THE UNIVERSITY OF CHICAGO

CONSUMER DEMAND FOR ENVIRONMENTAL GOODS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY

ASHTON RICHARD PALLOTTINI

CHICAGO, ILLINOIS

JUNE 2026

© 2026 by Ashton Richard Pallottini

ORCID iD: https://orcid.org/0009-0005-9734-8023

TABLE OF CONTENTS

LIST OF FIGURES vi

LIST OF TABLES viii

ACKNOWLEDGMENTS ix

ABSTRACT x

1 ONSET OF OFFSETS: THE EVOLUTION OF SOCIAL REWARDS IN DEMAND FOR ENVIRONMENTAL GOODS 1
- 1.1 Introduction 1
- 1.2 Model 5
  - 1.2.1 Private and Social Rewards 5
  - 1.2.2 The Evolution of Social Rewards 6
  - 1.2.3 Additional Model Predictions 9
- 1.3 Experimental Design 10
  - 1.3.1 Recruitment 11
  - 1.3.2 Measuring Commonness: Perceived Market Penetrations 13
  - 1.3.3 The Decision to Buy 15
- 1.4 Reduced Form Analysis 19
  - 1.4.1 Empirical Specification 19
  - 1.4.2 Demand for Carbon Offsets 20
  - 1.4.3 Impact of Social Rewards 22
  - 1.4.4 Additional Mechanisms 26
- 1.5 Structural Estimation 29
  - 1.5.1 The Value Added by Structural Estimation 29
  - 1.5.2 Estimation Procedure 30
  - 1.5.3 Willingness-to-pay for Social Rewards 31
  - 1.5.4 Optimal Dynamic Subsidies Under Social Rewards 32
- 1.6 Conclusion 37

2 CHANGING GEARS: THE WELFARE IMPACTS OF BRAND IDEOLOGY . . 40
- 2.1 Introduction 40
- 2.2 Background 45
- 2.3 Data 48
  - 2.3.1 Vehicle Data 50
  - 2.3.2 ZIP Characteristics 52
  - 2.3.3 Market Share Trends 53
  - 2.3.4 External Validity 55
- 2.4 Empirical Analysis 57
  - 2.4.1 Model of Demand 59
  - 2.4.2 Instrument for Price 62

2.4.3 Results . . . 63
2.4.4 Alternate Mechanisms . . . 71
2.5 Welfare Impacts . . . 73
2.5.1 Firm Revenues . . . 73
2.5.2 Social Surplus Loss from Emissions . . . 74
2.5.3 Private Surplus Loss . . . 76
2.6 Conclusion . . . 77

3 URBAN EXPANSION, DROUGHT RISK, AND WILLINGNESS-TO-PAY FOR PIPED WATER . . . 80
3.1 Introduction . . . 80
3.2 Background . . . 84
3.2.1 Bangalore Water Supply . . . 84
3.2.2 2024 Bangalore Water Crisis . . . 85
3.3 Data . . . 88
3.3.1 Housing Demand . . . 88
3.3.2 Long-Run Housing Supply . . . 90
3.3.3 Water Infrastructure & Other Public Amenities . . . 91
3.3.4 Survey Data . . . 92
3.4 Results . . . 93
3.4.1 Demand for Housing . . . 94
3.4.2 Health as a Driver of WTP . . . 100
3.4.3 Short-run Housing Supply Effects . . . 101
3.4.4 Long-run Housing Supply Effects . . . 104
3.5 Adaptation Via the Housing Market . . . 106
3.6 Conclusion . . . 110

APPENDIX A APPENDIX FOR CHAPTER 1 . . . 112
A.1 Model Appendix . . . 112
A.1.1 Derivations . . . 112
A.1.2 Extensions . . . 115
A.2 Extension Appendix: Electric Vehicle Subsidies . . . 118
A.3 Robustness Appendix . . . 124
A.3.1 Balance Tables . . . 124
A.3.2 Demographic Re-weighting . . . 124
A.4 Additional Tables and Figures . . . 125

APPENDIX B APPENDIX FOR CHAPTER 2 . . . 144
B.1 Robustness Checks . . . 144
B.1.1 Varying Cutoffs of Political Classifications . . . 144
B.1.2 Using 2024 Vote Shares . . . 144
B.1.3 Instrument Measurement Error . . . 146
B.2 Sentiment Analysis . . . 147
B.2.1 Reddit Comment Selection . . . 147

B.2.2 Analyzing Sentiment of Reddit Comments 148
B.3 Additional Tables and Figures 150

APPENDIX C APPENDIX FOR CHAPTER 3 169
C.1 Discussion of Missing Data Imputation 169
C.2 Additional Survey Findings 171
C.2.1 Information Frictions on Water Sources 171
C.2.2 Risk Perceptions Pre- and Post-Crisis 171
C.2.3 Micro-level Migration Patterns 172
C.3 Additional Tables and Figures 174

REFERENCES 180

LIST OF FIGURES

1.1 Simulated Social Rewards Functions Under Normality 8
1.2 Histograms of Perceived Market Penetrations 16
1.3 Scatterplot of Perceived Market Penetrations 17
1.4 Reduced Form Estimated Impact of Visibility on Demand 23
1.5 Reduced Form Impact of Visibility by Market Penetration 24
1.6 Structurally Estimated WTP for Visibility by Market Penetration 32
1.7 Optimal Subsidies to Mitigate Carbon by Market Penetration 36

2.1 Google Trends Search Interest for Elon Musk Over Time 46
2.2 Sentiment Analysis 49
2.3 Tesla Market Share Over Time 54
2.4 Trends in Non-Tesla Market Shares 56
2.5 Distribution of ZIP Code Level Ideology for Texas 58
2.6 Brand Fixed Effects for non-American Brands 66
2.7 Brand Fixed Effects for American Brands 67
2.8 Trends in Fixed Effects for Tesla and Other Brands Selling EVs 68
2.9 Trends in WTP for Teslas Relative to the Median Brand 70

3.1 Water Pipe Locations 86
3.2 Listing Distances from Nearest Water Pipe 87
3.3 Locations of Scraped Rental Listings by Month 90
3.4 Average Price Trends by Distance from Water Pipes 95
3.5 The Impact of an Urban Water Crisis on Housing Rents 98
3.6 The Impact of an Urban Water Crisis on Housing Rents by Lease Length 99
3.7 The Impact of an Urban Water Crisis on Housing Rents by Size 100
3.8 Health Issues During the Water Crisis 102
3.9 The Impact of an Urban Water Crisis on Residential Construction 105

A.1 Reduced Form Impact of Visibility by Second-order Belief Including Privacy Concerns 119
A.2 Structurally Estimated WTP for Visibility by Market Penetration Including Privacy Concerns 120
A.3 Optimal EV Subsidies by Market Penetration 122
A.4 Pigouvian Subsidies for EVs Under Counterfactual Parameterizations 123
A.5 Histograms of Beliefs by Political Ideology 128
A.6 First- Versus Second-order Beliefs 129
A.7 Posterior Perceptions of Quality by Randomly Assigned Commonness Signal . . 130
A.8 The Impact of Price on Uptake 131
A.9 Reduced Form Impacts to Uptake by Second-order Belief Separated by Visibility Treatment 132
A.10 Reduced Form Impact of Visibility by First-order Belief 133
A.11 Reduced Form Impacts to Uptake by First-order Belief Separated by Visibility Treatment 134

A.12 Reduced Form Impact of Visibility by Second-order Belief Omitting Referrers . 135
A.13 Reduced Form Impact of Visibility by Second-order Belief Omitting Virtue Signaling Concerns 136
A.14 Reduced Form Impact of Visibility by Second-order Belief Omitting Experienced Survey-Takers 137
A.15 Nonparametric Estimates of Visibility's Impact by Market Penetration 138
A.16 Two-Bin Estimates of Visibility's Impact by Market Penetration 139
A.17 Structurally Estimated Demand for Mitigating 145lbs. of Carbon 140
A.18 Structurally Estimated WTP for Visibility by Market Penetration for Full Sample 141
A.19 Structurally Estimated WTP for Visibility by Market Penetration Omitting Conservatives 142
A.20 Structural Estimates of the Model's Assumption 143

B.1 Alternative Cut-offs for Democrat and Republican ZIPs 145
B.2 Trends in WTP for Teslas Using 2024 Vote Shares 146
B.3 Google Trends Search Interest for Elon Musk's Politics 153
B.4 Trends in Tesla Stock Price 154
B.5 Trends in WTP for Teslas Relative to the Average Brand 155
B.6 Brand Fixed Effect for Rivian 156
B.7 Trends in News Stories Mentioning Tesla 157
B.8 Trends in Charging Station Availability 158
B.9 Sentiment Analysis - Alternative Definition of Positive and Negative Comments 159
B.10 Sentiment Analysis - Alternative Definition of Positive and Negative Comments 160
B.11 Permutation Test on Estimated Tesla WTP Changes 161
B.12 Tesla Delivery Times 162
B.13 Trends in WTP for Teslas Relative to the Median Brand with Two Political Classifications 163
B.14 Texan Vote Shares for Donald Trump in 2024 164
B.15 Trends in Total Vehicles 165
B.16 Trends in Competitor Entry Into Vehicle Markets 166
B.17 Supply Chain Constraints in Tesla Deliveries Over Time 167

C.1 Hedonic Estimates with $m = 1$ and $m = 10$ PMM Imputation 170
C.2 Self-reported Pipe Access by Distance to Nearest Pipe 172
C.3 Perceptions of Drought Risk 173
C.4 Surveyed Distance from Nearest Pipe 174
C.5 Trends in Listing Characteristics and Count by Distance 176
C.6 The Impact of an Urban Water Crisis on Housing Rents: Building-level Panel . 177
C.7 The Grouping of Listings More Than 0.2 km Away from Piped Water 178
C.8 Estimated Price Gradient Before and After Water Crisis 179

LIST OF TABLES

1.1 Demographics in the Sample Compared to the U.S. Population 12
1.2 Reduced Form Estimates of Visibility's Impact by Market Penetration 21

2.1 Summary Statistics of New Vehicle Registrations 51
2.2 Demand Estimation Coefficient Results 64
2.3 Annual Welfare Impacts of Musk's Involvement in the 2024 Election 75

3.1 Summary Statistics of Key Amenities 89
3.2 Summary Statistics of Key Characteristics of Construction Projects 91
3.3 Summary Statistics of Survey Respondents 93
3.4 The Impacts of Urban Drought on Health Outcomes 103
3.5 The Impacts of Urban Drought on New Residential Construction 107

A.1 Balance by Visibility Treatment 124
A.2 Balance by Price 125
A.3 Balance by Second-Order Belief Signal 126
A.4 Re-weighted Reduced Form Estimates of Visibility's Impact by Market Penetration 127
A.5 First Stage Regressions 138

B.1 Descriptive Statistics for Texas, California, and the U.S.A. 151
B.2 First Stage Results 152
B.3 Summary Statistics of Nests 168

C.1 Summary of Other Water Crises 175
C.2 Proportional Coverage Estimates on New Residential Construction 179

ACKNOWLEDGMENTS

I thank the many advisors who have helped me during my time as a Ph.D. student. I am especially grateful to Leonardo Bursztyn, Michael Greenstone, and Koichiro Ito for their excellent guidance. My growth as a scholar is due to your clever advice and unwavering confidence.

My decision to pursue a Ph.D. was inspired by excellent courses I enrolled in at Michigan State University. I particularly would like to offer my gratitude to Prabhat Barnwal, Chih-Li Sung, and Jeffrey Wooldridge for their fantastic teaching, mentorship, and support.

I have learned so much from my coauthors, without whom this dissertation would be significantly lessened. I am grateful to have worked with Sofia Shchukina on the second chapter of this dissertation and with Jun Wong and Emma Zhang on the third. I am likewise thankful to have collaborated in work outside this dissertation with Soren Anderson, Prabhat Barnwal, Arik Levinson, and Yixin Sun. Thank you all for joining me in 'making things work.'

I am thankful to my peers at The University of Chicago for creating a community which is intellectual, creative, and stimulating to engage with. My fondest memories of my Ph.D. include the cherished friends I have made over the years. We have come a long way from playing Among Us and making small talk in virtual chat rooms during our first year.

I also thank the undergraduate and EMBA students who I have had the great opportunity to teach as a Ph.D. student. There are no greater students at any university across the world. It is trite, but I genuinely feel that I learned more from you than you did from me.

I also send my thanks and my love to my family, especially my mother and father, brother, aunt, uncle, and grandparents. It takes a village, and I thank you for being the core of mine.

ABSTRACT

This dissertation includes three essays on the underlying motives of consumers demanding 'eco-friendly' environmental goods.

The first essay estimates how the visibility of such environmental goods influences demand via 'social rewards' that are garnered by consumers and change as a good's market evolves. We run an online RCT that generates variation in price, purchase visibility, and perceived market penetrations of carbon offsets – an environmental good. We show that social rewards increase willingness-to-pay (WTP) by 87% for environmental goods which are seen as having near-0% market penetration. This effect dissipates quickly for goods which are seen as more common, dropping below 15% for market penetrations of 12-50%. Consequently, We show that subsidies for carbon mitigation should be set $10 below the social cost of carbon for uncommon goods, sloping upwards in market penetration as social rewards disappear and consumer WTP for carbon mitigation decreases.

The second essay examines the extent to which ideology influences consumer WTP for an environmental good and, consequently, private and social welfare. Using Texas vehicle registrations, we estimate that WTP for the Tesla brand falls by $12,000 among Democrat ZIPs and moderate ZIPs from 2021Q1 through 2025Q1. We show that the steepest falls in WTP for Teslas correspond to CEO Elon Musk's Republican political involvement. We observe a particularly sharp fall in WTP for Teslas of $8,000 during Musk's involvement in the 2024 election. Consequently, we estimate that Musk's 2024 election involvement leads Tesla revenues to fall $2.3 billion per year in the U.S.A. Democrat and moderate consumers substituted primarily to gas vehicles, leading annual damages from vehicular carbon emissions to rise by $32 million and private surplus to fall by $1.6 billion per year.

The final essay studies rising groundwater drought risk to investigate how a shock to the quality of private groundwater wells impacts demand for publicly supplied piped water. We estimate that an unexpected major groundwater drought in Bangalore permanently increased

residents' willingness-to-pay to live near the pipe network by 0.4% of monthly rents per 100m. We show consequently that new residential construction projects that are close to the piped water network increased by 25-51% relative to those that are farther. We find that this housing market adjustment recovers 20% of the welfare gains of expanded pipe access at just 2.4% the cost.

CHAPTER 1

ONSET OF OFFSETS: THE EVOLUTION OF SOCIAL REWARDS IN DEMAND FOR ENVIRONMENTAL GOODS

1.1 Introduction

Markets for goods that have an environmental benefit are rapidly expanding. American expenditures on retail environmental goods are projected to nearly double to $400 billion by 2032, with consumers also paying hefty sums for durable goods like electric vehicles and solar panels.[1] Policymakers view this widespread expansion of demand for environmental goods as crucial in improving environmental quality, motivating wide-ranging applications of subsidies [U.S. White House, 2024]. Unfortunately, the importance of monetary considerations versus other components in driving demand for environmental goods is unclear.

Academic literature has speculated that demand for environmental goods is driven by 'social rewards' that consumers gain from visibly buying an environmental good [Griskevicius et al., 2010]. However, the impact of social rewards on willingness-to-pay (WTP) for environmental goods is largely unquantified. It is also unclear whether WTP for these social rewards is constant as environmental goods become more common, or whether social rewards will disappear as the novelty of environmental goods fades. Many environmental goods are transitioning from uncommon to common, so understanding how social rewards evolve is crucial.

In this paper, I use an online randomized controlled trial (RCT) on Prolific users to estimate the changing role of social rewards in driving consumer WTP for environmental goods.[2] I experimentally analyze consumer demand for goods offsetting carbon emissions linked to climate change, an externality which is estimated to have *annual* welfare impacts

1. See here: `https://capitaloneshopping.com/research/eco-conscious-consumer-statistics/`

2. The RCT is pre-registered on the AEA RCT registry with RCT ID AEARCTR-0012632. It also is approved by University of Chicago's IRB under approval number IRB23-1827.

worth trillions of dollars [U.S. EPA, 2023]. My RCT manipulates purchase visibility as well as how common the good is perceived to be. Using this induced variation in visibility and commonness, I provide the first estimates of how social rewards influence demand for environmental goods dynamically.

To show that social rewards may change as environmental goods become more common, I first present a model of consumer demand. Consumers in the model are classified by their privately-known 'types' which dictate their private rewards from buying the good [Bénabou and Tirole, 2011]. Part of the consumer's type is of interest to individuals who observe the consumer's purchase – for instance, the consumer's wealth, prosociality, or how 'cool' the consumer is. When this environmental good is visible, other individuals who see the consumer's decision to buy or not buy are able to infer these things about the consumer. This generates 'social rewards' for the consumer who enjoys being seen as person who is wealthy, prosocial, or cool. The model predicts that the WTP for these social rewards varies with the commonness of the good in a way that is theoretically ambiguous.

I then leverage an online RCT on a sample of Prolific users to measure WTP for social rewards among environmental goods with varying levels of commonness. The RCT elicits participants' beliefs about commonness, where commonness is measured by market penetration. I induce exogenous variation in these beliefs by sending consumers random signals about market penetrations, where signals are generated from prior small-scale surveys. From here, I offer participants a choice to purchase a carbon offset – an environmental good – at a randomly discounted price. Half of participants are randomly assigned to the visible group which offers them the choice to have their name published on a website if they buy the offset, with this website being promoted to thousands of individuals on Instagram. The other half do not receive this option and are treated as the 'invisible' group.

I estimate from my RCT that consumer WTP for environmental goods is influenced by social rewards, especially when the good is uncommon. Among consumers with perceived

market penetrations nearly equal to 0, I estimate that WTP for the environmental good increases by 87% when the purchase is made visible. This increase in WTP is less than 15% for penetrations of 12-50%, indicating that consumers primarily value social rewards in environmental goods that are seen as uncommon. I show that these estimates are robust to using both structural and reduced form estimation strategies. I in turn use my estimates to calculate optimal Ramsey and Pigouvian subsidy paths for environmental goods. Focusing on my setting of carbon mitigation, I estimate that carbon subsidies should be set $10 per metric ton lower than the social cost of carbon (SCC) for goods with near-0% market penetration and roughly equal to the SCC for goods with market penetrations between 10-50%.

Such findings are important because they show that social rewards play a large role in driving demand for environmental goods, particularly among goods that are uncommon. This study thus highlights the potential of purchase visibility enhancement as a policy tool to increase demand for environmental goods. These findings also directly inform optimal dynamic subsidy policy for environmental damage mitigation. When considering consumer demand, I find that dynamic subsidies for environmental goods should initially be set small and ramp up as the good becomes more common. This contrasts supply-side considerations of learning-by-doing and innovation spillovers that have led policymakers to set extant subsidies largest among nascent technologies and slope downwards [Spence, 1984, Bloom et al., 2013, Covert and Sweeney, 2024].[3] My findings thus highlight the importance in also considering demand-side mechanisms when setting dynamic subsidies for environmental goods.

I contribute to the environmental economics literature by quantifying how social rewards impact demand for environmental goods, including a focus on the dynamics of social rewards. Previous work has examined behavioral components in demand for environmental goods [Bollinger and Gillingham, 2012, Allcott and Rogers, 2014b, Allcott and Kessler,

3. For instance, subsidies for electric vehicles in the U.S. before August 2022 were set such that firms selling more than 200,000 vehicles would lose eligibility for the subsidy. Similarly, the California Solar Initiative and many other clean energy subsidies have sloped downwards over time.

2019b, Bollinger et al., 2020a, Gillingham and Bollinger, 2021a, Bollinger et al., 2022, Rodemeier, 2024]. Several works have specifically examined social rewards in non-causal or stated preference settings [Sexton and Sexton, 2014a, Gosnell et al., 2021, Carattini et al., 2022]. I thus contribute to this literature in three ways. First, I contribute by causally identifying the role of social rewards in driving WTP for environmental goods. Second, I contribute by analyzing social rewards as a *dynamic* influence in WTP for environmental goods, an analysis which helps explain upwards sloping optimal dynamic subsidies [Langer and Lemoine, 2022]. Finally, I contribute by showing that visibility enhancement in many cases can be a useful tool for policymakers in increasing demand for environmental goods.

I also contribute to the behavioral economics literature by estimating how social rewards vary as goods become more common. Many articles have shown that social rewards are important in high-stakes settings, usually focusing on one level of a good's commonness [Bénabou and Tirole, 2006, Gerber et al., 2008, DellaVigna et al., 2012, Ashraf et al., 2014b,a, Bursztyn and Jensen, 2017, DellaVigna et al., 2017, Perez-Truglia and Cruces, 2017, Karing, 2024]. However, only one paper has examined how social rewards change as a good becomes more common [Jee et al., 2024]. I contribute two things beyond this paper. First, I contribute by explicitly estimating WTP for social rewards, as my setting is able to include price variation. Second, I expand the range over which the evolution of social rewards is examined, including the crucial earliest stages of market development in which I show WTP for social rewards fall drastically.

The paper proceeds as follows: Section 2 provides my theoretical framework, Section 3 details my RCT, Section 4 summarizes my reduced form empirical results, Section 5 enumerates results from my structural estimation, and Section 6 concludes.

1.2 Model

In this section, I discuss the model I use to evaluate the importance of social rewards in driving willingness-to-pay for environmental goods [Bénabou and Tirole, 2011, Jee et al., 2024].[4] Consumers are modeled as receiving both private and social rewards from visibly buying environmental goods. The model yields that the importance of social rewards varies in ways that are theoretically ambiguous as markets develop from uncommon to common.

1.2.1 Private and Social Rewards

A consumer i makes a decision $a_i \in \{0, 1\}$ about whether to buy an environmental good or not.[5] Based on this decision, the consumer receives utility U_i. This utility depends on the good's price p and person i's privately-known idiosyncratic tastes for the good u_i. Utility also depends on person i's privately-known v_i which captures the components of preferences held by i which are of interest to individuals apart from i (e.g. how prosocial, wealthy, or 'cool' i is). Utility for person i thus also depends on the social rewards from being seen by others as having a high v_i. I thus model utility as

$$U_i = \big(\beta_p(y - p) + v_i + u_i\big)\, a_i + \mu E_{-i}[v_i|a_i]$$

with y representing a subsidy applied to the good and so β_p is a marginal utility of money. The coefficient $\mu = \lambda * x$ is a combined measure of how much people care about social pressure ($\lambda \geq 0$) and how visible the good is ($x \geq 0$). Social rewards are contained in the term $E_{-i}[v_i|a_i]$, which measures the expectations formed by people who are not i about person i's level of v_i, given person i's decision to buy the good or not. I assume that u_i and v_i are distributed independently from each other and combine them into one value $w_i := u_i + v_i$

4. I follow Jee et al. [2024] in including a term for idiosyncratic preferences, but refrain from endogenizing the visibility of goods as a function of cost.

5. The results of this model generalize to a continuous action space [Bénabou and Tirole, 2006, 2011]

which completely encompasses an individual's privately-known 'type.'

Equilibrium in this model is characterized by a cutoff type $w^*(y)$ who is ambivalent between choosing $a_i = 1$ and $a_i = 0$. Every agent with type above $w^*(y)$ buys the good, whereas every agent with type below $w^*(y)$ does not. This cutoff type solves the following fixed point equation:

$$w^*(y) + \mu \{E(v|w > w^*(y)) - E(v|w \leq w^*(y))\} = \beta_p(p - y)$$

The left-hand side of this equation captures the benefits for the marginal type of buying the good. The right-hand side of the equation captures the costs of buying the good. The marginal cutoff type $w^*(y)$ thus has benefits equal to costs.

Of interest in this study are the social rewards which are a function of how common a good is. For brevity, I define these social rewards as

$$\Delta(w^*(y)) := E(v|w > w^*(y)) - E(v|w \leq w^*(y))$$

This is the difference between what others infer as the v of the good buyers ($a = 1$) and what others infer as the v of the good non-buyers ($a = 0$). The social rewards are a function of the value of $w^*(y)$, indicating that social rewards vary by market penetration. The focus of this paper is understanding how social rewards vary as a good becomes more common, directly mapping out this $\Delta(*)$ as a function of $w^*(y)$.

1.2.2 The Evolution of Social Rewards

The shape of $\Delta(*)$ is theoretically ambiguous as a function of good commonness, depending on the unobservable underlying distribution of w. In other words, the model predicts that social rewards ambiguously vary as a good evolves from uncommon to common, motivating empirical estimation of the social rewards function. This follows directly from the social

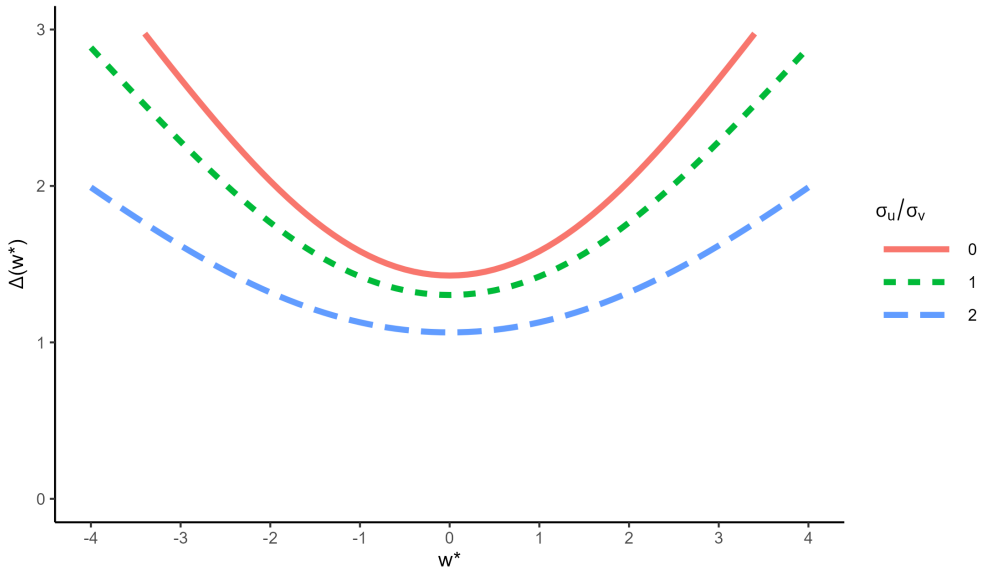
rewards being dependent on the conditional expectations $E(v|w > w^*(y))$ and $E(v|w \leq w^*(y))$.

One interesting social rewards function, a U-shape, is formed when w is distributed such that its probability density function (PDF) has one interior mode (e.g. a Normal distribution). I plot the social rewards function under a normally distributed w in Figure 1.1, with the top panel depicting the x-axis as w^* and the bottom panel depicting the x-axis as the market penetration.[6] Under a normally distributed w, social rewards are high when few people buy the good (w^* is large/positive), as those who buy the good reap significant "honor" from having been one of the few people to do so. As more individuals buy the good, this honor quickly dissipates as it becomes more difficult to infer which of the good buyers are at the highest values of v. At low/negative values of w^*, social rewards begin to increase again due to the "shame" felt by those few not buying the good.

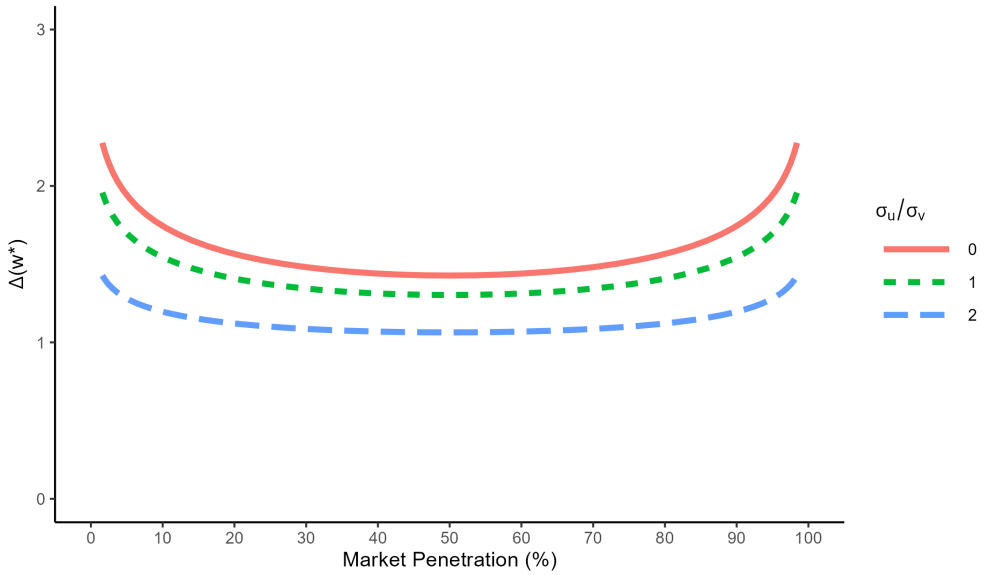
The dependence of the shape of social rewards on the unobservable distribution of w makes empirical estimation of social rewards important. Distributions other than the normal distribution will present different forms of social rewards, some of which differ vastly from the U-shape under normality.[7] Figure 1.1 shows further that even within the class of normal distributions the shape of social rewards can vary much. This figure depicts that social rewards depend on the standard deviation of idiosyncratic preferences, σ_u, where larger values of σ_u yield flatter social rewards curves. A larger value of σ_u indicates more variation in idiosyncratic preferences that drive purchases, flattening the social rewards function as purchases becomes less informative signals about v.

6. A normal distribution of w comes from e.g. assuming that u and v are both independently normally distributed, so that their sum w is as well.

7. Generally, social rewards will be subjected to 'honor' and 'shame' regions for any distribution of w with a unique interior maximum. See Bénabou and Tirole [2011] for further generalized results.



Panel A: Social Rewards by Cutoff Type



Panel B: Social Rewards by Market Penetration

Figure 1.1: Simulated Social Rewards Functions Under Normality

Notes: Figure depicts simulated social rewards by cutoff type w^* (top) and percent market penetration (bottom). All lines assume $\sigma_v = 0.5$ and $\mu = 1$. The solid red lines on top assume that $\sigma_u = 0$, the short-dash green lines in the middle that $\sigma_u = 0.5$, and the long-dash blue lines on the bottom that $\sigma_u = 1$.

1.2.3 Additional Model Predictions

The model generates three predictions, all of which can be tested in data to measure goodness-of-fit.

I define the cumulative density function (CDF) of w as $F_w(*)$, yielding the following expression which gives the proportion of people who buy the good:

$$\bar{a}(y) = 1 - F_w(w^*(y))$$

The first prediction is then:

Prediction 1 (Impact of μ)**.** *Increasing the value of $\mu = \lambda * x$ increases uptake of environmental goods at all values of w^*:*

$$\frac{\partial \bar{a}}{\partial \mu} = f_w(w^*(y)) \frac{\Delta(w^*(y))}{1 + \mu\Delta'(w^*(y))} > 0$$

with $f_w()$ representing the PDF of w.

This prediction states that increasing μ by making a good more visible or increasing the amount that people care about social signaling will increase the proportion of individuals buying the good.[8]

The second prediction is:

Prediction 2 (Impact of p)**.** *Increasing the price p decreases uptake of environmental goods at all values of w^*:*

$$\frac{\partial \bar{a}}{\partial p} = -f_w(w^*(y)) \frac{1}{1 + \mu\Delta'(w^*(y))} < 0$$

8. Note that this holds if and only if $1 + \mu\Delta'(w^*) > 0$. This inequality ensures a unique equilibrium in the model and prevents the indirect effect of visibility on $w^*(y)$ from outweighing the direct effect. This inequality holds for μ which is sufficiently small and $\Delta'(*)$ which is not sufficiently negative. I test this assumption in Appendix Figure A.20 and show that it holds for all observations in my data.

with $f_w()$ representing the PDF of w.

Consumers receive disutility from paying a higher price, and so increasing the price directly decreases the proportion of individuals buying the good.[9]

The final prediction is:

Prediction 3 (Relative Price Sensitivities)**.** *Consumers are less price sensitive in the visible case ($\mu > 0$) than in the invisible case ($\mu = 0$) if and only if $\Delta'(w^*(y)) > 0$.*

$$\left.\frac{\partial \bar{a}}{\partial p}\right|_{\mu \neq 0} = -f_w(w^*(y))\frac{1}{1+\mu\Delta'(w^*(y))} > -f_w(w^*(y)) = \left.\frac{\partial \bar{a}}{\partial p}\right|_{\mu=0} \iff \Delta'(w^*(y)) > 0$$

The term $(1+\mu\Delta'(w^*(y)))^{-1}$ is referred to as the "social multiplier," as it dictates social rewards' multiplicative impact on consumer price sensitivity. In the case with $\Delta'(w^*(y)) > 0$, the social multiplier is between 0 and 1 and so the price sensitivity is attenuated towards 0 in the visible case as compared to the invisible case. This is because increasing the price has an indirect effect on uptake – it lowers the number of people buying the good, which because $\Delta'(w^*(y)) > 0$ means that social rewards in turn increase. This increase in the social rewards serves to push people back into buying the good, an effect which is not present when $\mu = 0$.

1.3 Experimental Design

In this section, I describe a RCT on Prolific users which generates random variation in prices, visibility, and commonness for one specific environmental good – a carbon offset. The RCT captures invisible and visible purchases of this environmental good, allowing for estimation of willingness-to-pay for social rewards. Commonness is measured by perceived market penetrations, which are exogenously varied with randomly assigned signals to allow analysis of WTP for social rewards in goods that are common and uncommon.

9. Crucially, I enforce that social signaling value does not depend on price paid (i.e. that μ is not a function of p). In my RCT, the fact that prices vary is not privately nor publicly known.

1.3.1 Recruitment

To enable measurement and exogenous manipulation of visibility and commonness, I recruit participants in an online survey. I follow best practices to improve internal and external validity, yielding a fairly representative sample of individuals that were attentive while taking the survey.

Participants are recruited for my survey through the Prolific platform, which has been shown to produce results that are similar to Pew probabilistic samples under best practices [Tang et al., 2022, Douglas et al., 2023]. I recruit 3,111 participants located in the U.S. and offer $2 to take an approximately 8-minute survey. Participants are also made aware that they will receive bonus payments of up to $21.50, with an anticipated average bonus payment of approximately $1.60. Average payouts in my survey are thus listed as $3.60. This equates to an hourly wage of $27 per hour, which is more than twice the norm of $12 per hour set by major platforms such as Prolific. A relatively large payment was chosen to increase representativity of the sample by reducing selection into the survey on willingness-to-accept (WTA) taking surveys.

I further reduce concerns of a selected sample by following recruitment best practices [Tang et al., 2022, Douglas et al., 2023]. I set participation quotas on sex, race, and political views, which are important demographic components driving selection onto the Prolific platform. This makes my sample more representative of the U.S. population along these observable demographics, though it does not fully address selection on unobservables.[10] Notably, survey participants on the Prolific platform must also actively choose to take my survey when it is offered to them. To assuage concerns about selection based on interest in survey material, the survey recruitment script was left in general terms.

The practices I adhere to yield a sample which is largely representative of the U.S. across

10. Note that randomizing incentives can in part aid with this [Dutz et al., 2025]. However, the nature of online survey platforms makes it difficult to offer individuals different base pay rates for the same survey during recruitment.

key demographics. Table 1.1 depicts observable demographic variables of my sample and compares them to the general U.S. population. The sample is representative along several key demographic variables, including age, sex, race, migrant status, and political leanings. However, the sample overrepresents those with at least a Bachelor's degree from college and underrepresents low-earning households. I present results In Appendix Table A.4 which are re-weighted by demographics and show that this does not alter my estimates.

	N	Mean	SD	P25	P75	U.S. Mean
Age	2,963	40.0	12.9	30	48	38.9
Male	2,963	0.50	0.50	0	1	0.50
White	2,963	0.65	0.48	0	1	0.61
Democrat	2,963	0.27	0.44	0	1	0.28
Republican	2,963	0.26	0.44	0	1	0.28
Born in U.S.	2,963	0.89	0.31	1	1	0.86
College Degree	2,963	0.59	0.49	0	1	0.38
HH Income >$25k	2,963	0.90	0.31	1	1	0.82
HH Income >$50k	2,963	0.71	0.45	0	1	0.62
HH Income >$100k	2,963	0.35	0.48	0	1	0.33

Table 1.1: Demographics in the Sample Compared to the U.S. Population

Notes: Table shows participant demographics summarized by their mean, standard deviation, 25^{th} percentile, and 75^{th} percentile. Observations missing data for any row are omitted. The column on the far right depicts U.S. averages for each variable. The U.S. averages for the variables Democrat and Republican are estimated from a 2022 Gallup poll. The U.S. averages for the HH Income variables come from World Economics Forum. All other U.S. average data comes from the U.S. Census Bureau and American Community Survey.

In addition to representativity, I follow practices which enhance effort among my sample. First, I make a large portion of the pay based on actions in the survey. This incentivizes participants to put in more effort and pay attention during the survey, as not doing so could lower their pay. Second, I include multiple attention checks and a RECAPTCHA exercise within the survey, meant to catch individuals who are not paying close attention as they move through the survey. I remove a total of 218 participants that fail (7.0%) one of these checks. The end result is a sample of individuals that were putting in careful effort throughout the

survey.

1.3.2 Measuring Commonness: Perceived Market Penetrations

I measure commonness by collecting individual perceptions of market penetrations for carbon offsets. These perceptions are shocked with random signals to create exogenous variation, allowing for an analysis of social rewards for environmental goods at various levels of commonness.

To ensure accurate perceptions are collected, I address two major concerns about carbon offsets. The first concern is that people are generally unaware of what carbon offsets are, which I handle by providing participants a brief page of information about carbon offsets.[11] The second concern is that consumers do not believe that carbon offsets effectively mitigate carbon emissions [Aspelund and Russo, 2024, Calel et al., 2025], which if true would make offsets an inappropriate setting to study social rewards. I address this by offering carbon offsets from 501(c)(3) non-profit Climate Vault, a vendor which mechanically decreases carbon emissions by buying cap-and-trade permits in order to prevent firms from using them. After the RCT, I ask participants whether they trust carbon offsets, finding that only 15% of individuals report having strong feelings that carbon offsets are ineffective.

I then collect perceived market penetrations as my measure of commonness. I first ask people to self-report if they have previously bought a carbon offset.[12] I in turn use this information to incentivize participants to correctly guess what proportion of people taking the survey have reported previously buying a carbon offset. This is a measure of how common each individual thinks carbon offsets are, which I thus refer to as a "first-order belief". Finally, I ask each participant to guess what the average guess on carbon offset

11. This includes a screenshot of a major airline company offering carbon offsets bundled with the airline tickets. This is included to capture that offsets are often sold to consumers bundled with other goods.

12. I find that 23% of participants have previously bought a carbon offset and 28% know someone who has.

market penetrations was. This is asking participants to provide their beliefs about others' first-order beliefs, and so I refer to it as a "second-order belief." Both first- and second-order beliefs may be important in determining demand for carbon offsets, but the social rewards mechanism is governed solely by second-order beliefs.[13]

To create exogenous variation in commonness, I shock perceived market penetrations with randomly assigned signals.[14] I induce exogenous variation in second-order beliefs by randomly sending signals about the correct answer which are constructed from previously collected survey data as in Akesson et al. [2022].[15] This previous survey data collected responses from 1,000 individuals, each of whom provided their first-order beliefs about carbon offset market penetrations. Using these 1,000 first-order beliefs, I construct 6 signals by taking the average guess of two participants each, yielding signals of 1.4%, 5.6%, 12.8%, 21.7%, 31.5%, and 43.8%. I then present these signals to current survey participants as: "We previously surveyed 1,000 United States adults. In this previous survey, we found an average guess of (Signal) out of 1,000 people buying a carbon offset."[16] After showing these signals, I collect participants' posterior second-order beliefs.

My procedure produces a wide support of perceived market penetrations. Figure 1.2 depicts histograms for participants' first-, prior second-, and posterior second-order beliefs. As can be seen, there is a wide dispersion of beliefs around the correct answers which are depicted with vertical black lines. First-order beliefs are heavily skewed right, with support between 0% and 80%. Prior second-order beliefs are still skewed right, but to a lessened

13. For instance, a person driving a Tesla will only reap social rewards from driving the Tesla if they think that everyone else views the Tesla positively. The individual's own first-order beliefs would instead govern private rewards of driving the Tesla.

14. For instance, endogeneity is a concern if liberal people think that more people buy carbon offsets and are also more likely to buy offsets themselves.

15. To avoid confusing survey participants and as they are less important for social rewards, no such treatments are used for first-order beliefs.

16. This message is phrased so as it is not deceptive. In addition, the stakes are sufficiently low that I opted against debriefing participants after the study. Debriefing participants would have corrected their beliefs, but may have had negative impacts to their participation in future studies.

extent. Posterior second-order beliefs depict a significant amount of updating, as peaks shift towards the signals being sent to participants. I find on average that participants shift 71% of the way between their prior second-order belief and the signal they saw. This indicates that the signals induced a non-trivial amount of exogenous variation in second-order beliefs.

