdle model; the second column does include the violation in the first stage of the hurdle model. Standard errors are clustered at the zip code level.

households with at most a high school degree respond to *Other* violations. In the case of no violations in the hurdle model, the response is 2.80 additional ounces of bottled water, while when violations are included in the hurdle the response is a less precise but larger 5.24 additional ounces purchased.

The culmination of these results suggests surprising heterogeneity in household responses to these violations. An important conclusion from these results is that different types of households choose to respond to different violations. For example, it is predominantly black households that respond to *LCR* violations, while both black and white households respond to *Bacteria* violations. Lower-income households exhibit a strong response to *LCR* violations as well. These results suggest that household racial makeup is an important determinant for avoidance behavior as well as highlighting which violations are responded to. Furthermore, these results show that lower-income households also undertake avoidance behavior; however they are less able to afford the burden of additional avoidance costs from the violations relative to higher-income households. These differential responses are important because some households may not be responding to these violations or are selecting which violations to respond to, and it implies that an additional financial burden is placed upon households that respond to the violations.

Regression

Our results begin with a non-interacted specification at the weekly level. Here we use indicators for the occurrence of a violation as the treatment multiplied by the population ratio as described in the previous section. The results can be found in in Table 2.5. We run the regression with two specifications. The first specification has unit fixed-effects at the county level with standard errors clustered at the county level. The second specification has fixed-effects at the ZCTA level and is clustered at the ZCTA level.

The table shows that individuals react most strongly to *Bacteria* violations. Both specifications show that individuals have a statistically significant reaction to the *Bacteria* violations.

Table 2.5. Effect of Violation on Bottled Water Purchases

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Violation	County	ZCTA
	(1)	(2)
1(Bacteria)	4.940**	5.428***
	(2.273)	(1.579)
1(Chem/Element)	-0.173	-1.279
	(1.359)	(0.933)
1(LCR)	1.526	2.012
	(3.408)	(2.026)
1(Nitrate)	-3.956	-2.584
	(3.033)	(2.484)
1(Other)	-0.0832	0.247
	(1.533)	(1.276)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group. The first column has fixed effects at the county level and the standard errors are clustered at the county level; the second column has fixed effects at the ZCTA level and the standard errors are clustered at the ZCTA level.

When county FEs are used, the magnitude is 4.94 additional ounces of bottled water purchased per week, while when ZCTA FEs are used the reaction is 5.43 additional ounces of bottled water purchased per week. For the specification at the ZCTA level, the mean duration for a *Bacteria* violation is approximately 34 weeks, suggesting that the average consumer purchases about 185 extra ounces of bottled water for the duration of the violation. At approximately \$0.08 per fluid ounce of water, this suggests consumers spend an additional \$14.80 per *Bacteria* violation. As a comparison, the mean duration for an *LCR* violation is about 39 weeks, suggesting that the average consumer purchases about 78 extra ounces of bottled water.

These results broadly suggest that consumers react to these water quality violations. Our investigation now turns to which groups of individuals react. As discussed in a previous section, it is not obvious which group will react, as there are plausible arguments either way.

Table 2.6 displays the results of interacting the treatment with being above or below the median level of income. The results show that households above and below the median income both respond to *Bacteria* violations, and the magnitudes are nearly equal. This is encouraging, suggesting that there is not a discrepancy in reaction to these violation across income levels. However, given bottled water is a luxury good, it is surprising to see a behavioral response of the same magnitude from lower-income households compared to higher-income households. The lower-income households will be spending a larger portion of their income on avoidance behavior, which is concerning. This result is largely consistent with the tables in the hurdle section.

Table 2.6. Regression with Interaction of Income Level with Treatment

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Income, Violation	County	ZCTA
	(1)	(2)
High, Bacteria	4.011	5.254**
	(3.029)	(2.440)
Low, Bacteria	5.940**	5.657**
	(2.607)	(2.541)
High, Chem/Ele	1.298	-1.165
	(1.948)	(1.883)
Low, Chem/Ele	-1.330	-1.381
	(1.700)	(1.411)
High, LCR	-0.417	0.596
	(3.760)	(2.807)
Low, LCR	3.334	3.337
	(4.135)	(3.153)
High, Nitrate	-3.127	-0.0679
	(2.860)	(3.321)
Low, Nitrate	-4.503	-4.371
	(4.040)	(3.387)
High, Other	-2.812	-1.571
	(2.621)	(2.495)
Low, Other	1.977	1.607
	(2.244)	(1.983)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group, interacted with a binary variable for being above the median income level. The first column has fixed effects at the county level and the standard errors are clustered at the county level; the second column has fixed effects at the ZCTA level and the standard errors are clustered at the ZCTA level.

Next, Table 2.7 shows the interaction of the treatment variable with three coarse race categories: white, black, or "other". The category of other includes all households whose head of household does not identify as white or black.[45] The results show that white households have a statistically significant response at approximately 6 additional ounces of bottled water per week to *Bacteria* violation, almost identical to both tables described above. This implies an additional 204 ounces of bottled water purchased for the duration of the violation, or almost a 19% increase in consumption. These results suggest that white households are the primary group responding, while non-white households either do not respond at all or do not respond nearly as strongly. Assuming that non-white households are not engaging in some type of permanent avoidance behavior, this implies that the notifications that are dispersed when a violation occurs may only be reaching white households, or white households are the only ones with the capacity or willingness to respond.

Table 2.8 displays the results of interacting the treatment with education, where education is either "up to a high school degree" or "some college and above". The results show both groups responding to bacteria violations. However, there households with an education level of "up to a high school degree" respond to *Bacteria* violation with a somewhat imprecise 9.86 ounces of additional bottles of water per week – 335 ounces of water for the duration of the violation – while "some college and up" has a response of 4.40 – nearly 150 extra ounces of bottled water. Both groups respond, however the less educated households have the much stronger response than the more educated households.

2.6 Discussion

There are three possible behavioral explanations from particular households not responding: Either the households have already permanently altered their behavior, they are simply not concerned about the violations, or they do not receive the violation notifications. Investigating how various demographic groups are responding (if at all) is very important, as it suggests that

[45] The coarse categories are chosen because it exemplifies the same results when the racial groups are finer, however this is less cluttered.

Table 2.7. Interaction of Race with Treatment

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Race, Violation	County	ZCTA
	(1)	(2)
White, Bacteria	5.350**	6.011***
	(2.404)	(1.875)
Black, Bacteria	1.814	1.534
	(3.670)	(5.492)
Other, Bacteria	4.278	4.173
	(9.253)	(6.230)
White, Chem/Ele	0.140	-0.685
	(1.368)	(1.019)
Black, Chem/Ele	2.790	-0.444
	(5.994)	(4.734)
Other, Chem/Ele	-6.214	-8.204**
	(3.936)	(4.066)
White, LCR	1.755	1.843
	(4.393)	(2.319)
Black, LCR	7.301	7.958
	(6.312)	(6.633)
Other, LCR	-7.925	-5.148
	(5.282)	(6.421)
White, Nitrate	-4.062	-2.107
	(3.485)	(2.584)
Black, Nitrate	2.959	-10.75
	(23.11)	(21.47)
Other, Nitrate	-4.237	-3.465
	(5.741)	(6.140)
White, Other	0.425	0.450
	(1.834)	(1.520)
Black, Other	-0.969	1.227
	(3.978)	(4.643)
Other, Other	-2.781	-2.902
	(3.774)	(4.460)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group, interacted with a binary variable for a household head identifying as white or not. The first column has fixed effects at the county level and the standard errors are clustered at the county level; the second column has fixed effects at the ZCTA level and the standard errors are clustered at the ZCTA level.

Table 2.8. Regression with Interaction of Education and Treatment

	Dependent variable:	
	Ounces of Bottled Water Purchased Per Week	
Education, Violation	County	ZCTA
	(1)	(2)
Some College and Up, Bacteria	3.906*	4.397**
	(2.171)	(1.729)
Up to HS Deg., Bacteria	9.368*	9.855**
	(4.977)	(4.717)
Some College and Up,, Chem/Ele	0.513	-0.840
	(1.571)	(1.056)
Up to HS Deg., Chem/Ele	-3.027	-3.221
	(2.990)	(2.831)
Some College and Up, LCR	2.828	2.959
	(3.234)	(2.334)
Up to HS Deg., LCR	-4.645	-2.544
	(8.034)	(4.967)
Some College and Up, Nitrate	-3.346	-2.571
	(2.774)	(2.727)
Up to HS Deg., Nitrate	-6.975	-2.523
	(8.273)	(8.463)
Some College and Up, Other	-1.272	-1.022
	(1.629)	(1.543)
Up to HS Deg., Other	4.866	5.553
	(3.836)	(3.619)
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The dependent variable is ounces of bottled water purchased per week per household. The primary variable of interest is a set of binary variables indicating if a household has experienced a violation within a contaminant group, interacted with a binary variable for education status. The first column has fixed effects at the county level and the standard errors are clustered at the county level; the second column has fixed effects at the ZCTA level and the standard errors are clustered at the ZCTA level.

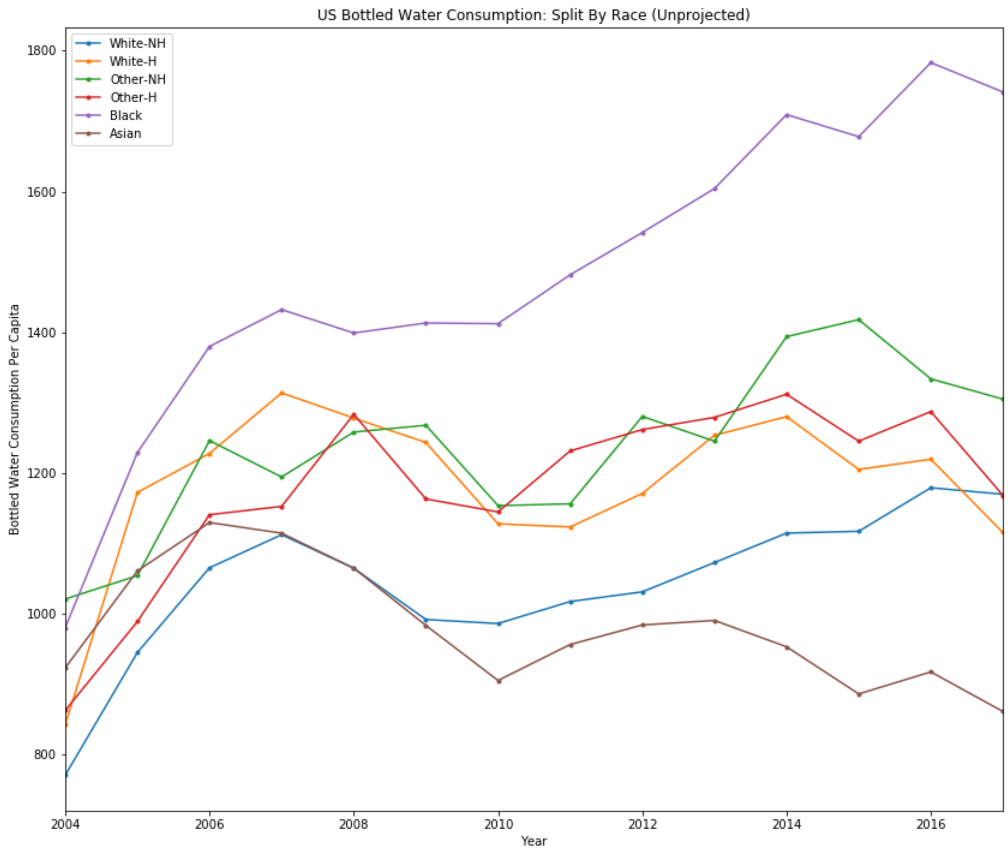

Figure 2.3. Ounces of Bottled Water Consumption in the US, by Racial Category

Note: This plot shows ounces of bottled water consumed per capita in the United States, broken down by racial category. "NH" stands for "non-hispanic" and "H" stands for "hispanic".

certain groups may be financially burdened by their permanent avoidance behavior (this includes permanent higher baseline bottled water purchases) or that some groups could not be switching at all. In the latter case, these groups are ingesting tap water that is considered unhealthy or dangerous and this will yield negative health effects.