I show that first- and second-order beliefs about market penetrations do not perfectly coincide, which is needed to estimate social rewards – driven by second-order beliefs – separately from other mechanisms governed by first-order beliefs. Figure 1.3 provides a scatterplot with participants' prior first-order beliefs on the x-axis and prior second-order beliefs on the y-axis.[17] If participants' first- and second-order beliefs coincide, then this figure should depict points which are tightly following the red 45° line. I show in the figure that there is a significant amount of dispersion around this 45° line. In fact, a locally estimated regression lies close to this line, but overall indicates that second-order beliefs tend to exceed first-order beliefs across most of the data's support. This implies that consumers tend to think that others overestimate how common carbon offsets are relative to the truth.[18]

1.3.3 The Decision to Buy

To measure demand for carbon offsets, I offer participants the opportunity to buy a carbon offset of a random price. Some individuals have visible purchases, while others have invisible purchases, allowing me to estimate social rewards.

I present participants with a decision between receiving a full bonus payment or instead allocating part of their bonus payment to go towards purchasing a carbon offset. Only consumers with a WTP above the price threshold will opt for the offset instead of receiving

17. Appendix Figure A.6 plots prior second-order beliefs against posterior second-order beliefs, divided by signal received. It shows that respondents largely trusted the signals they were sent and updated beliefs accordingly.

18. This may suggest that the social rewards people believe they should get do not match with the social rewards that they do get, on average. However, dispersions in beliefs are likely smaller for goods which are naturally visible to the public (e.g. EVs and solar panels) and thus allow individuals to form more accurate beliefs.

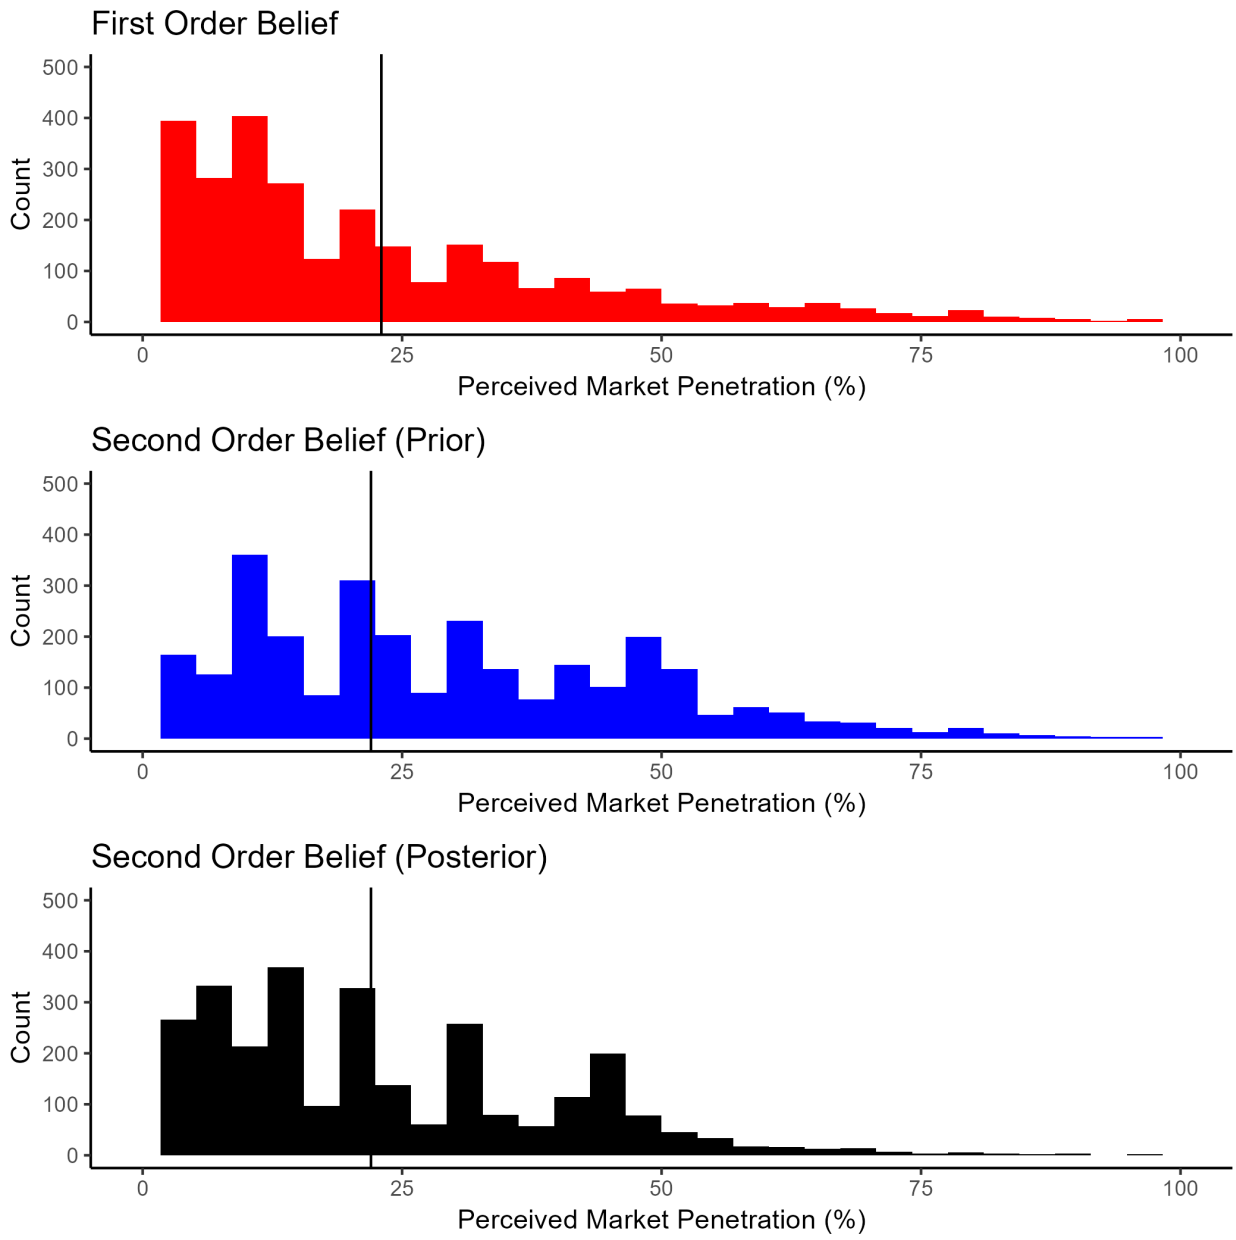


Figure 1.2: Histograms of Perceived Market Penetrations

Notes: Figure depicts histograms of first-order (top), prior second-order (middle), and posterior second-order (bottom) beliefs within the sample. All individuals failing attention checks or reCAPTCHA are dropped from this figure. First-order beliefs are an individual's beliefs on what percent of people have bought a carbon offset. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are depicted before individuals receive a signal (prior) and after they see a signal (posterior). Vertical black lines depict correct answers for each respective graph.

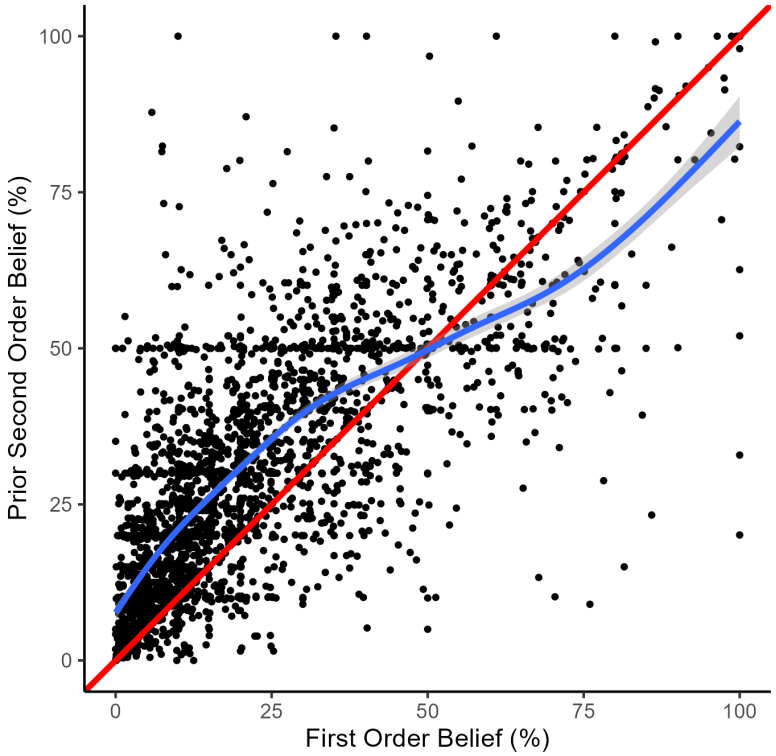


Figure 1.3: Scatterplot of Perceived Market Penetrations

Notes: Figure depicts a scatterplot of each individual's first-order and prior (i.e. before any signal was sent) second-order beliefs within the sample. First-order beliefs are an individual's beliefs on what percent of people have bought a carbon offset. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. The red 45-degree line indicates points at which first- and second-order beliefs perfectly coincide. The blue line depicts a local linear regression with corresponding 90% confidence interval shaded.

the bonus payment in full.[19] This exercise thus measures demand for carbon mitigation. The carbon offset in question offsets 145 lbs. of carbon, which is roughly equivalent to the average carbon released per individual on a flight from New York City to Chicago. The market price of such an offset is $1.50.

I induce two sources of experimental variation in this decision in order to allow for demand estimation in visible and invisible cases. First, the prices are randomly assigned with equal probability to be $0.10, $0.75, or $1.40.[20] Second, visibility is varied by randomly assigning half of participants to be in the "Visible" group and half in the "Invisible" group. The visible group is told that they will have the opportunity to have their name posted on a website if they choose to buy the offset and provide their name.[21] Participants are also told that this website is highly visible, being promoted on an Instagram account with thousands of followers.[22] The "Invisible" group does not receive this option to have their purchase visibly posted.[23]

19. This is true insofar as there is no "endowment effect" driving differences between decisions with money in the survey and the participants' own money [Thaler, 1980]. Endowment effects may increase donation rates in my RCT, but likely do not drive the shape of the social rewards function because endowment effects are not a function of perceived market penetrations.

20. Appendix Figure A.8 depicts bar graphs of the proportion of consumers who chose to buy the offset by price. As expected, higher prices lead to lower uptake of the carbon offset. As can be seen, approximately 55%, 29%, and 20% of individuals across all treatments buy the offset at prices of $0.10, $0.75, and $1.40, respectively.

21. Individuals are of course always able to self-promote their purchases. I designed this visibility treatment to enhance self-promotion by increasing credibility, reducing effort required, and removing concerns that people would be viewed as 'virtue signaling.'

22. The website was linked in the survey and is available here: `https://www.onset-of-offsets.com/`. The Instagram page is available here: `https://www.instagram.com/onsetofoffsets/`

23. If participants in the visible group choose not to provide their name, they are told that they will be listed as an anonymous donor. Only 16% of the sample chose to provide no name at all. However, I show in Section 4 that privacy concerns to some extent do reduce the impact of visibility on uptake in this setting.

1.4 Reduced Form Analysis

To examine the evolution of social rewards, I present parametric and nonparametric regression estimates. These regressions confirm model predictions, supporting that the model is a good fit for the data. They also provide estimates of social rewards that are high for uncommon goods and lower for common goods. I conclude by showing that virtue signaling, quality signaling, and referral effects are not driving my results.

1.4.1 Empirical Specification

My preferred reduced form specification is a linear probability regression, which models the probability of individual i buying an offset a_i:

$$Pr(a_i = 1) = \beta_0 + \beta_1 \mathrm{P}_i + \beta_2 f(\mathrm{SB}_i) + \beta_3 \mathrm{Pub}_i + \beta_4 \mathrm{P}_i \mathrm{Pub}_i + \beta_5 f(\mathrm{SB}_i)\mathrm{Pub}_i + \alpha X_i$$

in which P_i is price, SB_i is second-order beliefs about market penetration, Pub_i is an indicator variable for being in the visible treatment, and X_i is a vector of controls. I include controls for political views, income, education, self-reported eco-friendliness, race, sex, age, previous offset purchases, and a quadratic polynomial of first-order beliefs.

The form of $f(y)$ is assumed to be quartic to allow for social rewards which flexibly vary over perceived market penetrations. This assumption plays a large role in determining the estimated shape of the social rewards function. Introducing higher order terms may induce noisy estimates, particularly on the edges of the dataset's support. On the other hand, using only low order polynomials may understate key patterns. Using $f(y) := \gamma_1 y + \gamma_2 y^2 + \gamma_3 y^3 + \gamma_4 y^4$ allows for fine nonlinearity in reduced form estimates of the social rewards function while balancing concerns of noisy estimates.

To handle endogeneity in perceived market penetrations, I instrument for beliefs SB_i with

the random signals individuals were sent. This yields the following first stage regression[24]:

$$\text{SB}_i = \eta_1 \text{Signal}_i + \eta_2 X_i + \varepsilon_i$$

As discussed in Murphy and Topel [1985] and Wooldridge [2010], two-staged least squares provides incorrect standard errors when applied to this regression. I address this by presenting bootstrapped standard errors and confidence intervals in all that follows.[25] Usage of a linear probability model likewise avoids issues of running 'forbidden regressions' and generates consistent estimates of average treatment effects scaled by the first-stage coefficient, akin to LATE [Li et al., 2022].

1.4.2 Demand for Carbon Offsets

Estimates in Table 1.2 suggest that distrust in researchers, privacy concerns, and political leanings all play some role in governing desirability of the social rewards in my RCT. Column (1) depicts estimates for the overall sample.[26] Column (2) subsets to the sample which reported trusting researchers that the website posting would be visible. Column (3) further subsets to the sample which did not report having privacy concerns from the online name posting. Column (4) additionally subsets to those in the sample who did not identify as being politically very conservative. My preferred specification is in column (3), which I argue best reflects demand for visible goods in real world settings.[27]

24. The exclusion restriction is that my signals about other peoples' beliefs only impact demand for carbon offsets through their impacts on peoples' own second-order beliefs.

25. Another valid approach is to perform a Murphy-Topel correction. However, my endogenous variable is included in the second stage (i) in a quartic polynomial and (ii) interacted with Pub_i. This complex inclusion of the endogenous variable in the second stage motivates my usage of bootstrapped standard errors, which automatically account for this complexity.

26. I survey 3,111 individuals. I drop 218 who fail attention checks, 73 with outlier beliefs, and 35 with missing demographic data. This yields $N = 2,685$.

27. Omitting researcher trust reflects that researchers are not present in consumers' real-world decisions. Similarly, privacy concerns are limited in real-world signaling of e.g. EVs in which name anonymity is maintained.

	(1)	(2)	(3)	(4)
Public	0.209*	0.249**	0.292**	0.308**
	(0.107)	(0.114)	(0.125)	(0.132)
SB X Public	-0.0688***	-0.0716**	-0.0744**	-0.0855**
	(0.0267)	(0.0283)	(0.0312)	(0.0333)
$(SB)^2$ X Public	0.00506**	0.00512**	0.00550**	0.00676**
	(0.00213)	(0.00224)	(0.00245)	(0.00266)
$(SB)^3$ X Public	-0.000134**	-0.000133*	-0.000149**	-0.000192**
	(0.0000654)	(0.0000684)	(0.0000741)	(0.0000816)
$(SB)^4$ X Public	0.00000120*	0.00000115	0.00000134*	0.00000181**
	(0.000000672)	(0.000000699)	(0.000000752)	(0.000000840)
Price	-0.280***	-0.280***	-0.281***	-0.292***
	(0.0230)	(0.0230)	(0.0231)	(0.0239)
Price X Public	-0.00760	-0.0180	-0.0398	-0.0205
	(0.0324)	(0.0335)	(0.0355)	(0.0371)
Drop Visibility Doubters?	No	Yes	Yes	Yes
Drop Those With Privacy Concerns?	No	No	Yes	Yes
Drop Far Right?	No	No	No	Yes
Observations	2685	2525	2305	2107
Dep. Var. Mean	0.34	0.35	0.36	0.37

Table 1.2: Reduced Form Estimates of Visibility's Impact by Market Penetration

Notes: Table shows how the impacts of visibility on uptake vary by second-order beliefs about market penetration of carbon offsets. Estimates are generated by a two-stage linear probability model regression. In the first stage, second-order beliefs are instrumented with a randomly-generated signal. Column (1) depicts estimates for the full sample. Column (2) depicts estimates after dropping individuals who did not believe the website treatment was actually publicly visible. Column (3) depicts estimates after additionally dropping those who reported having strong privacy concerns about their name being posted. Column (4) depicts estimates additionally dropping those reported being politically very conservative. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Standard errors are depicted in parentheses.
$^*\ p < 0.10,\ ^{**}\ p < 0.05,\ ^{***}\ p < 0.01.$

The role of visibility seems to depend on individuals' second-order beliefs, as depicted in Figure 1.4. Figure 1.4 depicts demand curves estimated from column (3) of Table 1.2, separately presented for individuals with second-order beliefs of 30% and 0%. Visibility shifts demand outwards in both cases, but the shift is much larger when second-order beliefs are 0%, indicating that social rewards are higher when perceived market penetrations are lower. The demand curves also indicate that (i) demand decreases in price and (ii) demand increases in visibility, matching predictions 1 and 2 from the model.

Importantly, Appendix Table A.5 displays the signal I sent was a relevant instrument. The coefficient on Signal is significant at all conventional levels, with a corresponding t-statistic of 42.6. The estimated coefficient on Signal of 0.636 indicates that an increase to the signal of 1 led individuals to increase their beliefs by 0.636. Overall, instrument relevance appears to hold as the smallest F statistic is 84.41, far exceeding standard cut-offs for weak instruments (e.g. 10).

1.4.3 Impact of Social Rewards

The regression specification allows me to estimate social rewards within column (3) of Table 1.2. Here, visibility is estimated to increase uptake by 29.2 percentage points when second-order beliefs on market penetration are equal to 0. The effect varies as second-order beliefs increase. In particular, coefficients on all higher-order terms of the polynomial significantly differ from 0.

The linear probability model regression estimates highlight that social rewards are large for uncommon goods and quickly decrease as goods become more common, as depicted in Figure 1.5. At the smallest market penetration, social rewards increase uptake by 29.2 percentage points. This effect decreases sharply and remains below 10 percentage points at perceived market penetrations above 5%. Despite the presence of cubic and quartic terms in the linear probability regression, the effect of visibility largely appears similar to the left half

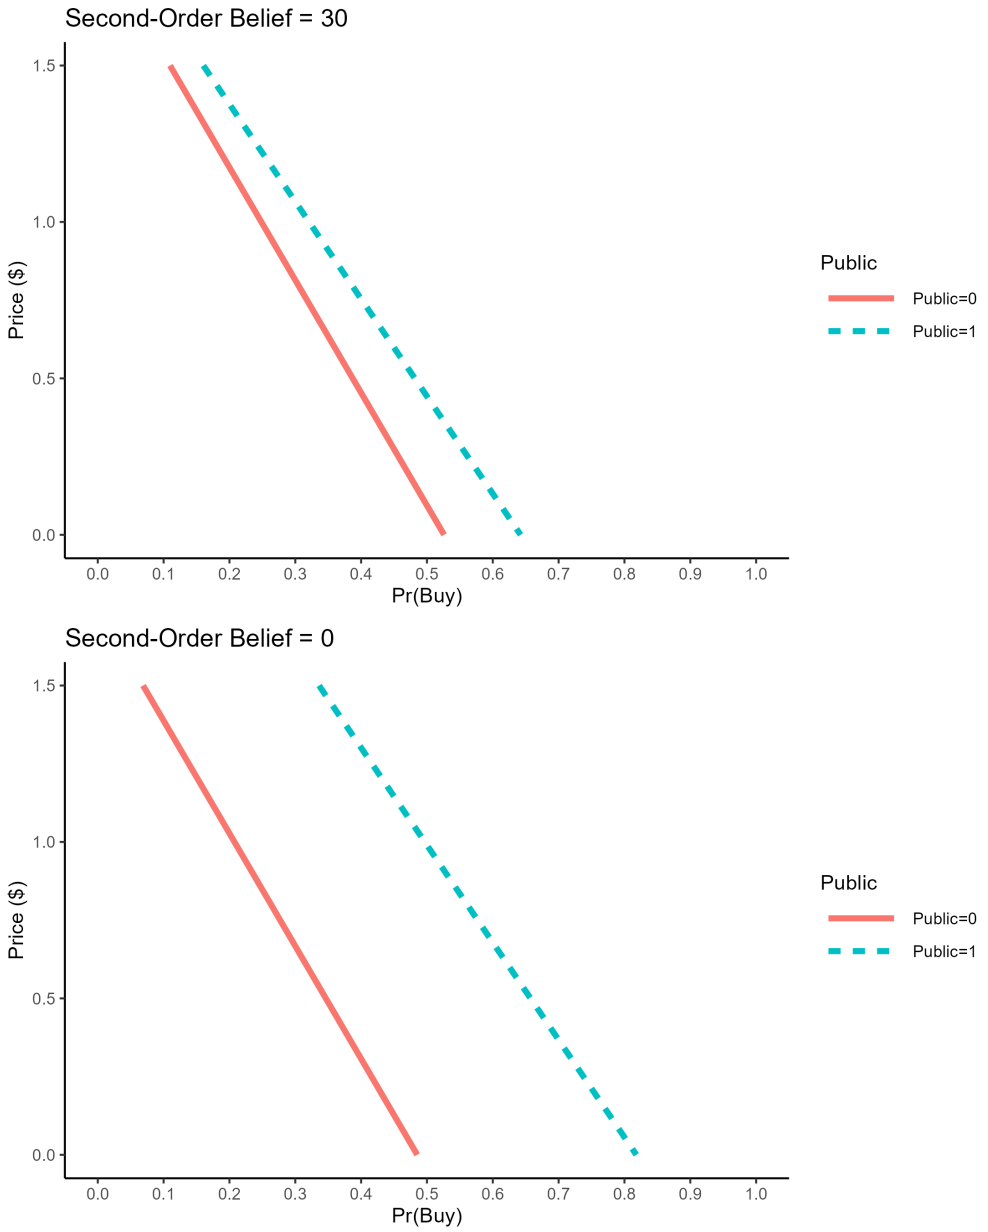


Figure 1.4: Reduced Form Estimated Impact of Visibility on Demand

Notes: Figure depicts demand curves when second-order beliefs are 30% (top) and 0% (bottom). First-order beliefs are an individual's beliefs on what percent of people have bought a carbon offset. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Blue dashed lines are in the visible case and red solid lines are in the private case. Predictions are generated by setting all control variables to their means if they are numerical (e.g. age) or their modes if they are categorical (e.g. political views). Second-order beliefs are instrumented for with a randomly sent signal.

of a quadratic function.[28] This mirrors the behavior of social rewards when the underlying distribution of consumer types w is normal.[29]

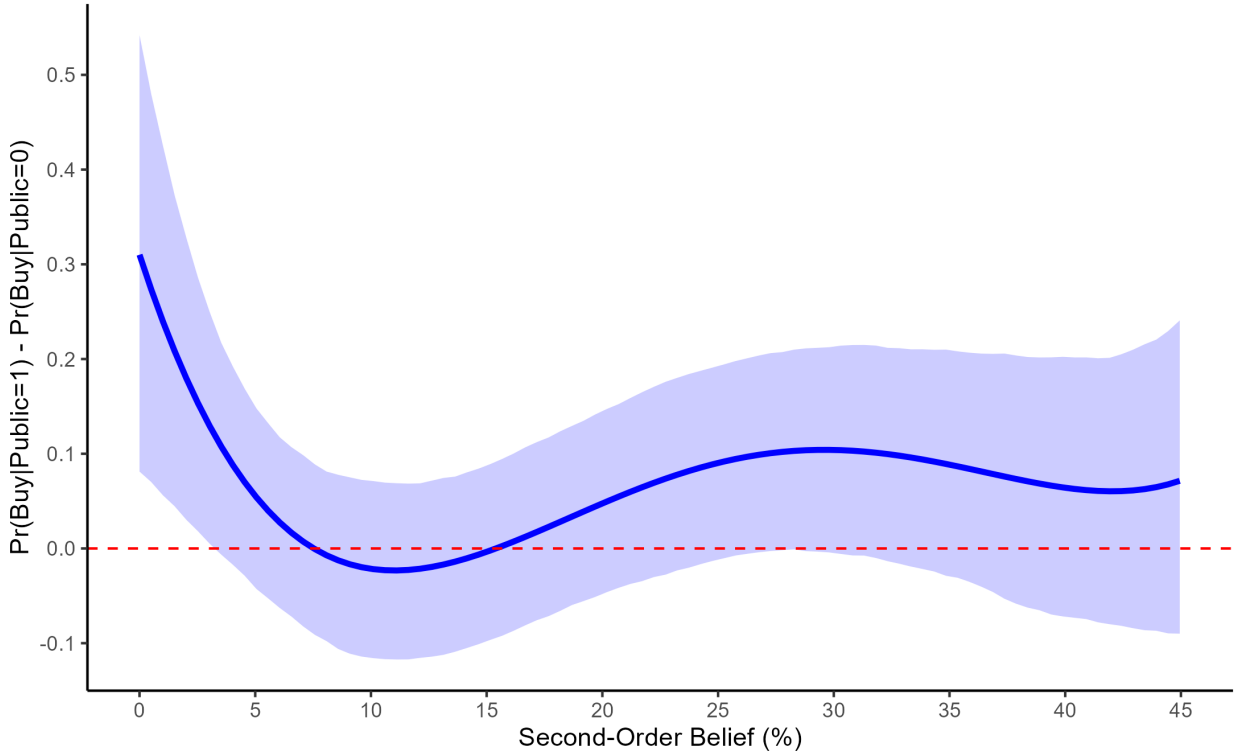


Figure 1.5: Reduced Form Impact of Visibility by Market Penetration

Notes: Figure depicts linear probability model regression estimates of uptake of offsets when randomly visible by second-order beliefs of market penetration of carbon offsets. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

Crucially, the shape of the estimated social rewards function in Figure 1.5 is robust across

28. While social rewards are only present at small market penetrations, these low levels of market penetrations can often exist for years. For instance, a 2023 Pew Research survey found that only 7% of homeowners have solar panels, despite residential solar being available since the 1980s.

29. In Appendix Figure A.9, I depict the probabilities of uptake separately by visible treatment. The variation in the difference between probabilities seems to be largely driven by variation in uptake probability for the visible group, rather than the control group.

multiple additional reduced form estimation strategies. In Appendix Figure A.15, I present multiple nonparametric results that confirm the robustness of social rewards estimates. This appendix figure provides local linear and binned regression estimates supporting my findings that social rewards are large at low market penetrations and quickly dissipate. Appendix Figure A.16 in particular provides a direct statistical comparison of social rewards from 0%-3% versus 3%-50%, showing that their estimates differ ($p < 0.05$). Taken together, these figures suggest that my findings are robust to using alternative estimation strategies.

The presence of social rewards which are large at small market penetrations and quickly decrease as a quadratic is akin to Panel A of Figure 1.1, supporting a normality assumption later used to structurally estimate the model. To statistically test this normality assumption, I perform a bootstrapped likelihood ratio test in which the 'true' specification is assumed to set $(\text{SB})^3 \times \text{Public}$ and $(\text{SB})^4 \times \text{Public}$ equal to 0. This restricted specification is used to generate data for a bootstrap in which each iteration yields a likelihood ratio comparing the full specification to the reduced specification. A p value is then generated by comparing the likelihood ratio observed in my data to the likelihood ratios generated by each bootstrap iteration. This process yields a p value of 0.15, supporting the notion that only the quadratic terms – corresponding to the normal distribution – are important.

Social Rewards and Demand Elasticity

I show that column (3) of Table 1.2 also aligns with Prediction 3 generated by the model. I estimate values of -0.281 and -0.321 for the price sensitivities in the private and visible groups: β_1 and $\beta_1 + \beta_4$. Prediction 3 stated that price sensitivity should be closer to 0 when social rewards slope downwards. Table 1.2 depicts that the interaction between Price and Public is not statistically significant, and so Prediction 3 would suggest that the social rewards function should either be dominated by regions that are flat or have both upwards-sloping and downwards-sloping regions. In Figure 1.5, I show that the estimated social rewards

slopes upwards between beliefs of 10 and 30 – roughly half of the support – and downwards elsewhere, which is why the average price sensitivities do not vary by public treatment.

1.4.4 Additional Mechanisms

To establish that social rewards drive my estimates, I discuss competing mechanisms. The impacts of referral effects, quality signaling, and virtue signaling are shown to be small.

Referral Effects

One alternative explanation to social rewards is that visibility increases demand for environmental goods through a referral effect. Participants know that many individuals in the U.S. are not aware of carbon offsets. Thus, they may see visibility as a chance to increase the effect of their own purchase by encouraging others in the U.S. to also buy offsets. This would lead to visibility increasing demand for carbon offsets in my RCT by causing individuals to think that their purchase will have a greater overall impact on the environment, rather than through yielding social rewards. The presence of this mechanism could explain why visibility's impact decreases in perceived market penetration if, for instance, individuals do not feel that the referral effect is as large when more people are already purchasing carbon offsets.

Evidence suggests that referral effects do not drive demand in my setting for two reasons. First, individuals at the end of the survey were directly asked whether they felt their name being posted would encourage others to donate. I find that only approximately 26% agreed that the posting of their name would potentially lead others to donate. I present results in Appendix Figure A.12 which estimate social rewards using only individuals in the treated group that report being motivated by social rewards but not by referrals. The results are largely unchanged in this figure compared to my main results.

The second reason to believe referral effects are unimportant is that the impact of public-

ity does not vary by *first-order* beliefs, but first-order beliefs are what govern referral effects. This is because a person's own beliefs about market penetration determine how much they think their purchase would spread knowledge of carbon offsets. All results presented in this section have controlled for first-order beliefs, ensuring that my estimates capture social rewards net of any referral effects. In Appendix Figure A.10, I examine the importance of referral effects by plotting the impact of visibility across first-order beliefs.[30] The plot depicts minimal variation in impacts of visibility across first-order beliefs, which seems to indicate that referrals are not driving consumer behavior in my setting.

Quality Signals

Another important mechanism which may impact demand is signal of quality. Consumer demand may increase in market penetration if goods with a higher market penetration are seen as higher quality than goods with a low market penetration. This could in turn interact with the visibility treatment if, for instance, individual consumers like to be seen consuming high quality goods and dislike being seen consuming low quality goods.

Evidence suggests that quality signals are of limited importance in my setting for two reasons. First, I directly asked individuals whether they had positive or negative feelings on carbon offsets at the end of the study. Conditioning on demographics, responses to this questions are positively correlated with *first-order* beliefs but uncorrelated with second-order beliefs. As discussed in the previous subsection and depicted in Appendix Figure A.12, first-order beliefs do not drive the impacts from visibility in my setting. Perceived quality is likewise uncorrelated with signals that individuals were assigned, as depicted in Appendix Figure A.7.

The second reason is that I observe that first-order beliefs of carbon offset market penetrations do not impact the probability of buying an offset in the invisible group. This is

30. Note that there is no instrument for first-order beliefs, though I include many demographic controls which in part mitigate endogeneity concerns.

depicted in Appendix Figure A.11, which plots the impacts of first-order beliefs on purchase probability by visibility treatment. If large market penetrations were associated with higher quality and this in turn impacted demand, it would then be expected that purchase probability slopes upwards across first-order beliefs. Appendix Figure A.11 shows no such upwards slope for either the visible or invisible group. Perceived market penetrations appear to be correlated with perceived quality, but this correlation does not seem to impact demand nor the magnitude of visibility as a driver in demand.

Virtue Signaling Backlash

A final mechanism which could play a role in my estimates of social rewards is concern over backlash to being seen as 'virtue signaling.' Virtue signaling is defined as the public expression of one's sentiments in order to demonstrate one's character. Oftentimes, 'virtue signaling' is used as a dismissive term to argue that an individual who signals prosociality is not actually virtuous. If participants are worried that their name being posted online will be dismissed by others, it may depress my estimated values of social rewards.

I show that virtue signaling plays little role in my setting. To measure the importance of concerns over being seen as virtue signaling, I directly ask individuals in the survey whether concerns about being seen as a virtue signaler played a role in their decision. Overall, I find that only 10% of the sample reported having some concerns about virtue signaling. To test whether virtue signaling significantly depresses my estimates, I present estimates in Appendix Figure A.13 that measure social rewards using only individuals in the treated group that report not having concerns about being seen as a virtue signaler. Overall, results are essentially unchanged when compared to Figure 1.5, indicating that concerns over being seen as a virtue signaler played only a minimal role in the estimation of social rewards.

1.5 Structural Estimation

Structural estimates which complement the reduced form analysis are generated using maximum likelihood estimation in order to calculate optimal dynamic subsidies. Structural estimates of willingness-to-pay for social rewards are high for uncommon goods and small for common goods, mirroring reduced form results. As a consequence, my estimates suggest that subsidies should slope upwards in market penetration, contradicting conventional knowledge that suggests the opposite.

1.5.1 The Value Added by Structural Estimation

The structural estimation in this section imposes additional assumptions on the data, but enables optimal dynamic subsidies to be calculated.[31] More specifically, the structure of this model enables calculations of Ramsey optimal subsidies, whereas reduced form estimates can only be used to calculate Pigouvian subsidies. This is because estimating Ramsey subsidies requires knowledge of what proportion of individuals are marginal versus inframarginal to the subsidy, which depends on the distribution of individuals' types w. As policymakers weigh the benefits of increasing demand for environmental goods against the costs of raising subsidy funds, understanding how Ramsey subsidies compare to Pigouvian subsidies is useful.

The calculation of how social rewards influence optimal dynamic subsidies is important, as previous dynamic subsidies have not accounted for social rewards. For logistical reasons, policymakers frequently must commit to subsidies in advance of administering them. Doing so in a way which maximizes welfare requires considering how the costs and benefits of such subsidies change as time passes and a good becomes more common. Past dynamic subsidies for environmental goods – including Federal EV subsidies, the California Solar Initiative,

31. An additional benefit of the model is that it allows me to project how the WTP for social rewards varies in settings which have differing levels of visibility and of idiosyncratic preferences when compared to my setting. For instance, purchases of EVs are likely more visible (higher x) and more idiosyncratic (larger σ_u) than the purchases of offsets in my setting are. I examine EV subsidies in the Appendix.

the Solar Investment Tax Credit, and many state renewable energy initiatives – have been set to slope downwards over time. This downwards slope accounts for technological growth, but fails to account for the reduction in consumer WTP as social rewards dissipate, which I show puts pressure on such subsidies to slope upwards as goods become more common.

1.5.2 Estimation Procedure

The main objective is to estimate social rewards $\Delta(w^*)$, which is a function of the cutoff type w^* determined in the equilibrium fixed-point equation[32]

$$w^* + \lambda x_i \Delta(w^*) = \beta_p p$$

In order to structurally estimate this model, it is necessary to make a distributional assumption on individuals' underlying types w. This assumption allows the writing of a likelihood function and provides an analytic form of $\Delta(w^*)$. This assumption is not trivial, as it plays a large role in determining the shape of the estimated social rewards function. In Section 4, I showed that assuming a normal distribution of w is reasonable.

Assuming that $w_i = u_i + v_i$ is distributed independently and identically as $Normal(0, \sigma_w^2)$, $\Delta(w)$ can be written as

$$\Delta(w) = \sqrt{\frac{\sigma_w^2}{\sigma_w^2 + 1}} \times \frac{\exp(-\frac{1}{2}\frac{w^2}{\sigma_w^2+1})}{\sqrt{2\pi}\sigma_w} \times \frac{1}{\Phi\left(\frac{w}{\sqrt{\sigma_w^2+=1}}\right)\left[1 - \Phi\left(\frac{w}{\sqrt{\sigma_w^2+=1}}\right)\right]}$$

in which $\Phi(*)$ is the standard normal CDF.[33] The normality assumption thus allows me to

32. Individuals in the RCT are not told that they are receiving a subsidized price for the offset, so y is dropped from these equations.

33. Note that the distribution of w has been re-centered to 0.

estimate the parameters $\{\lambda, \beta_p, \sigma_w\}$ by maximizing the likelihood

$$\prod_i \Phi(w_i^*)^{1-\text{Buy}_i} \left(1 - \Phi(w_i^*)\right)^{\text{Buy}_i}$$

subject to the fixed point equation $w_i^* + \lambda x_i \Delta(w_i^*) = \beta_p p_i$.

The sources of exogenous variation observed in my experiment allow for estimation of the model's parameters under likelihood maximization. For each participant, I observe their exogenously assigned price p_i and binary visibility treatment assignment x_i. I additionally see each individual's perceived market penetration, PMP_i, which can be inverted to an individual's perceived w_i^* using the inverse normal CDF.[34] To deal with endogeneity of beliefs, I instrument for PMP_i using exogenously assigned signals.

1.5.3 Willingness-to-pay for Social Rewards

Figure 1.6 plots the estimated WTP for social rewards across market penetrations, depicting that social rewards are again important for uncommon goods and less important for more common goods.[35] As can be seen, the WTP for visibility is initially large, reaching the same \$1 threshold estimated in Section 4. This represents a roughly 87% increase in WTP for carbon mitigation on a per-ton basis.[36] After initially being large, this WTP drops quickly, reaching under \$0.15 (a 13% per-ton increase) by the time a perceived market penetration of 12% is met. The model thus estimates that social rewards are concentrated primarily within

34. i.e. $w_i^* = \Phi^{-1}(\frac{PMP_i}{\sigma_w})$

35. Figure A.17 plots structural estimates of demand for carbon mitigation across the visible and private cases under the assumption that consumers hold accurate beliefs about market penetrations.

36. Specifically, WTP per metric ton of carbon mitigated is estimated to be \$18 in (Rodemeier, 2024). For 145lbs. of carbon, this scales to $\frac{145}{2204} \times 18 = \$1.18$. I estimate consumers are WTP \$1.03 for social rewards, roughly 87% of \$1.18.

uncommon goods that have very low market penetrations.[37]

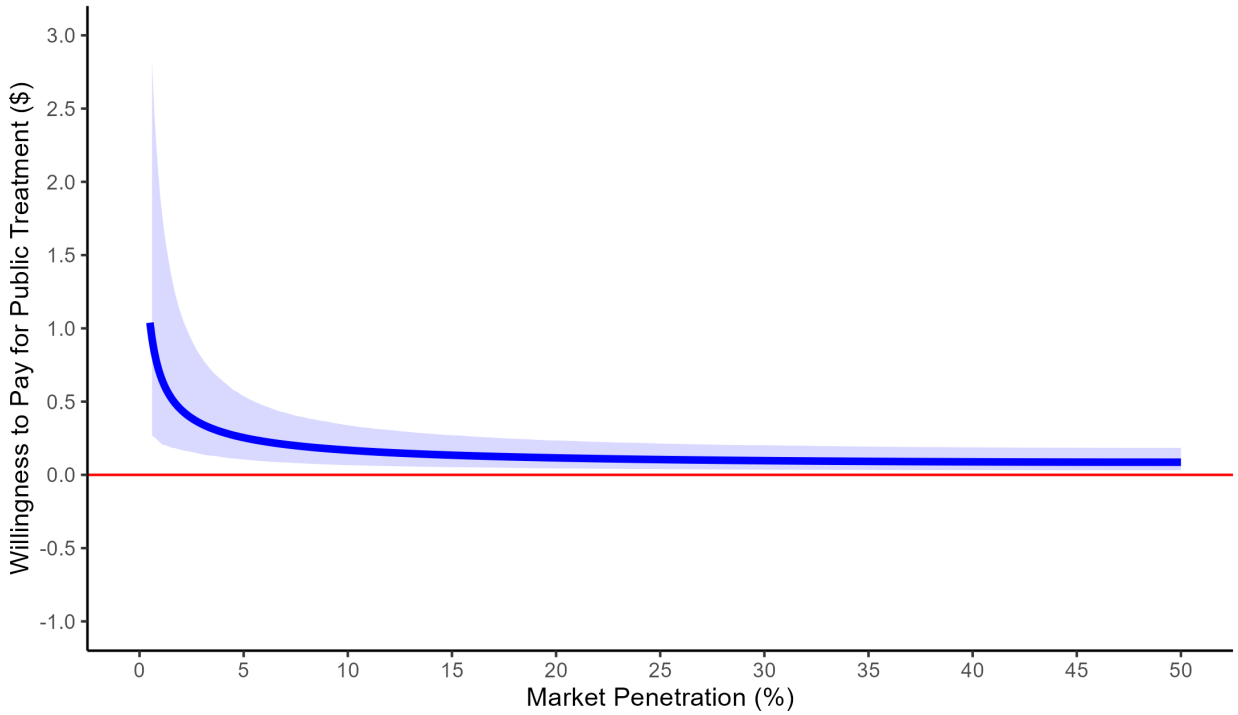


Figure 1.6: Structurally Estimated WTP for Visibility by Market Penetration

Notes: Figure depicts structural estimates of the willingness-to-pay (WTP) for visibility from purchasing an offset by second-order belief of market penetration of carbon offsets. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

1.5.4 Optimal Dynamic Subsidies Under Social Rewards

I consider a social planner who is setting a dynamic subsidy for a publicly visible environmental good in order to maximize social welfare. In maximizing social welfare, the social

37. Appendix Figures A.18 through A.19 plot this same graph except for columns (1), (3), and (4) from Table 1.2, respectively.

planner will consider individual utility

$$U = [w + \beta_p(y - p)]a + \mu E[v|a]$$

The planner will also consider the externality (in utils) $e \geq 0$ from each good as well as the shadow value of raising and administering subsidy funds $\eta \geq 0$. As such, social welfare has the form

$$SW(y) = \bar{U} + eN - (1 + \eta)\beta_p y N$$

in which N is the number of individuals buying the good. The number of individuals buying the good is dependent on the proportion of individuals whose types exceed the cutoff type $w^*(y)$. As such, the planner solves the following problem[38]

$$\begin{aligned} y^* &= \text{argmax}_y SW(y) \\ &= \text{argmax}_y \int_{w^*(y)}^{\infty} (e + w - \beta_p p - \eta\beta_p y) f(w) dw + \mu\bar{w} \end{aligned}$$

The social planner chooses a subsidy y which is lower when WTP for social rewards is

38. Note that $\bar{w} := (1 - F(w^*(y)))E[v|w > w^*(y)] + F(w^*(y))E[v|w < w^*(y)]$. Here, the values of honor and stigma are zero-sum across the pool of agents due to the assumption that beliefs are on average correct. The zero-sum nature of social rewards allows for social rewards to effectively be ignored in welfare calculations. This is not necessarily the case if beliefs are on average incorrect. In this case, perturbations to subsidies may be met with differing changes to beliefs among non-buyers and buyers, causing social welfare to increase or decrease.

higher.[39] In the Ramsey case in which $\eta > 0$, it can be shown that the optimal solution is

$$y = \underbrace{\frac{e}{(\eta+1)\beta_p}}_{\text{Externality}} - \underbrace{\frac{\mu\Delta(w^*)}{(\eta+1)\beta_p}}_{\text{Soc. Rewards}} - \underbrace{\frac{\eta(1+\mu\Delta'(w^*(y))[1-F(w^*(y))]}{(\eta+1)\beta_p f(w^*(y))}}_{\text{Marginal-Inframarginal Trade-off}}$$

Intuitively, the first term captures that the Ramsey subsidy increases as the monetary value of the positive externality $\frac{e}{\beta_p}$ increases. The second term captures the monetary value of social rewards, which the social planner deducts from the subsidy to avoid overcompensation of consumers. Finally, the third term captures the trade-off between subsidizing inframarginal consumers (measured by $1 - F(w^*(y))$) and marginal consumers (measured by $f(w^*(y))$). When more inframarginal consumers are present, the social planner decreases the subsidy because funds are expensive. When there are more consumers marginal to the subsidy, the subsidy is increased as the subsidy is more cost effective.

The Pigouvian subsidy ($\eta = 0$) also yields optimal dynamic subsidies that are lower when WTP for social rewards is higher. For the case in which $\eta = 0$, the above subsidy simplifies to

$$y^{Pigou} = \frac{e - \mu\Delta(w^*)}{\beta_p}$$

In the case of invisibility ($\mu = 0$), the Pigouvian subsidy further simplifies to the classic solution $y^{Pigou} = \frac{e}{\beta_p}$. When social rewards are not present the social planner chooses a Pigouvian subsidy equal to the dollar value of the externality from buying the good. When the good is visible ($\mu > 0$), the social planner instead chooses to shade the subsidy by the marginal WTP for social rewards.

Estimating Pigouvian subsidies can be done using both structural and reduced form

39. This subsidy calculation supposes that individuals hold beliefs which are on average correct about the market penetration of a good. This assumption is supported by Figure 1.2. Moreover, individuals are more likely to hold accurate beliefs for goods which are naturally publicly visible. The public visibility of goods such as EVs and solar panels likely aids individuals in estimating market penetrations.

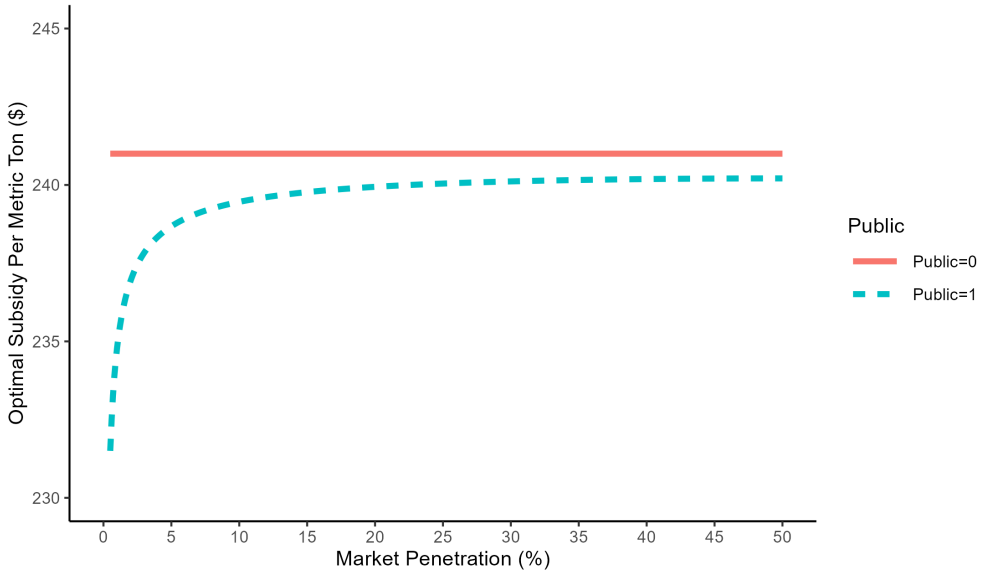
results, whereas Ramsey subsidies can only be done using structural estimation. Pigouvian subsidies require only knowledge of e and estimates of the social rewards $\mu\Delta(w^*)/\beta_p$, the latter of which can be estimated with reduced form strategies as in Figure 1.5. Ramsey subsidies, however, require knowledge of what proportion of individuals at any given market penetration are marginal and inframarginal to the subsidy. This knowledge involves imposing a distribution of private valuations of the good on individuals, as is done in the structural model.

In Figure 1.7, I plot structural estimates of the optimal dynamic subsidy per metric ton of carbon mitigated and show that they slope upwards for both Pigouvian and Ramsey subsidies. Panel A depicts Pigouvian subsidies in which the subsidy is set equal to marginal benefits at all market penetrations for invisible goods.[40] When the good is visible, the Pigouvian optimal subsidy is initially set \$10 smaller than the externality to account for the presence of large social rewards. As the good becomes more common and social rewards dissipate, the planner chooses to increase the subsidy in order to compensate individuals for the smaller social rewards. Thus, the presence of social rewards suggests that uncommon environmental goods should experience relatively small subsidies. Similar conclusions are drawn from examining Panel B of Figure 1.7, which calculates the Ramsey subsidy using $\eta = 0.25$ as listed in O'Callaghan and Prior [2018].

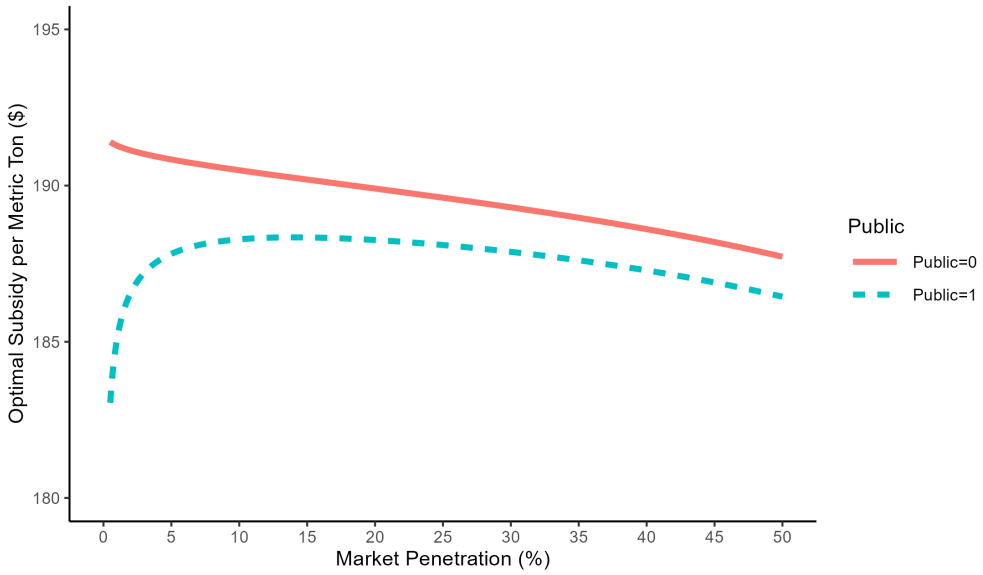
In Appendix Figure A.3, I apply these estimates to EVs to quantify the role of social rewards in governing optimal EV subsidies. Such estimates should be taken with caution because EVs represent a different context than carbon offsetting.[41] In particular, EVs may be subject to noisier demand and higher levels of visibility than the context I focus on, both of which impact social rewards as seen in Appendix Figure A.4. It should also be noted that

40. Marginal benefits are calculated here using a social cost of carbon of \$241 per metric ton.

41. It is well established that social rewards are important in high-stakes contexts, including EVs [Pallottini and Shchukina, 2025]. For my findings on the *evolution* of social rewards to generalize to these high-stakes settings is for demand in these other settings to vastly noisier without the action also being more visible.



Panel A: Pigouvian Optimal Subsidies for One Metric Ton of Carbon



Panel B: Ramsey Optimal Subsidies for One Metric Ton of Carbon

Figure 1.7: Optimal Subsidies to Mitigate Carbon by Market Penetration

Notes: Figure depicts estimates of the Pigouvian (Panel A) and Ramsey (Panel B) optimal subsidy for one metric ton of carbon mitigation by market penetration of carbon offsets, assuming individuals' second-order beliefs are on average correct. The optimal subsidy is depicted in the private case as the red solid line and in the visible case as the blue dashed line. The social cost of carbon is assumed to be $241 per metric ton [U.S. EPA, 2023]. Estimates of the shadow value of public funds ($\eta = 0.25$) are taken from O'Callaghan and Prior [2018].

my model captures only a demand-side mechanism which shifts the optimal dynamic subsidy path. In particular, my model of consumers does not capture supply-side mechanisms such as learning-by-doing, innovation spillovers, and other sources of technology growth that support downwards sloping subsidies. Finding the overall balance between supply-side and demand-side mechanisms in driving these optimal dynamic subsidy paths remains a rich area for future research.

1.6 Conclusion

I have leveraged an RCT which provides three main findings in an analysis of demand for environmental goods. Using exogenous variation in price and visibility, I have first shown that consumers are willing to pay a price premium for their environmental good purchases to be publicly visible. The presence of this premium indicates that consumers are not purely driven by price and environmental quality in their decisions to buy environmental goods. Using purchases as a signal to others ('social rewards') plays an important role in decisions to buy goods like EVs, solar panels, and heat pumps. This highlights that enhancing visibility of similar actions could be a useful tool for policymakers moving forward and explains in part why certain environmental goods (e.g. Priuses or Teslas) are designed by suppliers to visually stand out.

Second, I have shown that the WTP for social rewards varies greatly by how common the environmental good is. By creating exogenous variation in the beliefs about market penetration of a good, I estimate that WTP for social rewards is concentrated primarily among uncommon goods. This result is robust to using reduced form analysis as well as structural estimation. I estimate that visibility increases WTP by 87% for goods that have a near-zero market penetration. The value of social signaling quickly dissipates, falling to below a 15% increase in WTP at market penetrations exceeding 12%. This highlights that consumers have a desire to be seen as one of the first adopters of a environmental good.

Finally, I have used my estimates of social rewards to show that optimal dynamic subsidy paths for environmental goods should start lower at near-zero market penetrations and increase as the goods become more common. I show that the Pigouvian optimal per-ton subsidy of carbon should be set $10 lower than the SCC for near-0% market penetrations and ramp up to roughly equal the SCC at market penetrations from 10-50%. This directly contradicts supply-side research and actual extant subsidies which have favored the opposite – targeting subsidies towards uncommon goods. I believe that this study is of particular use to policymakers, as many environmental goods are transitioning from uncommon to common.

Such results enhance our knowledge around environmental goods in particular, but I believe they are of use also to other literatures. To the best of my knowledge, this is the first paper to directly analyze how WTP for social rewards for a visible good varies across market penetrations of 0% to 50%. Similar patterns may hold in other settings in which consumption is visible, particularly among other prosocial goods or status goods. This represents a starting point for a large amount of future work which could be of interest to economists in various settings.

Several other open questions may be interesting for future research. First, my paper is unable to address how social rewards vary as consumption of goods becomes near unanimous. While I am able to plot social rewards from 0 to 50% uptake, everything above 50% must be examined in future analysis. This could be particularly interesting to understand, as 'shame' may begin to play a larger role in driving visible decisions at these higher market penetrations. If shame exists as a mechanism, it would imply that social rewards are higher at near-unanimous uptakes than at moderate uptakes, which could indicate further non-linearity in optimal dynamic subsidy paths.

Second, future work could seek to understand further dimensions along which social rewards vary. In this setting, I have focused solely on one good with a fixed level of environ-

mental impact. However, the magnitude of impact likely plays a role in determining social rewards. Social rewards may be larger when an good has a more positive impact. Similarly, the cost borne by individuals to buy a good likely has implications for how they are perceived by peers. Analyzing how social rewards vary by price or societal benefit could thus unveil interesting findings.

Finally, expanding this work to field testing or observational data, while difficult, could potentially prove fruitful. I have chosen a lab setting for this RCT. While this makes generating exogenous variation in price, visibility, and commonness feasible, it comes with the usual drawbacks. Experimenter demand, Hawthorne effects, and selection bias are all potential threats to the validity of my findings. I have taken appropriate steps to assuage these concerns, focusing on constructing a high quality sample. Nevertheless, expanding on my results in a more natural setting – even at the expense of measuring strictly causal effects – would likely be of use.

CHAPTER 2

CHANGING GEARS: THE WELFARE IMPACTS OF BRAND IDEOLOGY

2.1 Introduction

Ideological polarization is increasingly pervasive in the day-to-day American life [Iyengar et al., 2019]. Ideology has been shown to influence the relationships people form, their labor market decisions, and the places they want to live [Gift and Gift, 2015, Gimpel and Hui, 2015, Chen and Rohla, 2018]. Moreover, this rise in "affective polarization" appears to be influencing individual consumption decisions. Recent ideological boycotts have targeted Amazon, American Eagle, Bud Light, Disney, General Mills, Jaguar, Nike, Pfizer, Target, Tesla, Twitter/X, and Walmart, among numerous other brands.[1]

Amid this rise in polarization, policymakers in the U.S.A. are increasingly focusing on mitigating environmental damages by encouraging consumers to buy environmental goods that may be ideologically polarized. Despite widespread subsidies at state and federal levels, goods such as electric vehicles (EVs) are far more likely to be purchased by Democrats than by Republicans [NORC, 2024, Davis et al., 2025]. It is unclear what drives Democrats to buy environmental goods more often, be it higher valuation of environmental quality or the ideological implications of buying such goods.

In this paper, we examine EVs to estimate the extent to which ideology drives demand for environmental goods and, consequently, private and social welfare. We do so by using varying perceptions of the ideology linked to Tesla via its association with CEO Elon Musk. Musk has increasingly expressed Republican viewpoints in the 2020s, leading to a wave of

. This chapter is in equal collaboration with Sofia Shchukina.

1. See for instance `https://www.ethicalconsumer.org/ethicalcampaigns/boycotts` for summaries of existing boycotts and `https://www.usatoday.com/story/money/2025/04/17/consumer-shopper-boycott-blackout-target-walmart-amazon/82776080007/` for an article discussing recent retailer boycotts.

speculation that his Republican shift is impacting Tesla's market share. However, Teslas are also among the cleanest vehicles consumers can purchase, which creates a friction for consumers choosing between ideology and environmental impact.[2]

We first use data from Reddit and Google Trends to establish that Musk's shift in ideology makes Tesla EVs an appropriate setting to study how ideology influences demand for environmental goods. We show that spikes in Google search interest for Musk coincide with his Republican-leaning actions, particularly his 2022 Twitter acquisition and 2024 election involvement. We then conduct a sentiment analysis on both Musk and Tesla using over 100,000 comments from the social media app Reddit. We find that such actions have led to a rise in negative sentiment towards both Musk and Tesla among Democrats and moderates. By contrast, we find spikes in positive sentiment towards Musk and *negative* sentiment towards Tesla among Republicans, possibly highlighting Republican hesitance to buy EVs.

We then describe a registration-level dataset on passenger vehicles in Texas which allows us to study how ideology impacts EV demand. This dataset includes three key components that are often missing in other datasets: registrant ZIP code, registration date, and transaction price. The presence of registrant ZIP codes in the dataset enables us to link each registration to ideology based on the voting patterns of the registrant's ZIP code in the 2020 presidential election. Because we observe registration dates, we are able to analyze how registrations from ZIPs with varying ideologies behave immediately before and after Musk's Republican-leaning actions. Finally, observing transaction prices allows us to estimate demand for EVs without relying on imperfect proxies like manufacturer suggested retail prices (MSRP).

To understand how ideology drives demand for EVs, we employ nested logit demand estimation which allows willingness-to-pay (WTP) for the Tesla brand to vary by ideology and

2. While Musk's role in Tesla's falling market share has been speculated on, this paper provides a causal analysis. Much media coverage has pointed to Musk's ideology as a main driver in Tesla Motors' falling market shares. Such coverage has so far been unable to separate the role of Musk from broader trends in the EV market.

quarter-year. This reflects that Musk's ideological shift likely impacted Democrat, moderate, and Republican uptake heterogeneously. Our estimation controls for a variety of vehicle and ZIP code characteristics in order to isolate the role of ideology. We also instrument for price by using federal subsidy variation generated by the 2022 Inflation Reduction Act, which we argue is exogenous conditional on our controls. Finally, we show that anticipated policy shocks to EVs, past changes in subsidy policy, competitor entry, charging infrastructure expansion, and supply chain constraints are unlikely to drive our estimated trends in WTP for Tesla by ideology. Our estimation strategy thus generates estimates of how ideology influences consumer WTP for Teslas while addressing concerns over many confounds.

We estimate that Musk's ideological shifts have corresponded to large falls in WTP for Teslas among Democrats and moderates, with no change among Republicans. At the beginning of our study period, WTPs for Teslas relative to the median vehicle brand are respectively -$7,000, -$8,000, and -$9,000 for Democrat, moderate, and Republican ZIPs. These WTPs trend in parallel across the three ideologies until the 2024 presidential election. During the 2024 election, we observe an absolute decrease in WTP for Teslas of $8,000 in Democrat ZIP codes and $7,000 in moderate ZIP codes. No increase in WTP for Teslas occurs in Republican ZIP codes over our study period, reflecting that Republicans consumption of EVs has not increased. Moreover, such WTP drops are not present among other EV brands, which suggests that anticipated changes to EV policy from the Trump administration does not drive the fall in WTP for Teslas.

This fall in WTP for Teslas reduced private surplus and resulted in consumers substituting to vehicles which are more polluting than Teslas. Examining Musk's involvement in the 2024 presidential election, we find that Tesla Motors' revenues fall by $2.3 billion per year in the U.S.A. Competing firms see annual revenue gains of $860 million during this period, with the largest gains accruing to internal combustion engine vehicle (ICEV, i.e. 'gas vehicle') manufacturers. Many individuals also select the outside option of not buying

a new vehicle at all, opting instead to buy a used vehicle or keep their old vehicle – both of which are overwhelmingly likely to be ICEVs. Driven by this substitution towards ICEVs, we estimate that social damages from vehicular carbon emissions increase by $32 million per year. Consumers also lose significant value from Musk's election involvement. We find that private surplus in the vehicle market falls by $1.6 billion per year due to Tesla's falling appeal.

Such findings are important because they show that ideology can crowd out other considerations among consumers, including in the context of environmental goods. With affective polarization rising, our estimates highlight the extent to which ideology drives demand. This is particularly important for environmental goods, as we show social surplus falling when Democrat and moderate consumers substitute to *more polluting* goods which better match their ideology. The detrimental impacts to private surplus are even larger, highlighting that consumers can lose out as a result of brand ideological activity. Notably, brands that produce non-environmental goods (e.g. Pfizer, Target, and Walmart) have also experienced ideological boycotting – many of which may also have large impacts to private welfare. In examining Tesla Motors, we link rising polarization to consumers buying goods based on ideology, often at the expense of themselves or others.

This paper contributes to the environmental economics literature by establishing how important ideology is in driving demand for environmental goods. Previous papers have analyzed the role of peers in driving demand for environmental goods [Bollinger and Gillingham, 2012, Allcott and Rogers, 2014a, Sexton and Sexton, 2014b, Allcott and Kessler, 2019a, Bollinger et al., 2020b, Gillingham and Bollinger, 2021b, Bollinger et al., 2022, Liu, 2022]. Recent work has also established a strong correlation between ideology and demand for EVs, including examining the role of Musk [Davis et al., 2025, Gillingham et al., 2025]. Our contributions beyond these papers are three-fold. First, our empirical strategy includes both Musk as a natural experiment *and* controls for subsidy policy, competitor entry, and charging

infrastructure, enhancing the causal interpretation of our estimates. Second, we employ a demand estimation strategy which uncovers consumer WTP for the ideological associations of a good and the consequent private surplus loss upon Tesla's brand devaluation. Third, we quantify substitution patterns when Tesla's ideological association shifts Republican, showing that many consumers substitute to more polluting ICEVs.

This paper also contributes to the political economy literature by examining the welfare impacts of ideological consumption. Economic studies have previously established that ideology impacts individual decision-making [Bishop, 2009, Iyengar et al., 2012, Gift and Gift, 2015, Gimpel and Hui, 2015, Huber and Malhotra, 2017, Mummolo and Nall, 2017, Chen and Rohla, 2018, McConnell et al., 2018]. Fewer studies exist that link ideology to consumption [Khan et al., 2013, Chan and Ilicic, 2019, Panagopoulos et al., 2020, Conway and Boxell, 2024]. To our knowledge, no economic studies exist that estimate how ideology influences WTP for a good and, consequently, welfare. We contribute to this literature a demand estimation exercise in which we quantify the welfare impacts of consumption that is driven by ideology. We do so in a setting with higher stakes relative to previous papers which study how ideology influences consumption.

This paper finally contributes to the industrial organization and marketing literatures by more thoroughly examining the role of brand fixed effects in demand. Many different inputs have been shown to be important in driving brand value [Schmalansee, 1983, Clark, 2007, Doraszelski and Markovich, 2007, Bronnenberg et al., 2009, 2012, Shin et al., 2012, Kamenica et al., 2013, Bronnenberg et al., 2015, Bronnenberg and Dubè, 2017]. Recent examinations highlight the role of figureheads in driving WTP for brands [Bendisch et al., 2013, Cohen, 2013, Fetscherin, 2015, Mkrtchyan et al., 2023, Rudeloff and Damms, 2023]. We contribute two things to this literature. First, we quantify the link between brand fixed effects and welfare, showing that falling brand value can decrease private and social surplus. Second, we provide estimates of consumer demand driven by CEO political activity, which is important

as many brands have now been boycotted due to the political involvement of their CEO.

The paper proceeds as follows: Section 2 details background information pertinent to our setting, Section 3 enumerates our data sources, Section 4 highlights our estimation strategy and results, Section 5 discusses welfare impacts, and Section 6 concludes.

2.2 Background

Tesla Motors ("Tesla") is an American EV manufacturer which has recently seen rapidly falling market shares. Tesla historically enjoyed a majority of the U.S.A. EV market, making up more than 75% of EV market share prior to 2022. Since 2022, Tesla's EV market share has fallen dramatically, dropping below 50% in mid-2024.[3] This period of rapid transition in the market for Teslas has been marked by alterations in federal EV subsidy policy, increased competition, expansions in public EV charging infrastructure, and supply chain constraints.

Media speculation has posited that CEO Elon Musk's highly visible shift towards Republican ideology has caused Tesla's falling market share [Mahtesian and McHugh, 2025]. This is plausible in part because of Musk's high level of visibility. In Figure 2.1, we plot Google Trends search frequencies for "Elon Musk" and show that they outpace searches for several other prominent American CEOs. The largest spikes in searches for Musk coincide with his political actions. In particular, the first large spike in search interest occurs in April 2022 when Musk initiated the acquisition of Twitter (now 'X') with the stated purpose of reducing censorship of Republican viewpoints on the platform. The second spike coincides with Musk's involvement in the 2024 Presidential Election cycle, initiated by Musk's endorsement of Republican candidate Donald Trump in July 2024.

Musk's Republican swing in the 2020s altered peoples' views of Musk's political leanings. Prior to the 2020s, Musk often donated to Democratic candidates and publicly asserted

3. See here for Tesla EV market shares from 2019-2025: `https://www.axios.com/2025/09/08/tesla-market-share-electric-vehicles-musk`

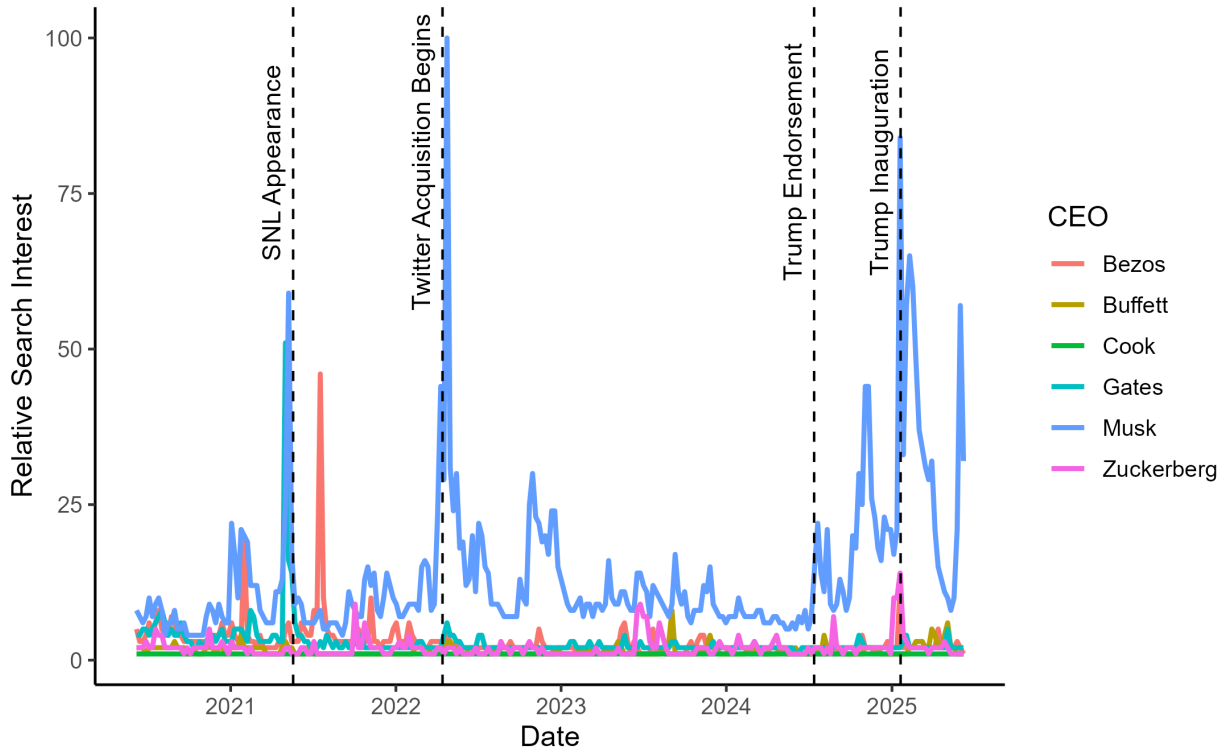


Figure 2.1: Google Trends Search Interest for Elon Musk Over Time

Notes: Figure depicts Google Trends data for the term 'Elon Musk.' The four vertical dashed lines, from left to right, depict the dates Musk went on SNL, Musk initialized the Twitter acquisition, Musk endorsed Donald Trump for President, and Musk attended Trump's inauguration. The y-axis is normalized such that the maximum point is equal to 100 and all other points within the graph are relative to this maximum.

support for moderate or Democrat causes. This, and his chairing of an EV company, led to the general public perception that Musk was not a Republican. We show that this view changes in the 2020s in Appendix Figure B.3, which plots Google search trends for "Elon Musk Republican" and "Elon Musk Trump". Searches for the former term peaked during the Twitter acquisition, whereas the 2024 election cycle was the peak of the latter. This highlights that Musk's actions were unexpected at the times of their occurrence and thus shifted public perception of Musk.

Musk's Republican swing also altered consumer sentiment about Tesla. To show this, we conduct a sentiment analysis of comments about Elon Musk and Tesla posted on the social media site Reddit. We collect over 100,000 individual comments that were posted between January 2021 and December 2024 on 'threads' discussing Elon Musk, Tesla, or SpaceX [Watchful1, 2025].[4] To examine whether sentiment changed differently by ideology, we specifically gather comments from three popular 'subreddits': r/liberal, r/conservative, and r/centrist.[5] We then use ChatGPT's GPT-4o to classify the sentiment for each comment as being negative or positive on a scale of -1 to 1 separately for both Musk and Tesla. We define all comments with a score of 0.3 or greater as positive, and all comments with a score of -0.3 or less as negative.[6] We then add up the total number of positive and negative comments expressed towards Musk and Tesla in each month.[7]

We show in Figure 2.2 that sentiment for Tesla and Musk are linked, with large changes

4. A thread on Reddit is a discussion, which comprises of an original post and associated comments. We conduct sentiment analysis on both the original post and all associated comments for all relevant threads.

5. A subreddit is a specific page on Reddit which serves as a collection of posts and comments on one particular broad topic. For instance, r/conservative describes their page as of August 12, 2025 as "a place on Reddit for conservatives, both fiscal and social, to read and discuss political and cultural issues from a distinctly conservative point of view."

6. Comments with scores in the interval (-0.3,0.3) are omitted to eliminate noise from weak, unclear, and uncertain scores. Appendix Figures B.9 and B.10 show that our results are robust to using alternative cutoffs ($>0/<0$, and $\geq 0.5/\leq -0.5$ respectively) to define positive and negative comments.

7. Further details about this analysis are provided in Appendix Section B.2.

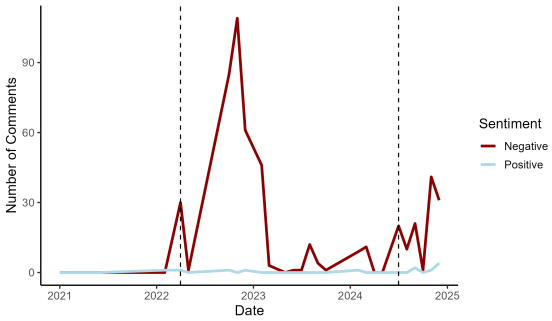
in both corresponding to Musk's political activity.[8] We observe marked increases in negative sentiment about both Musk and Tesla in r/liberal following the Twitter acquisition, and to a lesser extent also during the 2024 election cycle. The r/centrist subreddit sees similar spikes in negative sentiment about Musk and Tesla corresponding to those two events, although the spike in negative sentiment is stronger during the 2024 election cycle. The fact that sentiment for Tesla and Musk move in parallel supports the notion that Musk's political activity directly influenced liberal and moderate valuations of Teslas. On r/conservative, there is a sharp increase in positive sentiment about Musk following the Twitter acquisition in, but sentiment about Tesla during this time is predominantly negative. This suggests that conservatives' increasingly positive views of Musk do not necessarily translate into increasingly positive views of Tesla.[9]

2.3 Data

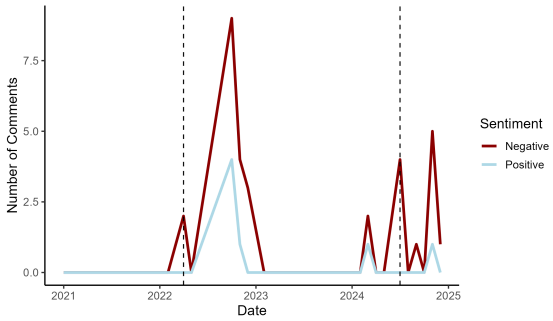
In this section, we describe the data which allows us to estimate consumer WTP for Teslas separately by ideology. We use the universe of vehicle registrations from the Texas Department of Motor Vehicles (DMV) from January 2021 to April 2025. The Texas DMV data is uniquely apt for our study because it includes registrant ZIP, registration date, and transaction price, allowing demand to be linked to political views over time at a spatially granular level.

8. Our findings on sentiment largely match NBC News survey data on ideological views about Musk and Tesla [NBC, 2025].

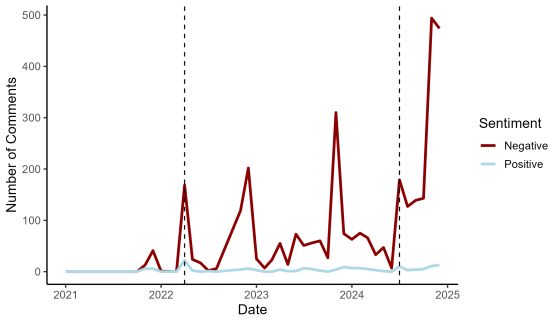
9. Our conclusion that Musk's political activity influences consumer valuation is further reinforced by Appendix Figure B.4, which depicts Tesla's stock price trends. The Twitter acquisition immediately preceded a large fall in Tesla's stock price. Tesla stock additionally saw a precipitous increase in share price upon Trump winning the 2024 election, an effect which reversed following Musk's involvement with the administration.



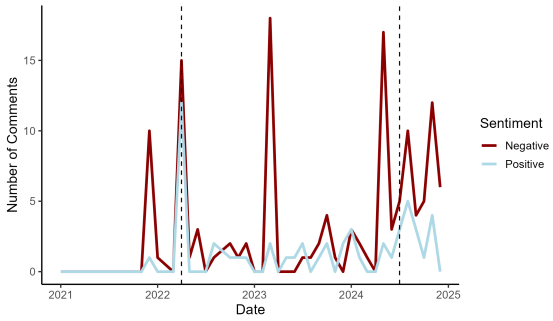
r/liberal: Elon Musk Sentiment



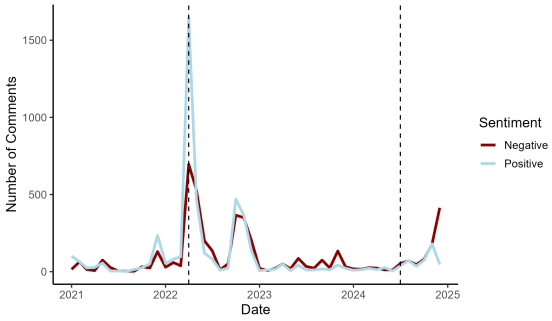
r/liberal: Tesla Sentiment



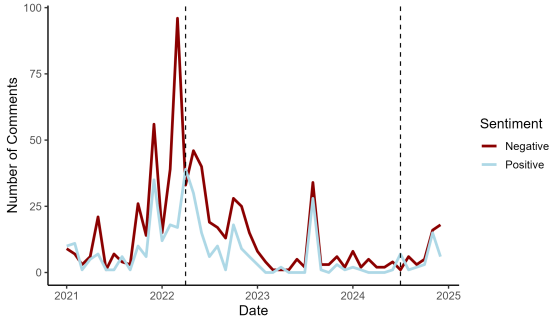
r/centrist: Elon Musk Sentiment



r/centrist: Tesla Sentiment



r/conservative: Elon Musk Sentiment



r/conservative: Tesla Sentiment

Figure 2.2: Sentiment Analysis

Notes: This figure depicts the sentiment analysis results from the subreddits r/liberal, r/centrist, and r/conservative. We analyze all comments on threads where the title contains one of the following keywords: "Elon", "Musk", "Tesla", "SpaceX". For each comment, we produce two separate sentiment scores: one measuring sentiment about Elon Musk, and one measuring sentiment about Tesla. This sentiment analysis was performed with ChatGPT's GPT-4o model. Positive comments are defined as comments with a score greater than or equal to 0.3, and negative comments are defined as comments with a score less than or equal to -0.3. The figures on the left depict sentiment about Elon Musk, while the figures on the right depict sentiment about Tesla. In each subfigure, the first dashed vertical line corresponds to the month immediately prior to the month of Musk's Twitter acquisition, and the second dashed vertical line corresponds to the month immediately prior to the month of Musk's endorsement of Trump (June 2024).

2.3.1 Vehicle Data

To analyze trends in EV demand, we make use of data from the Texas DMV which includes all private vehicle registrations in Texas from January 2021 through April 2025. For each registration, we observe the associated unique vehicle identification number (VIN), transaction price of the vehicle purchase, odometer reading at the time of registration, and ZIP code where the registrant resides. As vehicle registration is required by state law in Texas, this dataset encompasses *all* vehicle owners across the entire state.

The Texas DMV data includes VINs for each registration, allowing us to observe key characteristics of the vehicles being registered. We decode the VINs using the U.S. National Highway Traffic Safety Administration's (NHTSA) bulk VIN decoder, yielding the make, model, year, vehicle body type, and several other characteristics for each observed registration.[10] We remove all registrations of non-passenger vehicles and additionally drop all registrations in which the registrant resides outsides of Texas.[11] Finally, to analyze new vehicle registrations, we restrict our sample to include only vehicles with a model year that is at most one year old at the time of registration.[12]. Our final sample includes the roughly 6.2 million new passenger vehicle registrations occurring in Texas from January 2021 to April 2025.[13]

To capture important characteristics which are not observed in the DMV data, we gather

10. See here: `https://vpic.nhtsa.dot.gov/decoder/`. Throughout this paper, EVs are treated as including both battery EVs and plug-in hybrid EVs. All other vehicles are classified as ICEVs, following the NHTSA's definition of primary fuel as well as federal subsidy classifications which count only battery EVs and plug-in hybrids as eligible.

11. Vehicle bodies that are dropped include buses, fire apparatuses, incomplete vehicles, limousines, motor homes, freight trucks, and delivery vans

12. This fits with norms on vehicle listing sites such as Kelley Blue Book and Edmunds which list all vehicles of the current or previous year as "new" and all vehicles from two years prior or older as "used". An alternative to this approach is to consider vehicles with low odometer readings at the time of registration This approach yields many vehicles which are no longer being manufactured, possibly reflecting underlying issues with odometer data collection.

13. This count is in line with registration counts from other sources which estimate annual new vehicle registrations at the state level [F&I Tools USA, 2025].

information on body length, fuel efficiency, and horsepower from various sources. Data on fuel efficiency is taken primarily from the U.S.A. Environmental Protection Agency (EPA), which formally tests fuel efficiency for light duty vehicles sold in the U.S.A. Fuel efficiency for heavy duty vehicles (e.g. the Ford F-250, F-350, and many Ram trucks) is taken from various Edmunds review pages when available, and from self-reported fuel efficiency on the app Fuelly when not. Horsepower and length for all vehicles are gathered primarily from Edmunds review pages, with Kelley Blue Book and Car and Driver being used for vehicles that do not list specification on Edmunds. With this diverse set of sources, we are able to collect length, fuel efficiency, and horsepower data for nearly every vehicle registration in our data (>99.99%).[14] We present statistical descriptors of vehicle characteristics in Table 2.1.

Variable	Mean	SD	Min	Max	N
Price (USD)	48,173	25,008	0	6,162,023	6,240,435
Combined MPG	28.94	20.29	10.00	132.70	6,239,522
Horsepower	285.41	101.26	76.00	1020.00	6,239,967
Length (in.)	199.96	22.48	143.00	257.03	6,240,418
Diesel dummy	0.06	0.23	0.00	1.00	6,240,435
EV dummy	0.05	0.22	0.00	1.00	6,240,435
Passenger Truck dummy	0.25	0.43	0.00	1.00	6,240,435
Class 1 Weight dummy	0.61	0.49	0.00	1.00	6,240,435
Class 3 Weight dummy	0.04	0.19	0.00	1.00	6,240,435

Table 2.1: Summary Statistics of New Vehicle Registrations

Notes: Table depicts summary statistics of newly registered passenger vehicles in Texas from January 2021 to April 2025.