It is possible to look for signs of permanent avoidance by particular groups. One way is to look for higher levels of consumption in the groups who do not respond before the violation occurs. Figure 2.3 shows the breakdown in ounces of bottled water consumption by race. The

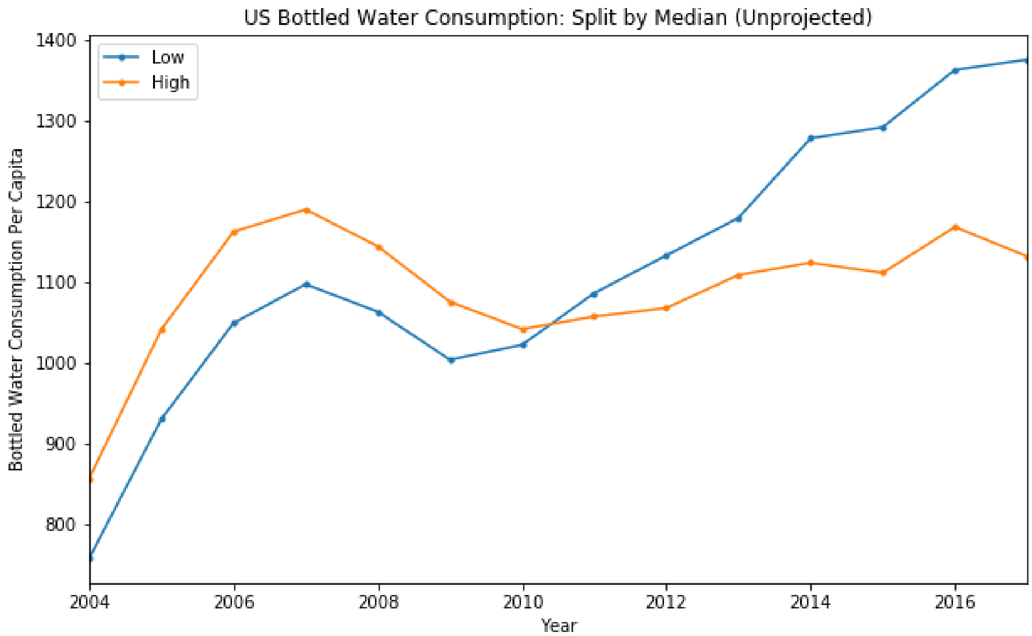

Figure 2.4. Ounces of Bottled Water Consumption in the US, by Income
Note: This plot shows ounces of bottled water consumed per capita in the United States, split by being above the median income and below the median income.

graph shows that black households consume the most bottled water, which is important because they also respond to violations by purchasing even more bottled water. Combined with the results from Table 2.2, this describes black households not only having higher baseline levels of bottled water consumption but also stronger reactions to these violations in magnitude. The remaining groups consume much less bottled water. However, even with white households consuming approximately 400 to 600 ounces less of bottled water per capita, our results show they respond to *Bacteria* violations at a similar rate compared to black households.

Related to this, Figure 2.4 shows that for the early part of the sample, households above the median income consume more ounces of bottled water than households below the median. But, beginning in 2011 this trend shifts and households below the median income begin consuming substantially more bottled water than households above the median income. Lower-income households began to have higher baseline consumption than higher-income households, implying either some shift to more permanent avoidance behavior or a mistrust in their tap water. Whatever the reason, this imposes an additional financial burden upon them. In particular, with this higher baseline, lower-income households respond to *LCR* violations quite strongly, implying that these violations exercise an even heavier burden.

These two graphs show that certain types of households consume more bottled water at baseline levels. However, these households also have strong responses to water quality violations, suggesting that they are at least aware of the violations. Assessing this higher baseline as permanent behavioral response to violations or as a larger trend in consumption is not possible in this analysis. Yet, it is significant because water bottles are a luxury good and households below the median income now purchase more bottled water than households above the median. This could indicate that households below the median income do not have the means to engage in permanent avoidance behavior.

2.7 Conclusion

The goal of this paper was to investigate the heterogeneity in household responses to bottled water consumption. After presenting baseline results, we presented our main results. We used a hurdle model due to excess zeroes and estimated the responses. Broadly, white and black households both respond to *Bacteria* violations, while both black households and households below the median income respond to *LCR* violations. A striking result is that black households increase their bottled water consumption by 23% when faced with a *LCR* violation. Further, we replicate the more educated households respond to these violations, but only find them responding to *Bacteria* violations. We also ran our analyses using only OLS sans the hurdle model, however our conclusions are largely unchanged.

Our results suggest that household characteristics are an important consideration in response to water quality violations. Further, while we are unable to identify why some groups do not respond, it is an important concern for policy makers because it might be that these households are accessing their drinking water from other sources or that they are simply ignoring the violations. The households that we do observe responding are facing an additional cost by avoiding the contaminated water which is also problematic. Not only do certain groups express skepticism in the cleanliness of their tap water, but some they face an additional financial burden and are not necessarily readily able to afford the burden of avoidance. Remedying these issues is important for an equitable water delivery system, as is identifying what is needed to allow households the opportunity to respond to any water quality violations.

Chapter 2, in part is currently being prepared for submission for publication of the material. The dissertation author was the sole author of this chapter.

3. Gasoline Demand Response to Taxes Versus Embedded Cap and Trade Prices

3.1 Introduction

In 2019, the United States Environmental Protection Agency estimated that the US transportation sector is the largest emitter of greenhouse gas emissions, making up twenty-nine percent of all GHG emissions (EPA (2021a)), while thirty-six percent of end-use carbon dioxide emissions result from transportation. Gasoline taxes are a ubiquitous policy tool that reduce vehicle miles, thereby lowering pollution.[46] There is evidence suggesting gas taxes elicit strong reactions from consumers relative to gasoline input price changes. Carbon cap-and-trade (CCT) and low-carbon fuel standards (LCFS) are alternative pollution control policies many countries have implemented. How consumers will react to the price signal sent by CCT permits or LCFS credits compared with taxes is not immediately obvious and is investigated in this paper.

This paper examines the consumer response of gasoline consumption to the cap-and-trade and low-carbon fuel standards[47] policies implemented in the state of California, comparing this reaction to the reaction from gasoline tax changes and tax-exclusive gas price[48] changes. These permits/credits are both an upstream cost and a tax, and recent literature[49] shows consumers have different reactions to these two costs. CCT permits[50] are an upstream cost bourne (and

[46]They also generate significant revenues – the US federal gasoline tax is 18.4 cents per gallon, which in 2017 accrued \$26.6 billion dollars in revenue FHWA (2017). California's gasoline excise tax is 41.7 cents, which generated \$6.1 billion in revenue in 2017.

[47]A LCFS credit program operates very similarly to a CCT permit system; this is explained in the appendix.

[48]For example, oil price changes, labor costs, etc.

[49]See Li et al. (2014), Rivers and Schaufele (2015), Tiezzi and Verde (2016).

[50]For conciseness, I will refer to CCT permits and LCFS credits solely as CCT permits. When the distinction is

likely passed down) by firms; the literature suggests consumers don't react to tax-exclusive (i.e. upstream) price changes nearly as much as taxes, so it is plausible that consumers treat CCT permits similarly to other upstream price changes. On the other hand, these CCT permits are implicitly a tax on gasoline levied by the government, so a consumer might react to them similarly to a gasoline tax. My results show that consumers react to CCT permit price changes very similarly to tax changes; both reactions are significantly stronger than tax-exclusive price changes.

Neoclassical economics suggests that, when considering price, consumers are only concerned with the final price. Indeed, the consumer only sees the posted price (which includes all taxes and costs) when making a gasoline purchase, so gasoline taxes, CCT permit, and LCFS credit prices all impact the consumer solely through the final posted price. In general, the long-run gasoline price elasticity is around -0.8 (Sterner (2012)). Both Levin et al. (2017) and Knittel and Tanaka (2019) use consumer microdata and find consumers have larger short-run elasticities, very similar to those suggested in Sterner (2012) of between -0.2 and -0.3. A pair of papers Davis and Kilian (2011) and Coglianese et al. (2017) suggest that consumers anticipate tax increases and purchase gasoline just prior to the tax increase.

One mechanism put forth to explain the different reaction between gas taxes and other gasoline input costs is salience. Gasoline tax changes usually generate significant media attention; thus making gasoline taxes "salient" to the consumer, inducing a strong reaction. Applying this mechanism to consumers' reaction to CCT permit/LCFS credit prices seems reasonable. In particular, given permit/credit prices change more frequently than taxes, on average they might be more salient than gasoline tax changes. This should cause consumers to have relatively high CCT permit/LCFS credit price elasticities of demand.

The salience explanation is undermined by the data presented in Figure 3.1, displaying article counts discussing gasoline price impacts from gasoline taxes, the CCT program, and the LCFS program within California. As can be seen, there is much more frequent discussion of

necessary, it will be made clear.

gasoline taxes in the news than the CCT program and the LCFS program. This pokes a hole in the above salience explanation, as it suggests consumers might not be as aware of (the impact of) these greenhouse gas reducing programs.

The next section provides a brief background on the policies and the data used. The section after presents the model and the results, followed by a section discussing salience as a mechanism, with the conclusion being last.