To generate exogenous variation in the price of each registration, we measure variations in federal subsidy eligibility generated by the 2022 Inflation Reduction Act. Prior to August 2022, vehicles from makes which had sold more than 200,000 EVs in the U.S.A. (Tesla and Chevrolet) were ineligible to receive the $7,500 subsidy. In August 2022, foreign-made EVs

14. Data on fuel efficiency is not consistently available for all trims of a make-model-year. Similarly, many trims have data on length and horsepower which is unavailable or not easily matched due to inconsistent naming conventions across data sources. To maintain consistency, we thus match data at the make-model-year level and use the average value across available trims.

were also made ineligible to receive the federal subsidy. In January 2023, the 200,000 vehicle sales limit was removed and replaced with restrictions on subsidy eligibility for EV cars with manufacturer suggested retail prices (MSRPs) above $55,000 and EV SUVs with MSRPs above $80,000.[15] Additional alterations to subsidy eligibility occurred for some Rivian, Jeep, Audi, BMW, and Nissan models in April 2023 based on battery material sourcing. We merge the time-variant subsidy amount for each vehicle based on its MSRP (gathered from Kelley Blue Book pages), make, model, vehicle body, and time of purchase.

2.3.2 ZIP Characteristics

We measure ZIP code characteristics which may be relevant drivers of demand for Teslas and other vehicles. Specifically, we account for a ZIP's ideology, demographics, and tax filing rates.

We measure ZIP code ideology in order to capture how WTP for Teslas shifts heterogeneously across the ideological spectrum. To do so, we use precinct-level 2020 Presidential election votes from MIT Election Lab which are assigned to ZIP codes according to spatial overlap. We then classify each ZIP code as being Republican, moderate, or Democrat based on the percent of inhabitants who voted for 2020 Republican presidential candidate Donald Trump. ZIP codes which had more than 60% of individuals voting for Trump are classified as Republican. ZIP codes with less than 40% of voters choosing Trump are classified as Democrat. All ZIP codes between these cut-offs are thus moderate.[16]

We collect ZIP code median income, population, and proportion of the ZIP code classified as rural in order to capture important demographic differences across ZIP codes which drive

15. See Allcott et al. [2024a] for a more thorough discussion of the Inflation Reduction Act's impact to EV subsidy policy.

16. These cutoffs are arbitrary, but are chosen to balance a trade-off between exclusivity (i.e. "too many" moderate ZIP codes being counted as Democrat or Republican) and statistical power (too few ZIPs in each category to detect effects). In Appendix Figure B.1, we show that the choice of cutoffs matters little for our conclusions.

demand for vehicles. All three variables are collected from the 2023 American Community Survey. We collect income to allow certain brands to be more appealing to individuals with different levels of income – such as luxury brands among higher earners. We collect population and rural proportion to account for the fact that individuals living in less urban environments are likely to rely on privately owned vehicles to a greater extent for commuting.

Finally, we collect tax filing rates to measure the extent to which subsidies for EVs are likely to have historically actually been realized in specific ZIP codes. Prior to January 1, 2024, federal subsidy allocations were administered as a rebate received only when filing taxes. This implies that individuals who did not file their taxes could not have received any subsidy funds for buying an EV. We do not observe actual subsidy amounts received in our data, instead relying on subsidy *eligibility* as our instrument. We thus capture tax filing rates to mitigate measurement error as a potential violation of the exclusion restriction when using subsidy eligibility as an instrument for price.

2.3.3 Market Share Trends

We present raw trends in Tesla market share in Figure 2.3 separately by a ZIP code's ideological classification, yielding three main findings which motivate our work. First, Teslas are more common in Democrat and moderate ZIP codes than they are in Republican ZIP codes throughout our study period. Second, Tesla's market share in Democrat ZIP codes falls relative to moderate ZIP codes following Musk's April 2022 acquisition of the social media platform Twitter (now 'X'). As we show in Figure 2.2, this coincides with the largest spike in negative sentiment about Musk on the subreddit r/liberal. Third, the 2024 election cycle coincides with a dramatic decrease in Tesla's market share in Democrat and moderate ZIP codes. This again reflects precisely what we find in the Reddit data in Figure 2.2: rising negative sentiment among *both* Democrats and moderates.

We additionally present raw trends in the market shares of non-Tesla vehicles in Fig-

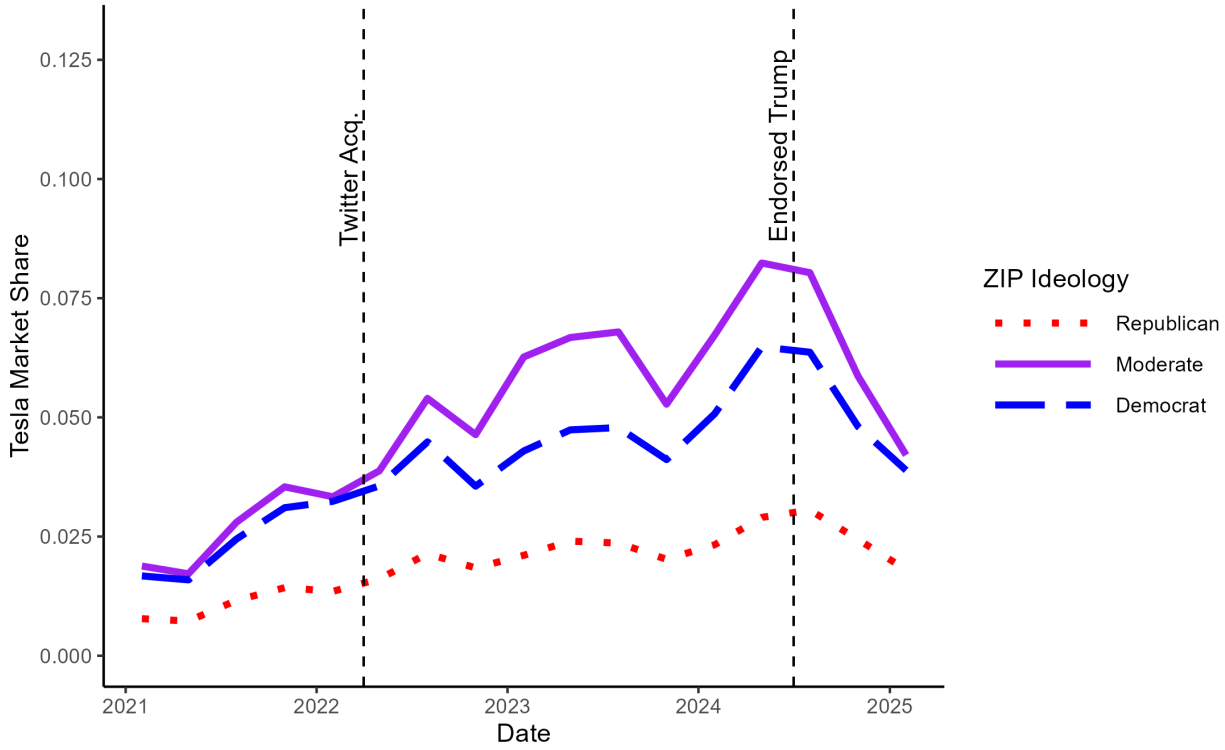


Figure 2.3: Tesla Market Share Over Time

Notes: Figure depicts trends in Tesla market share, defined as the proportion of new registrations in a given month that are Tesla makes. Figure depicts these trends separately for Republican ZIP codes (red dotted), moderate ZIP codes (purple solid), and Democrat ZIP codes (blue dashed). Republican ZIP codes are defined as those for whom less than 40% of votes in the 2020 election were for Joe Biden, moderate ZIP codes between 40% and 60%, and Democrat ZIP codes greater than 60%. Vertical dashed lines indicate the beginning of the Twitter acquisition (left) and Elon Musk's endorsement of Donald Trump in the 2024 Presidential Election cycle (right).

ure 2.4, suggesting that many would-be Tesla buyers substitute towards ICEVs. Despite the presence of many EV alternatives to Tesla, we find that the market shares of EVs not manufactured by Tesla stagnate between 2023-2025. Of particular note, we observe that the proportion of vehicle registrations comprising of EVs manufactured by Tesla's competitors remains stagnant following the Trump endorsement. This contrasts with a sharp increase in the prevalence of ICEV registrations among both moderate and Democrat ZIP codes.[17]

The raw trends presented in Figure 2.3 and Figure 2.4 are compelling, but motivate further analysis to capture WTP for the Tesla brand net of confounding mechanisms. In particular, market share trends may reflect variations in vehicle characteristics, prices, subsidy policy, competition, charging infrastructure, or supply chain constraints. This motivates an analysis which controls for or otherwise examines these confounding mechanisms.

2.3.4 External Validity

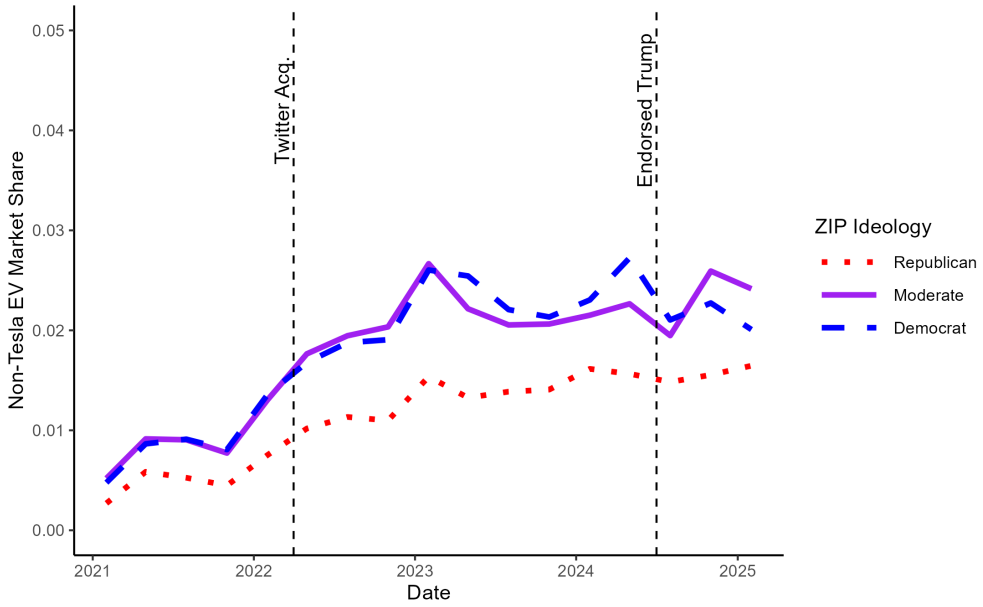
As our approach uses one state's data to describe a nationwide phenomenon, a discussion of external validity is pertinent.[18] We believe that Texas differs from other areas in the U.S.A. along two important axes, both of which likely make our estimates lower bounds of patterns across the nation.[19]

First, Texas is more Republican-leaning than many other states, which may depress our estimates of the falling WTP for Teslas. In Figure 2.5, we present a map and a histogram of vote shares for Trump in 2020 within Texas. The vast majority of the state by land area voted in favor of Trump, though there is still a significant portion of ZIP codes that aligned

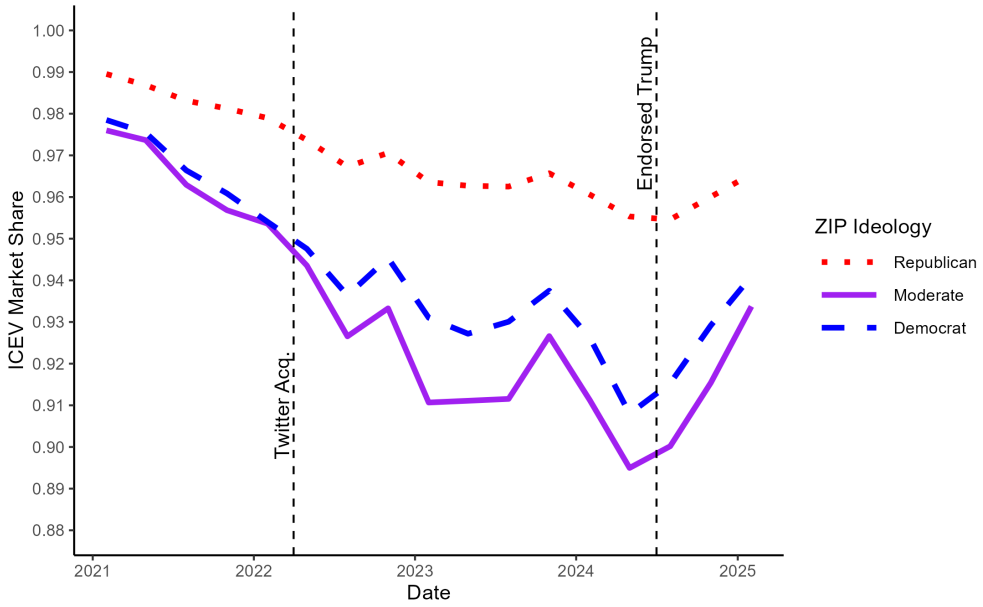
17. We present total registrations in Appendix Figure B.15 which seems to suggest that the late 2024 vehicle market also observed fewer overall registrations than the late 2023 vehicle market. This is consistent with individuals substituting away from Teslas and towards not buying a new vehicle.

18. Data for other states tend to not include all three of registrant ZIP, registration date, and transaction price.

19. Appendix Table B.1 depicts summary statistics for Texas compared to the whole U.S.A. and to California, the largest EV market in the U.S.A. Texas is more representative of the U.S.A. in terms of racial demographics, median income, population density, and EV market shares.



Panel A: Non-Tesla EV Market Shares



Panel B: ICEV Market Shares

Figure 2.4: Trends in Non-Tesla Market Shares

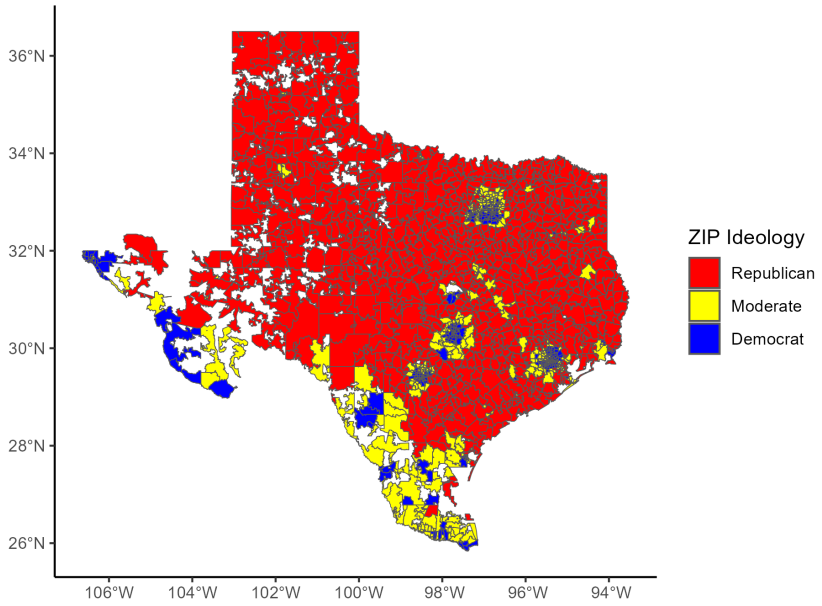
Notes: Figure depicts in Panel A trends in the market share of non-Tesla EVs, defined as the proportion of new registrations in a given month that are EVs not made by Tesla. Panel B depicts in Panel B the trends in market share of ICEVs, defined similarly. Figure depicts these trends separately for Republican ZIP codes (red dotted), moderate ZIP codes (purple solid), and Democrat ZIP codes (blue dashed). Republican ZIP codes are defined as those for whom less than 40% of votes in the 2020 election were for Joe Biden, moderate ZIP codes between 40% and 60%, and Democrat ZIP codes greater than 60%. Vertical dashed lines indicate the beginning of the Twitter acquisition (left) and Elon Musk's endorsement of Donald Trump in the 2024 Presidential Election cycle (right).

with Democratic presidential candidate Joe Biden. This means that the support of our data in terms of political leanings does not have the same breadth as it would if we also had data for other states. In particular, our data is not able to speak to the effects that exist in relatively Democratic rural ZIP codes that exist in other large EV markets like New York and California. The fact that Texas is more Republican than most other states may indicate that our estimates of push-back against Musk's shift towards the Republican Party are lower than the average effects in the U.S.A., even among ZIP codes which we classify as Democrat.

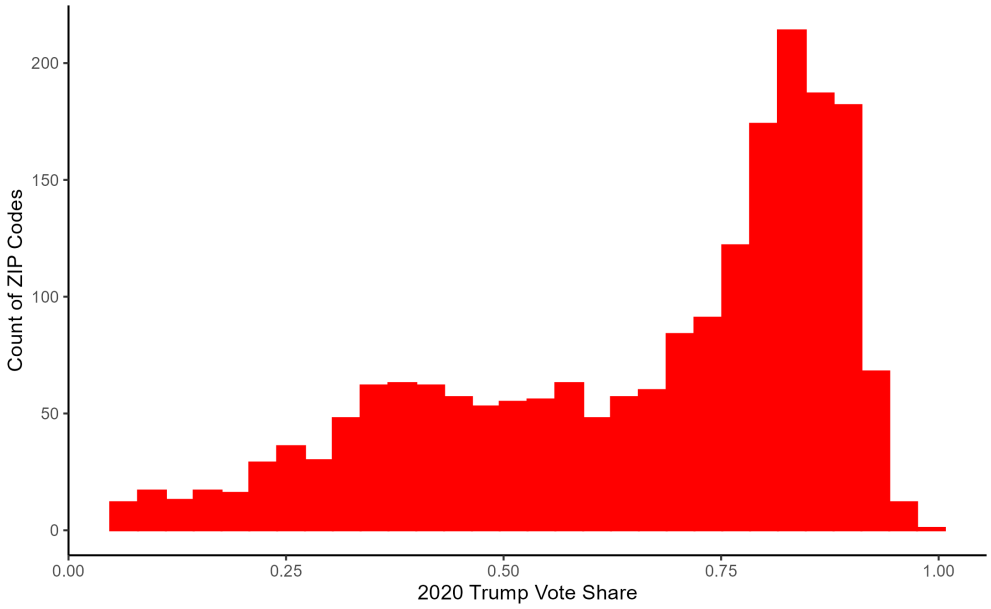
Second, Tesla is headquartered in Texas and has a production plant in the city of Austin, which may depress our estimates of the falling WTP for Teslas. Geography has previously been shown to impact brand loyalty [Bronnenberg et al., 2009, 2012]. Many residents of Texas may thus be particularly loyal to Tesla even if they find Musk's political involvement to be undesirable. If true, this could reduce estimated effect sizes of Musk's political involvement on moderate and Democrat WTP for Teslas, again meaning our estimates could potentially be lower bounds of what occurs on average in the U.S.A.

2.4 Empirical Analysis

In this section, we use nested logit to estimate a model of demand with flexible substitution patterns among vehicles. We include brand fixed effects that vary by ZIP ideology and quarter-year in this model, which allows trends in WTP for Teslas to vary by ideology. Our estimates reveal that WTP for Teslas dropped among Democrats and moderates during the 2024 election cycle. We show that anticipated EV policies, changes in subsidy policy under the Inflation Reduction Act, competitor entry, charging infrastructure, and supply chain constraints do not explain these results.



Panel A: Spatial Distribution of Texan Ideology



Panel B: Histogram of ZIP Code Vote Shares

Figure 2.5: Distribution of ZIP Code Level Ideology for Texas

Notes: Figure depicts a map of vote shares for Donald Trump in 2020 (Panel A) and a corresponding ZIP-code level histogram of vote shares (Panel B). Missing values on vote shares are omitted in the map (depicted as empty white space). Republican ZIP codes are defined as those for whom less than 40% of votes in the 2020 election were for Joe Biden, moderate ZIP codes between 40% and 60%, and Democrat ZIP codes greater than 60%.

2.4.1 Model of Demand

Model Setup

In order for estimation to be tractable, we assume that the utility of Texan households over vehicles is characterized by a discrete choice model. We model the indirect utility for individual in ZIP code z buying a vehicle with model-year j (henceforth 'product j') in year-month t. This decision is modeled as depending on transaction price p_{jt} and a vector of ZIP- and product-specific characteristics x_{zjt}. We include also a term ξ_{jt} which captures month-specific characteristics of product j which are unknown to the econometrician but known to consumers and firms. Finally, we include ε_{zjt} which represents idiosyncratic tastes for products among consumers. Utility is thus modeled as

$$u_{zjt} = \alpha p_{jt} + \beta x_{zjt} + \xi_{jt} + \zeta_{zg(j)t} + (1-\sigma)\varepsilon_{zjt} \tag{2.1}$$

The vector of characteristics x_{zjt} allows brand values to vary by ideology and income. We do so for ideology by including a fixed effect for brand interacted with ZIP ideology and quarter. We also allow brand values to vary by income by interacting a ZIP's decile of median income with a brand-by-quarter fixed effect. This reflects the fact that certain brands may predominantly appeal to different income groups, such as luxury brands appealing to high-income consumers.

To capture other characteristics that consumers care about, x_{zjt} also includes a variety of controls. These include vehicle body, fuel efficiency, horsepower, weight class, and length. We additionally control for county-quarter fixed effects, ZIP rural percentage, and ZIP population. This allows the outside option (not buying a new vehicle) to have varying attractiveness over space and across time. Finally, we include controls for competitor entry and public EV charging stations at the ZIP-month level. The former is measured as the number of products being registered for the first time ever in a ZIP in any month, and the

latter as the overall number of open public charging stations in each ZIP in a given month.

To allow flexible substitution patterns, we model demand to follow the structure of a nested logit [McFadden, 1981, Berry, 1994]. This structure allows for correlation in unobserved utility across products that belong to the same nest. We denote in equation (2.1) each product j belonging to a nest $g(j)$, with nests indexed by $g \in \mathcal{G}$. The term $\zeta_{ig(j)t}$ is a nest-specific unobservable that is common to all products within nest $g(j)$, and ε_{ijt} are idiosyncratic preferences which are assumed to be i.i.d. Type-1 extreme value. The parameter $\sigma \in [0, 1)$ captures the degree of within-nest correlation in idiosyncratic preferences. When $\sigma = 0$, the model collapses to the simple logit case.

We normalize the utility from the outside option as $u_{i0t} = \varepsilon_{i0t}$ for all t, which yields a closed form solution for market shares.[20] Denoting by $\mathcal{J}_{gt}$ the set of products belonging to nest g in period t, the market share of product j can be decomposed into two components: the conditional share of product j within its nest and the share of the nest within the overall market. Specifically,

$$s_{zjt} = s_{j|zg(j)t} \cdot s_{zg(j)t} \tag{2.2}$$

where

$$s_{j|zg(j)t} = \frac{\exp\left(\frac{V_{zjt}}{1-\sigma}\right)}{\sum_{j' \in \mathcal{J}g(j)t} \exp\left(\frac{V_{zj't}}{1-\sigma}\right)}$$

$$s_{zg(j)t} = \frac{\left[\sum_{j' \in \mathcal{J}g(j)t} \exp\left(\frac{V_{zj't}}{1-\sigma}\right)\right]^{1-\sigma}}{1 + \sum_{g' \in \mathcal{G}} \left[\sum_{j' \in \mathcal{J}g't} \exp\left(\frac{V_{zj't}}{1-\sigma}\right)\right]^{1-\sigma}}$$

and $V_{zjt} = \alpha p_{jt} + \beta x_{zjt} + \xi_{jt}$ represents the mean utility component of product j.

20. This normalization and the type 1 extreme value assumption are both standard in the literature on logit demand estimation [Berry, 1994].

Taking logarithms and rearranging, we obtain the estimating equation:

$$\ln(s_{zjt}) - \ln(s_{z0t}) = \alpha p_{jt} + \beta x_{zjt} + \sigma \ln(s_{j|zg(j)t}) + \xi_{jt} \tag{2.3}$$

The left-hand side and the right-hand side variables of this equation are all observable in data (except for the residual ξ_{jt}). This enables estimation of the nested logit parameters under ordinary least squares so long as price and within-nest share are uncorrelated with ξ_{jt}. This is unlikely to hold because firms set prices and produce quantitites based on ξ_{jt}, motivating the usage of an instrumental variables approach which we describe in the following subsections.

Nest Structure and Instrument

We model demand to capture the fact that individuals who are likely to purchase Teslas may also be likely to purchase other cleaner vehicles such as competing EVs or high-performing ICEVs. To capture this substitution, we nest demand by fuel efficiency, measured in miles-per-gallon for ICEVs and miles-per-gallon-equivalent for EVs. We discretize fuel efficiency into deciles, creating a total of ten nests. Appendix Table B.3 depicts summary statistics for these nests, showing in particular that the nest occupied by Tesla (nest 10) primarily includes other EVs and high-performing ICEVs. We instrument for nest share using the number of products in each nest, as is common practice in demand estimation exercises [Gandhi and Houde, 2019].[21]

21. This instrument is relevant because a greater number of alternatives is likely to reduce the market share of any given product. It is argued to be exclusive under an assumption that the drivers of entry into a nest are uncorrelated with unobserved demand shocks (i.e. that entry is driven by supply-side influences conditional on x_{jt}).

2.4.2 Instrument for Price

To deal with concerns of price endogeneity in equation 2.3, we employ an instrumental variable approach. A primary concern with estimating equation 2.3 is that price is likely correlated with unobservable characteristics contained in ξ_{jt}. This means that a least squares regression produces a biased estimate of α. To deal with this issue, we instrument for price using variation in federal subsidy eligibility over time for each vehicle model. During our study period, the Inflation Reduction Act restricted eligibility for receiving subsidies in some models and enhanced eligibility among others.

Two assumptions are required for federal subsidy eligibility to be a valid instrument. First, federal subsidies need to be correlated with the price that individuals pay (instrument relevance). This is to likely to hold because subsidies directly reduce the price that consumers pay for EVs, assuming non-zero pass-through. Indeed, we show a strong first-stage regression with F-stats ranging from 221-250 in Appendix Table B.2. The estimated coefficient on subsidy amount is -0.62, indicating that a $1 increase in subsidy allocation is associated with a fall in price paid of $0.62.[22]

Second, federal subsidies need to impact market shares of vehicles only through their impact to prices (instrument exclusivity). This is violated if subsidies shift non-price characteristics of vehicles, which may occur if subsidy eligibility is dependent upon threshold values for unobservables or eligibility changes coincide with product alterations. The Inflation Reduction Act eligibility depends only on vehicles being American made, of certain body types, and of a sufficiently low Manufacturer's Suggested Retail Price (MSRP). The former two of these are controlled for in our regression in the brand dummy variables and

22. One-for-one changes in subsidy eligibility and price paid is likely not observed for two reasons. First, many suppliers may absorb some of the subsidy value and not pass it on to consumers. Second, many ineligible EVs are technically able to receive the subsidy if they are leased, which may diminish first stage estimates. However, Texas requires all leased vehicles to pay sales tax on the entire value of a vehicle, diminishing the rates of leasing in Texas relative to other states. It is estimated that 12.8% of financed vehicles in Texas were leased in 2024Q1 `https://flexpathcapital.com/images/uploads/experian-safm-q1-2024.pdf`.

body dummy variables, while the latter likely does not impact demand conditional on price. We alleviate additional concerns over endogeneity by controlling for a variety of other crucial characteristics in driving demand.[23]

2.4.3 Results

Vehicle Characteristics

We show that our estimated coefficients of model parameters are robust to several specifications that control for general trends in vehicle markets in Table 2.2. Column (1) of Table 2.2 includes fixed effects for a market's make-quarter-ideology, make-quarter-income, and county-quarter. Column (2) controls also for the count of EV charging stations present in a ZIP, capturing increasing public charging infrastructure. Finally, column (3) controls additionally for competitor entry into ZIPs, capturing trends in when vehicles enter markets.

Price has a negative and statistically significant coefficient which is robust across all specifications of Table 2.2. Our preferred specification in column (3) controlling for both charging infrastructure and competitor entry yields a price coefficient of -0.018. This corresponds to a market-price elasticity of approximately -0.79, which falls within the range of estimates among papers examining demand for new vehicles [Berry et al., 1995, Leard et al., 2023, Leard and Wu, 2023, Grieco et al., 2023].

Further validating our estimates is the fact that coefficients on fuel efficiency, horsepower, Class 3 weight, and length are all positive. This indicates, as expected, that individuals prefer larger and heavier vehicles with efficient engines. Finally, we note that our nest parameter is approximately 0.76. Values closer to 1 collapse to simple logit demand in which goods within a nests do not have correlated demand. An estimated value of 0.76 thus indicates

23. The prices we observe prior to January 2024 are gross of subsidy and we do not explicitly observe subsidies being applied to vehicles. There is thus potential for measurement error which would violate the exclusion restriction if correlated with the error term ξ_{jt}. We discuss measurement error and rule out this concern in Appendix Section B.1.3.

	(1)	(2)	(3)
Price (1000 USD)	-0.015***	-0.016***	-0.018***
	(0.006)	(0.006)	(0.006)
Nest Parameter (σ)	0.735***	0.742***	0.759***
	(0.013)	(0.013)	(0.013)
MPG	0.001***	0.001***	0.000**
	(0.000)	(0.000)	(0.000)
Horsepower	0.002**	0.002**	0.002***
	(0.001)	(0.001)	(0.001)
Class 3 Weight dummy	0.111**	0.118***	0.133***
	(0.045)	(0.046)	(0.047)
Class 1 Weight dummy	-0.102***	-0.107***	-0.120***
	(0.039)	(0.040)	(0.042)
Length (in.)	0.003***	0.004***	0.004***
	(0.001)	(0.001)	(0.001)
Observations	4,136,625	4,136,625	4,136,625
Charging Controls	No	Yes	Yes
Entry Controls	No	No	Yes
Make x Quarter x Ideology FE	Yes	Yes	Yes
Make x Quarter x Income FE	Yes	Yes	Yes
County x Quarter FE	Yes	Yes	Yes

Table 2.2: Demand Estimation Coefficient Results

Notes: Table depicts estimated marginal utilities for characteristics that drive vehicle demand. Price is instrumented for using federal subsidy variation and nest share by number of competing vehicles in the nest. All estimates include make-quarter-ideology, make-quarter-income and county-quarter fixed effects. Column (1) includes only these fixed effects. Column (2) additionally controls for number of EV charging stations in each ZIP code in each month. Finally, Column (3) additionally controls for number of EVs and total of vehicles entering markets in each ZIP in each month. Standard errors are clustered at the market level and depicted below estimates in parentheses. All specifications additionally control for vehicle body type, percentage of a ZIP which is rural, and percentage of a ZIP which files taxes.

that unobserved preferences within nests are to some extent correlated with each other.

Fixed Effects for Brands

Estimates of brand fixed effects are shown to vary by ideology in Figure 2.6 and Figure 2.7. We find that the largest non-American brands are more attractive among Democrat and moderate ZIPs. On the other hand, the largest American brands tend to be more attractive among Republican ZIPs. These brand fixed effects tend to trend together with only minor variations in the gap between ideologies. None of these brands depict large trend breaks across ideology. Similarly, no brands experience sharp and sudden declines corresponding to the 2024 election.

Figure 2.8 depicts fixed effects for Tesla, which show that Musk's involvement in the 2024 election leads to a relative devaluation of Tesla among moderate and Democrat ZIPs. The fixed effects across ideology for Tesla trend together until Musk's involvement in the 2024 election, at which point moderate and Democrat fixed effects fall while Republican fixed effects stay flat. At the end of our data in 2025Q1, the fixed effects for Tesla are the lowest they have been across our study period for both moderates and Democrats.

WTP for Tesla Brand

We calculate WTP for Tesla relative to the median brand using our estimates. We do so by first subtracting the median fixed effect across brands from Tesla's fixed effect, where medians are calculated separately for each ideology-quarter pairing. We choose to use the median brand because brand values are skewed by high-end luxury vehicles such as Rolls-Royce and Lamborghini.[24] We then normalize this difference in fixed effects by the coefficient on price. This yields a measure of WTP for Teslas relative to the median brand.

24. In Appendix Figure B.5, we provide qualitatively similar estimates which instead compare Tesla to the mean.

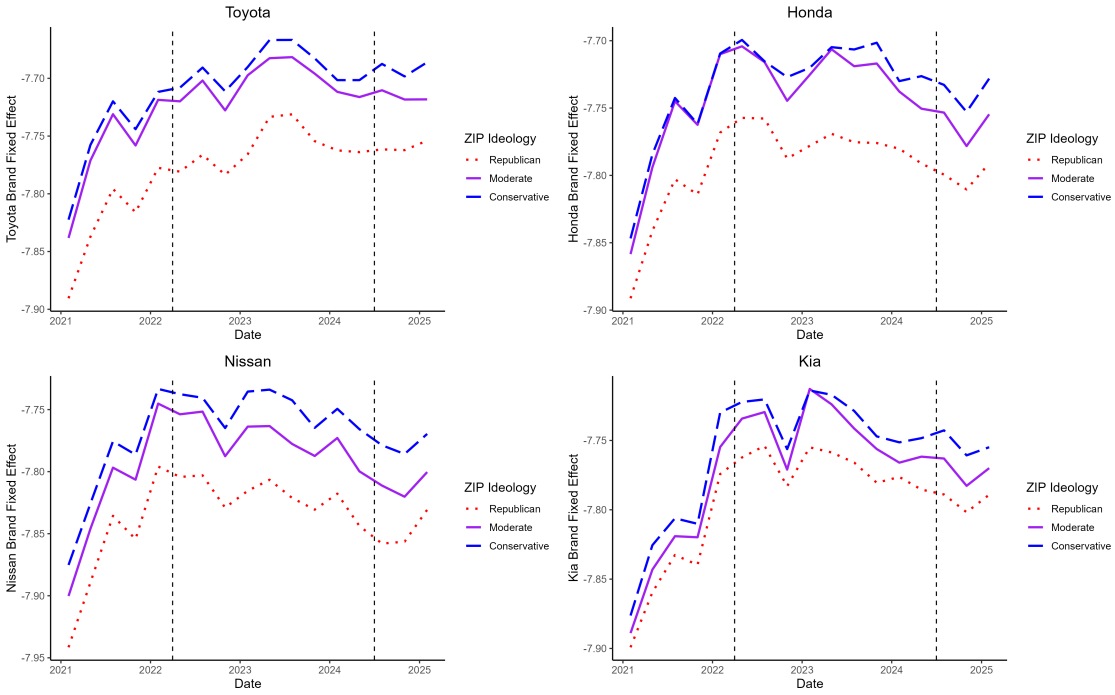


Figure 2.6: Brand Fixed Effects for non-American Brands

Notes: Figure depicts estimates of fixed effects for the top four non-American brands by observed new registrations. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Fixed effects are all relative to the excluded group, which is the outside option of not purchasing a new vehicle. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).

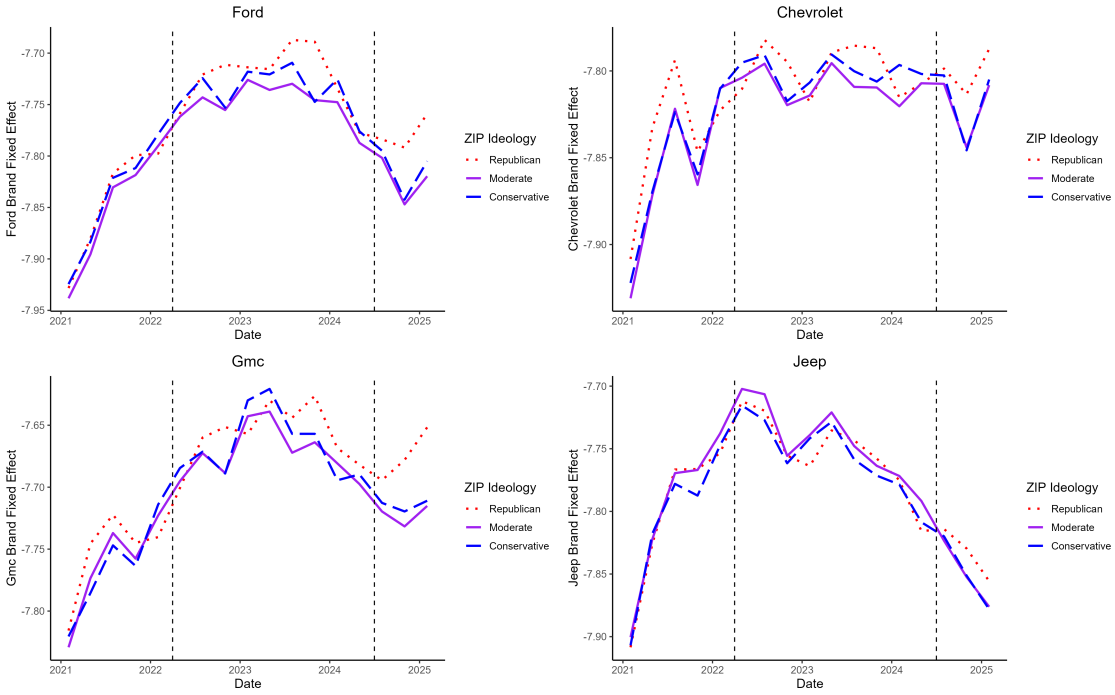
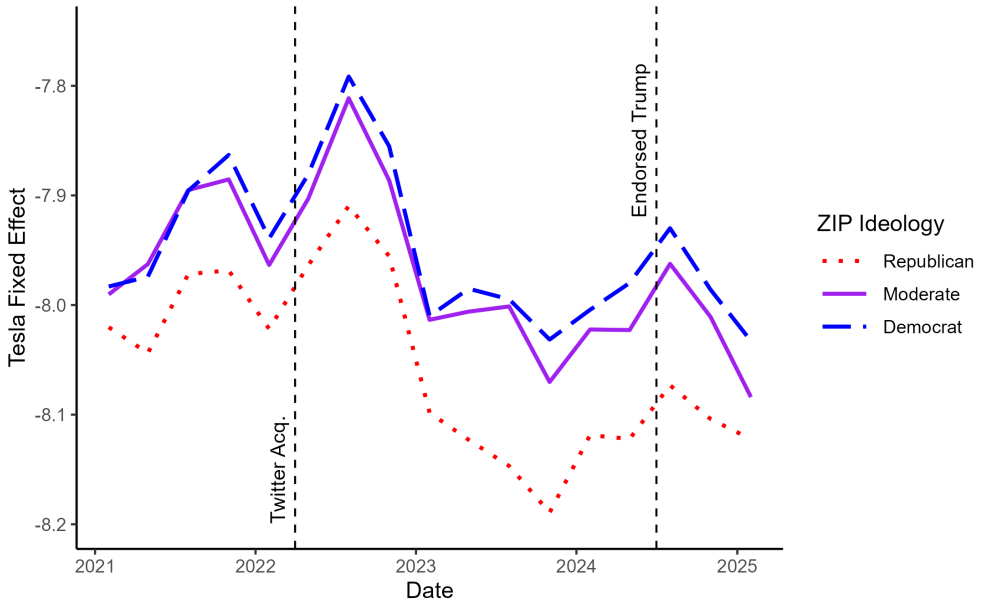
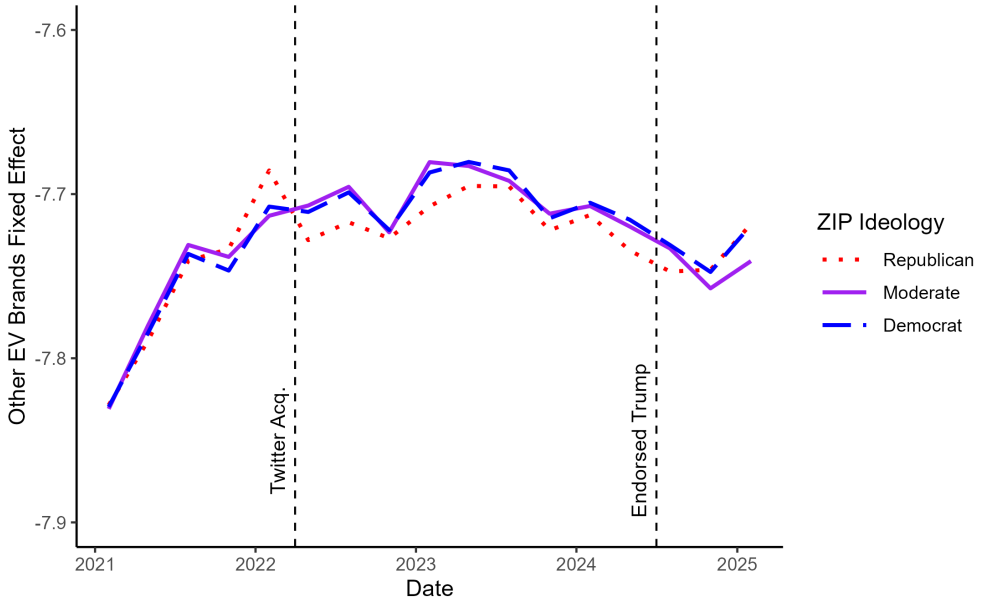


Figure 2.7: Brand Fixed Effects for American Brands

Notes: Figure depicts estimates of fixed effects for the top four American brands by observed new registrations. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Fixed effects are all relative to the excluded group, which is the outside option of not purchasing a new vehicle. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).