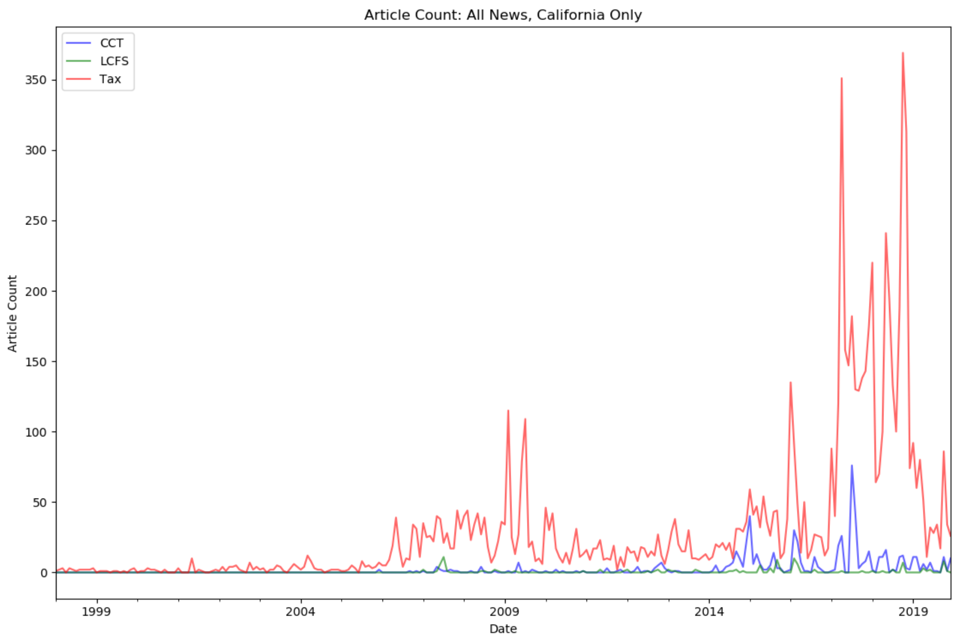

Figure 3.1. Article Counts: All News Sources, California Only

Note: This plot shows article counts for articles discussing "gas taxes", "cap-and-trade", or "low-carbon fuel standard" pertaining to California.

3.2 Policy and Data

In September of 2006, California passed AB-32, mandating that the California Air Resources Board (CARB) develop and implement a plan for greenhouse gas emissions. On December 17th, 2010, CARB implemented a cap-and-trade program. In 2012, the California Cap-and-Trade program began. In this program, a firm is required to have one permit for every metric ton of CO_2e emissions emitted for the emissions covered under the program. Some of these permits are allocated for free, with the remainder being purchasable at a quarterly auction. The price floor for permits increases every year by 5% plus inflation. The low-carbon fuel standard was enacted by Executive Order S-1-07 by then-Governor Schwarzenegger on January 19, 2007. The LCFS operates by lowering the allowed carbon intensity (CI) of the fuel blend produced by oil refineries and distributors over time.

Figure 3.2 displays the per-gallon cost to consumers from cap-and-trade and the low-carbon fuel standard in California. Before 2012, the programs were not yet phased in. The LCFS CI limit was flat for the first years and not particularly burdensome. Recently, the cost of LCFS compliance has greatly increased. In 2015, California's CCT program expanded to include transportation fuels and the LCFS program began to implement more stringent compliance CI levels each year. By 2020, the mandate is a 10% reduction in the CI of fuels.

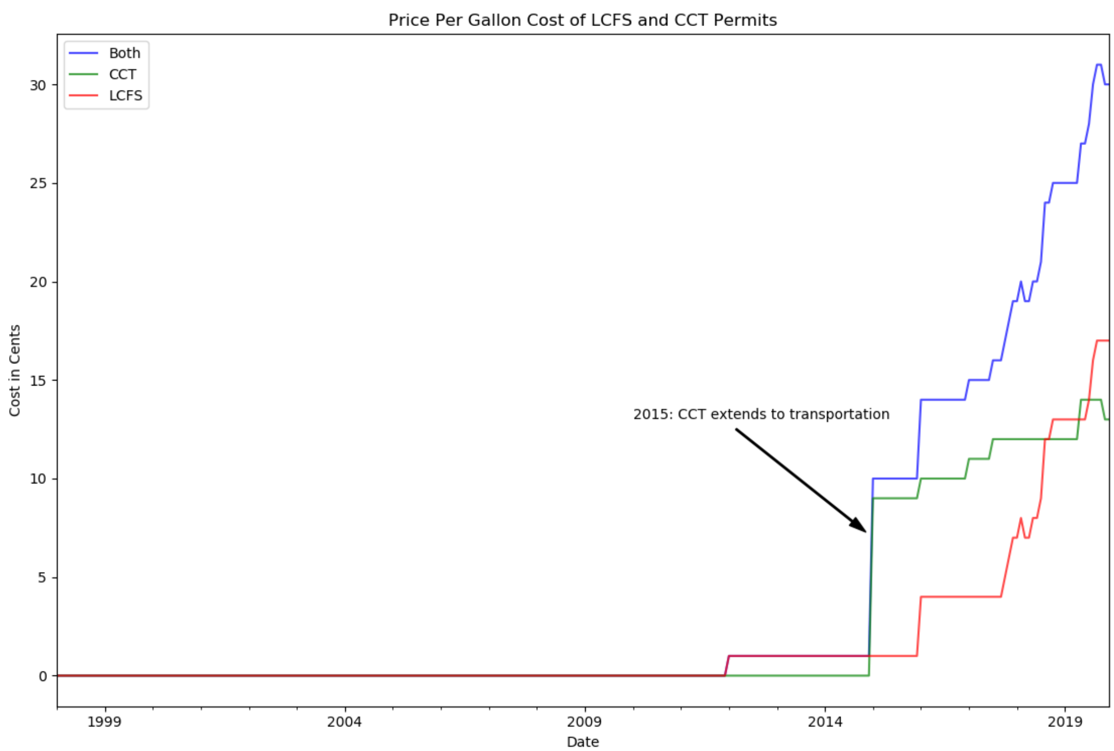

Figure 3.2. LCFS credit prices and CCT permit prices per-gallon cost

Note: This plot shows the at-the-pump prices of LCFS credit prices and CCT permit prices.

Every state in the US has an excise tax on gasoline, but most states do not have a sales tax on motor fuel. Figure 3.3 shows the highest, lowest, and mean state gasoline taxes, as well as the (unchanging) federal gasoline tax. The mean state gasoline tax rises slowly over time, while some state taxes rise much more quickly than others. In 1998, the difference between the highest and lowest taxes was 9.9 cents, while in 2019 the disparity was 29.4 cents.

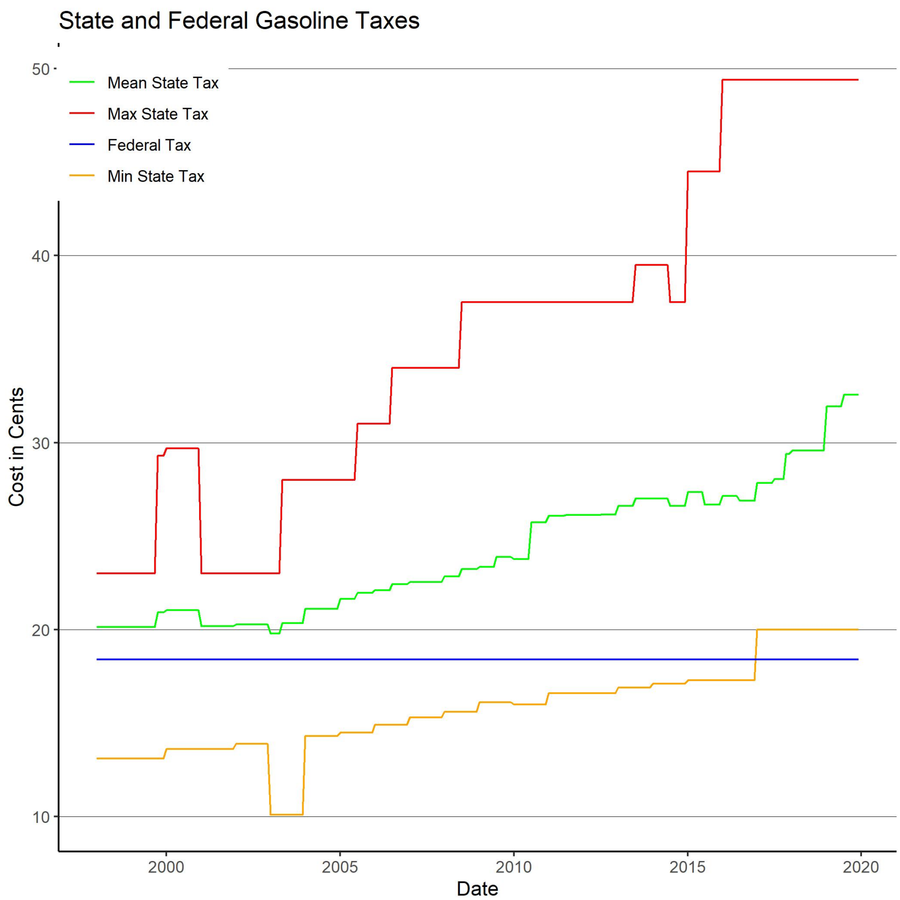

Figure 3.3. State and Federal Gas Taxes Over Time

Note: This plot shows the median, maximum, and minimum state gasoline taxes and the federal gasoline tax for the duration of the sample.

I use a monthly panel data set for nine US states[51] consisting of gasoline consumption, gasoline prices, gasoline taxes, carbon cap-and-trade permit prices, and low-carbon fuel standard credit prices. The data begins January 1998 and ends December 2019. The gasoline consumption and gasoline tax data come from the Highway Statistics published by the Federal Highway

[51]The nine states are California, Colorado, Florida, Massachusetts, Minnesota, New York, Ohio, Texas, and Washington.

Administration. Tax-inclusive gasoline prices come from the Energy Information Administration. Carbon cap-and-trade permit prices and LCFS prices both come from the California Air Resources Board. CARB provides conversion tables to obtain an average cost per gallon of gasoline for the relevant permits. I collected demographic variables of average family size, income, population, and proportion of population living in metropolitan statistical areas from the Current Population Survey. Further demographic variables include miles of public roads, number of cars and trucks, and number of licensed drivers are collected from the Federal Highway Administration Highway Statistics Series.

3.3 Model and Results

Recall that the goal is to estimate the consumer response to input price changes, tax changes, and CCT permit price changes. Due to the endogeneity problem inherent in sales taxes, they are left out of the analysis. Thus, taxes are taken only as excise taxes, both state and federal. This decision is standard in the literature.

The econometric model is:

$$\log q_{ts} = \alpha + \beta_p p_{ts} + \beta_\tau \tau_{ts} + \beta_c c_{ts} + \Theta X_{ts} + \gamma_t + \varepsilon_{ts} \tag{3.2}$$

Let t denote the month-year in the sample and s denote the state. The variables are: q is quantity of gasoline consumed, p is the tax-exclusive price of gasoline, τ is the gasoline tax, c is the adding of the per-gallon cost of CCT and LCFS permit prices, X is a vector of controls, and γ is a generic placeholder for various time fixed-effects which will vary by specification.