Panel A: Tesla Fixed Effects



Panel B: Average Fixed Effects for Other Brands Selling EVs

Figure 2.8: Trends in Fixed Effects for Tesla and Other Brands Selling EVs

Notes: Figure depicts estimates of fixed effects for Tesla (Panel A) and for other brands weighted by the number of EVs they sell (Panel B). Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Fixed effects are all relative to the outside option of not purchasing a new vehicle.

We plot WTP for Teslas in Figure 2.9, which shows that WTP for Teslas is initially largest in Democrat ZIP codes and smallest in Republican ZIP codes.[25] When compared to the median brand, Democrat ZIPs are initially willing to pay $6,000 less for a Tesla. This slightly exceeds the value of -$7,000 that moderate ZIPs place on Tesla. Republican ZIPs have the lowest WTP, equal to approximately -$8,000.[26]

We show that WTP for Teslas trend similarly by ideology until Musk's involvement in the 2024 election, which precedes a decrease in WTP for Teslas in Democrat and moderate ZIP codes. We estimate a decrease in WTP approximately equal to $8,000 for both Democrat and moderate ZIP codes, but no change in WTP among Republican ZIPs. The magnitude of these losses are such that WTP for Teslas is now less than WTP for the median brand by about $20,000 among both Democrats and moderates.[27] The figure shows also that WTP for Teslas did not increase among Republicans, highlighting that further barriers exist to universal EV adoption beyond one brand's ideology. This period represents the first break in WTP trends between ideologies, which is encouraging for our interpretation that the effect is driven by Musk's involvement in the 2024 election.[28]

25. We provide also Appendix Figure B.13, which plots WTP trends with only two political classifications based on being greater than or less than a 50% vote share for Trump in 2020.

26. In Appendix Figure B.1, we show that all results are robust to altering the cut-off for what is considered Democrat, moderate, and Republican ZIP codes. Results are qualitatively unchanged when considering cutoffs on 2020 Trump vote shares that vary between $\{45\%, 55\%\}$ and $\{30\%, 70\%\}$.

27. We use 2024 Trump vote shares in Appendix Figure B.2, showing that the results are largely robust. The primary difference is that the moderate trends in WTP are attenuated, which is driven largely by the fact that only very Democratic ZIPs had less than 40% vote share for Trump in 2024, as depicted in Appendix Figure B.14.

28. We run a permutation test in which vehicle market share is permuted within each month by ZIP code and we then run the logit demand estimation. This is designed to nonparametrically approximate p-values for the falls in WTP for Tesla that we observe. As depicted in Figure B.11, we observe p-values of less than 0.01 for falls in Democrat and moderate Tesla fixed effects and a p-value of 0.744 for the fall in Republican Tesla fixed effects.

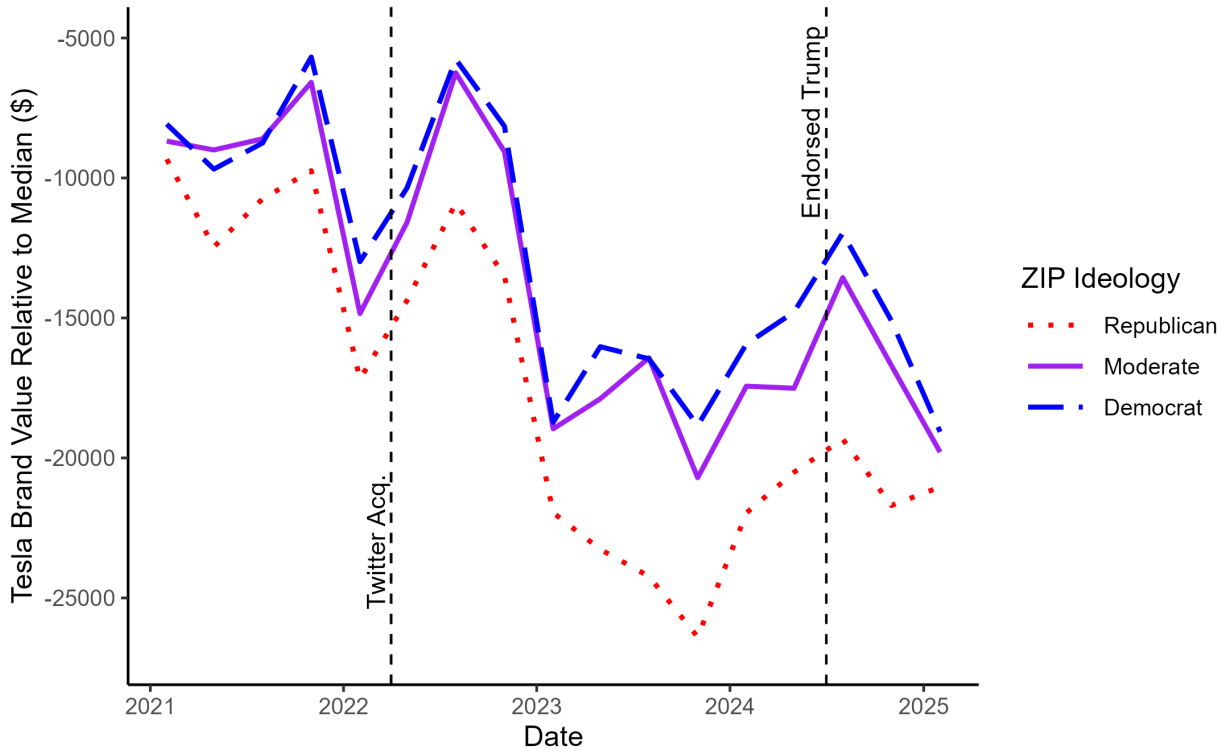


Figure 2.9: Trends in WTP for Teslas Relative to the Median Brand

Notes: Figure depicts estimates of consumer WTP for a Tesla vehicle relative to the median brand. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Tesla's quarter-year specific fixed effects are then compared to the median of quarter-year fixed effects of all other brands for each political leaning. This is then converted to WTP for a brand by normalizing the difference in fixed effects by an estimated price coefficient. Price is instrumented for by variation in eligibility for federal subsidies generated by the 2022 Inflation Reduction Act. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).

2.4.4 Alternate Mechanisms

Anticipated Policy Under Trump in 2024

We rule out anticipated EV policy changes under Trump as drivers of the fall in WTP for Teslas during the 2024 election. Figure 2.8 depicts average fixed effects for all other brands that sell EVs in Panel B, showing that no fall in WTP occurs for such brands. In Appendix Figure B.6, we also plot the fixed effects for Rivian, the only other brand with a significant presence in the U.S.A. which primarily sells EVs.[29] We again see no fall in the fixed effect for Rivian. This evidence suggests that the fall in WTP during the 2024 election is Tesla-specific, rather than reflecting a general anticipation of policies that will detrimentally impact EVs more broadly.

Subsidy Policy Changes

We rule out subsidy policy changes under the Inflation Reduction Act as a driver of our estimates of the fall in WTP during the 2024 election. Alterations which enhanced Tesla's eligibility or restricted the eligibility of Tesla's competitors occurred in August 2022, January 2023, and April 2023.[30] These alterations to subsidy eligibility did not drive the relative fall in WTP during the 2024 election, since they occurred well before 2024Q3.

However, these subsidy changes may play a role in driving our estimates of WTP across all ideologies. We observe a particularly large fall in brand value for Tesla for all three ideologies in 2023Q1, when Tesla gains subsidy eligibility contingent on being below an MSRP threshold. This MSRP threshold may drive the fall in brand value that we observe

29. Several other brands sell primarily EVs, including Lucid, Polestar, and Vinfast. These brands have all sold fewer than 750 vehicles across our data, making their brand fixed effects difficult to estimate.

30. The August 2022 changes made foreign-made EVs ineligible to receive the federal subsidy. The January 2023 changes removed a previous 200,000 sales limit and replaced it with restrictions on subsidy eligibility for EV cars with MSRPSs above $55,000 and EV SUVs with MSRPs above $80,000. The April 2023 alterations impacted eligibility for Rivian, Jeep, Audi, BMW, and Nissan models based on battery material sourcing.

across all ideologies in 2023Q1 because many Tesla models in late 2022 have MSRPs several thousand dollars above the MSRP thresholds, making them ineligible for the subsidy. As such, we observe that Tesla in January 2023 reduces prices by a greater amount than other brands when they gain eligibility.[31]

Competitor Entry

We control for when competitors enter vehicle markets in our regressions, mitigating the concern that competitor entry drives our estimates of WTP for Teslas. Table 2.2 shows that the inclusion of these controls does little to change our coefficient estimates. Moreover, Appendix Figure B.16 presents trends in the number of first-time make-model-year registrations (excluding Teslas) within a ZIP code, aggregated by date and ZIP ideology. We find that there is no spike in either EVs nor ICEVs entering the market during the 2024 election.

Charging Infrastructure Expansion

We rule out expansion to public charging infrastructure as a confounding mechanism by controlling for ZIP level charging station availability. Table 2.2 again shows that there is minimal impact to estimates. We present also charging station data from Alternative Fuels Data Center [2025] in Appendix Figure B.8. Examining infrastructure trends, it is clear that charging stations both for Tesla and non-Tesla EVs are more prevalent in Democrat ZIPs and less prevalent in Republican ZIPs.[32] However, the expansion of charging stations appears to largely be parallel since January 2023, and there do not appear to be discrete changes in charging availability during the 2024 election.

31. See, for instance, this article on the January 2023 Tesla price cuts exceeding subsidy amount: `https://www.reuters.com/business/autos-transportation/tesla-cuts-prices-electric-vehicles-us-market-2023-01-13/`.

32. We observe that the overwhelming majority of public charging stations (74%) over our study period are accessible both to Tesla and Non-Tesla EVs.

Supply Chain Constraints

We address concerns about Tesla's supply chain constraints in two ways. First, we measure fixed effects at the quarter level which to an extent alleviates concerns about shortages which last less than three months. Second, we plot in Figure B.17 the measure of Teslas in inventory as scraped from Tesla's website by independent data collectors [Tesla-Info.com, 2025, Internet Archive, 2025]. This figure shows that Tesla had many vehicles in inventory for the 2024Q3 through 2025Q1 drop in WTP for Teslas. This suggests that supply chain constraints do not drive the drop in WTP for Teslas during the 2024 election, though unfortunately historical ZIP level data on wait times is to our knowledge unavailable.[33]

2.5 Welfare Impacts

In this section, we use our demand estimates to calculate the welfare impacts of ideological preferences over Teslas and the ensuing substitution to other vehicles. We show that Democrat and moderate consumers substituted primarily towards ICEVs or not purchasing a vehicle during the 2024 election. This led Tesla's revenue to decrease, the social damages from vehicular emissions to increase, and private surplus to decrease.

2.5.1 Firm Revenues

We use a counterfactual simulation from our demand estimates to quantify how firm revenues shifted as a result of Musk's 2024 political activity. We consider as a baseline the Texas vehicle market in the second quarter of 2024, right before Musk became involved in the 2024 election cycle. We then simulate how vehicle sales would change in 2024Q2 if WTP for Teslas had fallen to the same extent as it did following Musk's involvement in the 2024 election. We hold the set of products and characteristics (including prices) fixed in order to perform this

33. Appendix Figure B.12 depicts wait times as gathered from Tesla.com and the Wayback Machine for January 2021 - December 2022. These wait times can not be disaggregated to ZIP levels.

exercise. We are thus essentially estimating the extent to which demand shifts in a simple model of supply and demand, which is distinct from an analysis of equilibrium impacts.

We show that consumers substituted away from Teslas and towards Tesla's competitors in Table 2.3, leading Tesla's annual revenues to fall by $2.3 billion.[34] The effects are presented separately for Democrat and moderate ZIPs, with Republican ZIPs having zero effect. We estimate that Tesla's revenue fell by $955 million in moderate ZIPs. The effect is larger in Democrat ZIPs, $1.3 billion, reflecting that the change in WTP for Tesla is slightly larger among Democrat ZIPs. Consumers of both political views appear to substitute largely towards either buying new ICEVs or choosing the outside option (i.e. buying a used vehicle or keeping their old vehicle). Overall revenues for ICEV manufacturers increase by $544 million, as compared to a smaller $316 million increase for manufacturers of EVs.

2.5.2 Social Surplus Loss from Emissions

As depicted in the top panel of Table 2.3, Tesla's newfound ideological association led many consumers to substitute to ICEVs which emit greenhouse gases from their tailpipe. Many other consumers substituted to EVs not manufactured by Tesla. As Tesla makes are among the most energy-efficient EVs on the market, this likely increased electricity consumption and raised emissions from upstream electricity generation.[35]

In the second panel of Table 2.3, we estimate that emissions of carbon dioxide equivalents increased as a result of Musk's involvement in the 2024 election cycle. We assume that all vehicles drive approximately 13,500 miles per year in order to generate estimates of annual

34. All estimates are scaled to the U.S.A. population of moderates and Democrats using 2024 population estimates. We assume that the overall population of the U.S.A. and Texas are 340.1 million and 31.3 million, respectively. To account for Texas being more Republican, we then additionally scale estimates based on the vote shares for Trump in the 2024 election: 49.8% in the U.S.A. and 56.1% in Texas.

35. These substitution patterns also likely played a role in traffic mortality (via vehicle weight) and air pollution. Because Tesla's most popular models are relatively light vehicles, the social welfare impacts we present are likely lower bounds.

Panel	Vehicle Group	Democrat ZIPs	Moderate ZIPs	Total
Firm Revenues (USD Millions)				
Revenues	Tesla	-1,312	-955	-2,267
Revenues	Competing EVs	191	125	316
Revenues	Competing ICEVs	396	248	544
Social Welfare (USD Millions)				
CO_2 Increases	All	-18	-14	-32
Private Surplus (USD Millions)				
Consumer Losses	All	-839	-722	-1,561

Table 2.3: Annual Welfare Impacts of Musk's Involvement in the 2024 Election

Notes: Table depicts welfare results for a counterfactual simulation in which Elon Musk's involvement in the 2024 election and the ensuing Tesla brand damages occur in quarter 2 of 2024. All other vehicle characteristics and prices are held constant. The firms panel depicts changes in revenue. The consumers panel depicts changes in private surplus. The emissions panel depicts impacts to social welfare driven by carbon dioxide emission increases. All estimates cover one year worth of impacts scaled to the population of the U.S.A. To calculate emissions, all vehicles are assumed to drive 13,500 miles per year and the outside option of buying a used vehicle or keeping one's own vehicle is assumed to yield a 20 MPG vehicle. The category 'Competing EVs' includes battery EVs and plug-in hybrids. The category 'Competing ICEVs' includes vehicles powered solely by gasoline and diesel as well as hybrids which cannot run solely on electricity.

emissions.[36] We assume also that choosing the outside option of a used vehicle or keeping one's own vehicle results in a vehicle that obtains 20 miles-per-gallon.[37] Using a social cost of carbon of $241 per metric ton (EPA, 2024), we estimate increases to social damages from emissions that equate to $18 million dollars per year in Democrat ZIP codes and $14 million dollars per year in moderate ZIP codes.

2.5.3 Private Surplus Loss

Though consumers are able to buy substitutes for Teslas following Musk's political involvement, the devaluation of Tesla as their first choice led to losses in private surplus. These losses in private surplus are welfare-relevant and represent a measure of the impacts of politically active brands on consumers in an increasingly polarized society.

We estimate the impacts to private surplus from Musk's involvement in the 2024 election and present our estimates in the final panel of Table 2.3. We estimate losses in private surplus of $722 million for moderates and $839 million for Democrat ZIP codes. In total, this represents a total loss of $1.6 billion in private surplus for *one year's* worth of sales.[38] This effect is likely to be larger when considering that the effects of Musk's political activity may last longer than one year. This effect is also likely to be a lower bound for overall private surplus loss because it does not account for private surplus losses accrued by individuals who had previously bought a Tesla before the 2024 election.

36. Recent work suggests that EVs may be driven less than ICEVs [Burlig et al., 2021]. This indicates that our estimates of impacts to carbon emissions may be lower bounds if buying an EV induces individuals to drive less.

37. This assumption is made because used vehicle markets are primarily ICEVs and the typical used vehicle gets around 20 miles-per-gallon (`https://afdc.energy.gov/data/10310`).

38. This effect is reasonable compared to the estimated private surplus loss of $5.1 billion from removing EV subsidies, as seen in Allcott et al. [2024a].

2.6 Conclusion

We have used a natural experiment and a dataset of vehicle registrations to provide three main findings on demand for environmental goods. We have first shown that WTP for vehicle brands is heterogeneous by ideology. We find that Republicans have a relatively high valuation of American brands and Democrats foreign brands. More pertinent to the setting of EVs, Democrats and moderates historically value the Tesla brand to a greater extent than Republicans. This has larger implications not just for EV adoption, but for the adoption of other environmental goods, as most tend to be seen as aligned with Democrat values.

Second, we have leveraged quasi-experimental variation in ideological associations of the Tesla brand based on CEO Elon Musk's political activities to show that WTP for environmental goods is influenced by ideology. We establish that Musk's involvement in the 2024 Presidential Election cycle led to absolute decreases in WTP for Teslas of $8,000 among Democrat and moderate ZIP codes, with no change in WTP for Teslas occurring in Republican ZIP codes. We are able to rule out anticipated policy shocks, past subsidy policy variation, competitor entry, charging infrastructure expansion, and supply chain constraints as alternative mechanisms driving differential trends in our estimated WTP values.

Finally, we have shown that ideological preferences for brands have an impact on private and social surplus. We show that consumers substituted away from Teslas and towards other vehicles with higher emissions intensity. This redistributed revenues across vehicle manufacturers, with Tesla losing $2.3 billion dollars per year due to Musk's involvement in the 2024 election. A large portion of revenue increases accrued to competing ICEV producers, causing social damages from vehicular emissions to rise by an estimated $32 million per year. We estimate also that consumers were subjected to $1.6 billion dollars in lost private surplus per year, highlighting that consumers can lose from brand political activity. Our analysis of Texas may be a lower bound for effects on the rest of the U.S.A. because Texas has a relatively Republican political climate and because Tesla is headquartered within Texas.

Such results enhance our knowledge of consumer behavior in an increasingly ideologically polarized society. To the best of our knowledge, this is the first paper to estimate how consumers trade off between ideology and environmental impact. Similar decreases in welfare likely exist for other environmental goods, and consumer welfare is also likely impacted by ideological boycotts of Amazon, Target, Walmart, Nike, Twitter/X, Meta, Google, and an array of other brands.

Several other open questions are likely interesting for future work. First, our paper is unable to thoroughly examine impacts to consumer welfare beyond private surplus from product choice. Individuals may experience a significant amount of cognitive distress when polarization becomes so pervasive that it infiltrates their daily consumption purchases. Moreover, we have focused on a relatively high-end vehicle for which a significant number of substitutes are available. Consumer welfare impacts may be different for Amazon boycotts in which there are few close substitutes, or Pfizer boycotts which pertain to goods that directly impact human health. Future work could thus look to expand into other contexts in which brands are becoming increasingly political and determine what welfare impacts exist.

Second, our paper is limited in terms of its breadth of geographical coverage. As discussed in the paper, our usage of Texas data likely depresses our estimates of the impacts of Tesla's ideological rebranding due to the Republican lean of Texas and the fact that Tesla is headquartered in Texas. Other datasets, such as those provided by WardsAuto or Experian, are available for purchase at the ZIP-quarter level. The costs of acquiring such datasets at this level of granularity often ranges in the millions, likely beyond what is feasibly attainable for many individual researchers. We thus encourage any individuals or groups with access to this data to engage in a broader analysis of Musk's impacts on the EV market. Such exercises will provide a clearer understanding of the future of EV demand.

Finally, we believe that a thorough analysis of the precise behavioral mechanisms which drive substitution away from ideological goods would be useful. In particular, Tesla EVs are

highly visible, representing one of the most recognizable EV brands on the market [Davis et al., 2025]. The substitution away from Tesla following Musk's behavior may thus in part be driven by peer effects in addition to any individual distaste that consumers hold for the brand. Decomposing the boycotting of goods into a private tastes term versus a social utility term would be informative as to which goods in which scenarios are more or less likely to experience brand devaluation.

CHAPTER 3

URBAN EXPANSION, DROUGHT RISK, AND WILLINGNESS-TO-PAY FOR PIPED WATER

3.1 Introduction

The urban population in the Global South is projected to expand 81% from 2020 to 2050 [World Resources Institute, 2021]. While this urbanization is driving economic growth, it is also pushing housing to suburbanize faster than public infrastructure can expand [George et al., 2019]. Governments are faced with a decision between making costly expansions to infrastructure or encouraging settlement in areas that already have infrastructure. This decision has implications for a city's congestion, productivity, and climate resilience [Akbar et al., 2023b,a, Harari, 2020, Ahlfeldt et al., 2023, Rosenthal-Kay, 2025].

Suburban households without access to public infrastructure often turn to private substitutes. These private substitutes frequently are of higher risk or lower quality than their public counterparts [Filion and Keil, 2016]. Despite widespread reliance on privately-supplied roads, electricity, internet, transport, water, garbage collection, and security, the role of this quality disparity in driving urban expansion is largely unquantified.

In this paper, we study how consumers and real estate developers respond to a climate-driven change in the risk of relying on privately-supplied borewell groundwater.[1] We conduct our study using Bangalore's unanticipated 2024 water crisis. During this crisis, the half of Bangalore's 14 million residents without access to publicly-supplied piped water experienced groundwater shortages for 2 months. Residents without piped water were subjected to rations, increased prices, and detrimental health impacts during this period. Residents with

. This chapter is in equal collaboration with Jun Wong and Emma Zhang.

1. The risk of urban groundwater shortage is rapidly increasing worldwide. Experts estimate that 97% of the cities located in the global south have higher chances of experiencing droughts, which are also becoming 73- 88% more intense [Stolte et al., 2023].

piped water were largely unaffected during this crisis and are likewise insulated from the risk of future droughts.[2]

We first analyze a novel dataset of monthly rental listings across Bangalore to show that the water crisis increased demand for piped water access. We observe each listing's exact location, monthly price, and amenities. This allows us to link proximity to the piped water network with rental prices conditional on a set of other characteristics. Hedonic valuations find that residents' willingness-to-pay (WTP) to live 100 meters closer to the pipe network permanently increased by 0.4% of monthly rents following the drought.[3] This premium is permanent for long-term listings but lasts only during the drought for short-term leases that are less exposed to future drought risk. We find that the estimate is particularly large (0.5% per 0.1km) for units occupied by families, highlighting the uneven distribution of impacts from drought across households.

To explain this shift in demand, we then combine original household survey data with administrative data on hospitalizations to show that groundwater shortages have negative impacts to human health. Out-patient hospitalizations spike at the beginning of the drought and in-patient hospitalizations remain at elevated levels throughout the water crisis.[4] Heterogeneity analysis of the household survey finds that the health effects were attributable to the water crisis. We observe reported health issues that are 3 percentage points higher for those who reported relying on tanker water – which has high contamination risk – during the crisis.

With demand shifting to favor housing with access to piped water, we show that long-run

2. Households relying on piped water are insulated from drought risk in Bangalore because Bangalore's pipes draw from surface water sources outside the city that are resilient to depletion. This does not necessarily reflect the situation for individuals residing in cities where the piped water system draws from groundwater (e.g. Mexico City) who are not insulated from groundwater shortages.

3. Listings range in distance from under 0.1km to up to 3km, implying that the listings farthest from water pipes experience a predicted rent decrease of 12%.

4. The water crisis also corresponds to reports of cholera outbreak in the city at the time: `https://economictimes.indiatimes.com/small-biz/sustainability/from-dump-to-data-how-this-start-up-is-fixing-indias-waste-management-issue/articleshow/121205721.cms`

supply of housing adjusted to areas near water pipes. We divide Bangalore into 50 distinct geographic areas and classify each area by its access to water pipes. We use a difference-in-difference approach to compare areas close to the piped water network relative to those farther from the piped water network. We estimate with this approach that construction increased by 25-51% in areas near the piped water network relative to areas farther from the piped water network. This implies that adaptation via the housing market to some extent reduces the damages of urban droughts in the long-run, decreasing the necessity for utilities to make costly expansions to the pipe network.

Back-of-the-envelope calculations suggest that the housing market response is more cost-effective that pipe expansions. We conservatively estimate that the housing expansion increased consumer surplus by 20% as much as a counterfactual policy of expanding piped water access to be universal. The cost of the housing expansion, however, is only 2.4% of the cost of such a pipe expansion. This finding highlights that encouraging housing development in areas serviced by public utilities (e.g. via zoning policy, information campaigns, or subsidies) is a cost-effective policy approach for governments to improve social welfare for cities with limited public infrastructure.

Our paper contributes to the literature on urban expansion in developing countries by examining how the relative quality of public infrastructure influences such expansion. Many paper have previously studied the causes and effects of the way a city expands [Harari, 2020, Ahlfeldt et al., 2023, Akbar et al., 2023b,a, Balboni et al., 2020, Pratama et al., 2022, Rosenthal-Kay, 2025, Sorin Le Guével, 2025]. We contribute to this literature in two ways. First, we contribute by examining the role of a change in relative quality of a publicly-provided amenity in driving urban expansion. Second, we contribute as the first to study piped water access in driving urban expansion. In particular, we show that rising drought risk limits the extent to which cities will expand into the suburbs, a finding which has important implications for cities with similar water infrastructure to Bangalore, such as

Mumbai, Jakarta, New Delhi, Lima, and Nairobi.

We also contribute to the literature on adaptation to environmental disamenities by examining highly localized housing demand and supply responses. The conceptualization of consumer valuation of their surrounding environment dates back to Rosen [1974] and Roback [1982] and has been empirically examined in many settings [Chay and Greenstone, 2005, Greenstone and Gallagher, 2008, Haninger et al., 2017, Hornbeck, 2012, Hornbeck and Naidu, 2014, Desmet and Rossi-Hansberg, 2015, Desmet et al., 2018, Cruz and Rossi-Hansberg, 2024, Aron-Dine, 2024]. We contribute to this literature in two ways. First, we examine the extent to which a natural disaster with spatially differentiating impacts can lead consumers to weigh intra-city migration versus other available adaptation mechanisms. Second, we are among the first to study how housing *supply* responds to these environmental shocks.

Finally, we contribute to the literature on water access in low- and medium-income contexts by estimating how groundwater drought risk influences WTP for piped water. A large body of work has examined demand for and impacts of piped water and other clean water alternatives [Galiani et al., 2005, Gamper-Rabindran et al., 2010, Devoto et al., 2012, Kremer et al., 2011, Luoto et al., 2011, Szabo, 2015, Dupas et al., 2016, Null et al., 2018, Alsan and Goldin, 2019, Haushofer et al., 2021, dos Santos and Guidetti, 2024, Burlig et al., 2025]. These papers have little focus on the role of rising drought risk, excepting Abajian et al. [2025] who examine "Day Zero" where *piped* water was nearly entirely depleted in Cape Town. We contribute by studying the impacts of borewells drying up, as opposed to piped water. It is more difficult for households with drying wells to obtain new access to piped water (without migrating) than it is for households with drying pipes to drill new wells. We contribute also by analyzing WTP as an outcome, allowing us to quantify residential valuations of drought risk.

The rest of the paper proceeds as follows: Section 2 provides background details, Section

3 describes our data, Section 4 presents reduced form results, Section 5 conducts simple back-of-the-envelope calculations, and Section 6 concludes.

3.2 Background

3.2.1 Bangalore Water Supply

There are two sources of water used to meet the needs of Bangalore's residents: (1) piped water and (2) borewell groundwater. Piped water comes from the Cauvery river and is available in the center of Bangalore. Borewell water supplies groundwater to residents in Bangalore's suburban periphery.

The pipe system was largely developed around the historical city center of Bangalore [Castán Broto et al., 2021] and is managed by the Bangalore Water Supply and Sewage Board (BWSSB). Residents of single-family homes that have access to BWSSB pipes pay a metered fee each month for water usage. Residents of apartment buildings that have access to BWSSB pay a fixed building maintenance fee each month which encompasses the cost of water.[5] An alternate way to access river water is to utilize public pumps, which are connected to the pipe system. The BWSSB pipe system services about 50% of the city [Kulranjan et al., 2023].

Bangalore's suburban population relies on groundwater, which is obtained from privately-managed borewells. Groundwater reliance occurs in areas where housing developments outstripped the pace of pipe infrastructure expansion – typically in newly developed peripheral areas or in buildings commissioned by individual developers.[6] While BWSSB manages

5. The building maintenance fee can change month-to-month as water usage does change between the dry and wet seasons.

6. Like many cities in the Global South, Bangalore has experienced rapid urbanization over the past several decades. The city's population has surged from approximately 8.5 million in 2011 to an estimated 14 million in 2024. This surge in population has led to increased demand for housing, infrastructure, and amenities, particularly in the city's suburban outskirts [Bangalore Mirror Bureau, 2024].

around 6,000 public borewells, the overwhelming majority of borewells are drilled privately for personal use [Castán Broto et al., 2021].[7]

The presence of piped water only in the urban downtown of Bangalore creates disparities in water access between those residing downtown and those residing in the suburbs. Figure 3.1 overlays on a map of Bangalore the location of water pipes throughout the city, depicting that water pipes are densely populated throughout the urban center of the city. Many portions of Bangalore's suburban outskirts have minimal access to this piped water and instead must rely predominantly on groundwater. This yields many rental listings which are distant from water pipes. Figure 3.2 displays the number of rental listings' distances from the nearest pipe, highlighting the fact that thousands of listings are distant from the nearest water pipe.

Water pipes can be expanded, but these expansions are costly, difficult, and take a large amount of time. A recent expansion of water pipes into Bangalore's suburbs is depicted by the red lines in Figure 3.1. This expansion was announced in 2014 and only began supplying water in limited capacities from September 2024. The expansion is estimated to have cost approximately $500 million to provide water connections for a planned 500,000 individuals [Japan International Cooperation Agency, 2024].

3.2.2 2024 Bangalore Water Crisis

From late February to late April of 2024, Bangalore experienced its worst drought in decades. A combination of below-average monsoon rainfall, reductions in natural replenishment rates, and increasing suburban demand led to severe water shortages for those relying on groundwater. Groundwater reservoirs supplying Bangalore reached critically low levels, forcing BWSSB to impose water rationing measures. Despite these measures, around 7,000 of the cities 16,000 borewells ran dry [Mogul et al., 2024].

7. Some experts suspect there are up to 500,000 total borewells in Bangalore [Cave, 2024].

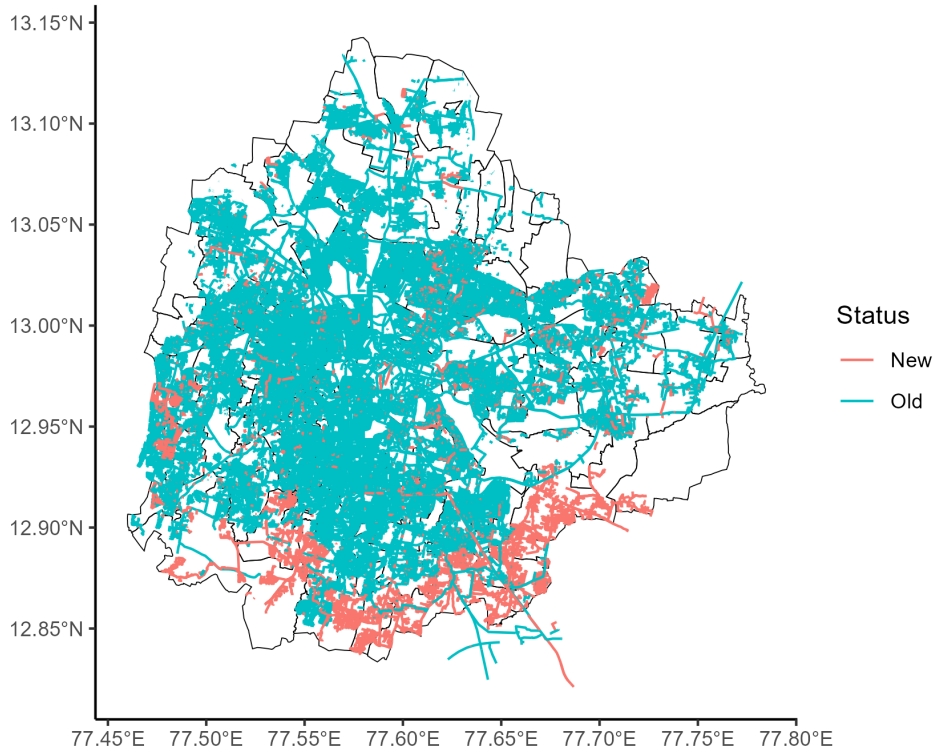


Figure 3.1: Water Pipe Locations

Notes: Figure depicts the map of all pipes in the publicly provided piped water network within Bangalore. Black outlines are utility wards in Bangalore. Blue and red lines depict water pipes, with red lines indicating pipes which were created in the September 2024 expansion and blue lines indicating which existed as of the prior expansion in 2014.

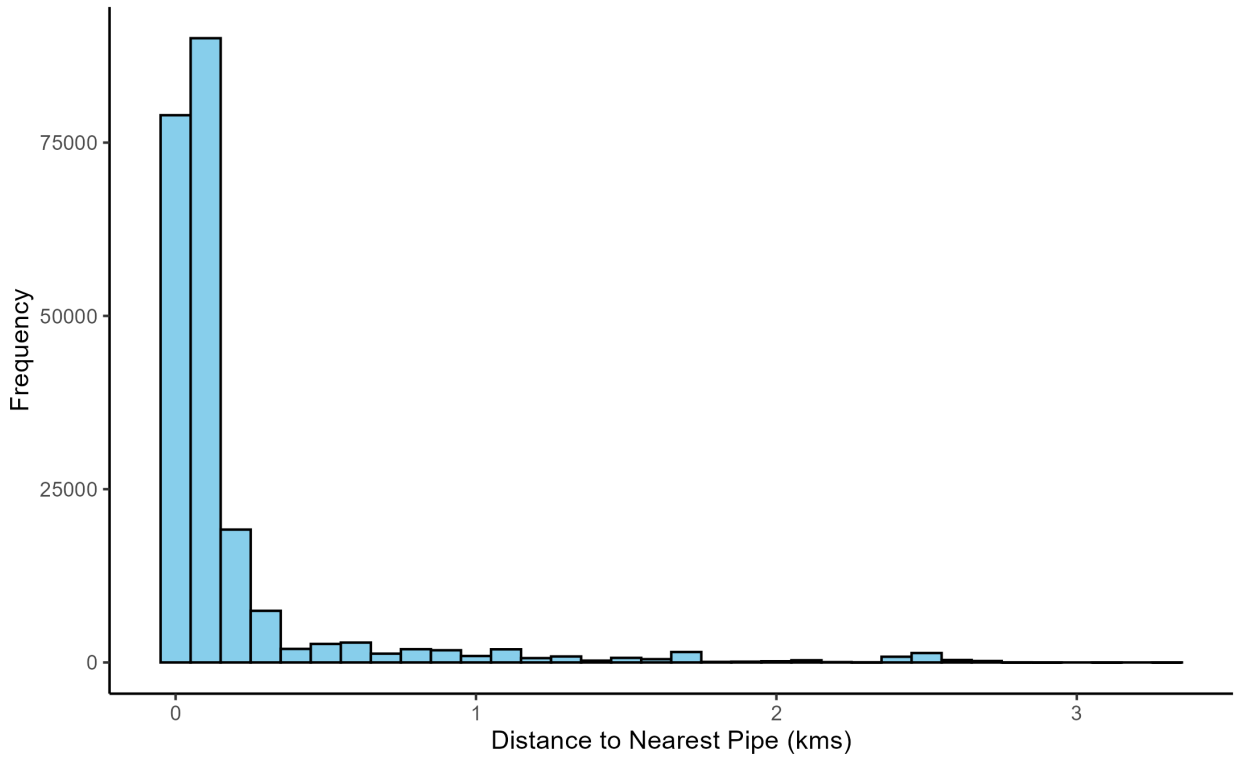


Figure 3.2: Listing Distances from Nearest Water Pipe

Notes: Figure depicts a histogram of distances of listings to the closest water pipe. Distance is measured in kilometers. The vertical axis is a count of listings. Listings across all months in our data are included.

This drought primarily affected households and businesses without access to piped water, as the Cauvery's river supply was unaffected.[8] Individuals without access to the pipe network were forced to ration water, with many reporting that they were unable to flush toilets, wash dishes, wash laundry, or take showers [Cave, 2024]. Impacted residents spent time waiting in long queues for tanker water, searching for public borewells that had not dried up, and purchasing bottled water for drinking and cooking. We estimate with survey data that median monthly water expenditures during the crisis increase by ₹642 for those reporting no access to piped water, with a more modest ₹250 increase among those reporting having a pipe connection.[9]

3.3 Data

We combine data from four sources to investigate housing demand, housing supply, water amenities, and micro-level impacts of the water crisis.

3.3.1 Housing Demand

We collect data on rental prices and housing amenities comes from Housing.com. Once each month, from May 2023 to April 2025, we scrape every active housing rental listing for Bangalore that is present on the site. The scrapes occur during the middle of each month and capture an average of 14,500 units each month. After omitting all listings outside of Bangalore's city limits, we observe a total of 347,426 listings across 24 months. Our analysis only includes listings which are residential.[10] In total, our analysis sample has 323,655 listings.

8. Appendix Table C.1 summarizes several other major water crises which have occurred worldwide.

9. The increase is driven by households that rely on water tanker water to supplement their pipe connection. Removing households that rely on water tankers yields an increase of ₹652 for those reporting not having a pipe connection and ₹0 for those reporting having a pipe connection.

10. Approximately 24,000 listings are meant for companies.

For each rental listing, we observe the exact location, monthly rent price, other monetary expenses associated with renting there (security deposit, brokerage, and maintenance), as well as a rich set of characteristics about the listing (furnishing, bathroom count, building age, etc.). Table 3.1 provides summary data on these amenities.[11]

Variable	Mean	SD	Min	Q1	Median	Q3	Max	N
Rent Price	36,007	29,701	200	17,000	28,000	45,000	270,000	323,655
Security Deposit	188,027	195,599	0	70,000	150,000	250,000	3,000,000	323,193
Brokerage	28,317	32,598	0	0.00	23,000	40,003	300,000	323,378
Maintenance	3,786	8,027	1	1,000	2,500	5,000	700,000	137,961
Square Footage	1,202	748	10	650	1,120	1,500	24,000	323,019
Furnished (yes = 1)	0.17	0.38	0	0.00	0.00	0.00	1.00	323,646
Bedrooms	2.01	0.90	1	1.00	2.00	3.00	10.00	323,536
Bathrooms	1.93	0.91	1	1.00	2.00	2.00	11.00	294,745
Balconies	1.14	1.01	0	0.00	1.00	2.00	9.00	295,792
Building Age	4.47	4.27	0	1.00	4.00	6.00	99.00	241,820
Gated Community (yes = 1)	0.21	0.41	0	0.00	0.00	0.00	1.00	323,655
Parking Spots	1.28	0.87	0	1.00	1.00	2.00	10.00	281,125
Building Height (Floors)	7.44	6.68	1	4.00	5.00	7.00	69.00	247,641
Nearest Park Dist. (km)	0.31	0.28	0	0.10	0.24	0.42	2.42	323,655
Nearest Office Dist. (km)	0.59	0.61	0	0.18	0.36	0.76	4.63	323,655
Nearest Transit Stop Dist. (km)	0.42	0.35	0	0.19	0.31	0.52	2.66	323,655
Nearest Shop Dist. (km)	0.26	0.29	0	0.08	0.16	0.32	3.14	323,655
Shops Within 500m	29.42	41.96	0	3.00	10.00	37.00	266.00	323,655

Table 3.1: Summary Statistics of Key Amenities

Notes: Table depicts key summary statistics for a variety of listing-specific variables in our data. All monetary characteristics are in nominal INR. The column 'N' reflects the number of listings for which the pertinent characteristic is not missing.

Figure 3.3 displays a map of wards in Bangalore along with dots indicating our scraped listings for four months – one well before the crisis (July 2023), one just before the crisis (December 2023), one during the crisis (March 2024), and one after the crisis (August 2024). Our data captures listings across Bangalore's urban downtown and suburban sprawl. The listings include both low-price and high-price units, reflecting that the site captures a range

11. For many listings, characteristics such as building height, building age, and balcony count are optional to include in the Housing.com posting. This leads to a large amount of missing data for important characteristics that are likely correlated with both piped water proximity and the price of a listing. Fortunately, we observe a rich set of characteristics for every single listing – including bedroom count, square footage, and listing type (e.g. "flat", "apartment", "independent builder floor", etc.). We thus handle missing characteristics data by employing a multiple imputation strategy that uses observed characteristics to predict missing characteristics within each listing. This strategy is discussed in more detail in Appendix for Chapter 3.

of listings for all types of renters.