This model allows for direct interpretation of the coefficient estimates as semi-elasticities, i.e. $\dfrac{\partial \log q}{\partial i} = \beta_i$ for $i \in \{p, \tau, c\}$. One may wonder whether elasticities or semi-elasticities are the proper object to consider for reactions. Elasticities seem improper here because we are concerned with an individual's reaction to specific price components which impact the total price. For example, an elasticity specific to a tax implicitly assumes that individuals react differently to taxes than other costs. Further, the concern here is about how a direct cost increase impacts consumption, not relative percent increases (e.g. \$0.05 increase in taxes versus \$0.05 increase from input prices). While the difference in reactions appears to be empirically founded, one reason elasticities are preferred is because they are unitless. The units here are the same, so this adjustment is not necessary. Therefore, semi-elasticities are the preferred measure.

An implicit assumption in the construction of c is that permit prices are passed through one-for-one. While this itself has not been empirically verified, it is a logical assumption given taxes have been shown to pass through one-for-one (Marion and Muehlegger (2011)). Further, recall c is comprised of both the per-gallon cost of CCT permits and LCFS credits.

The regressions presented in Table 3.1 have two primary specifications (each specification has 3 columns). The first specification uses state fixed-effects and month-of-sample fixed-effects. The second specification uses state, year, state-year, state-year trend, and state-month fixed-effects. As these two specifications include different controls, the sources of variation are different. The first specification uses time-varying monthly deviations from the average of the month sample and state average. The second relies on deviations from the state-average-month effect and state-year average & linear trend. The second specification contains the preferred estimates. The coefficients in the first specification are generally less precisely estimated, but not different from the second specification in an impactful way. Hypothesis tests for equality between coefficients within the same model can be found in Table 3.1 under the regression coefficients results.

The first columns of the two specifications have total price as the only variable of interest. The estimated coefficient is -0.00031 for column (1) and -0.00016 for column (4). For comparison, the implied elasticity of these two columns is -0.05 and -0.03, respectively. These values are within the standard range of estimates. The middle columns separate tax-exclusive prices and tax variables. Overall, columns (2) and (5) show that taxes generate much strong reactions than tax-exclusive gas price changes.

The final columns of the two specifications – which answer the question this paper pursues – separately estimate tax-and-CCT-exclusive [all-tax] price changes, tax changes, and CCT (and LCFS) price changes. For these variables, column (3) has values 0.00012, -0.00357, and -0.00141 respectively. The difference between tax changes and all-tax-exclusive gas price changes remains large. CCT price changes are of a similar magnitude to the tax change variable and a difference between the two is not statistically distinguishable, yet this is also the case for CCT prices and all-tax exclusive gas price changes. Column (6) has similar results with different (and significant) coefficients: the all-tax exclusive price coefficient is -0.00014, the tax coefficient is -0.00231, and the CCT coefficient is -0.00185. This time, the taxes and the CCT coefficient are statistically different from the all-tax-exclusive price change, but not statistically different from each other. This suggests that individuals do not react identically with all-tax-exclusive price changes and taxes or CCT prices. It cannot be determined whether taxes and CCT prices elicit the same reaction, so the null hypothesis of no difference remains intact.

Both estimated specifications show a large differential response by individuals to all-tax exclusive price changes compared with tax and CCT price changes. At the bottom of Table 3.1, the percent change in quantity for a 5 cent increase in the relevant variable is presented. A casual summary suggests that an all-tax-exclusive price increase affects consumption of gasoline by a little under 0.1%, while both a 5 cent tax increase and a 5 cent CCT price increase affect gasoline consumption by about 1%. In 2019, passenger vehicle transportation emissions in California totaled a little over 119 MMTCO2e, implying that a 5 cent increase in the CCT price would lead to about a 1.1 MMTCO2e reduction in transportation emissions, or approximately a 0.3% decrease in total California 2019 CO2e emissions.

3.4 Discussion of Salience As A Mechanism

The primary piece of evidence for salience uses polling conducted by the Public Policy Institute of California from their "Statewide Survey: Californians & the Environment", which asks questions about environmental policy knowledge in California. This survey is conducted

Table 3.1. Gas Prices, Taxes, and Cap and Trade Results

	Dependent variable:					
	Log Quantity of Gasoline					
	Month-of-year, State			State, State-Year, State-Month, Year		
	(1)	(2)	(3)	(4)	(5)	(6)
Gas Prices	-0.00031^{***}	0.00007	0.00012	-0.00016^{***}	-0.00015^{***}	-0.00014^{***}
	(0.00010)	(0.00016)	(0.00014)	(0.00004)	(0.00004)	(0.00004)
Taxes		-0.00355^{***}	-0.00357^{***}		-0.00215^{**}	-0.00231^{***}
		(0.00092)	(0.00093)		(0.00085)	(0.00082)
CCT			-0.00141			-0.00185^{***}
			(0.00101)			(0.00045)
	Hypothesis Tests for Equality					
$p = \tau$		0.0005	0.0002		0.0227	0.01
$p = c$			0.1403			<0.0001
$\tau = c$			0.2013			0.6464
	Percent Change in Quantity for a 5 cent change in...					
Input Prices	-0.15344	0.03503	0.06023	-0.07841	-0.07359	-0.06909
	(0.05229)	(0.07879)	(0.07065)	(0.01983)	(0.0221)	(0.02118)
Taxes		-1.77727	-1.78689		-1.0748	-1.15505
		(0.4601)	(0.46603)		(0.42713)	(0.41189)
CCT			-0.70299			-0.92573
			(0.50634)			(0.22306)

*p<0.1; **p<0.05; ***p<0.01

Note: Data series from January 1998 to December 2019. The first three columns correspond to month-of-year fixed-effects and state fixed-effects. The second three columns are run with state, state-year linear trend, state-month, and year fixed-effects. Both regressions include state-level covariates. They are percentage of population living in an MSA, family size, and the log of: automobiles per capita, trucks per capita, drivers licences per capita, public road miles per adult, real income per capita, and population.

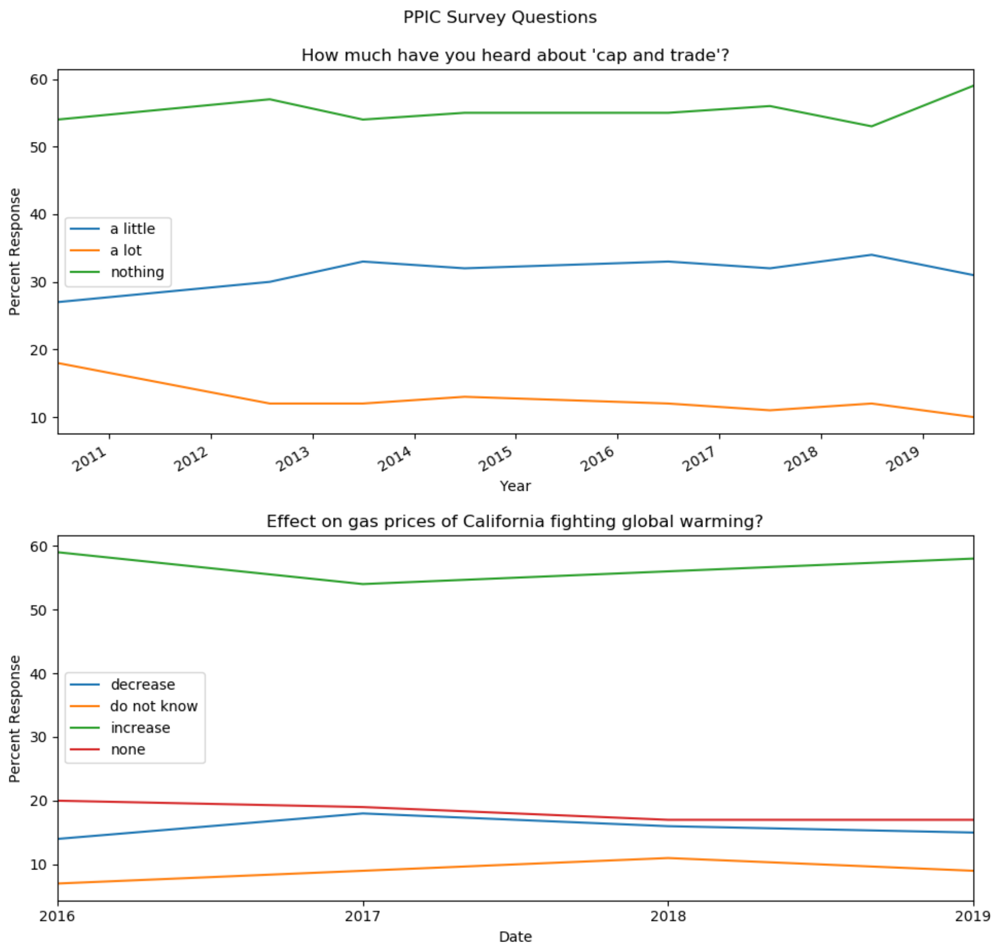

Figure 3.4. Data from ‘PPIC Statewide Survey: Californians and the Environment’.

once per year, with some questions phasing in and out. The survey queries a representative sample of California adults and likely voters. The results presented in Figure 3.4 are for “all adults”.

There two key questions to focus on from this survey are question 22 and question 26. The first notable question is “How much, if anything, have you heard about the state government policy called ‘cap and trade’ that sets limits on greenhouse gas emissions? Have you heard a lot, a little, or nothing at all?” The response of “nothing at all” is steady around 55%, while the response of “a little” is around 30%, suggesting that Californians are largely unaware of the cap-and-trade program in the state. The high percentage of those surveyed who know nothing is striking. These two responses, along with the general lack of news regarding cap-and-trade and its impact on gas prices in California, greatly diminish the likelihood that salience is a driving force in the heightened response of CCT prices, given most people seem to not be aware of the program at all.

The second question asked is this: “Do you think that California doing things to reduce global warming in the future would cause gasoline prices at the pump around the state to increase, or to decrease, or wouldn’t affect gasoline prices at the pump around the state?” A majority of respondents say that gasoline prices would “increase”, with the response rate hovering between 55% and 60%. That is, the respondents are saying they believe if a policy is implemented which will reduce the impact of global warming, gas prices will go up. This does not require specific policy knowledge, merely that there is some policy reducing global warming in effect. If it is true that consumers believe there are recent policies aimed at combating global warming, they may attribute that to observing gas price increases. Further, about an equal number of respondents, between 15% and 20%, believe the effect will be “nothing at all” or that gas prices will “decrease”.