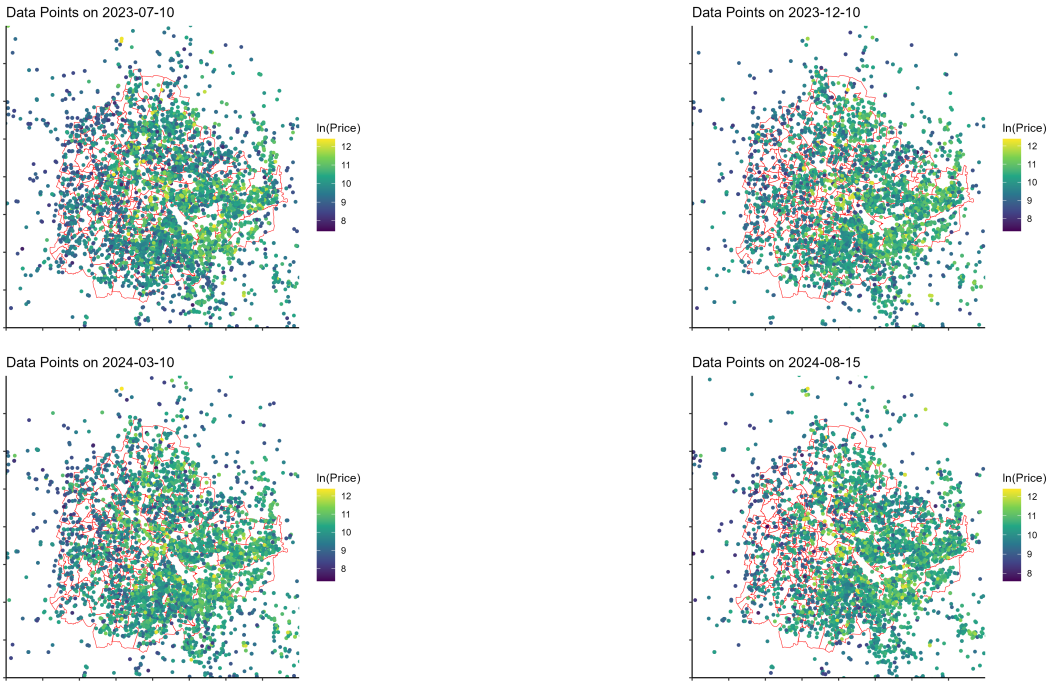


Figure 3.3: Locations of Scraped Rental Listings by Month

Notes: Figure depicts all scraped rental listings from Housing.com for four separate months. The water crisis occured from February to April of 2024. Red outlines delineate utility in Bangalore and each dot is a scraped listing for that month. Dots vary by color based off the log of the monthly rental price of the listing. All listings outside the outer edge of utility wards are removed from the analysis.

3.3.2 Long-Run Housing Supply

Our data on long-run housing supply comes from Karnataka's Department of Housing, which collects data on in-progress construction to adhere to the 2016 Real Estate Regulatory Authority (RERA) Act. This act, meant to ease grievances of housing stakeholders, requires that all real estate projects of at least 500 square meters or at least 8 apartments be registered with the respective state RERA authority. RERA requires that each real estate developer submit an application to the relevant state authority with estimated costs and duration of

construction.

We scrape data on each submitted real estate project application from January 2023 to April 2025. Each application includes details of the constructing enterprise, project start and proposed end dates, progress reports, building specifications, and the exact project location. We also observe cost of construction, area measures, and a description of the listings to be provided within the building. The data also includes details on whether the application was rejected or approved, though rejections are rare in practice.[12]

Table 3.2 provides summaries of key variables for all registered development sites within Bangalore's governing limits. The average construction duration is approximately 3.5 years, reflecting the importance of monitoring in-progress construction to understand long-run housing supply adjustments to the water crisis.

Variable	Mean	SD	Min	Q1	Median	Q3	Max	N
Est. Completion Time (Years)	3.58	1.63	0.24	2.42	3.41	4.68	11.77	645
Promoter Project Count	2.09	3.03	1.00	1.00	1.00	2.00	18.00	645
Total Land Area (sq. mtr.)	14,117	20,814	557	4,012	8,328	16,444	289,573	464
Total Land Cost (Million INR)	422	1,112	0	62	146	361	18,747	616
Est. Total Cost (Million INR)	1,731	3,978	0	230	569	1,655	59,690	596

Table 3.2: Summary Statistics of Key Characteristics of Construction Projects

Notes: Table depicts key summary statistics for a variety of variables in our construction projects data. All monetary characteristics are in nominal INR. The column 'N' reflects the number of listings for which the pertinent characteristic is not missing.

3.3.3 Water Infrastructure & Other Public Amenities

We generate measures of access to piped water by utilizing the geographic locations of all water pipes managed by the Bangalore Water Supply and Sewage Board (BWSSB). The

12. Data is available on the RERA website dating back to 2017. However, the data from 2017 through 2022 is inconsistent and often missing key data, particularly location. As such, we focus on the data from 2023 and later.

water pipe location data comes from a shapefile that was generated by Karnataka State Remote Sensing Applications Centre (KSRSAC) and updated in October of 2024.[13]

We use the distance from a residential building to proxy for the cost of receiving piped Cauvery River water during a groundwater shortage. Buildings that are close to the pipe network are more likely to be connected to the network. It is also lower cost for unconnected households to obtain a connection when closer to the pipe network. Additionally, pipes are also known to leak into groundwater reservoirs, potentially dampening the effects of the water crisis for households near the pipe network even for buildings without a BWSSB connection.

We also generate access to other public amenities such as shops, parks, office buildings, and transit stops by using data from OpenStreetMap.[14] Table 3.1 provides summary statistics for our key public-amenity control variables. In particular, we calculate the distance of each listing from the nearest transit stop as a measure of access to the public transit network. We additionally calculate distances of each listing from the nearest park, office building, and shop to gauge neighborhood amenities. Finally, as a measure of density we calculate the numbers of shops within 500 meters of each listing. These provide a rich set of amenities which we control for in our regression analysis.

3.3.4 Survey Data

Our final data source is a household survey of renters in urban and suburban Bangalore, which allows us to understand residential moving patterns as well as the crisis' impact to health, water expenditures, and adaptive measures.

We generate a geographically representative sample of renters by randomly selecting 300

13. We sourced the shapefile from Open City, a public repository that aggregates data compiled by other government agencies.

14. OpenStreetMap is a collaborative open-source mapping software which is considered as up-to-date or more up-to-date than other API tools such as Google Maps in urban and suburban areas. We access the data using the OverPass Turbo API.

out of 2,423 locales in which we observe apartments. Locales were stratified by proximity to the nearest water pipe and whether the September 2024 pipe expansion yielded their closest pipe. Each of these 300 locales was visited by survey enumerators. At each location, the survey enumerators approached individuals outside and asked if they were renters living within a 3-minute walk and who were residing in Bangalore continuously since January 2024.

The survey was conducted over the time frame August 4-18, 2025 and collected a total of 1,251 responses, targeting 4 responses per locale. For each individual respondent, we collect address history from January 2024 to August 2025, pre- and post-crisis perceptions of drought risk, self-reported health issues during the crisis, adaptation behaviors, and demographics. Summary statistics of survey respondents are presented in Table 3.3.

Variable	Mean	SD	Min	Q1	Median	Q3	Max	N
Usual Water Spend (INR)	2,588	2,867	0	608	1,738	4,345	21,726	1,229
Crisis Water Spend (INR)	2,691	2,371	0	826	2,173	4,345	8,951	1,225
Moved 2024-Aug 2025 (=1)	0.12	0.33	0	0	0.00	0.00	1.00	1,250
Annual Lease (=1)	0.55	0.50	0	0	1.00	1.00	1.00	1,250
BWSSB Water (=1)	0.58	0.49	0	0	1.00	1.00	1.00	1,250
Distance to Pipes (km)	0.13	0.37	0	0	0.01	0.05	2.54	1,250
Crisis Health Issue (=1)	0.04	0.20	0	0	0.00	0.00	1.00	1,250
Unit Occupants	3.98	1.14	1	3	4.00	5.00	8.00	1,250
Secondary School (=1)	0.93	0.26	0	1	1.00	1.00	1.00	1,250
Bachelor's Degree (=1)	0.67	0.47	0	0	1.00	1.00	1.00	1,250
Unit Bedrooms	3.70	1.31	0	3	4.00	4.00	10.00	1,248
Inc. > 50,000 INR (=1)	0.65	0.48	0	0	1.00	1.00	1.00	1,250
Inc. > 100,000 INR (=1)	0.26	0.44	0	0	0.00	1.00	1.00	1,250

Table 3.3: Summary Statistics of Survey Respondents

Notes: Table depicts key summary statistics for a variety of variables in our survey data. All monetary characteristics are in nominal INR and capture monthly values. The column 'N' reflects the number of respondents for which the pertinent characteristic is not missing.

3.4 Results

We use a hedonic approach to uncover the change in WTP to be closer to the piped water network. We estimate an increase in WTP of 0.4% of rents per 100m proximity to the nearest pipe. As a consequence, we show with a difference-in-difference that long-run supply

of housing increased by 25-51% in areas near pipes compared to areas far from pipes.

3.4.1 Demand for Housing

Raw Price Trends

Baseline trends in rental prices suggest that the water crisis made listings that are farther from the piped water network less appealing. Figure 3.4 plots trends in average rental price separately based on proximity of the listings to the piped water network before and after the 2024 pipe expansion. The average prices of listings that are farther from water pipes tend to be higher before the crisis, likely reflecting a difference in amenities. Prices appear to trend in parallel until the crisis begins, at which point there is sharp drop in prices for listings that were far from pipes. The duration of this gap is dependent on anticipation of future water pipe access, as locations that anticipate the September 2024 pipe expansion immediately rebound, whereas listings that remain far from pipes do not.

Figure 3.4 highlights three facts which motivate our analysis. First, absent of any crisis units that are farther from water pipes tend to be priced higher. This likely reflects preferences for other attributes that are correlated with distance to pipes, highlighting the importance of controlling for listing characteristics in our analysis. Second, the crisis was associated with an immediate and large decrease in price for units that did not have access to piped water. Third, this decrease in rent prices was permanent for areas not expecting to gain pipe access in the near future, but only temporary for those that were. This highlights that future drought risk seems to drive down rental prices, as units with new access to piped water are insulated from future water shortages.

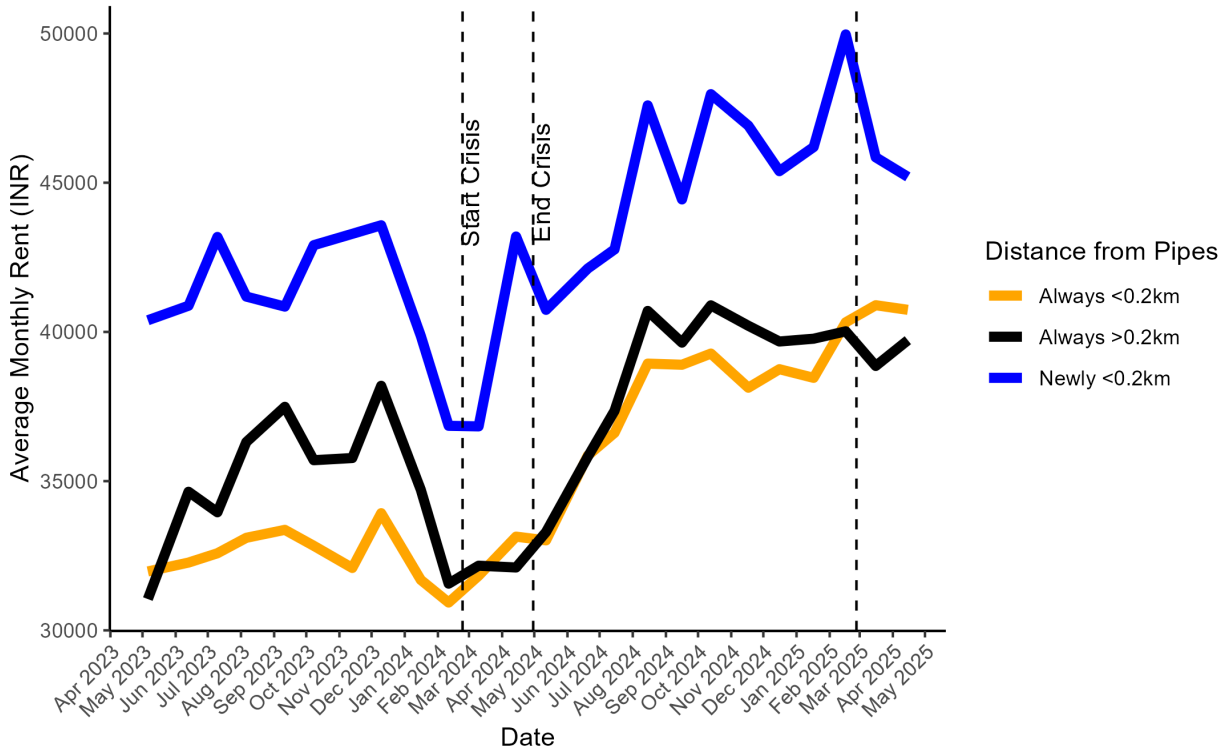


Figure 3.4: Average Price Trends by Distance from Water Pipes

Notes: Figure depicts raw trends in average monthly prices to rent listings in our data. The vertical axis is average monthly rent measured in nominal INR. The orange trend line is for listings less than 0.2km away from the nearest water pipe at all times before and after the September 2024 pipe expansion. The black trend line is for listings always at least 0.2km away from the nearest pipe both before and after the 2024 expansion. The blue line is for listings that were originally more than 0.2km away from the nearest pipe but moved within 0.2km of the nearest pipe due to the 2024 expansion. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

Empirical Specification

We run a hedonic analysis which isolates the impact of piped water access on rental prices of listed units before, during, and after the crisis [Rosen, 1974, Roback, 1982]. In particular, our estimates are interpretable as WTP to have piped water access in a given month.[15] In order to uncover this value, we run the following specification for listing i which resides in ward w during month t:

$$\ln(P_{it}) = \beta_0 + \left(\sum_{t'=1}^{T} \beta_{t'} \mathbb{I}_{[t=t']} KM_{it} \right) + \alpha A_{it} + \gamma_{wt} + \varepsilon_{it} \tag{3.1}$$

where P_{it} is the monthly rental price in INR, $\mathbb{I}_{[t=t']}$ is an indicator for being month t, KM_{it} is the distance (in kilometers) of a listing from the nearest BWSSB water pipe, A_{it} is a vector of listing-specific characteristics, and γ_{wt} is a ward-by-month fixed effect.[16]

Our coefficients of interest are $\beta_{t'}$ which capture WTP to be 0.1 kilometers closer to the nearest BWSSB water pipe separately each month. The inclusion of a variety of building-level control variables ensures that our estimates account for potential correlates of pipe distance that drive price, such as building age, building height, number of bedrooms, and other characteristics. Similarly, the inclusion of ward-by-month fixed effects ensures that we utilize only variation within a ward in a given month to estimate the effect.[17] This implies that all characteristics of a ward which do not vary more rapidly than on a monthly basis (e.g. infrastructure, public utilities, and other amenities) do not bias our results. Finally, the inclusion of other listing-specific amenities like nearby shop count and distance to the

15. Importantly, our data is not structured as a panel dataset because we do not observe listings once they are no longer posted. In Appendix Figure C.6, we present results that create a building-level panel and show that, while underpowered, the results are quantitatively similar.

16. For all listings prior to September 2024, this distance is calculated using distance from the nearest old pipe. From September 2024, all pipes – including newly expanded pipes – are included.

17. Note that there are 243 wards in Bangalore. Average ward size is approximately 9 square kilometers (about 3.5 square miles).

nearest park ensures that important within-ward correlates do not bias our estimates.

Average WTP Effects

We plot estimates of consumer WTP for the overall sample in Figure 3.5, showing that WTP for piped water increases as a result of the crisis. When the crisis begins, WTP for piped water increases by 0.4% of rents per 0.1 kilometers of distance from the piped water network. This implies that a listing that is directly on top of the piped water network is predicted to be approximately 1.0% cheaper than a listing that is the average distance (0.25km) from the nearest pipe. This premium appears relatively permanent, only attenuating slightly for a few months during Bangalore's wet season. The premium again sees a modest spike upon the advent of Bangalore's drought season in 2025, potentially reflecting fears of another crisis arising in early 2025.

WTP Effects by Lease Length

Figure 3.5 captures the preferences of all renters, including those on short-term leases, meaning it may underestimate the true impacts of long-term drought risk on WTP for piped water. Housing.com facilitates both short-term leases ($\sim$ 1-3 months) and long-term leases (up to a year), though we do not explicitly observe lease length in our data. As drought risk in Bangalore is seasonal (primarily high in February-April), many individuals signing short-term leases following the 2024 crisis may not have had any need to internalize future drought risk. This may attenuate our estimates in Figure 3.5 relative to individuals on long-term leases.

We present Figure 3.6 which depicts that individuals on longer-term leases have larger and lengthier WTP increases for piped water due to the crisis. In the left panel, we present the estimates only for listings which have security deposits worth at least 5 months of rent, the median in Bangalore. The right figure presents the results for listings which have security deposits of less than 5 months worth of rent. These security deposit sizes are likely highly

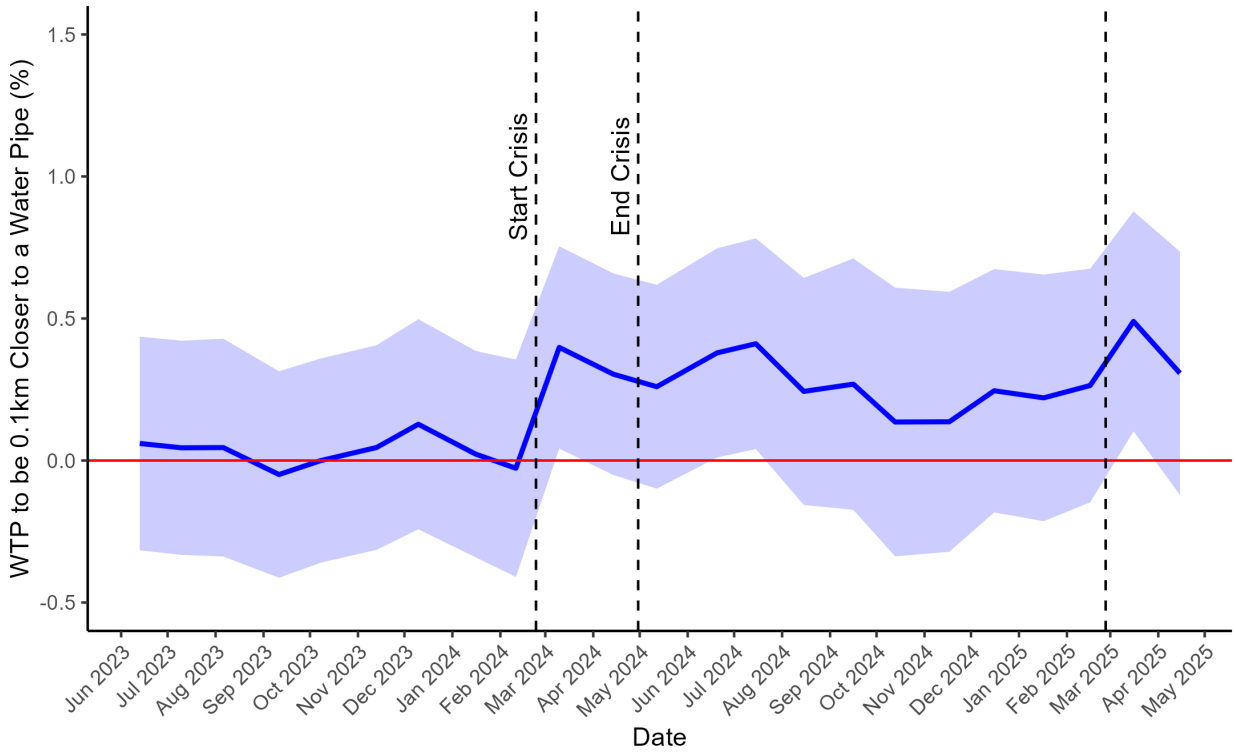


Figure 3.5: The Impact of an Urban Water Crisis on Housing Rents

Notes: Figure depicts estimated rent premium for rental listings being 0.1km closer to the nearest pipe in each month of our data. Estimates are generated by running regressions of the log of monthly rent prices on the distance (km) from the nearest water pipe, interacted by month. Regressions control for a rich set of listing-level characteristics and include ward and year-month fixed effects. Standard errors are clustered at the ward and year-month levels. The shaded blue area represents a 95% confidence interval, whereas the solid blue line represents point estimates. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

correlated with lease length, with the left figure thus portraying mainly long-term leases and the right figure mainly short-term leases. We show that the effects for long-term leases are relatively permanent, whereas short-term leases show rent premiums that exist only during drought season. This figure thus shows more concretely that future drought risk is a large driver in the sizes of observed WTP for piped water.

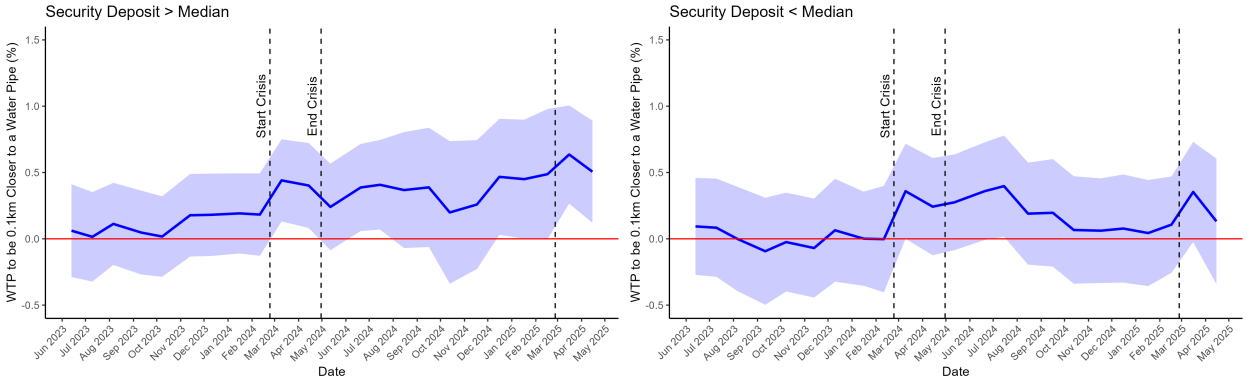


Figure 3.6: The Impact of an Urban Water Crisis on Housing Rents by Lease Length

Notes: Figure depicts estimated rent premium for rental listings being 0.1km closer to the nearest pipe in each month of our data separately for listings that have an above median security deposit size (left) and those that have a below median security deposit size (right). Median security deposit size is roughly 5 months worth of rent. Security deposit size is meant to proxy for lease length, as higher deposits correspond to longer leases. Estimates are generated by running regressions of the log of monthly rent prices on the distance (km) from the nearest water pipe, interacted by month. Regressions control for a rich set of listing-level characteristics and include ward and year-month fixed effects. Standard errors are clustered at the ward and year-month levels. The shaded blue area represents a 95% confidence interval, whereas the solid blue line represents point estimates. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

WTP Effects by Household Size

Figure 3.7 shows that WTP increases for piped water are also larger among large units than small units. In this figure, we depict the estimated WTP for piped water by month separately for listings that are at least three bedrooms (left) versus those that are two bedrooms or less

(right). The estimated effects are much larger (0.5% per 0.1km) for multi-bedroom units that tend to be occupied by families. Effects for small bachelor-style units appear to be muted, with notable spikes only occurring during drought months. This highlights that the impacts of drought risk are not borne evenly by the population and are particularly large for families more impacted by rations and price spikes.

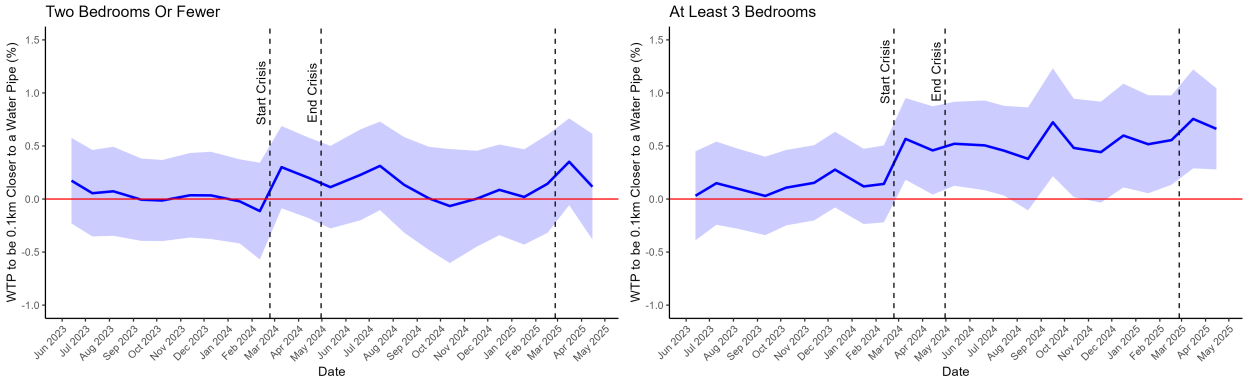


Figure 3.7: The Impact of an Urban Water Crisis on Housing Rents by Size

Notes: Figure depicts estimated rent premium for rental listings being 0.1km closer to the nearest pipe in each month of our data. Figure does so separately for listings with two bedrooms or fewer (left) and those with three bedrooms or more (right). Estimates are generated by running regressions of the log of monthly rent prices on the distance (km) from the nearest water pipe, interacted by month. Regressions control for a rich set of listing-level characteristics and include ward and year-month fixed effects. Standard errors are clustered at the ward and year-month levels. The shaded blue area represents a 95% confidence interval, whereas the solid blue line represents point estimates. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

3.4.2 Health as a Driver of WTP

To complement our WTP exercise and provide a more thorough understanding of the impacts from the water crisis, we show that the water crisis had detrimental health impacts. We do so using a combination of publicly available hospital visits data from the Government of Karnataka and self-reported health from our survey. Hospitals visit data is available on a

daily basis for two of Bangalore's nine major hospitals. For each day, we observe the total number of in-patient visits, out-patient visits, diagnostic tests, and surgeries.

In Figure 3.8, we show that hospital visits increase during the crisis. The crisis was accompanied by an immediate temporary spike in out-patient visits as well as a lasting increase in in-patient visits. Encouragingly, diagnostics and surgeries being performed do not change during this period, suggesting that these increases in in-patient and out-patient visits do not reflect general increases in demand for these two hospitals (e.g. from openings or closings in other hospitals). Moreover, we do not observe similar increases in in-patient nor out-patient visits during other dry seasons, suggesting that seasonality is not driving the increases during the crisis.

However, it is difficult to causally attribute these hospital visits to the water crisis, so we supplement these findings with our survey data. Two mechanisms primarily could explain health issues from a water crisis: decreased water usage (rationing) and substitution to using potentially contaminated water sources (water tankers). Table 3.4 presents regression estimates of these two mechanisms on health issues during the crisis. As can be seen, reliance upon often-contaminated tanker water is associated with more health issues during the crisis.[18] Rationing, while having a positive point estimate, is not significantly associated with increasing health issues. These findings suggests that climate-change induced groundwater shortages may lead households to substitute to less clean water sources (e.g. water tankers) from unregulated private markets that could have negative health impacts.

3.4.3 Short-run Housing Supply Effects

A crucial assumption of our WTP analysis is that short-run housing supply is not impacted by the crisis, which we find support of. In particular, our estimated relative increases in price could reflect either (i) an increase in the supply of housing with piped water at the start of

18. Water from tankers has been shown to contain heavy metals and bacteria, thus being unfit for human consumption in many cases [Sarkar, 2025]

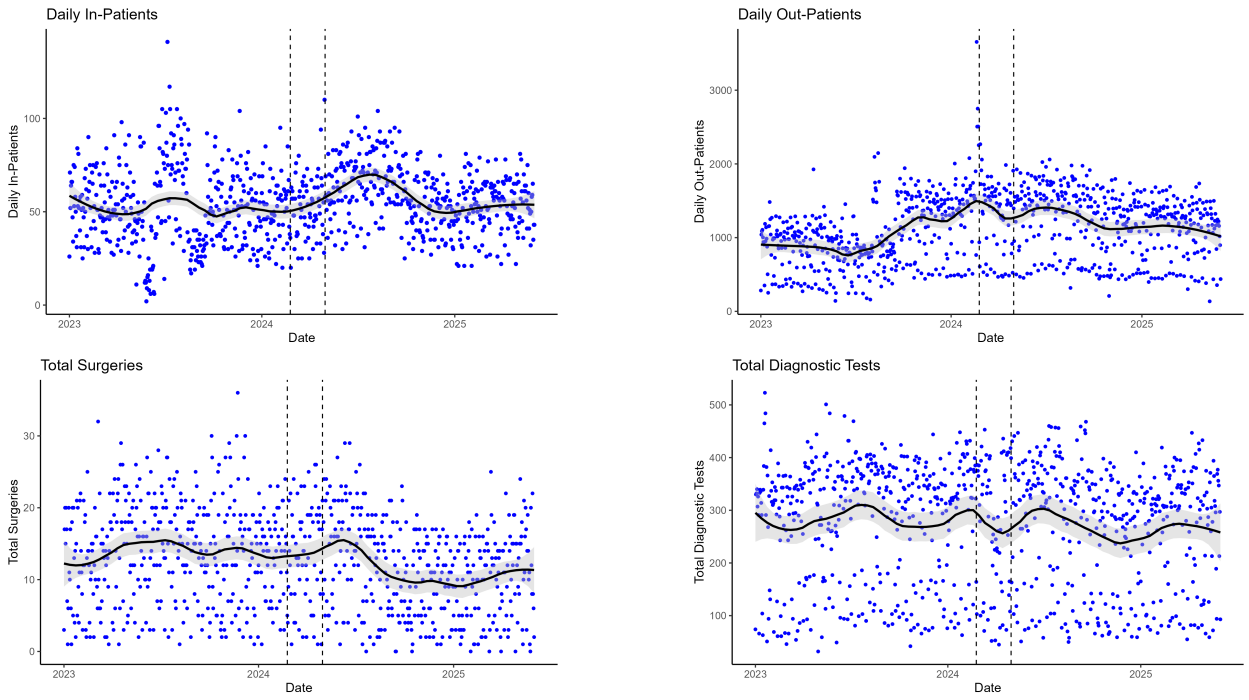


Figure 3.8: Health Issues During the Water Crisis

Notes: Figure depicts trends in hospital visits for two major hospitals in Bangalore. Figure does so for in-patient visits (top left), out-patient visits (top right), performed surgeries (bottom left), and diagnostic tests performed (bottom right). Note that y-axis scale varies for each panel. Blue dots indicate daily values, whereas black lines indicate locally smoothed trends with 95% confidence intervals.

	(1)	(2)	(3)
Annual Lease (=1)	0.011	0.019	0.009
	(0.010)	(0.012)	(0.008)
Secondary School (=1)	-0.013	-0.018	-0.013
	(0.065)	(0.023)	(0.040)
Bachelor's Degree (=1)	0.013	0.020	0.014
	(0.031)	(0.014)	(0.020)
Unit Bedrooms	-0.015	-0.004	-0.017
	(0.013)	(0.004)	(0.013)
Inc. >50,000 INR (=1)	-0.013***	-0.030**	-0.014***
	(0.005)	(0.013)	(0.005)
Independent House (=1)	-0.003	-0.006	-0.004
	(0.010)	(0.014)	(0.010)
Tanker Water (=1)	0.031	0.037***	0.029*
	(0.026)	(0.012)	(0.017)
Rationed (=1)	0.002	0.008	0.002
	(0.005)	(0.013)	(0.005)
Observations	914	914	914
Model	Logit	LPM	Probit

Table 3.4: The Impacts of Urban Drought on Health Outcomes

Notes: Table depicts regressions of health events during the 2024 water crisis on an array of demographics, including water tanker usage. Column (1) does so using logistic regression. Column (2) uses a linear probability model. Column (3) uses a probit regression. All effects are interpretable as average impacts to probability of a health event. For instance, the estimate on annual lease in column (1) is 0.011, which is interpretable as 1.1 percentage point higher prevalence of health events during the crisis among those with an annual lease. Health events include dizziness, nausea, diarrhea, and hospital visits. Standard errors are reported in parentheses below the estimated coefficient.

the crisis or (ii) landlords of groundwater-dependent residences strategically refraining from listing their most expensive units during the crisis.[19] In Appendix Figure C.5 we plot the total number of listings as well as trends in several key characteristics of these listings that determine pricing: building age, height, and bedroom count. We observe no sharp increase to the number of listings and no differential trends in any housing characteristics for listings. This seems to indicate that short-run supply is unchanged by the crisis.

3.4.4 Long-run Housing Supply Effects

We analyze the universe of in-progress residential construction listed on the RERA website and show that the crisis increased long-run housing supply in areas near the piped water network. Figure 3.9 depicts the trends of residential builds in the RERA system by month, separated by whether the build is within 100 meters of the nearest pipe or not. The beginning of the water crisis immediately preceded an increase in the number of residential projects being built within 100 meters of the nearest pipe. In contrast, the number of residential projects being built more than 100 meters from the nearest pipe appears unchanged following the crisis.

To causally estimate the change in residential construction for areas with piped water access, we convert the entirety of Bangalore into a set of 50 subsections of equal land area. We then calculate the distance of each subsection's centroid from the nearest water pipe.[20] Our estimation utilizes the following difference-in-difference regression to compare sub-sections b

19. This may be done if, for instance, landlords are worried about temporary decreases to demand during the crisis and are worried about the Housing.com algorithm de-prioritizing listings that have been unfilled for multiple months.

20. In Appendix Table C.2, we consider an alternative measure of pipe access which captures the proportion of a bubble's area within 100m, 200m, and 300m of a pipe. We show that our estimates are robust to this alternative measure.

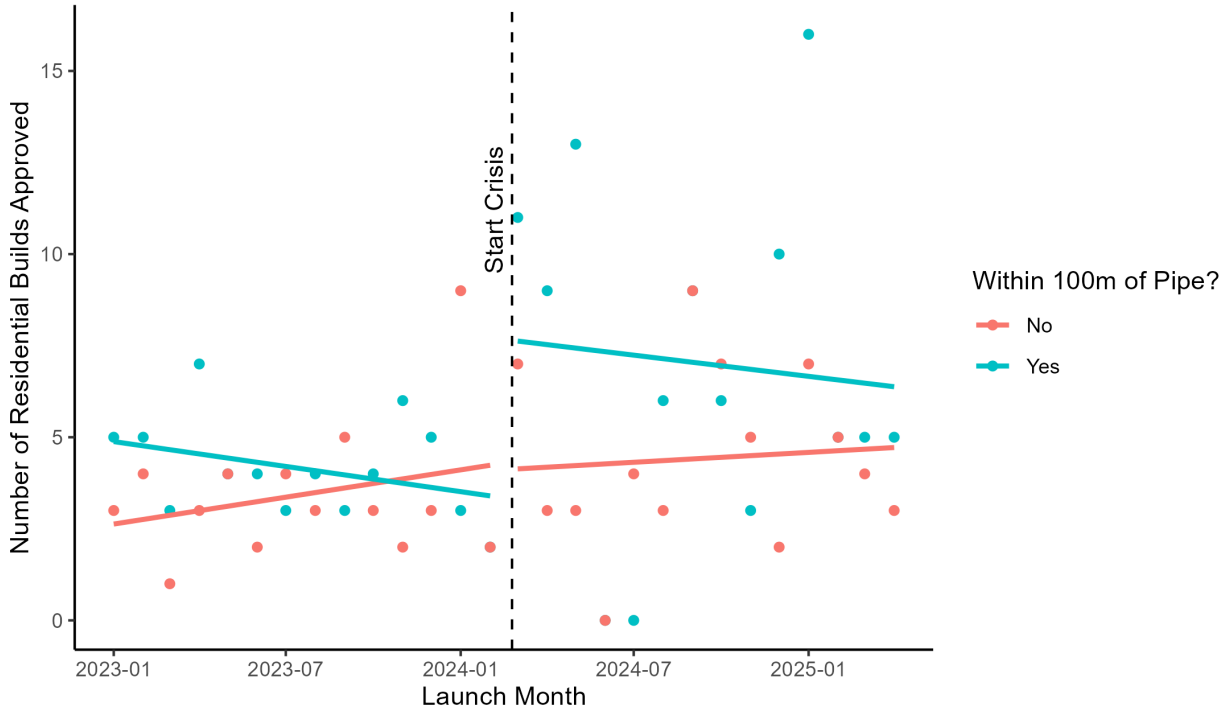


Figure 3.9: The Impact of an Urban Water Crisis on Residential Construction

Notes: Figure depicts trends in monthly new residential construction projects for areas that are within 100m (blue) of the nearest water pipe versus those that are not (red). The black vertical dashed line depicts the beginning month of the water crisis. The horizontal axis corresponds with approval/launch month.

at time t that are close to pipes to those that are far:

$$y_{bt} = g(\eta_0 + \eta_1 D_{bt} E_t + \lambda_b + \lambda_t) \tag{3.2}$$

where y_{bt} is the count of new builds that began construction in subsection b during month-year t, D_{bt} is an indicator taking on 1 if the subsection is within a certain range of the nearest water pipe, E_t is an indicator taking on one for months past the beginning of the water crisis, λ_b is a subsection fixed effect, and λ_t is a month-year fixed effect. As y_{bt} is a count variable, we assume that the link function is $g(x) = \exp(x)$, equivalent to Poisson regression.[21]

We persistently estimate values of η_1 that indicate increases in the rate of construction among units close to water pipes which are statistically significant at conventional levels. The estimates for these coefficients are presented in Table 3.5, which includes difference-in-difference estimates on construction for ranges of 100 meters, 200 meters, and 300 meters. The magnitude of the estimated effect ranges from a 25% increase for 100 meters to a 51% increase for 300 meters.

3.5 Adaptation Via the Housing Market

To estimate counterfactual welfare measures under increased housing supply near pipes compared to pipe expansions, we make use of a structural model. We argue that housing supply expansions are cost-effective means of enhancing welfare.

We have established that demand and long-run supply for housing near water pipes increased following the 2024 crisis, but an exercise which benchmarks welfare impacts of pipe expansions versus this new construction is useful. We note that we are unable to directly evaluate welfare effects of an increase in the supply of housing closer to the piped

21. An overdispersion test for the suitability of Poisson regression yields a p value of 0.57, supporting that Poisson regression is suitable for the data.

	100m cutoff	200m cutoff	300m cutoff
Closer than 100m × Post-Drought	0.250**		
	(0.107)		
Closer than 200m × Post-Drought		0.391**	
		(0.182)	
Closer than 300m × Post-Drought			0.513***
			(0.193)
Observations	1,620	1,620	1,620
Mean of build	0.12	0.12	0.12
Fixed effects	Subsection, Month	Subsection, Month	Subsection, Month

Table 3.5: The Impacts of Urban Drought on New Residential Construction

Notes: Table depicts estimates of the change in construction frequency for areas close to piped water following the 2024 drought. Column (1) does so comparing areas within 100m of the nearest pipe to those outside of 100m. Column (2) does the same except with a 200m cutoff. Column (3) does the same except with a 300m cutoff. Estimates are interpreted as percent increases after being scaled by 100. For instance, column (3) finds an estimate of 0.513, which is a 51.3% increase in construction in areas closer to water pipes than 300m following the drought. Standard errors are clustered at the subsection and date levels and reported in parentheses below the pertinent estimate.

water network because these units are not yet listed.[22] Our estimates suggest though that the shift in long-run housing is non-marginal, as construction increased by up to 51% within 300 meters of the piped network. Evaluating the welfare effects of this long-run supply shift and other counterfactuals requires evaluating non-marginal shifts in the distance to the piped water network, where the hedonic price schedule along is not sufficient by itself [Greenstone, 2017].

We estimate the hedonic price schedule using difference-in-differences estimates which are sufficient to construct a lower bound of the welfare effects of housing supply expansions [Banzhaf, 2021b].[23] We construct a lower bound on the counterfactual consumer surplus for non-marginal changes in a consumer's distance to the piped water network. We consider a

22. According to the RERA data, the average construction completion time in Bangalore is roughly 3.6 years.

23. These results complement our monthly estimates WTP from Section 3.4.1, but are a difference-in-difference rather than event study.

two-way fixed effect estimator:

$$\ln(P_{it}) = \beta_1 KM_i + \beta_2 Crisis_t + \beta_3 KM_i \times Crisis_t + \alpha A_{it} + \gamma_w + \delta_t + \varepsilon_{it} \tag{3.3}$$

Equation 3.3 is equivalent to a difference-in-difference with a continuous treatment. The estimated coefficient β_3 represent the "capitalization effect," which Banzhaf [2021b] shows is a lower bound on the Hicksian equivalent surplus (ES).

Our lower bound estimates of welfare echo the estimates which were found in Figure 3.5. Prior to the crisis, we do not observe any effects of distance to the piped water network on rents. After the crisis, listings that are 1 km closer to the piped network are roughly 2.8% more expensive (roughly ₹1008 over the average).[24] We use this elasticity to construct bounds on the welfare effects of counterfactual shifts.