3.5 Conclusion

This paper opened by discussing how consumers react differently to tax changes versus tax-exclusive price changes. The natural extension of this is how other, equivalent policies (those which implicitly tax individuals) affect consumption. I examined the effect of California’s cap-and-trade program and the LCFS program, showing that individuals appear to react to tax changes and CCT price changes similarly. Salience does not appear to be a plausible explanation, but rather that consumers expect prices to increase if a policy intended to fight global warming is implemented (if they are aware of it).

An important consideration for this literature is to better understand the process undertaken by individuals in their gasoline purchase decisions. Until we understand whether individuals consider their decision to purchase gasoline on a daily basis, or, for example, when their gasoline tank is below some threshold, is important. Further, the extent to which an individual is aware of all prices of gasoline stations readily accessible should be considered. Knowing more about the decision process allows better insight into additional mechanisms not considered here or in the previous literature. Yet, recent papers such as Levin et al. (2017) and Knittel and Tanaka (2019) are using higher-frequency microdata to explore some of these issues.

Determining how consumers react to taxes and other similar policies is important for deciding which policies will push consumers' behavior as the policy intends and which policy is ultimately the most beneficial. As this paper shows, economic policies that are equivalent in theory might result in markedly different observable impacts.

3.6 Appendix

Gasoline Prices and Gasoline Consumption

Figure 3.5 shows gasoline prices and gallons consumed over the period. The states follow similar changes across the sample. In particular, the graph shows that the same drops and rises occur across states. Even with the incidence of the CCT program in California, there is a trend upward in accordance with unaffected states and no noticeable decline.

Background on Cap-and-Trade and the Low-Carbon Fuel Standard

The state of California passed assembly bill AB-32 in 2006. The specified goal was to reduce greenhouse gas emissions. The bill gave the California Air Resources Board (CARB) control of designing and operating a carbon cap-and-trade program in California. This came into place in 2012, where the regulated industries were large manufacturing and electricity. In 2015, the program expanded (this was known ahead of time) to cover transportation fuels. The CCT program allocates some permits to firms every year, while holding quarterly auctions for the

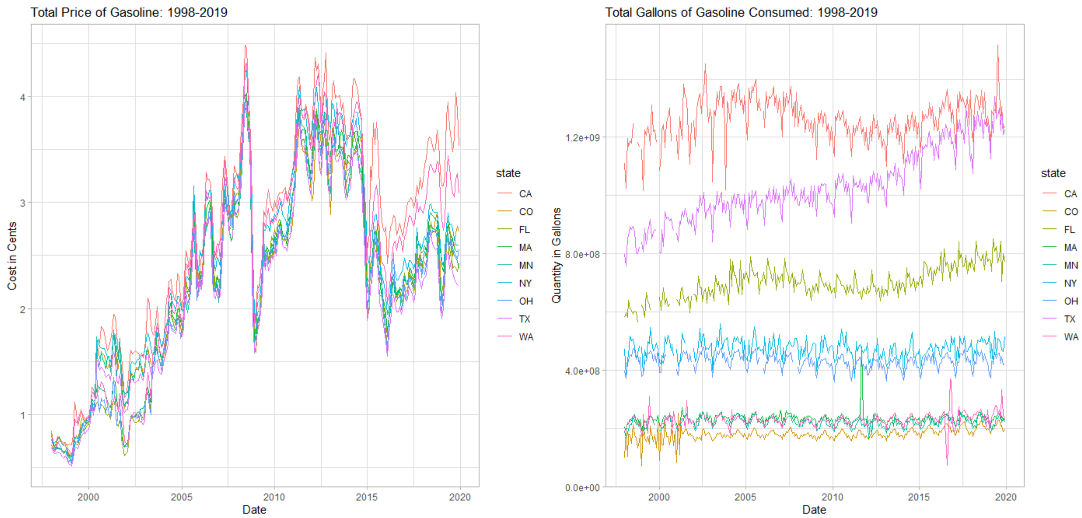

Figure 3.5. Gasoline Prices and Gallons Consumed

Note: This plot shows the gasoline prices and gallons of gasoline consumed for each state in the sample over the sample period.

remainder. The auction has a price floor which increases yearly by a known rate: 5% plus the inflation rate.

The low carbon fuel standard (LCFS) was included as a program in AB-32. The goal of the LCFS is to reduce the carbon intensity (CI) of the fuel pool of California, thereby reducing emissions (all else being equal). The LCFS sets a ceiling for the carbon intensity of the fuel pool, which declines every year until 2030. This forces firms to decrease the CI of their fuel stock or to purchase credits on the market. Firms with CIs below the ceiling earn credits, those above the ceiling create a deficit. Firms need to have a balance of credits at the end of the credit-year, thereby creating a market for credits.

Both of these programs operate by firms writing up individual contracts amongst themselves and determining the necessary number of permits/credits to trade as required under California State Law. These are also submitted to CARB, who oversees the bank of permits/credits. At the end of the permit cycle, firms will either receive or disburse the remaining credits/permits required determined by the firm's GHG emissions. One primary difference, though not important for this analysis, is the CCT program (not LCFS) allows permits to be paid by either entity (the upstream supplier or the downstream demander), whereas LCFS requires the supplier to pay.

Many countries, states, or groups of states have implemented cap and trade programs. The breadth of these programs varies by governed body. In particular, some cap and trade programs cover transportation fuels, while others do not, solely covering some larger entity such as electricity or large manufacturing. California was the first jurisdiction to implement an LCFS program in 2007, though it didn't take effect until 2011. In 2016, Oregon officially began an LCFS-like program of their own, called the Clean Fuels Program. Its creation and implementation deliberately mimics California's. Unfortunately, data for Oregon is not readily available and so is not included. British Columbia (in Canada) and the EU also adopted LCFS programs in 2008.

Additional Results Tables

A natural doubt to hold is that California is very different from other states, hence the effect of a CCT program in California may not have external validity. Table 3.2 should allay these concerns. The coefficients on both the tax-exclusive price variable and tax variable are close, yet estimated very imprecisely. Thus, one should expect that the reaction to the CCT program in California will be not far from how other states react.

Table 3.2. Comparison of Non-California states and California alone

	Dependent variable:	
	Log Quantity of Gasoline Per Capita	
	Not California	California
	(1)	(2)
Gas Prices	-0.00017^{**}	-0.00016^{*}
	(0.00007)	(0.00008)
Taxes	-0.00029	-0.00016
	(0.00380)	(0.00080)

*p<0.1; **p<0.05; ***p<0.01

Data series from January 1998 to December 2019. The first columns regresses log quantity of gasoline per capita on tax-exclusive gas prices and taxes for California only. The second column does the same but with all states excluding California.

One concern might be that there are not enough states for either the standard errors to be correct or the estimate is biased due to lacking other states. Here I replicate a table from Li et al. (2014). They present a table using monthly data from 1983 to 2008 with all 50 states. In their paper, it is "Table 5" which I have replicated here as Table 3.3, with an important modification. Their regression model is log-log, while the model estimated here is log-linear. The first column is the regression specified above where there is only the total price, while the second columns of both specifications are when total price is split into tax-exclusive price (i.e. input prices only) and taxes.

The coefficients presented in this table closely match the regression estimates from the data I have collected, relieving potential worries about biased estimates. Columns (1) and (2) should be compared from both tables, while columns (3) and (4) from Table 3.3 should be compared to columns (4) and (5) in Table 3.1. Comparing the first specification, the coefficients are nearly identical. In the second specification, the coefficients from Table 3.3 are larger, but the standard errors easily include the coefficients from the other table. This suggests concerns regarding insufficient data are not a problem.

Table 3.3. Li, Linn, and Muehlegger (2014) Replication

	Dependent variable:			
	Log Quantity of Gasoline			
	Month-of-year, State		State, State-Year, State-Month, Year	
	(1)	(2)	(3)	(4)
Input Prices	−0.00045***	−0.00042***	−0.00060	0.00061
	(0.00005)	(0.00006)	(0.00107)	(0.00125)
Taxes		−0.00300		−0.00651*
		(0.00186)		(0.00354)

*p<0.1; **p<0.05; ***p<0.01

Note: Data series from January 1983 to December 2008 for all 50 states in the USA. This table replicates the "Table 5" found in Li, Linn, and Muehlegger (2014). This table shows the result of regressing log quantity of gasoline per capita on (non-logged) tax-inclusive gasoline prices in columns (1) and (3), while columns (2) and (4) use tax-exclusive gas prices and gasoline excise taxes. The fixed effects are displayed above the columns; in particular, the first two columns use monthly time fixed effects and state fixed effects, while the latter two use state fixed effects, a state-year fixed effect, a state-month (12 for each state) fixed-effect, and year fixed effects.

It remains to investigate the impact of various assumptions made in the primary analysis. Naturally, there are many assumptions made, but two are of significance. The first assumption is simply including CCT prices in the tax-exclusive price variable as opposed to tax variable when initially separating the tax and tax-exclusive price changes. This is easily tested by instead including CCT prices in the tax variable in columns (2) and (5) of Table 3.1. The second is that cap-and-trade permit prices and LCFS credit prices engender the same response. The per-gallon costs are added together in the primary specifications, but the same estimation can run the variables separately to test for equal effects.

Challenging the assumption that CCT prices are more similar to tax-exclusive prices is simple: assume that individuals respond to CCT prices changes as they do to tax changes. To do this, CCT prices are added into the tax for columns (2) and (5) of Table 3.1. The results are contained in Table 3.4. The only differences between the two tables lies in columns (2) and (5). The coefficients reported in column (2) of Table 3.4 are 0.00014 for (all-)tax-exclusive gas prices and -0.00357 for tax (and CCT) changes. Compare this with -0.00007 and -0.00356, respectively, in Table 3.1 and it is seen the impact of this inclusion is identical. Column (5) leads to similar conclusions. The coefficients are -0.00014 and -0.00223, compared to the original table of -0.00015 and -0.00215. These results are identical as well.

The most interesting observation is the inclusion of CCT prices in taxes causes the two specifications to adjust their coefficients in opposite directions. Comparing column (2) in both tables shows the coefficients have decreased in Table 3.4, while examining column (5) shows both coefficients increasing. These changes are not statistically significant, but the change going in opposite directions is particularly interesting. As I cannot test this further, the baseline hypothesis remains that it is noise due to the differing sources of variation.

Overall, the interpretation to these imperceptible differences may be that CCT prices don't make a difference. Yet, an alternative explanation is that CCT prices act on a different source of information than taxes and gas prices. This will be investigated in the next section.

The prior tables estimate the effect of CCT price changes by effectively taking a weighted

Table 3.4. CCT Prices Included in Taxes

	Dependent variable:					
	Log Quantity of Gasoline					
	Month-of-year, State			State, State-Year, State-Month, Year		
	(1)	(2)	(3)	(4)	(5)	(6)
Gas Prices	-0.00031^{***}	0.00014	0.00012	-0.00016^{***}	-0.00014^{***}	-0.00014^{***}
	(0.00010)	(0.00014)	(0.00014)	(0.00004)	(0.00004)	(0.00004)
Taxes		-0.00323^{***}	-0.00357^{***}		-0.00223^{***}	-0.00231^{***}
		(0.00086)	(0.00093)		(0.00069)	(0.00082)
Cap and Trade			-0.00141			-0.00185^{***}
			(0.00101)			(0.00045)

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Data series from January 1998 to December 2019. This table differs from Table 1 in that columns (2) and (5) include CCT prices in the tax variable instead of the price variable; otherwise, the remaining columns are identical. The first three columns correspond to month-of-year fixed-effects and state fixed-effects. The second three columns are run with state, state-year linear trend, state-month, and year fixed-effects. Both regressions include state-level covariates. They are percentage of population living in an MSA, family size, and the log of: automobiles per capita, trucks per capita, drivers licences per capita, public road miles per adult, real income per capita, and population.

average of the impact of LCFS credit prices and cap-and-trade permit prices.[52] A question of direct interest is whether individuals react to these two implicit taxes differently.

The same specification is run as in Equation 3.2, except the c_{ts} component is split into an LCFS credit price variable and a cap-and-trade permit price variable. The regression runs these components as separate variables. The results of this regression can be found in Table 3.5. This table has the same columns as $(1), (2), (4)$, and (5) in Table 3.1. Columns (3) and (6) show the result of splitting up cap-and-trade permit prices and LCFS credit prices.

Column (3) estimates a positive CCT coefficient of 0.00918, while LCFS has a statistically significant coefficient of -0.00485, while the tax change coefficient increases a little to -0.00374. This coefficient is much larger than the coefficient on tax changes, but is not statistically significantly different. In column (6), the coefficient on CCT is much smaller, yet insignificant and positive at 0.00156. The coefficient on LCFS is slightly smaller, still significant, and negative at -0.00287. Both of these columns together suggest that individuals are reacting to the changes in the LCFS market. Further, this reaction remains of a similar size to the reaction to tax changes.

An important issue raised in Coglianese et al. (2017) is that anticipation can bias the estimate of the tax change coefficient because consumers might purchase gasoline just prior to the price increase due to the tax. To control for this, the authors suggest and use 1-month lead and lag coefficients to control for this anticipation. While their model is different than the typical one in this literature, their point stands.

What should be expected from the inclusion of these variables? Conventional economics says that a contemporaneous price increase should cause quantity to decrease. Now, consider a price increase tomorrow, about which the consumer is aware. In general, we should expect a consumer to purchase more gasoline today as the price will be higher tomorrow. That is, consumers will fill their gas tanks in anticipation of the tax increase.[53]

[52] A clarifying point: my reference to "cap-and-trade permit prices" means *only* those permit prices; it does not include LCFS credit prices.

[53] As an aside, it is somewhat peculiar that consumers do this. Most tax increases are at most \$0.10 per gallon

Table 3.5. Separating LCFS and CCT Prices

	Dependent variable:					
	Log Quantity of Gasoline					
	Month-of-year, State			State, State-Year, State-Month, Year		
	(1)	(2)	(3)	(4)	(5)	(6)
Gas Prices	-0.00031^{***}	0.00007	0.00015	-0.00016^{***}	-0.00015^{***}	-0.00014^{***}
	(0.00010)	(0.00016)	(0.00013)	(0.00004)	(0.00004)	(0.00004)
Taxes		-0.00355^{***}	-0.00374^{***}		-0.00215^{**}	-0.00238^{***}
		(0.00092)	(0.00091)		(0.00085)	(0.00085)
CCT			0.00918^{*}			0.00156
			(0.00499)			(0.00285)
LCFS			-0.00485^{***}			-0.00287^{***}
			(0.00111)			(0.00066)

*p<0.1; **p<0.05; ***p<0.01

Note: Data series from January 1998 to December 2019. This regression splits the CCT variable from Table 1 into a CCT *only* at-the-pump price and an LCFS *only* at-the-pump credit price in columns (3) and (6). The first three columns correspond to month-of-year fixed-effects and state fixed-effects. The second three columns are run with state, state-year linear trend, state-month, and year fixed-effects. Both regressions include state-level covariates. They are percentage of population living in an MSA, family size, and the log of: automobiles per capita, trucks per capita, drivers licences per capita, public road miles per adult, real income per capita, and population.

To test for the existence of this effect, model 3.2 is augmented to include both a lead and lag of the relevant price change. This can be found in Table 3.6. The inclusion of these variables changes the point estimates of the relevant contemporaneous variables. The all-tax-exclusive gas price variable increases slightly in magnitude, while the tax and CCT variables decrease. In general, the lead and lag coefficients do not have the expected pattern. In both cases, hypothesis tests reject at the 1% level equality between the sum of tax coefficients and the sum of all-tax-exclusive price coefficients. In the second specification, equality between the sum of CCT price coefficients and the sum of all-tax-exclusive price coefficients is rejected at the 1% level as well, but not in the first specification. This is consistent with results described above.

It has not yet been considered whether the tax-exclusive price might be endogenous within this article. If there are relevant demand or supply shocks to the economy which correlate with prices or consumption, then gasoline consumption might not be exogenous. To control for this, an instrument using crude oil prices is constructed to control for potential endogeneity of the tax-exclusive price of gasoline. The instrument uses gasoline monthly prices from 1996 interacted with monthly WTI oil prices spanning the timescale of the sample. This is a similar instrument construction as used in Li et al. (2014). The assumption is that any demand shocks after 1996 are not correlated with demand shocks occurring in 1998 and beyond. Similarly, CCT or LCFS prices might be endogenous. To control for this, a one-year lagged price of CCT prices is used. The resting assumption here is that any contemporaneous demand shocks will not be correlated with the previous year's CCT price. To control for possible endogeneity of taxes, I use the inflation-adjusted tax level.

The results can be found in Table 3.7. Overall, the instruments do not change the values too much. It is suggestive that the regression with monthly FEs are absorbing much of the variation. Comparing columns (1) and (2), the tax-exclusive gas price coefficient decreases in magnitude, while the tax and CCT price coefficients decrease in magnitude. Comparing columns

(many are much less), meaning a standard vehicle carrying capacity of 14 gallons costs a consumer at most $1.40 per gas station visit.

Table 3.6. Gas Prices, Taxes, and Cap and Trade

	Dependent variable:					
	Log Quantity of Gasoline					
	Month-of-year, State			State, State-Year, State-Month, Year		
	(1)	(2)	(3)	(4)	(5)	(6)
Gas Prices	0.00031	0.00036	0.00036	-0.00014^{**}	-0.00014^{**}	-0.00014^{**}
	(0.00028)	(0.00026)	(0.00026)	(0.00007)	(0.00007)	(0.00007)
Taxes		-0.00185	-0.00187		-0.00009	-0.00006
		(0.00221)	(0.00221)		(0.00135)	(0.00133)
Cap and Trade			-0.00183			0.00068
			(0.00230)			(0.00113)
Price lag	-0.00029^{*}	-0.00010	-0.00007	-0.00010	-0.00009	-0.00008
	(0.00017)	(0.00019)	(0.00018)	(0.00008)	(0.00008)	(0.00008)
Tax lag		0.00025	0.00008		-0.00120	-0.00142
		(0.00237)	(0.00233)		(0.00151)	(0.00147)
CCT lag			-0.00116			-0.00124^{**}
			(0.00168)			(0.00051)
Price lead	-0.00042^{**}	-0.00024	-0.00021	0.00007^{**}	0.00008^{**}	0.00009^{**}
	(0.00020)	(0.00019)	(0.00018)	(0.00003)	(0.00004)	(0.00004)
Tax lead		-0.00181^{***}	-0.00166^{***}		-0.00096	-0.00097
		(0.00058)	(0.00059)		(0.00084)	(0.00084)
CCT lead			0.00168			-0.00136^{*}
			(0.00124)			(0.00082)

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Note: Data series from January 1998 to December 2019. This regression adds a one-month lag and lead to the regression for each of tax-and-CCT-exclusive gas prices, gas taxes, and CCT prices. The first three columns correspond to month-of-year fixed-effects and state fixed-effects. The second three columns are run with state, state-year linear trend, state-month, and year fixed-effects. Both regressions include state-level covariates. They are percentage of population living in an MSA, family size, and the log of: automobiles per capita, trucks per capita, drivers licences per capita, public road miles per adult, real income per capita, and population.

Table 3.7. Instrumental Variables of Gas Prices, Taxes, and Cap and Trade

	Dependent variable:			
	Log Quantity of Gasoline Per Capita			
	Month-of-year, State		State, State-Year, State-Month, Year	
	(1)	(2)	(3)	(4)
Gas Prices	−0.00094	−0.00018	−0.00031**	−0.00029**
	(0.00120)	(0.00064)	(0.00012)	(0.00012)
Taxes	−0.00299	−0.00316**	−0.00249***	−0.00221***
	(0.00203)	(0.00148)	(0.00082)	(0.00085)
Cap and Trade	0.00080	−0.00103	−0.00227***	−0.00309***
	(0.00208)	(0.00103)	(0.00086)	(0.00080)
Covariates?	No	Yes	No	Yes

*p<0.1; **p<0.05; ***p<0.01

Note: Data series from January 1998 to December 2019. This table shows the result of an IV regression of Equation 1. To instrument for gas prices, the monthly WTI oil price is used; for taxes, the inflation-adjusted tax level is used; for CCT prices, a one-year lagged CCT price is used. The first three columns correspond to month-of-year fixed-effects and state fixed-effects. The second three columns are run with state, state-year linear trend, state-month, and year fixed-effects. Both regressions include state-level covariates. They are percentage of population living in an MSA, family size, and the log of: automobiles per capita, trucks per capita, drivers licences per capita, public road miles per adult, real income per capita, and population.

(3) and (4), the tax-exclusive gas price and CCT price coefficients increase in magnitude, while the tax coefficient decreases.

Analysis with News Articles

I gathered data from NexisUni on all articles mentioning gasoline prices, gasoline taxes, the California cap-and-trade program and gasoline prices, and the California low-carbon fuel standards program and gasoline prices. The latter two are not merely articles about the cap-and-trade program in California, but they must have included some key words pertaining to gasoline prices as well. An article discussing the cap-and-trade program without mentioning that gas prices change does not necessarily impart information about current gas prices.

The graph in the main paper displays article counts discussing gasoline taxes, cap-and-trade, and the low-carbon fuel standard. The graphs here also include article counts of gasoline prices. Figure 3.6 shows that discussion about gasoline prices and gasoline taxes in California is significantly more than articles about CCT or LCFS.

Figure 3.7 shows articles counts for the four searches mentioned above in NexisUni "Major US Newspapers". This set of newspapers does not discuss California's cap-and-trade program as it pertains to gasoline prices. The highest article count in a month is five. Even less discussed is the low-carbon fuel standard.