To capture sizable expansions to the piped water network, we consider a counterfactual that expands the piped water network such that every listing that is more than 0.2 kms away from the piped water network is exactly 0.2 kms away. We estimate that this implies a capitalization effect of roughly ₹17 million ($\sim$ \$195,000) for the 12 months after the start of the crisis. This suggests an average capitalization effect of roughly \$10 per listing (among the 20,000 listings that are further than 0.2kms away).

We next consider our estimate that there is a 39% increase in new construction builds within 0.2 km of the piped water network. This is equivalent to 25 additional construction builds in the year after the crisis within 0.2 km of the piped water network, which translates into 3,625 additional listings.[25] We assume that households that move into this increased number of listings are randomly drawn from the distribution of distances outside of the

24. As Figure 3.5 shows, we find no evidence of differential pre-trends between "more exposed" vs. "less exposed" listings.

25. The average construction in the pre-period has an average of 145 units per construction. We make the conservative assumption that the 39% is the growth over the baseline number of construction in areas within 0.2km of the piped water network, which is 64.

0.2km radius, and that the listings are distributed uniformly within the 0.2km. This implies that the lower bound of the gains in consumer surplus is roughly ₹3.4 million ($\sim$ \$39,000). Note that since this approach does not take the heterogeneity in MWTP into account, the lower bound on the average surplus is the same as above. This represents 20% of the total gain in surplus compared to the counterfactual of universal pipe expansion.

To assess the cost-effectiveness of the market-adjustment relative to universal piped water access, we consider the costs of both alternatives. The most recent Cauvery stage V project cost ₹43 billion (\$490 million) to lay roughly 1,500 km of pipes, which translates to ₹28 million per km of pipe. To ensure universal pipe connection, we estimate that an additional 250km of pipes will be required.[26] This implies that universal piped water access would cost an additional ₹7 billion in expenditures. We use the reported total project costs in the RERA data to estimate the cost of constructing *closer* to the piped water network. The average listing is more expensive to build within 0.2kms by ₹47,000, so the total additional incurred cost is ₹170 million.

These back-of-the-envelope calculations suggest that market adjustments from private developers could represent an efficient manner of adaptation compared to the large-scale policy intervention of expanding the piped water networks. We find that the lower bound on the consumer surplus from the increase in housing supply near the piped water network is 20% that of the counterfactual of universal piped water expansion—but we estimate the costs to be 2.4% that expanding universal piped water access. These figures suggest that subsidies for developers to build closer to piped water networks might be a cost-efficient way of ensuring piped water access, rather than the physical expansion of the network footprint.

26. To estimate this, we first group all 26,457 listings that are further than 0.2km away from the current piped water network into clusters that are within 0.2km from each other, which implies that each listing within this cluster would be served by the same pipe extension. We estimate 521 such clusters. We calculate the length of the pipe extension that would be needed to reach the centroid of each cluster. Appendix Figure C.7 shows the grouped listings.

3.6 Conclusion

We have leveraged a natural experiment which provides three main findings in an analysis in residential WTP for publicly-supplied urban piped water. We have first shown that consumer WTP for piped water increases when groundwater drought risk increases. We show that rental listings experienced a 0.4% increase in price per every 0.1km closer to the piped water network following the drought. This effect is permanent for long-term leases, emphasizing the role of future drought risk in driving residential sorting. The effect is particularly large (0.5%) for units housing families. We show that the shift in demand for housing is likely driven in part by health impacts from the drought, as the crisis corresponded to increased hospital visits and self-reported health episodes.

We have then analyzed the impacts to housing supply from rising drought risk. We observe that the drought had minimal changes in short-run supply of housing, measured as the number of rental listings posted. However, using a difference-in-difference approach we estimate an immediate 25-51% increase in construction of rental listings in areas near the piped water network relative to areas far from the piped water network. This result implies that the long-run supply of rental listings is elastic, indicating that the demand shock presented by the drought increased quantity of apartments near a water pipe. As a result, the housing market in part allows individuals to adapt to future drought risk by moving to areas closer to water pipes, mitigating the need for piped water expansion into previously unserviced suburbs.

Our reduced-form estimates have a natural structural interpretation as they form lower bounds on consumer surplus [Banzhaf, 2021b]. We find that the long-run response from the housing market delivers 20% of the consumer surplus relative to a counterfactual of universal piped water expansion for only 2.4% of the cost. These findings suggest that policies that encourage housing market responses could be cost-effective in settings with limited public infrastructure.

Such results enhance our knowledge of urban expansion in the Global South. In particular, many Southern cities are experiencing an influx of migrants. Consequently, Southern cities have tended to expand into suburban outskirts with minimal access to publicly-provided goods, including piped water. The privately-provided substitutes for many of these utilities are of low quality or come with poorly understood risks. Our estimates help illuminate the extent to which risk surrounding these substitutes - in our case groundwater - drive outwards expansion of these cities.

Several other open questions may be interesting for future research. First, we analyze only one public good, whereas many others may be of interest. In particular, electricity, phone/internet connectivity, waste collection, transportation, and sewerage, are often differentially accessible within a city. Understanding the role of quality or perceptions about quality of their substitutes is of interest. Analyzing the role of formal markets, informal markets, and other forms of adaptation will likely illuminate the extent to which each utility requires expansion.

Second, our analysis has focused primarily on renters, who face different incentives for movement than do homeowners. Homeowners likely face higher expenses to move to a new residence, but also likely internalize future drought risk to a greater extent than do renters. As such, it is ambiguous whether homeowner demand for housing will shift more or less than does demand for renters. Future analysis that considers homeowners is likely important, as they make up the majority of residents in Bangalore and many other cities.

Finally, we believe a more robust analysis of the health impacts of acute water shortages is necessary. To the best of our knowledge, relatively little literature analyzes the mortality or morbidity impacts of water shortages in economics or other fields. A thorough causal analysis of water shortages on these health outcomes, along with linking these estimates to adaptation, would provide policymakers and planners with a more clear understanding of the overall social benefits of expanded pipe provision.

APPENDIX A

APPENDIX FOR CHAPTER 1

A.1 Model Appendix

A.1.1 Derivations

Social Rewards Analytical Form Derivation

Here I derive the analytical form of the social rewards function. This derivation is provided directly in Bénabou and Tirole [2011] and Jee et al. [2024], but I provide it here as well for transparency. In this case, we want to derive an analytical form of

$$\Delta^*(y) = \frac{-1}{F_w(w^*(y))[1 - F_w(w^*(y))]} \int_{-\infty}^{\infty} vF(w^*(y) - v)f(v)dv$$

under the assumptions that $v \sim N(0,1)$, $u \sim N(0,\sigma)$, and that u and v are independently distributed. Consider only the second term of the above, and rewrite $w^* = w^*(y)$. We have that

$$\begin{aligned}
\int vF_u(w^*(y) - v)f_v(v)dv &= \int v\phi(v)\Phi\left(\frac{w-v}{\sigma}\right)dv \\
&= \int_{d(-\phi(v))} -\phi\left(\frac{w-v}{\sigma}\right)d\Phi(v) \\
&= \left[-\phi\left(\frac{w-v}{\sigma}\right)\right]_{-\infty}^{\infty} + \int_{-\infty}^{\infty} \frac{1}{\sigma}\phi\left(\frac{w-v}{\sigma}\right)(-\phi(v))dv \\
&= \left[-\phi\left(\frac{w-v}{\sigma}\right)\right]_{-\infty}^{\infty} - \frac{1}{\sigma}\int \phi\left(\frac{w-v}{\sigma}\right)\phi(v)dv
\end{aligned}$$

Consider only the portion of the above that remains within an integral:

$$\begin{aligned}
\int \phi\left(\frac{w-v}{\sigma}\right)\phi(v)dv &= \frac{1}{\sqrt{2\pi}\sqrt{2\pi}}\int \exp\left[-\frac{1}{2}\left(\frac{w-v}{\sigma}\right)^2 + v^2\right]dv \\
&= \frac{1}{\sqrt{2\pi}\sqrt{2\pi}}\int \exp\left[-\frac{1}{2}\left(\left(\frac{w}{\sigma}\right)^2 + v^2\frac{\sigma^2+1}{\sigma^2} - \frac{2wv}{\sigma^2}\right)\right]dv \\
&= \frac{1}{\sqrt{2\pi}\sqrt{2\pi}}\int \exp\left[-\frac{1}{2}\left(\frac{w^2}{\sigma^2+1} + \frac{\sigma^2+1}{\sigma^2}\left(v - \frac{w}{\sigma^2+1}\right)^2\right)\right]dv \\
&= \frac{1}{\sqrt{2\pi}\sqrt{2\pi}}\exp\left[-\frac{1}{2}\frac{w^2}{\sigma^2+1}\right]\int \exp\left[-\frac{1}{2}\frac{\sigma^2+1}{\sigma^2}\left(v - \frac{w}{\sigma^2+1}\right)^2\right]dv
\end{aligned}$$

The first exponential term is just a function of w and σ. The second term resembles the density of a Normal distribution. Indeed, it can be seen that:

$$\begin{aligned}
\frac{1}{\sqrt{2\pi}}\int \exp\left[-\frac{1}{2}\frac{\sigma^2+1}{\sigma^2}\left(v - \frac{w}{\sigma^2+1}\right)^2\right]dv &= \sqrt{\frac{\sigma^2}{\sigma^2+1}}\frac{1}{\sqrt{2\pi}}\int \exp\left[-\frac{1}{2}\left(\frac{v - \frac{w}{\sigma^2+1}}{\sqrt{\frac{\sigma^2}{\sigma^2+1}}}\right)^2\right]dv \\
&= \sqrt{\frac{\sigma^2}{\sigma^2+1}}
\end{aligned}$$

where the integral evaluates to one as is the property of any probability density evaluated across its full support. Thus, we have derived that

$$\int \phi\left(\frac{w-v}{\sigma}\right)\phi(v)dv = \frac{1}{\sqrt{2\pi}}\sqrt{\frac{\sigma^2}{\sigma^2+1}}\exp\left[-\frac{1}{2}\frac{w^2}{\sigma^2+1}\right]$$

We thus see that

$$\begin{aligned}\int \phi\left(\frac{w-v}{\sigma}\right)\phi(v)dv &= \left[-\phi\left(\frac{w-v}{\sigma}\right)\right]_{-\infty}^{\infty} - \frac{1}{\sigma}\int \phi\left(\frac{w-v}{\sigma}\right)\phi(v)dv \\ &= \left[-\phi(v)\phi\left(\frac{w-v}{\sigma}\right)\right]_{-\infty}^{\infty} - \frac{1}{\sigma}\frac{1}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\frac{w^2}{\sigma^2+1}\right]\sqrt{\frac{\sigma^2}{\sigma^2+1}} \\ &= -\frac{1}{\sigma}\frac{1}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\frac{w^2}{\sigma^2+1}\right]\sqrt{\frac{\sigma^2}{\sigma^2+1}}\end{aligned}$$

From plugging in the cumulative distribution function of a Normal distribution, it thus follows that

$$\begin{aligned}\Delta^*(y) &= \frac{-1}{F_w(w^*(y))[1-F_w(w^*(y))]}\int_{-\infty}^{\infty} vF(w^*(y)-v)f(v)dv \\ &= \frac{1}{\Phi\left(\frac{w}{\sqrt{\sigma^2+1}}\right)\left(1-\Phi\left(\frac{w}{\sqrt{\sigma^2+1}}\right)\right)} \times \frac{1}{\sqrt{2\pi}}\exp\left[-\frac{1}{2}\frac{w^2}{\sigma^2+1}\right] \times \sqrt{\frac{\sigma^2}{\sigma^2+1}}\end{aligned}$$

Optimal Subsidy Derivation

We have the following form of social welfare:

$$\text{argmax}_y \int_{w^*(y)}^{\infty} (e + w - \beta_p p - \eta\beta_p y) f(w)dw + \mu\bar{w}$$

in which $e > 0$ is the externality, $\eta \geq 0$ is the shadow cost of subsidy funds, $f(*)$ is the PDF of w, and $\bar{w}$ is the total value of signaling across all agents. By Leibniz's rule, taking a derivative of $SW(y)$ with respect to y and equating to 0 yields:

$$(e + w^*(y) - \beta_p p - \eta\beta_p y)\left(\frac{-\partial w^*(y)}{\partial y}\right) f(w^*(y)) = \eta\beta_p[1 - F(w^*(y))]$$

Using the equilibrium condition, we can solve for $\frac{\partial w^*(y)}{\partial y}$:

$$\begin{aligned} w^*(y) &= \beta_p p - \beta_p y - \mu\Delta(w^*(y)) \\ \frac{\partial w^*(y)}{\partial y} &= -\beta_p - \mu\Delta'(w^*(y))\frac{\partial w^*(y)}{\partial y} \\ \frac{\partial w^*(y)}{\partial y}(1+\mu\Delta'(w^*(y))) &= -\beta_p \\ \frac{\partial w^*(y)}{\partial y} &= -\frac{\beta_p}{1+\mu\Delta'(w^*(y)} \end{aligned}$$

The optimal subsidy thus solves

$$(e + w^*(y) - \beta_p p - \eta\beta_p y)\left(\frac{\beta_p}{1+\mu\Delta'(w^*(y)}\right) f(w^*(y)) = \eta\beta_p[1 - F(w^*(y))]$$

Solving for y:

$$\begin{aligned} e + \{\beta_p p - \beta_p y - \mu\Delta(w^*)\} - \beta_p p - \eta\beta_p y &= \frac{\eta(1+\mu\Delta'(w^*(y))[1-F(w^*(y))]}{f(w^*(y))} \\ (\eta+1)\beta_p y &= e - \mu\Delta(w^*) - \frac{\eta(1+\mu\Delta'(w^*(y))[1-F(w^*(y))]}{f(w^*(y))} \\ y &= \frac{e-\mu\Delta(w^*)}{(\eta+1)\beta_p} - \frac{\eta(1+\mu\Delta'(w^*(y))[1-F(w^*(y))]}{(\eta+1)\beta_p f(w^*(y))} \end{aligned}$$

A.1.2 Extensions

Correlation Between Visibility and Price

I have supposed that the marginal value of signaling $\mu = \lambda x$ is not a function of the price p of the good. This is appropriate in my setting as p is not publicly signaled. However, in many settings μ may be endogenous to the price or cost of an good if costly goods are more visible. This arises, for instance, in cases in which agents care more about being seen buy a

good that was costly to them or in which agents are more likely to display the costly good. In this scenario, μ is a function of $p - \mu(p)$ – as described in Jee et al. [2024].

Under this version of the model, equilibrium is still categorized by a cutoff type w^* who is ambivalent between buying the good or not. However, the relationship between visible and invisible price sensitivity described in Prediction 3 differs, now depending on whether p increases or decreases the visibility μ. For cutoff types w^* for which social rewards $\Delta(*)$ are increasing in p but visibility μ decreases in p, it is unclear whether price sensitivity is higher in visible or invisible cases just from examining the social rewards. As in Jee et al. [2024], the prediction in this setting will depend on the direction and the extent to which $\mu(p)$ is elastic with respect to p.

Correlation Between Visibility and Types

I have also supposed that the marginal value of signaling $\mu = \lambda x$ is not correlated with the agents' v. In my setting, x is unlikely to be correlated with v as the intensity of visibility is determined by the researcher and not the agent. However, this assumption may not hold if agents who are more prosocial/cool care more or less about being seen as prosocial/cool. That is, it is possible for λ to be correlated with v.

Here, I consider two cases in which μ and v are correlated. First, μ and v will be positively correlated if those who care more about being seen as prosocial/cool also tend to be more prosocial/cool. In this case, the model's equilibrium will still be governed by a cutoff type w^* who is ambivalent between buying the good or not. To see this, consider holding price p and idiosyncratic preferences u_i fixed. As v_i increases (thus increasing w_i), so too does the social rewards $\lambda x \Delta(*)$, indicating that at some point a sufficiently large value of w_i will be reached such that the good is purchased for all larger realizations.

However, the model's equilibrium is no longer governed by a cutoff type in the case when less prosocial/cool people care more about social rewards (i.e. μ and v are negatively cor-

related). In this scenario, a tension arises between private and social utility. As a result, situations may arise in which individuals who are very prosocial/cool buy the good alongside those who are not at all prosocial/cool but wish to blend in with the prosocial/cool individuals. That is, individuals with a small value of w_i may buy the good to signal having a large w_i if their value of μ is sufficiently large.[1]

Privacy Concerns

In Section 4, I have shown that privacy concerns play a non-trivial role in determining my estimated social rewards function. These privacy concerns are not the main focus of my paper, and in fact to some extent represent a divergence from social rewards from purchasing solar panels or EVs where privacy concerns are lessened. As such, I have presented results which largely remove privacy concerns from the data. However, I am unable to do so entirely. As such, it is interesting to consider how privacy would be modeled and how thus they may impact my estimates.

Suppose that consumers receive disutility $\varsigma_i < 0$ from their name being publicized. Then their utility becomes

$$U_i = \left(\beta_p(y-p) + v_i + u_i + \varsigma_i x\right) a_i + \mu E_{-i}[v_i|a_i]$$

Suppose that ς_i is uncorrelated with both v_i and u_i. Then, three different cases can arise. The first case is that there may be little heterogeneity in ς_i (i.e. σ_ς is small). The second case is instead that there is sufficient heterogeneity in ς_i, but ς_i is not private information to consumers (e.g. people can observe when others tend to be more private with their doings). In both of these two cases, the presence of privacy concerns will decrease the value people place on being seen buy a good but will not necessarily alter the strength of the signal across

1. Notably, subsidy policy also changes when μ varies with w. In particular, signaling can be a negative or positive-sum game as determined by the underlying distribution of w.

market penetrations. That is, estimated values of μ and social rewards will shift downwards, but the shape estimated social rewards will not change. In the final case, ς_i is both private information and sufficiently noise relative to u_i and v_i. In this case, the presence of privacy concerns will not only shift downwards estimated social rewards, but also alter its shape by decreasing the strength of the signal from purchasing.

When comparing Figure A.1 to Figure 1.5 and Figure A.2 to Figure 1.6, it can be seen that including individuals that report having privacy concerns tends to shift the entire curve downwards with minimal impact to the curvature of estimated social rewards. As such, it seems likely that either privacy concerns are relatively noiseless ($\sigma_\varsigma << \sigma_w$) or that ς_i is relatively well-known among peers.

A.2 Extension Appendix: Electric Vehicle Subsidies

In Section 5.4, I showed how optimal Pigouvian and Ramsey subsidies vary across market penetrations per ton of carbon that is being mitigated. Here, I scale these estimates up to the amount of carbon reduced by consumers purchasing electric vehicles. To do so, I assume that WTP for the social rewards scales linearly in the tonnage of carbon offset, allowing me to analyze how social rewards impact optimal subsidies for EVs.

According to the U.S. EPA, new gas vehicles (GVs) and EVs with average fuel efficiency are associated with roughly 0.93lbs and 0.30lbs of emitted carbon dioxide per mile. Assuming the average vehicle lasts 200,000 miles, the decision to buy an EV instead of a GV thus reduces atmospheric carbon dioxide emissions by a total of 126,000lbs. (roughly 57.15 metric tons) over the lifetime of the vehicle. I use this amount to scale up WTP for social rewards of carbon mitigation.

In Figure A.3, I use my estimates to plot optimal subsidy paths for EVs by market penetration based on estimates in [Allcott et al., 2024a] that gross lifetime negative externalities of an average EV are \$3,200 lower than for an average GV. The figure suggests that EVs

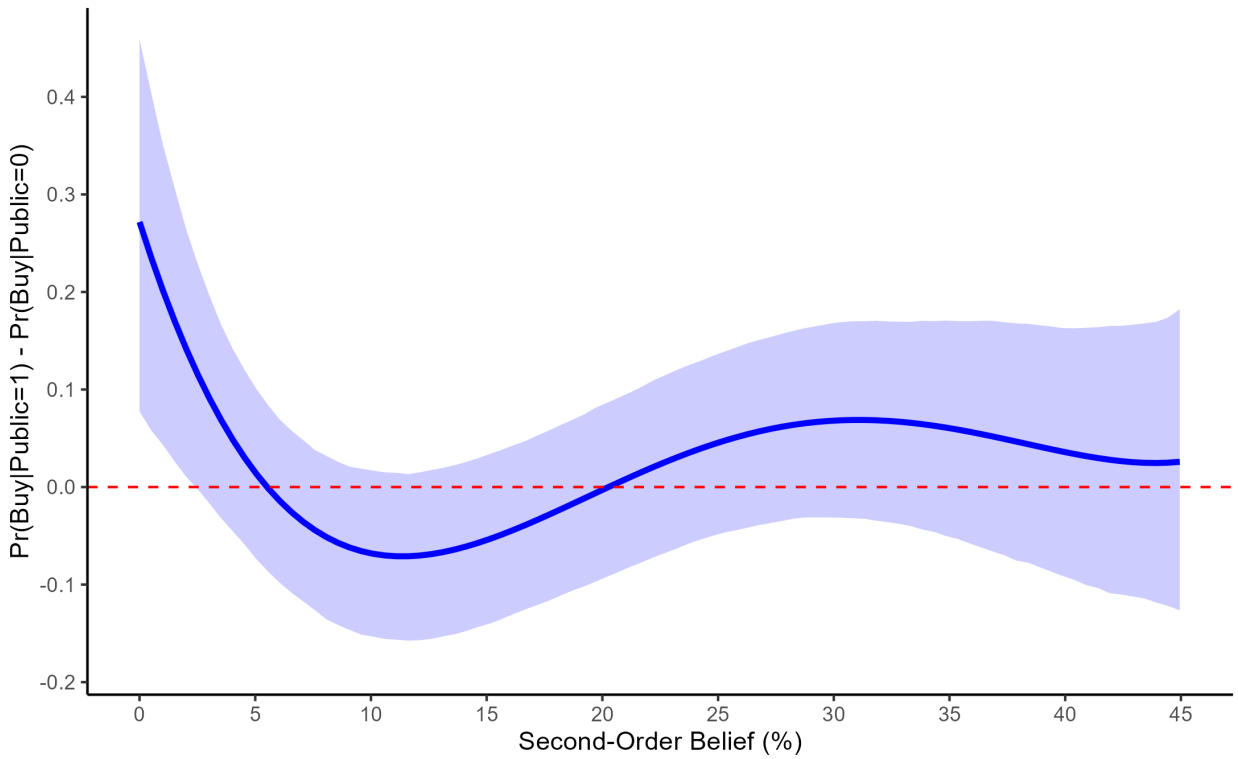


Figure A.1: Reduced Form Impact of Visibility by Second-order Belief Including Privacy Concerns

Notes: Figure depicts the linear probability estimated impact of visibility on purchasing an offset by second-order beliefs of market penetration of carbon offsets, including all those who reported that they "Strongly agree" that there are privacy concerns with being named online. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

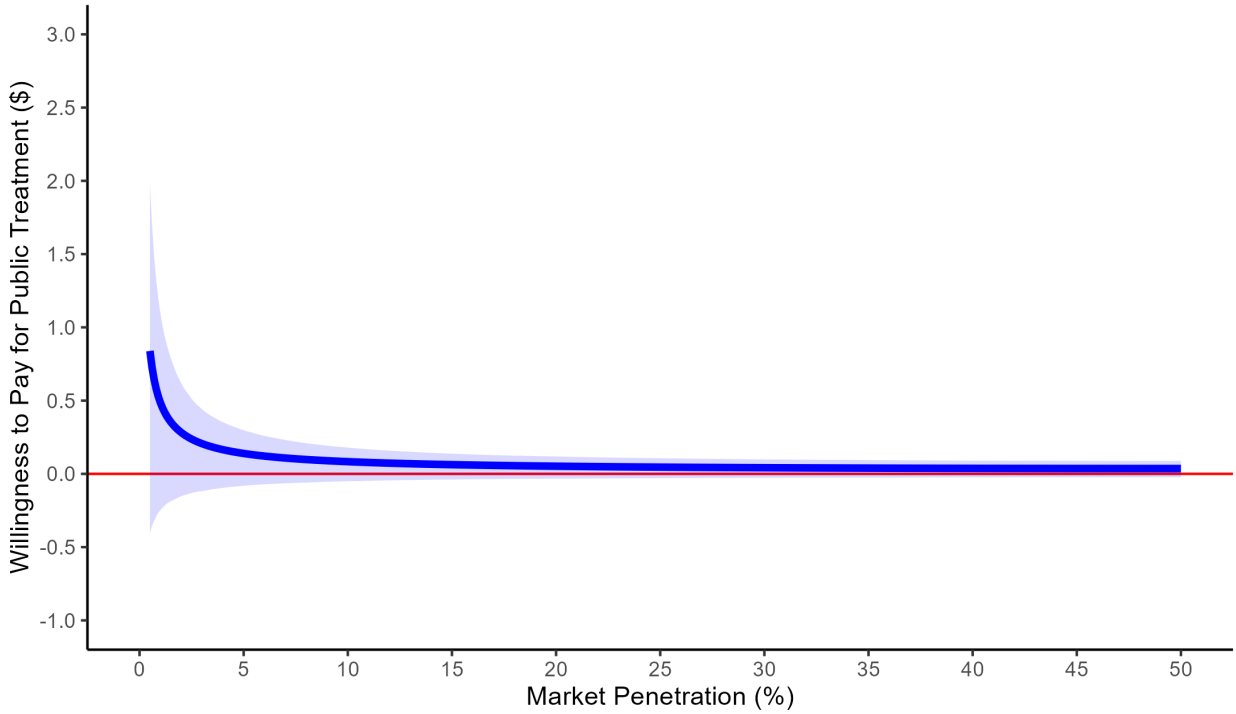


Figure A.2: Structurally Estimated WTP for Visibility by Market Penetration Including Privacy Concerns

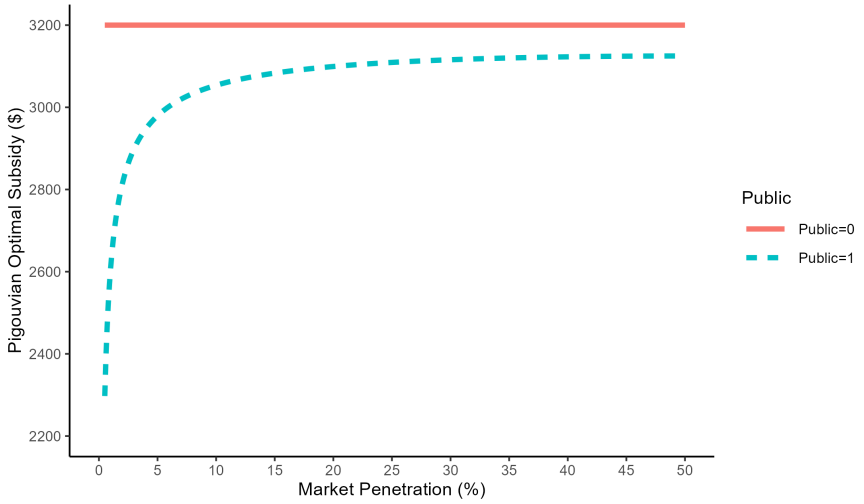
Notes: Figure depicts structural estimates of the willingness-to-pay (WTP) for visibility from purchasing an offset by second-order belief of market penetration of carbon offsets, including all those who reported that they "Strongly agree" that there are privacy concerns with being named online. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

should receive subsidies which are increasing at small market penetrations, before flattening in the Pigouvian setting or eventually decreasing in the Ramsey case. This is not reflected by subsidy policy in the U.S., which has set a fixed amount of $7,500 for the past decade. In fact, the Inflation Reduction Act has made subsidies *more exclusive* in recent years. To some extent, this subsidy path and amount likely reflects other supply-side externalities (e.g. learning-by-doing). However, my estimates suggest that demand-side considerations should be taken into account when setting optimal dynamic subsidy paths for goods such as EVs.

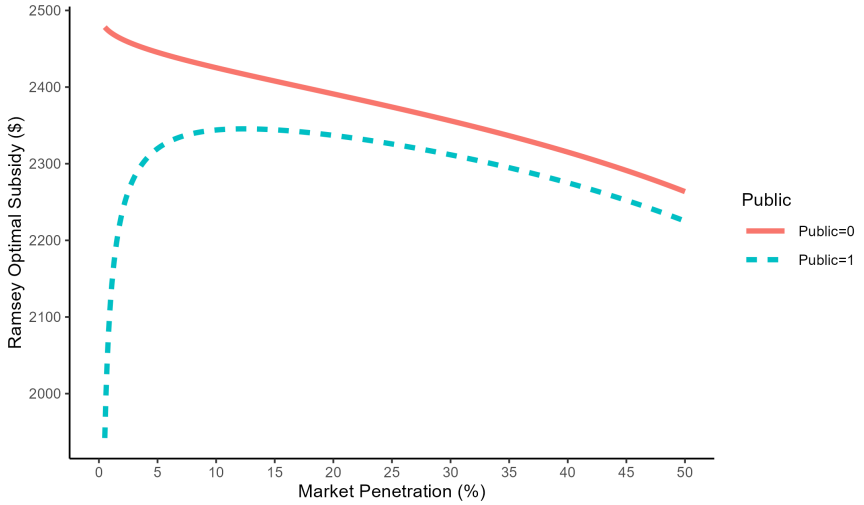
Counterfactual Optimal Subsidies for Electric Vehicles

In the above, optimal Pigouvian subsidies for EVs were calculated under the supposition that my estimates of the model's parameters translate from my setting into the setting of real world EV purchases. This assumption is strong, with violations occurring if EVs tend to be more visible (higher x) or being driven by higher variance in preferences (higher σ_w) than the carbon mitigation examined in my setting. To understand the importance of these differences in determining optimal subsidies, I consider counterfactuals under which goods are made more visible or peoples' types become more variant.

The optimal subsidies under such counterfactuals are presented in Figure A.4. In this Figure, I consider Pigouvian subsidies for combinations of x and σ_w being multiplied by 1, 2, and 3. As can be seen, enhancing the visibility of goods further emphasizes both the gap between visible and invisible subsidies as well as the upwards slope of the subsidy in the visible case. However, increasing the variability in peoples' types does the opposite: narrowing the gap between visible and invisible subsidies and flattening the publicly visible subsidy curve. These two forces oppose each other, leaving the impacts of increasing both visibility and preference variability ambiguous. In general, the figure shows that equivalent proportional increases to visibility and preference variability yields preference variability dominating. However, optimal subsidies slope steeply upwards at low market shares even



Panel A: Pigouvian Optimal Subsidies for EVs



Panel B: Ramsey Optimal Subsidies for EVs

Figure A.3: Optimal EV Subsidies by Market Penetration

Notes: Figure depicts structural estimates of the Pigouvian (Panel A) and Ramsey (Panel B) optimal subsidy for electric vehicles by market penetration of electric vehicles, assuming individuals' second-order beliefs are on average correct. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common electric vehicles are. Figure depicts optimal subsidies to mitigate the amount of carbon the buying of an average EV mitigates when compared to buying an average GV. The optimal subsidy is depicted in the invisible case as the red solid line and in the visible case as the blue dashed line. Estimates of EV benefits (i.e. the solid red line) are taken from Allcott et al. [2024b]. Estimates of WTP for social rewards are my own and are assumed to scale linearly in the amount of carbon mitigated.

in the case where both visibility and preference variability are doubled. As such, the conclusions from Figure A.3 seem to hold under reasonably large deviations from my estimated parameters.

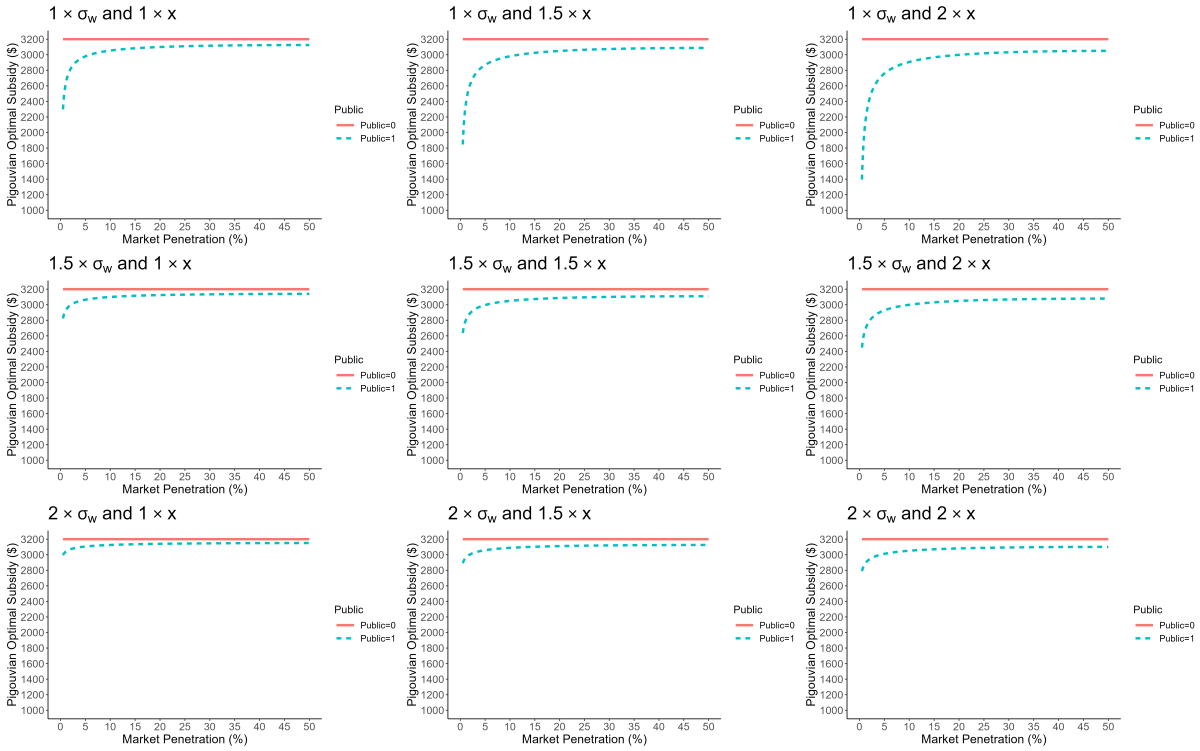


Figure A.4: Pigouvian Subsidies for EVs Under Counterfactual Parameterizations

Notes: Figure depicts structural estimates of the Pigouvian optimal subsidy electric vehicles by market penetration of electric vehicles, assuming individuals' second-order beliefs are on average correct. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common electric vehicles are. Figure depicts these subsidies across counterfactual regimes in which visibility increases (moving left to right) and preference variability increases (moving top to bottom). In all panels, optimal subsidy is depicted in the invisible case as the red solid line and in the visible case as the blue dashed line. Estimates of EV benefits are taken from Allcott et al. [2024b]. Estimates of WTP for social rewards are my own and are assumed to scale linearly in the amount of carbon mitigated.

A.3 Robustness Appendix

A.3.1 Balance Tables

Balance tables across the three elements of randomization from the RCT are presented in Table A.1, Table A.2, and Table A.3. The sample sizes are sufficiently large such that randomization should ensure that confounding variables are balanced across the different treatment arms. As can be seen, this is indeed the case. For all presented variables, no economically significant deviations occur. As expected, results presented in the main body of the paper are not driven by demographic confounds.

	White	Male	Dem.	Rep.	Born US	College	Inc>25k	Inc>50k	Inc>100k
Public=0									
Mean	0.65	0.51	0.28	0.27	0.90	0.60	0.90	0.72	0.35
SD	0.48	0.50	0.45	0.45	0.30	0.49	0.30	0.45	0.48
N	1,470	1,470	1,470	1,470	1,470	1,470	1,470	1,470	1,470
Public=1									
Mean	0.65	0.49	0.27	0.25	0.89	0.57	0.89	0.70	0.35
SD	0.48	0.50	0.44	0.44	0.32	0.50	0.31	0.46	0.48
N	1,485	1,485	1,485	1,485	1,485	1,485	1,485	1,485	1,485

Table A.1: Balance by Visibility Treatment

Notes: Table shows balance in demographic variables across assignments to the visibility treatment.

A.3.2 Demographic Re-weighting

The survey is fielded on an online platform Prolific, which is known among economists for providing quality human responses. However, selection into both the Prolific platform and taking each survey on the platform is non-random. Any survey fielded on Prolific thus likely is not representative of the U.S. population both along observable and unobservable demographics. In Table 1.1, I showed that my survey is representative of the U.S. in terms of age, sex, race, and political leanings. However, my sample is more likely to be college

	White	Male	Dem.	Rep.	Born US	College	Inc>25k	Inc>50k	Inc>100k
Price=$0.10									
Mean	0.64	0.49	0.30	0.26	0.88	0.57	0.90	0.69	0.34
SD	0.48	0.50	0.46	0.44	0.32	0.49	0.30	0.46	0.48
N	934	934	934	934	934	934	934	934	934
Price=$0.70									
Mean	0.67	0.51	0.26	0.26	0.91	0.59	0.89	0.72	0.37
SD	0.47	0.50	0.44	0.44	0.29	0.49	0.31	0.45	0.48
N	1,047	1,047	1,047	1,047	1,047	1,047	1,047	1,047	1,047
Price=$1.40									
Mean	0.64	0.49	0.26	0.26	0.89	0.59	0.89	0.71	0.33
SD	0.48	0.50	0.44	0.44	0.32	0.49	0.31	0.46	0.47
N	974	974	974	974	974	974	974	974	974

Table A.2: Balance by Price

Notes: Table shows balance in demographic variables across assignments of price.

educated and higher income than the general population. This leads to concerns of external validity.

In Table A.4, I present regression results which re-weight education and income to be more representative of the U.S. population. As can be seen, the results are virtually unchanged. This shows that results are robust to a re-weighting along observable demographics to better reflect the general population of the U.S.

A.4 Additional Tables and Figures

	White	Male	Dem.	Rep.	Born US	College	Inc>25k	Inc>50k	Inc>100k
Signal=1.4%									
Mean	0.66	0.53	0.25	0.29	0.88	0.58	0.91	0.71	0.38
SD	0.48	0.50	0.43	0.46	0.32	0.49	0.28	0.45	0.48
N	496	496	496	496	496	496	496	496	496
Signal=5.6%									
Mean	0.65	0.49	0.27	0.25	0.90	0.64	0.92	0.76	0.40
SD	0.48	0.50	0.45	0.43	0.31	0.48	0.27	0.43	0.49
N	491	491	491	491	491	491	491	491	491
Signal=12.8%									
Mean	0.64	0.48	0.31	0.25	0.93	0.58	0.89	0.71	0.34
SD	0.48	0.50	0.46	0.43	0.26	0.49	0.31	0.46	0.47
N	492	492	492	492	492	492	492	492	492
Signal=21.7%									
Mean	0.66	0.51	0.26	0.29	0.89	0.58	0.90	0.70	0.33
SD	0.47	0.50	0.44	0.45	0.32	0.49	0.30	0.46	0.47
N	481	481	481	481	481	481	481	481	481
Signal=31.5%									
Mean	0.64	0.50	0.24	0.29	0.89	0.60	0.89	0.71	0.34
SD	0.48	0.50	0.43	0.46	0.31	0.49	0.31	0.46	0.48
N	490	490	490	490	490	490	490	490	490
Signal=43.8%									
Mean	0.65	0.50	0.29	0.21	0.88	0.53	0.86	0.66	0.32
SD	0.48	0.50	0.45	0.41	0.32	0.50	0.35	0.47	0.47
N	511	511	511	511	511	511	511	5110	511

Table A.3: Balance by Second-Order Belief Signal

Notes: Table shows balance in demographic variables across assignments to signal varying second-order beliefs.