Finally, Figure 3.8 is all news articles as in Figure 3.6, but for all states in the dataset. It should be noted that only California has either a CCT or LCFS program, so the counts for gasoline prices or gasoline taxes will be much higher. This plot shows even more pronounced results.

Chapter 3, in part is currently being prepared for submission for publication of the material. The dissertation author was the sole author of this chapter.

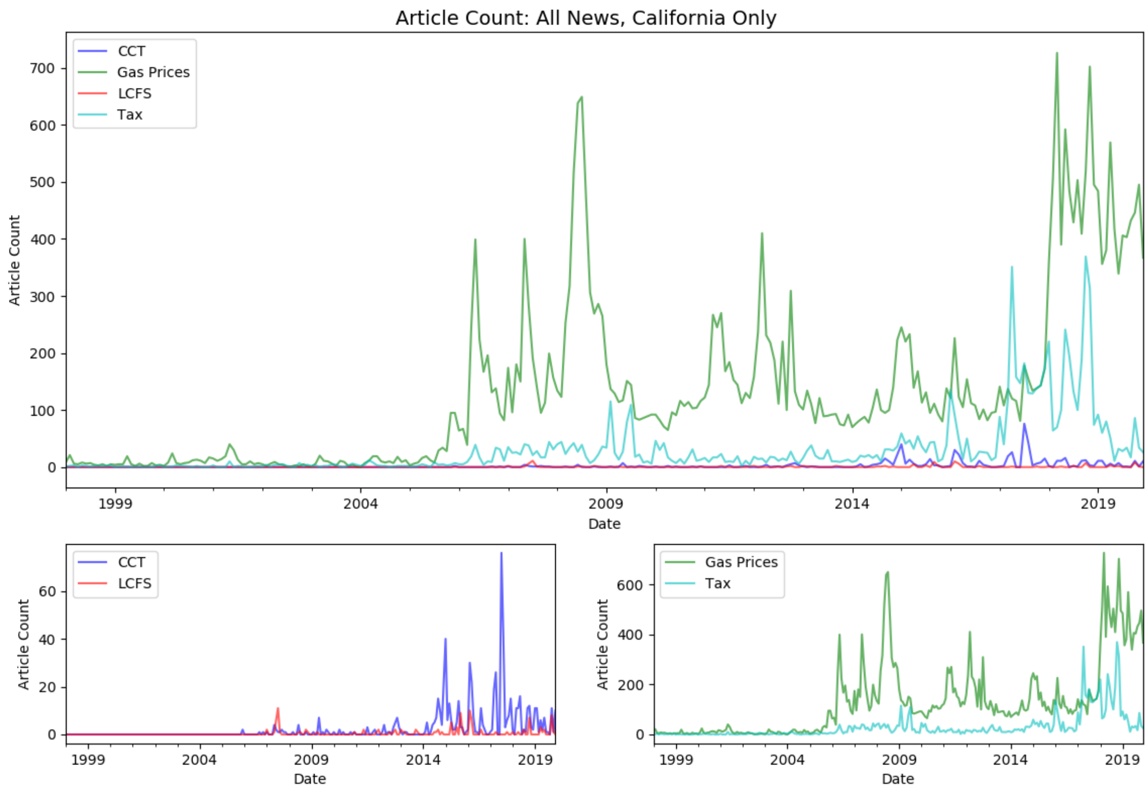

Figure 3.6. Article Counts: All News Publications, California Only

Note: The top graph shows monthly article counts in California for gasoline prices, gasoline taxes, cap-and-trade, and the low-carbon fuel standard. The lower two figures split this larger graph to better display the magnitude. The lower-left hand plot displays article counts for CCT and LCFS. The lower-right hand plot shows article counts for gas prices and gas taxes. Note the axes are different between the bottom two plots.

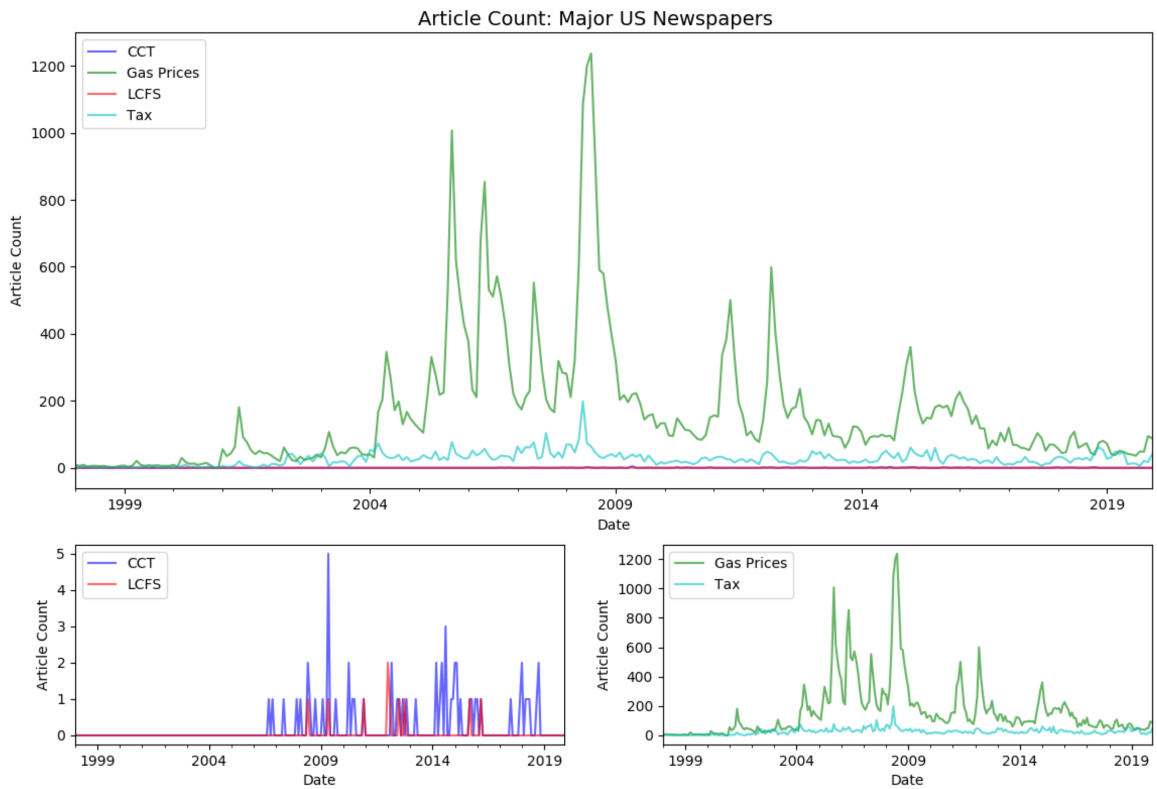

Figure 3.7. Article Counts: Major US Newspapers

Note: The top graph shows monthly article counts in NexisUni "Major US Newspapers" for gasoline prices, gasoline taxes, cap-and-trade, and the low-carbon fuel standard. See description notes in Figure 3.6.

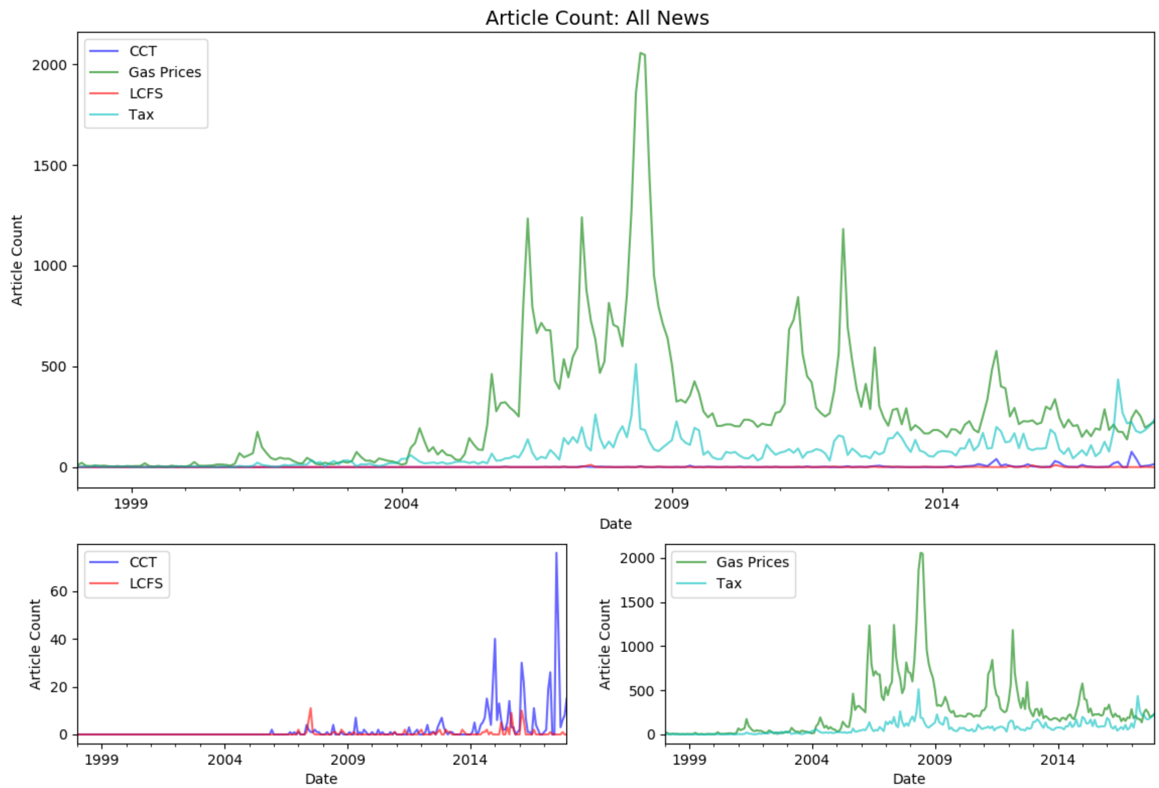

Figure 3.8. Article Counts: All News Publications

Note: The top graph shows monthly article counts for gasoline prices, gasoline taxes, cap-and-trade, and the low-carbon fuel standard. The lower two figures split this larger graph to better display the magnitude. The lower-left hand plot displays article counts for CCT and LCFS. The lower-right hand plot shows article counts for gas prices and gas taxes. Note the axes are different between the bottom two plots.

Bibliography

Alacevich, Caterina, Petyo Bonev, and Magnus Soderberg, "Pro-environmental interventions and behavioral spillovers: Evidence from organic waste sorting in Sweden," *Journal of Environmental Economics and Management*, 2021, *108.*

Allaire, Maura, Haowei Wu, and Upmanu Lall, "National trends in drinking water quality violations," *Proceedings of the National Academy of Sciences*, 2018, *115* (9), 2078–2083.

_ , Taylor Mackay, Shuyan Zheng, and Upmanu Lall, "Detecting community response to water quality violations using bottled water sales," *Proceedings of the National Academy of Sciences*, 2019, *116* (42), 20917–20922.

Allcott, Hunt and Dmitry Taubinsky, "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market," *American Economic Review*, 2015, *105* (8), 2501–38.

Andreoni, James, "Impure altruism and donations to public goods: A theory of warm-glow giving," *The economic journal*, 1990, *100* (401), 464–477.

_ , "Warm-glow versus cold-prickle: the effects of positive and negative framing on cooperation in experiments," *The Quarterly Journal of Economics*, 1995, *110* (1), 1–21.

Attanasio, Orazio P. and Luigi Pistaferri, "Consumption Inequality," *Journal of Economic Perspectives*, May 2016, *30* (2), 3–28.

Barlaz, Morton, Christoper Evans, Adam Brundage, Victoria Thompson, and Anne Choate, "WARM component-specific decay rate methods," *Memo to EPA*, 2009.

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

Busse, Meghan R, Christopher R Knittel, Jorge Silva-Risso, and Florian Zettelmeyer, "Did "cash for clunkers" deliver? The consumer effects of the car allowance rebate system," *Manuscript November*, 2012.

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

Cameron, A Colin and Pravin K Trivedi, *Microeconometrics: methods and applications*, Cambridge university press, 2005.

CARB, "SUMMARY OF CALIFORNIA-QUEBEC JOINT AUCTION SETTLEMENT PRICES AND RESULTS," 2021. [Online; accessed 10/27/21].

Coglianese, John, Lucas W Davis, Lutz Kilian, and James H Stock, "Anticipation, tax avoidance, and the price elasticity of gasoline demand," *Journal of Applied Econometrics*, 2017, *32* (1), 1–15.

Currie, Janet and Hannes Schwandt, "Mortality Inequality: The Good News from a County-Level Approach," *Journal of Economic Perspectives*, May 2016, *30* (2), 29–52.

Datta, Souvik and Sumeet Gulati, "Utility rebates for ENERGY STAR appliances: Are they effective?," *Journal of Environmental Economics and Management*, 2014, *68* (3), 480–506.

Davis, Lucas W, Alan Fuchs, and Paul Gertler, "Cash for coolers: evaluating a large-scale appliance replacement program in Mexico," *American Economic Journal: Economic Policy*, 2014, *6* (4), 207–38.

_ **and Lutz Kilian**, "Estimating the effect of a gasoline tax on carbon emissions," *Journal of Applied Econometrics*, 2011, *26* (7), 1187–1214.

de Coninck, Heleen, A. Revi, M. Babiker, P. Bertoldi, M. Buckeridge, A. Cartwright, W. Dong, J. Ford, S. Fuss, J.C. Hourcade, D. Ley, R. Mechler, P. Newman, A. Revokatova, S. Schultz, Linda Steg, and T. Sugiyama, "Strengthening and implementing the global response," in "Global warming of 1.5 C: Summary for policy makers," IPCC-The Intergovernmental Panel on Climate Change, 2018, pp. 313–443.

Duren, Riley M, Andrew K Thorpe, Kelsey T Foster, Talha Rafiq, Francesca M Hopkins, Vineet Yadav, Brian D Bue, David R Thompson, Stephen Conley, Nadia K Colombi, and C. Frankenberg, "California's methane super-emitters," *Nature*, 2019, *575* (7781), 180–184.

Ek, Claes and Jurate Miliute-Plepiene, "Behavioral spillovers from food-waste collection in Swedish municipalities," *Journal of Environmental Economics and Management*, 2018, *89*, 168–186.

EPA, "Waste Diversion as a Greenhouse Gas Control Strategy for Landfills," *United States Environmental Protection Agency Region 9 Air and Waste Divisions*, 2010.

EPA, US, "Documentation for Greenhouse Gas Emission and Energy Factors Used in the Waste

Reduction Model (WARM)," *United States Environmental Protection Agency*, 2020.

_ , "Inventory of U.S. Greenhouse Gas Emissions and Sinks 1990-2018," *United States Environmental Protection Agency*, 2021.

_ , "Inventory of U.S. Greenhouse Gas Emissions and Sinks 1990-2021," *United States Environmental Protection Agency*, 2021.

Errickson, Frank C, Klaus Keller, William D Collins, Vivek Srikrishnan, and David Anthoff, "Equity is more important for the social cost of methane than climate uncertainty," *Nature*, 2021, *592* (7855), 564–570.

FHWA, US, "Status of the Federal Highway Trust Fund 2017," 2017.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram, "Do energy efficiency investments deliver? Evidence from the weatherization assistance program," *The Quarterly Journal of Economics*, 2018, *133* (3), 1597–1644.

Fullerton, Don and Thomas C. Kinnaman, "Household Responses to Pricing Garbage by the Bag," *The American Economic Review*, 1996, *86* (4), 971–984.

GAO, (Government Accountability Office), "Drinking Water: Unreliable State Data Limit EPA's Ability to Target Enforcement Priorities and Communicate Water Systems' Performance. Report to Congressional Requesters," 2011.

Gillingham, Kenneth and James H Stock, "The cost of reducing greenhouse gas emissions," *Journal of Economic Perspectives*, 2018, *32* (4), 53–72.

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

Graff-Zivin, Joshua, Matthew Neidell, and Wolfram Schlenker, "Water quality violations and avoidance behavior: Evidence from bottled water consumption," *American Economic Review*, 2011, *101* (3), 448–53.

Hausman, Catherine and Lucija Muehlenbachs, "Price regulation and environmental externalities: Evidence from methane leaks," *Journal of the Association of Environmental and Resource Economists*, 2019, *6* (1), 73–109.

IPCC, "Volume 5: Waste, Chapter 3: Solid Waste Disposal," Jul 2006.

IWG, U.S.A., "Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide Interim Estimates under Executive Order 13990," February 2021.

Jacob, Daniel J, Alexander J Turner, Joannes D Maasakkers, Jianxiong Sheng, Kang Sun, Xiong Liu, Kelly Chance, Ilse Aben, Jason McKeever, and Christian Frankenberg, "Satellite observations of atmospheric methane and their value for quantifying methane emissions," *Atmospheric Chemistry and Physics*, 2016, *16* (22), 14371–14396.

Jacobsen, Mark R, "Evaluating US fuel economy standards in a model with producer and household heterogeneity," *American Economic Journal: Economic Policy*, 2013, *5* (2), 148–87.

_ **and Arthur A Van Benthem**, "Vehicle scrappage and gasoline policy," *American Economic Review*, 2015, *105* (3), 1312–38.

Knittel, Christopher R. and Shinsuke Tanaka, "Driving Behavior and the Price of Gasoline: Evidence From Fueling-level Micro Data," *NBER Working Paper No. 26488*, 2019.

la Cruz, Florentino B De and Morton A Barlaz, "Estimation of waste component-specific landfill decay rates using laboratory-scale decomposition data," *Environmental science & technology*, 2010, *44* (12), 4722–4728.

Levin, Laurence, Matthew S Lewis, and Frank A Wolak, "High frequency evidence on the demand for gasoline," *American Economic Journal: Economic Policy*, 2017, *9* (3), 314–47.

Li, Shanjun, Joshua Linn, and Erich Muehlegger, "Gasoline taxes and consumer behavior," *American Economic Journal: Economic Policy*, 2014, *6* (4), 302–42.

Marion, Justin and Erich Muehlegger, "Fuel tax incidence and supply conditions," *Journal of Public Economics*, 2011, *95* (9-10), 1202–1212.

Nordhaus, William D., "Revisiting the social cost of carbon," *Proceedings of the National Academy of Sciences*, 2017, *114* (7), 1518–1523.

of Austin, Texas The City, *Austin Resource Recovery Master Plan*, City of Austin, December 2011. [Online; accessed 05/10/21].

of Sciences, National Academies, *Improving characterization of anthropogenic methane emissions in the United States*, National Academies Press, 2018.

Pachauri, R. K., M. R. Allen, V. R. Barros, J. Broome, W. Cramer, R. Christ, J. A. Church, L. Clarke, Q. Dahe, P. Dasgupta, N. K. Dubash, O. Edenhofer, I. Elgizouli, C. B. Field, P. Forster, P. Friedlingstein, J. Fuglestvedt, L. Gomez-Echeverri, S. Hallegatte, G. Hegerl, M. Howden, K. Jiang, B. Jimenez Cisneroz, V. Kattsov, H. Lee, K. J. Mach, J. Marotzke, M. D. Mastrandrea, L. Meyer, J. Minx, Y. Mulugetta, K. O'Brien, M. Oppenheimer, J. J. Pereira, R. Pichs-Madruga, G.-K. Plattner, Hans-Otto Pörtner, S. B. Power, B. Preston, N. H. Ravindranath, A. Reisinger, K. Riahi, M. Rusticucci,

R. Scholes, K. Seyboth, Y. Sokona, R. Stavins, T. F. Stocker, P. Tschakert, D. van Vuuren, and J.-P. van Ypserle, *Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change*, IPCC, 2014.

Pai, Shantanu, Ning Ai, and Junjun Zheng, "Decentralized community composting feasibility analysis for residential food waste: A Chicago case study," *Sustainable Cities and Society*, 2019, *50*, 1016–83.

Pierce, Gregory and Silvia Gonzalez, "Mistrust at the tap? Factors contributing to public drinking water (mis)perception across US households," *Water Policy*, 2017, *19* (1), 1–12.

RGGI, "Allowance Prices and Volumes," 2021. [Online; accessed 10/27/21].

Rivers, Nicholas and Brandon Schaufele, "Salience of carbon taxes in the gasoline market," *Journal of Environmental Economics and Management*, 2015, *74*, 23–36.

Shahid, Kanwal and Eric Hittinger, "Technoeconomic optimization of food waste diversion to treatment facilities to determine cost effectiveness of policy incentives," *Journal of Cleaner Production*, 2021, *279*, 1226–34.

Shimshack, Jay P, Michael B Ward, and Timothy KM Beatty, "Mercury advisories: information, education, and fish consumption," *Journal of Environmental Economics and Management*, 2007, *53* (2), 158–179.

Steeter, Virginia and Brenda Platt, "Residential Food Waste Collection Access in the US," *BioCycle*, 2017.

Sterner, Thomas, "Distributional effects of taxing transport fuel," *Energy Policy*, 2012, *41*, 75–83.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Tiemann, Mary, *Safe drinking water act (SDWA): a summary of the act and its major requirements*, Congressional Research Service Washington, DC, 2017.

Tiezzi, Silvia and Stefano F Verde, "Differential demand response to gasoline taxes and gasoline prices in the US," *Resource and Energy Economics*, 2016, *44*, 71–91.