	(1)	(2)	(3)	(4)
Public	0.228*	0.251**	0.293**	0.299**
	(0.120)	(0.128)	(0.141)	(0.147)
SB X Public	-0.0696**	-0.0677**	-0.0672*	-0.0756**
	(0.0295)	(0.0313)	(0.0347)	(0.0367)
$(SB)^2$ X Public	0.00492**	0.00467*	0.00486*	0.00577**
	(0.00233)	(0.00244)	(0.00268)	(0.00289)
$(SB)^3$ X Public	-0.000127*	-0.000117	-0.000128	-0.000159*
	(0.0000712)	(0.0000740)	(0.0000804)	(0.0000883)
$(SB)^4$ X Public	0.00000112	0.000000994	0.00000116	0.00000146
	(0.000000729)	(0.000000753)	(0.000000813)	(0.000000908)
Price	-0.265***	-0.265***	-0.265***	-0.278***
	(0.0247)	(0.0247)	(0.0247)	(0.0256)
Price X Public	-0.0188	-0.0302	-0.0568	-0.0368
	(0.0342)	(0.0354)	(0.0373)	(0.0390)
Drop Visibility Doubters?	No	Yes	Yes	Yes
Drop Those With Privacy Concerns?	No	No	Yes	Yes
Drop Far Right?	No	No	No	Yes
Observations	2685	2525	2305	2107
Dep. Var. Mean	0.34	0.35	0.36	0.37

Table A.4: Re-weighted Reduced Form Estimates of Visibility's Impact by Market Penetration

Notes: Table shows how the impacts of visibility on uptake vary by second-order beliefs about market penetration of carbon offsets. Results are re-weighted by income and educational attainment using raking. Estimates are then generated by a two-stage linear probability regression. In the first stage, second-order beliefs are instrumented with a randomly-generated signal. Column (1) depicts estimates for the full sample. Column (2) depicts estimates after dropping individuals who did not believe the website treatment was actually publicly visible. Column (3) depicts estimates after additionally dropping those who reported having strong privacy concerns about their name being posted. Column (4) depicts estimates additionally dropping those reported being politically very conservative.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

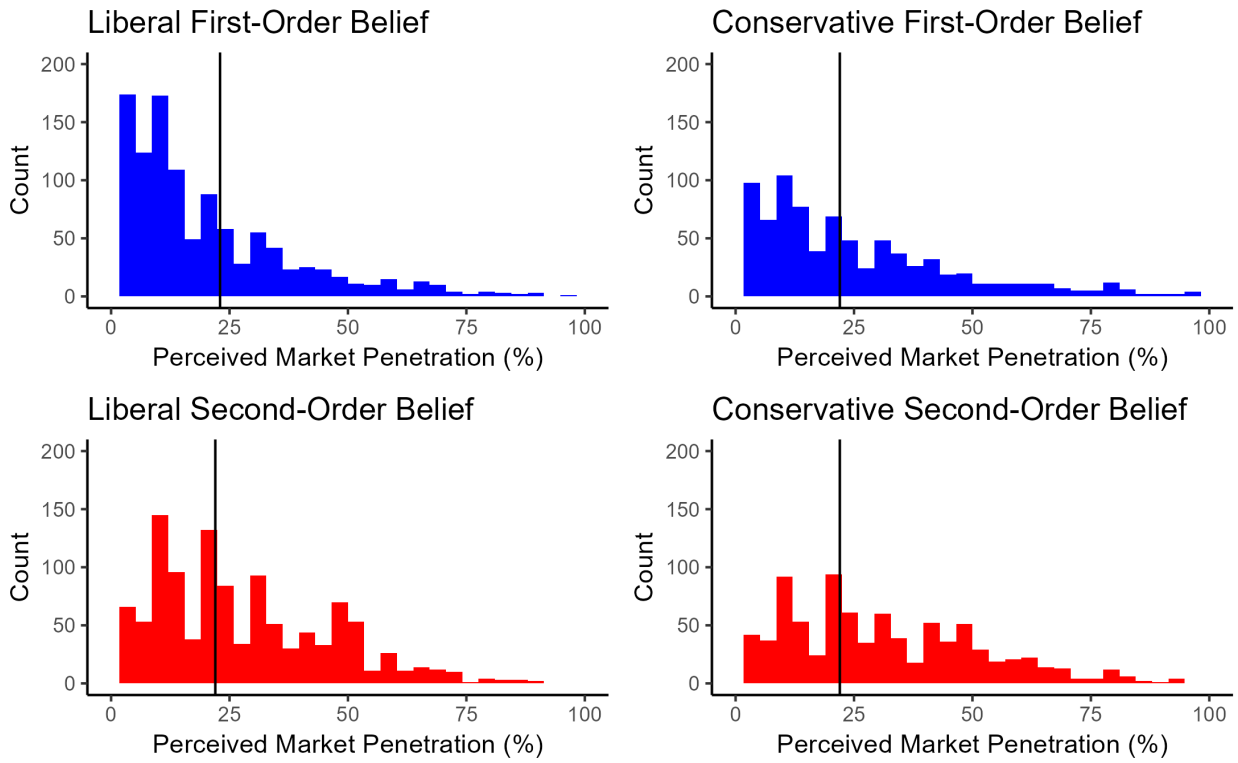


Figure A.5: Histograms of Beliefs by Political Ideology

Notes: Figure depicts histograms of first-order (top) and prior second-order beliefs (bottom) separately by political ideology. All individuals failing attention checks or reCAPTCHA are dropped from this figure. First-order beliefs are an individual's beliefs on what percent of people have bought a carbon offset. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Vertical black lines depict correct answers for each respective graph.

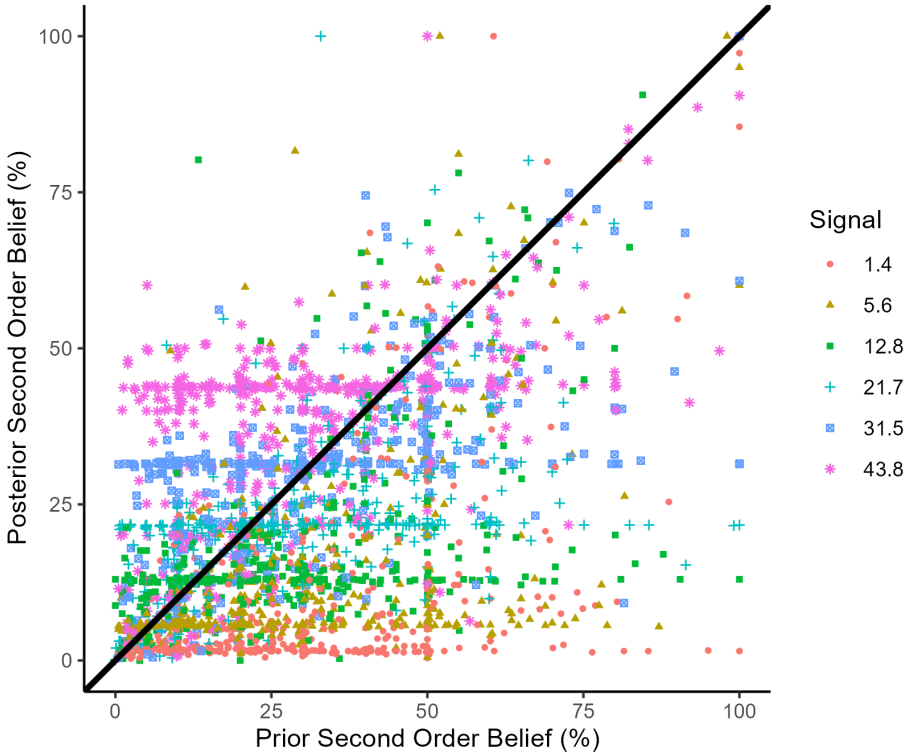


Figure A.6: First- Versus Second-order Beliefs

Notes: Figure depicts a scatterplot of individuals' prior second- and posterior second-order beliefs. The shape and color of each point is determined by the signal each individual was sent. The red 45∘ line indicates when priors and posteriors perfectly coincide. All individuals failing attention checks or reCAPTCHA are dropped from this figure. First-order beliefs are an individual's beliefs on what percent of people have bought a carbon offset. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are.

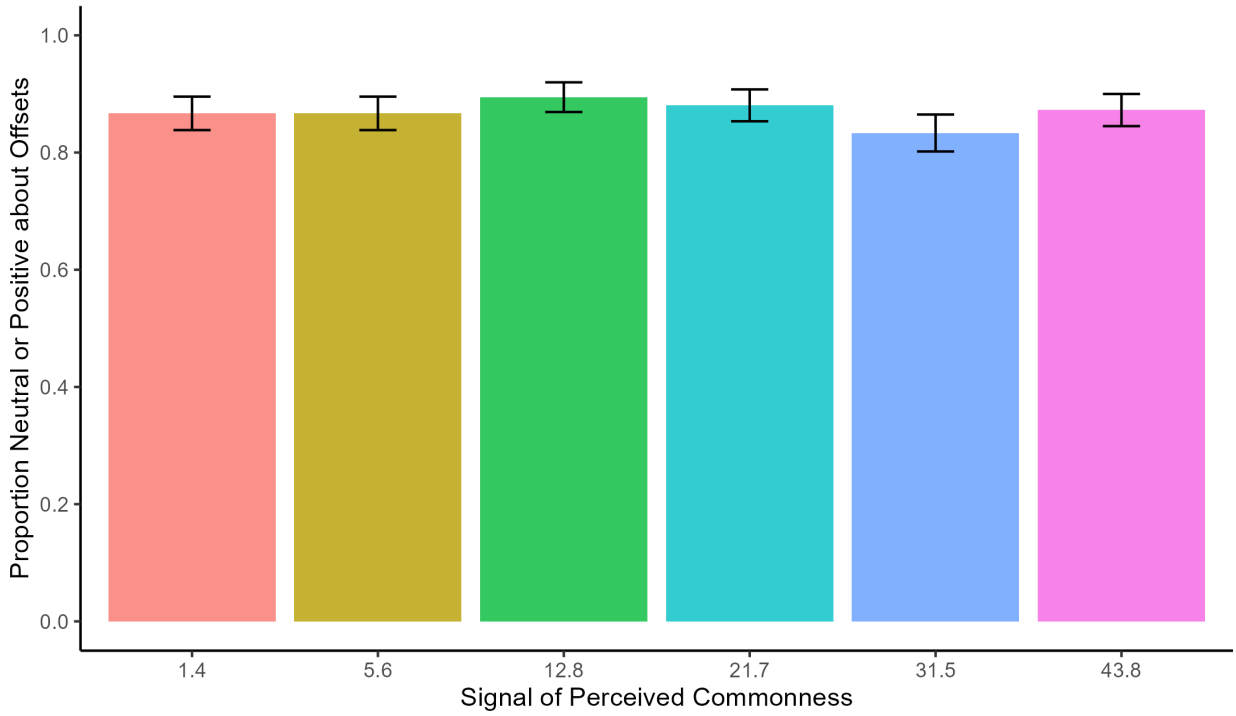


Figure A.7: Posterior Perceptions of Quality by Randomly Assigned Commonness Signal

Notes: Figure depicts bar charts of the proportion of individuals that report being positive or neutral about the effectiveness of carbon offsets by the signal of perceived commonness they received. Black intervals bars depict 90% confidence intervals.

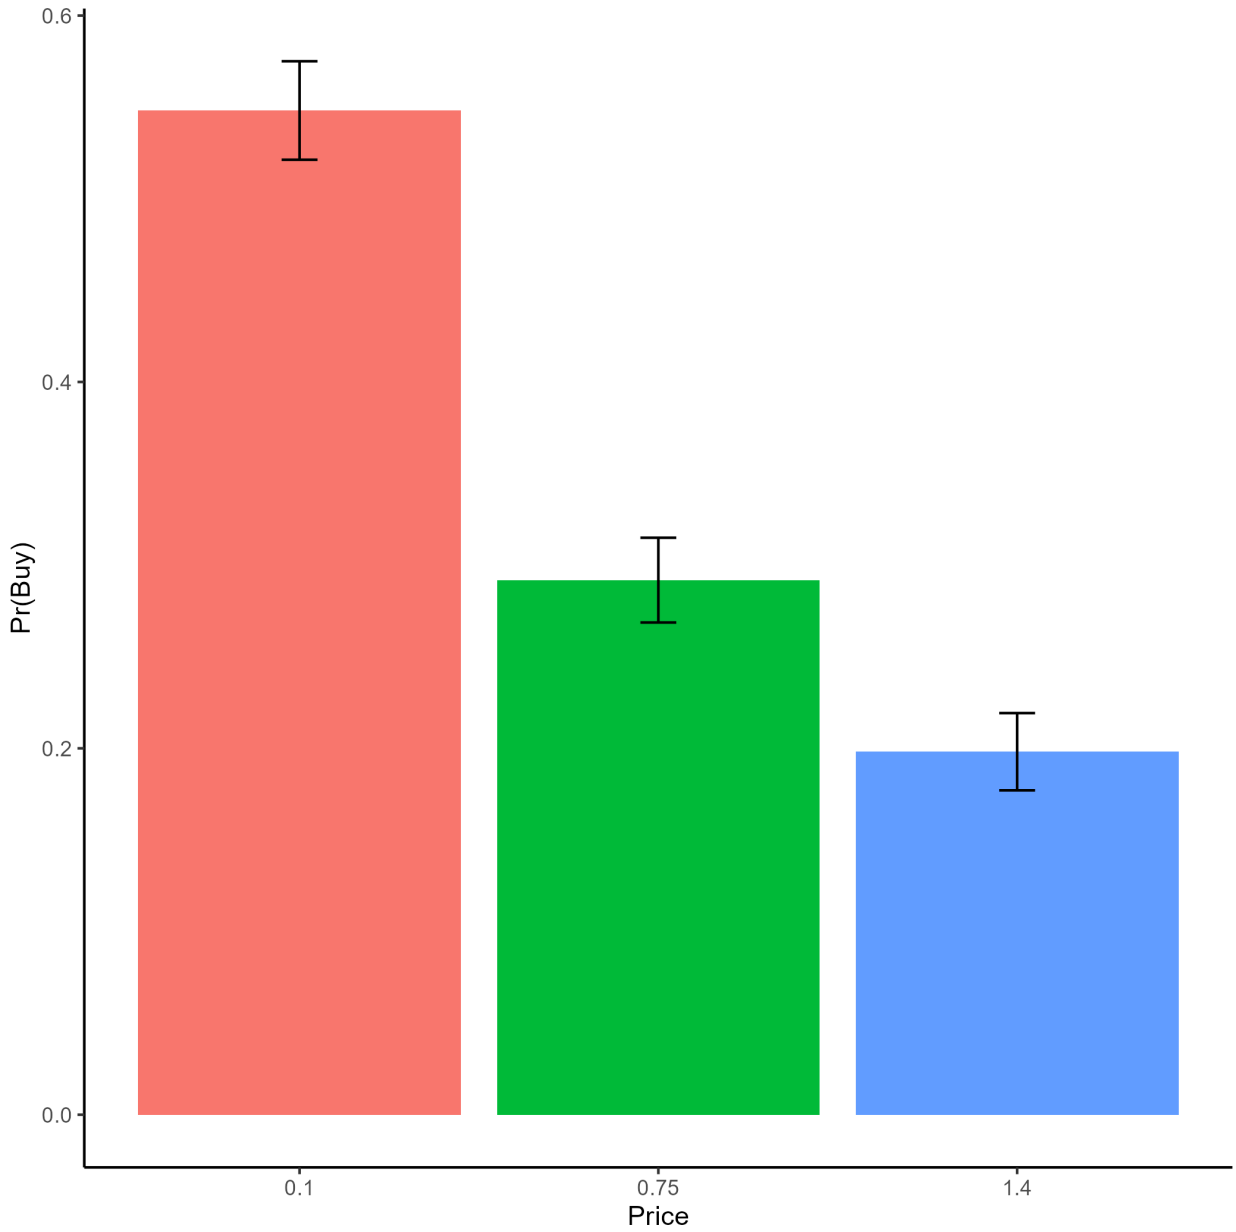


Figure A.8: The Impact of Price on Uptake

Notes: Figure depicts a bar graph of uptake of carbon offsets by price. Black horizontal lines indicate the borders of a 90% confidence interval.

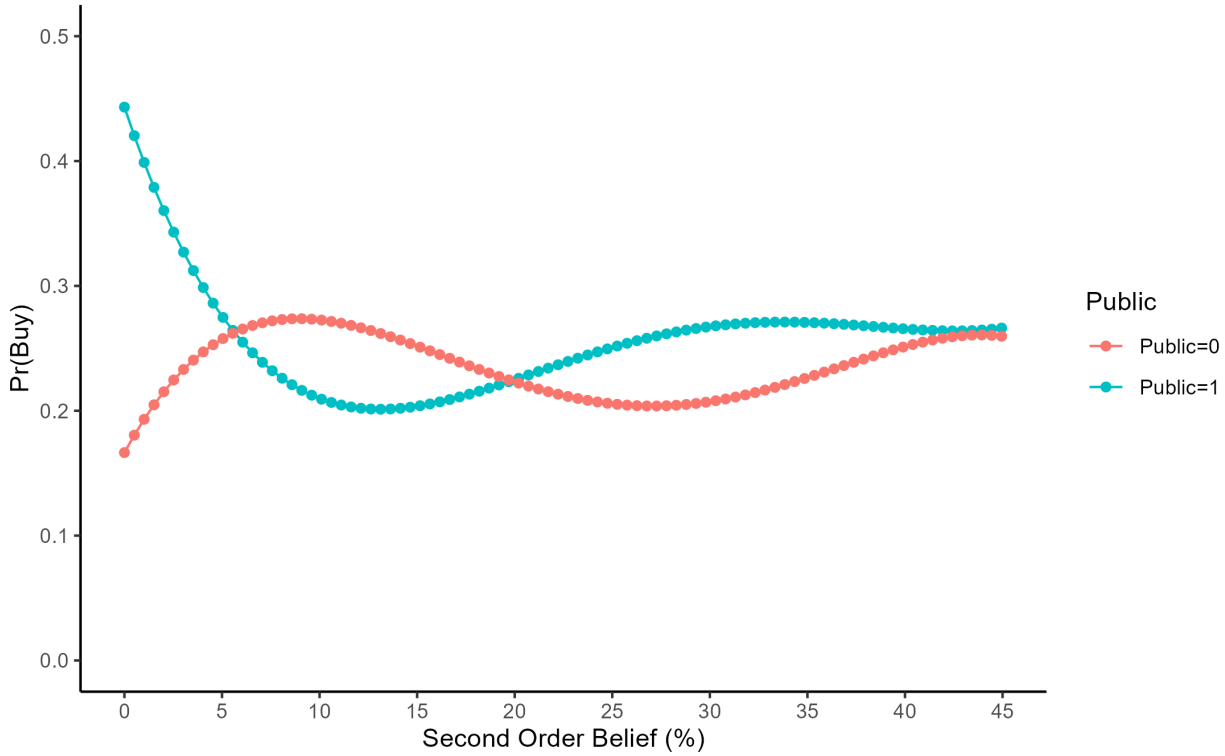


Figure A.9: Reduced Form Impacts to Uptake by Second-order Belief Separated by Visibility Treatment

Notes: Figure depicts the estimated probabilities of purchasing an offset by second-order beliefs of market penetration of carbon offsets, separated by whether an individual was in the visibility treatment or not. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue line is for the visible group and the red line is for the invisible control group.

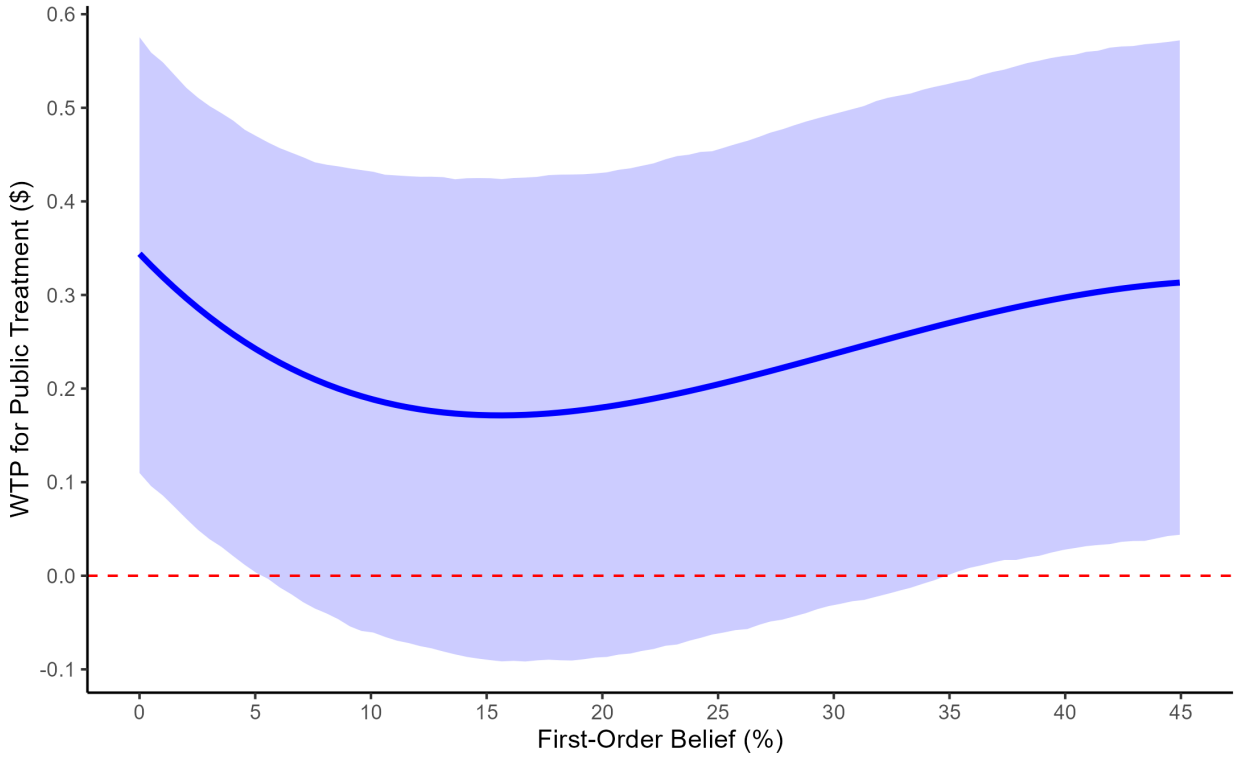


Figure A.10: Reduced Form Impact of Visibility by First-order Belief

Notes: Figure depicts the estimated impact of visibility on uptake probability by first-order beliefs of market penetration of carbon offsets. First-order beliefs are an individual's beliefs about how common carbon offsets are and are allowed to vary with publicity quadratically. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

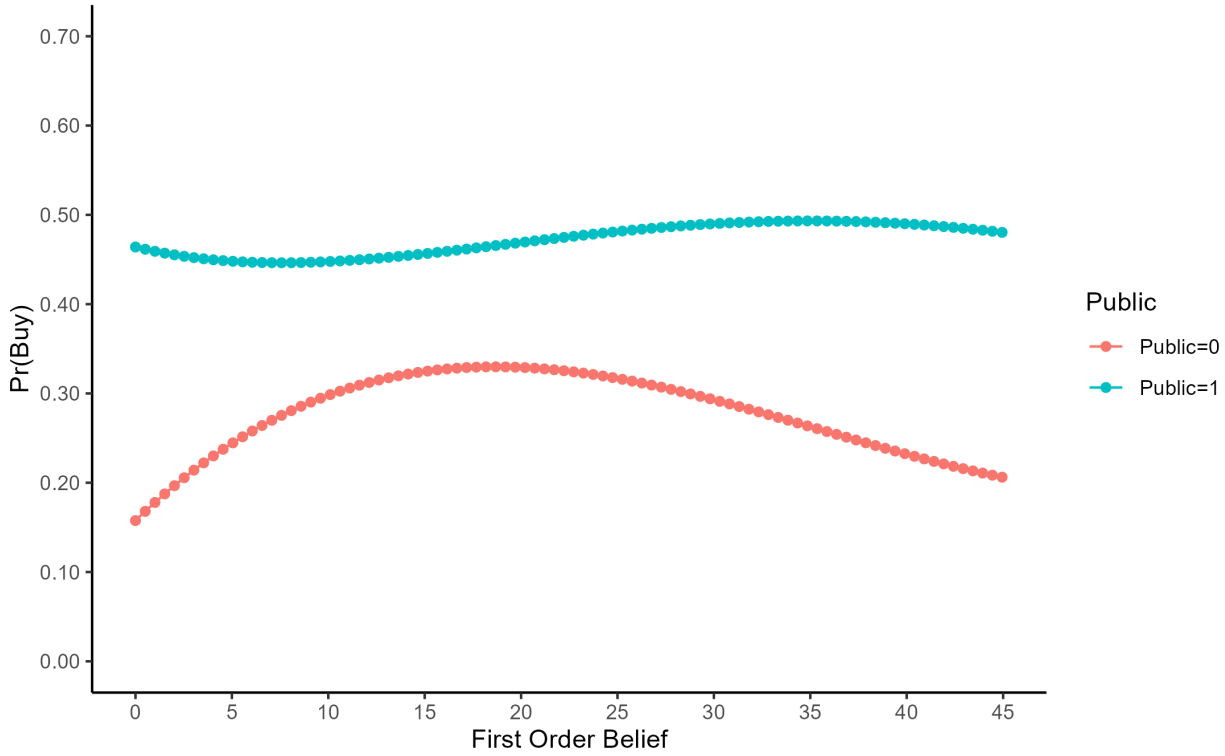


Figure A.11: Reduced Form Impacts to Uptake by First-order Belief Separated by Visibility Treatment

Notes: Figure depicts the estimated probabilities of purchasing an offset by first-order beliefs of market penetration of carbon offsets, separated by whether an individual was in the visible treatment or not. First-order beliefs are instrumented for with a randomly generated signal. The blue line is for the visible group and the red line is for the control group. The red and blue bands depict 90% confidence intervals.

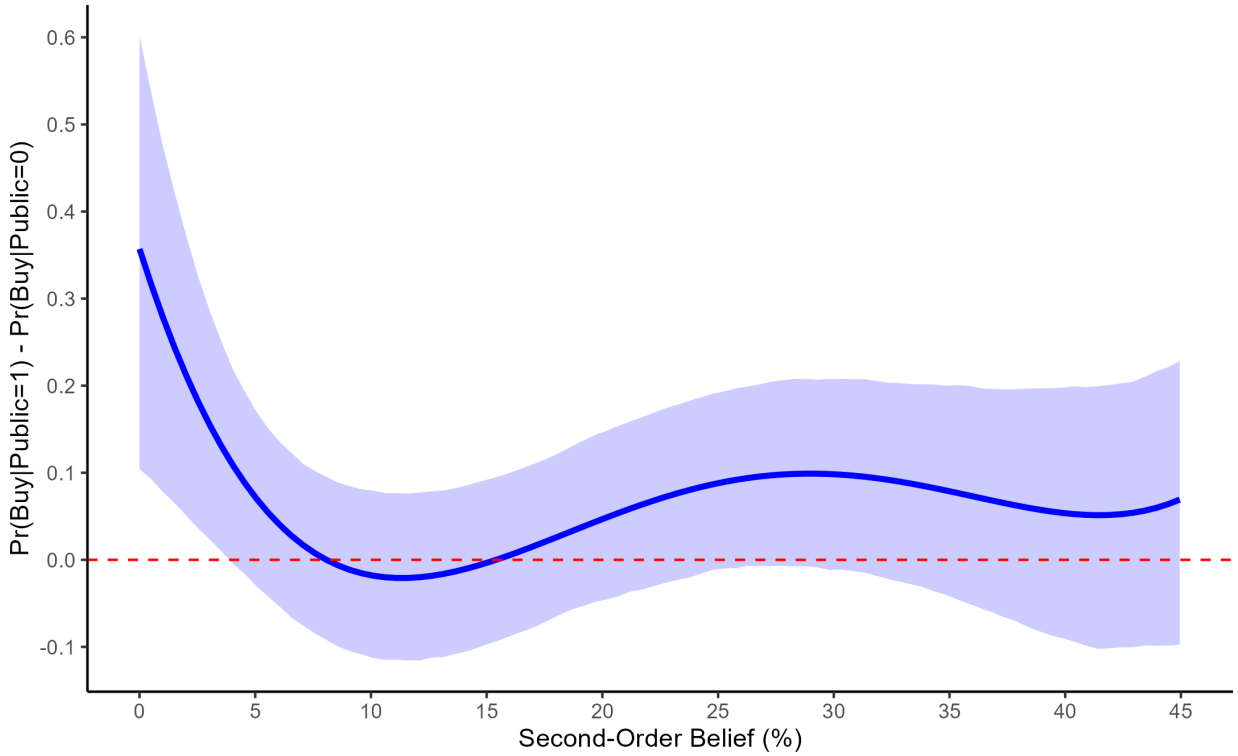


Figure A.12: Reduced Form Impact of Visibility by Second-order Belief Omitting Referrers

Notes: Figure depicts the linear probability estimated impact of visibility on purchasing an offset by second-order beliefs of market penetration of carbon offsets, omitting all those who reported that they "Strongly disagree" or "Somewhat disagree" that the website posting signaled eco-friendliness but were neutral or agreed that it would encourage others to donate. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

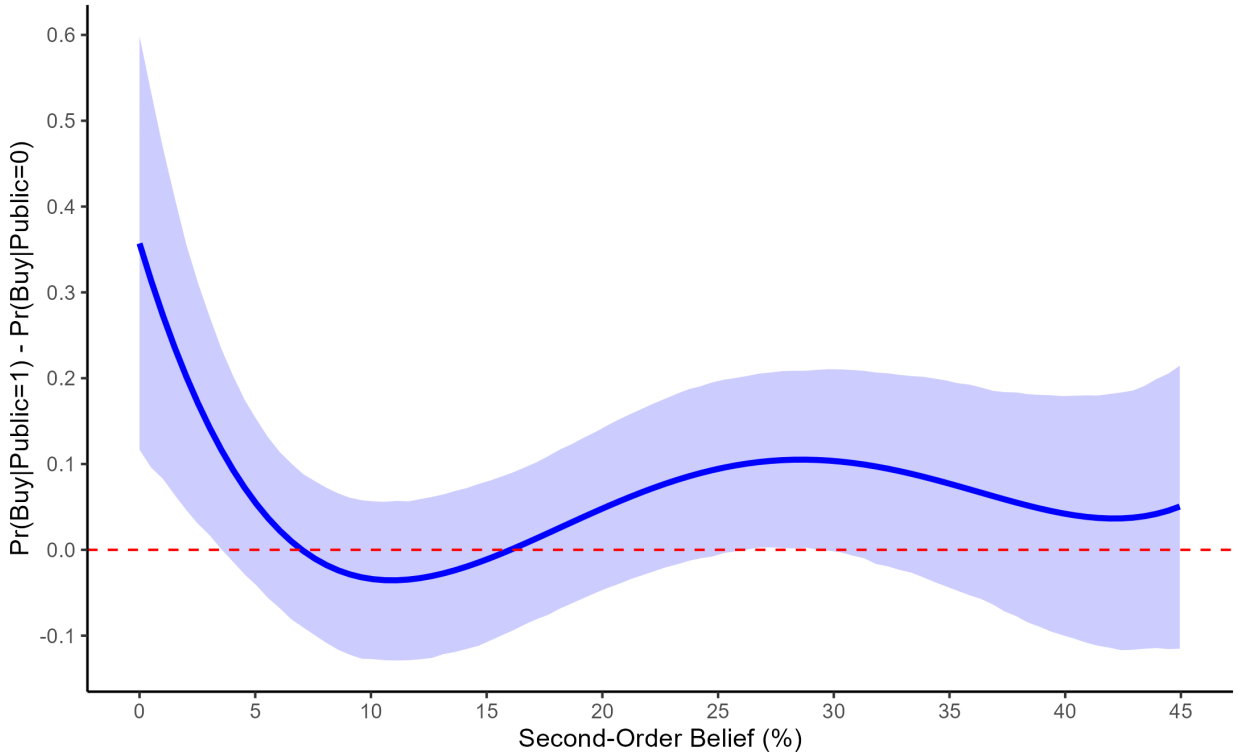


Figure A.13: Reduced Form Impact of Visibility by Second-order Belief Omitting Virtue Signaling Concerns

Notes: Figure depicts the linear probability estimated impact of visibility on purchasing an offset by second-order beliefs of market penetration of carbon offsets, excluding all those who reported that they "Strongly agree" that they were concerned about being seen as virtue signaling or bragging. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

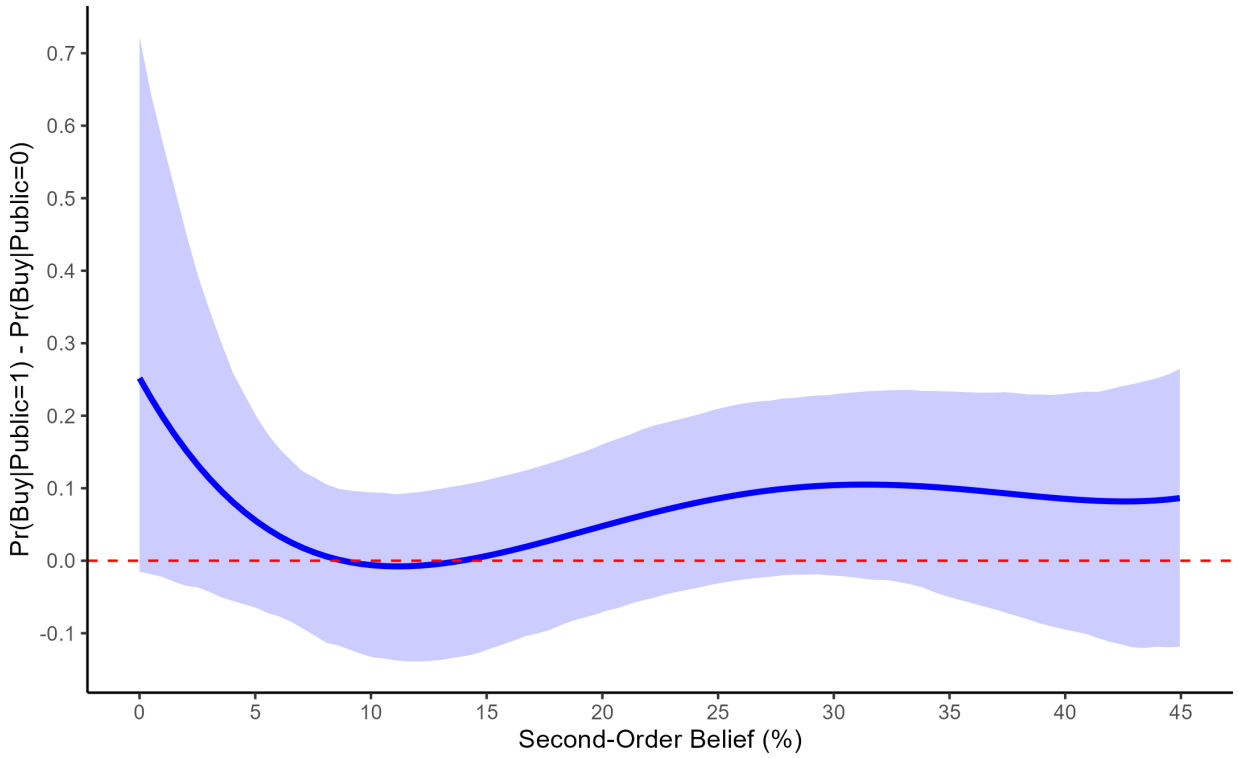


Figure A.14: Reduced Form Impact of Visibility by Second-order Belief Omitting Experienced Survey-Takers

Notes: Figure depicts the linear probability estimated impact of visibility on purchasing an offset by second-order beliefs of market penetration of carbon offsets, excluding individuals who have taken more than 2,000 surveys on Prolific (roughly the top 26%). Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

	(1)	(2)	(3)	(4)
Signal	0.627***	0.633***	0.636***	0.644***
	(0.0147)	(0.0152)	(0.0158)	(0.0163)
Drop Visibility Doubters?	No	Yes	Yes	Yes
Drop Those With Privacy Concerns?	No	No	Yes	Yes
Drop Far Right?	No	No	No	Yes
Observations	2,685	2,525	2,305	2,107
F Stat	100.00	93.88	87.05	84.41

Table A.5: First Stage Regressions

Notes: Table shows results from the first stage regression of second-order beliefs on signals. Column (1) depicts estimates for the full sample. Column (2) depicts estimates after dropping individuals who did not believe the website treatment was actually publicly visible. Column (3) depicts estimates after additionally dropping those who reported having strong privacy concerns about their name being posted. Column (4) depicts estimates additionally dropping those reported being politically very conservative.
$^{*}\, p < 0.10,\ ^{**}\, p < 0.05,\ ^{***}\, p < 0.01.$

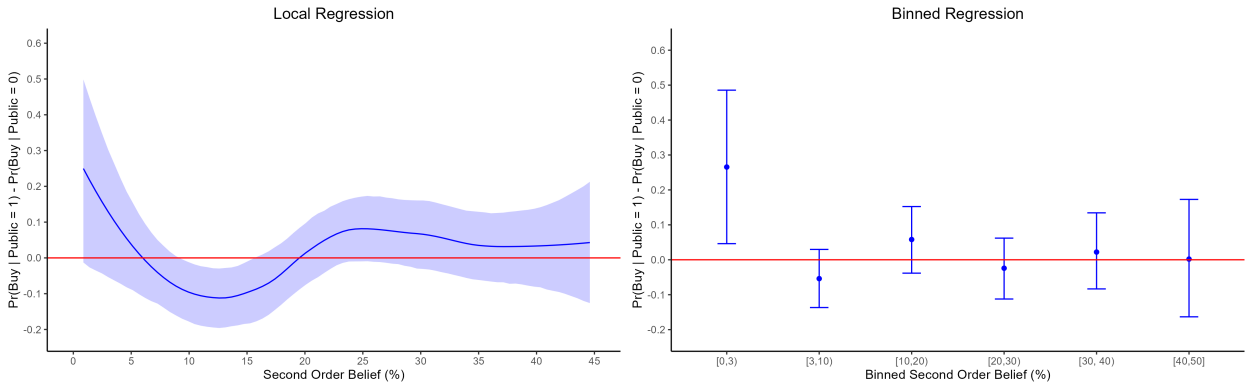


Figure A.15: Nonparametric Estimates of Visibility's Impact by Market Penetration

Notes: Figure depicts reduced form nonparametric estimates of the impact to uptake from visibility. The left panel does so with a local linear regression. The right panel does so with a linear probability model regression that interacts the visibility treatment with bins of fitted second-order beliefs. In both panels, second-order beliefs are instrumented for by a randomly shown signal. The blue ribbon in the left panel and blue bands in the right panel depict 90% confidence intervals. In both panels, confidence intervals are calculated via a bootstrap with 2,000 iterations.

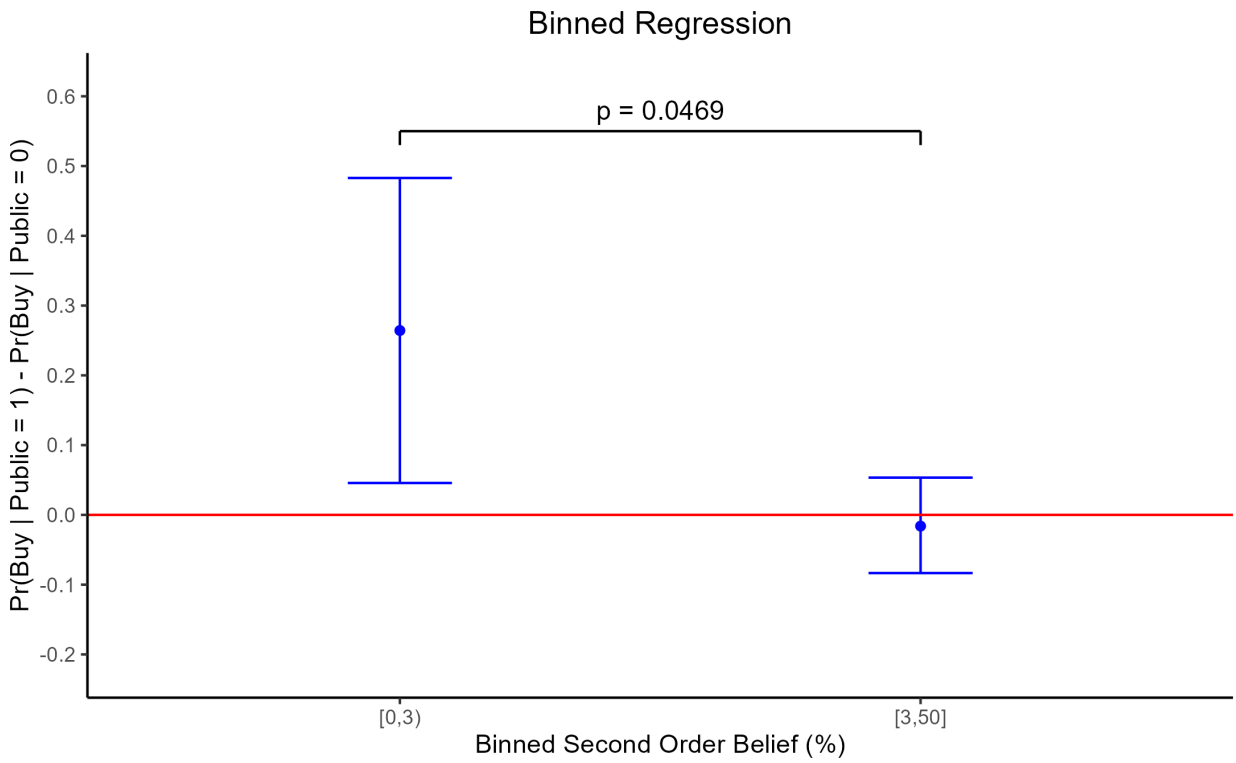


Figure A.16: Two-Bin Estimates of Visibility's Impact by Market Penetration

Notes: Figure depicts reduced form nonparametric estimates of the impact to uptake from visibility. Figure does so comparing binned fitted beliefs of 0-3% versus 3-50%. Second-order beliefs are instrumented for by a randomly shown signal. Blue bands depict 90% confidence intervals which are calculated via a bootstrap with 2,000 iterations.

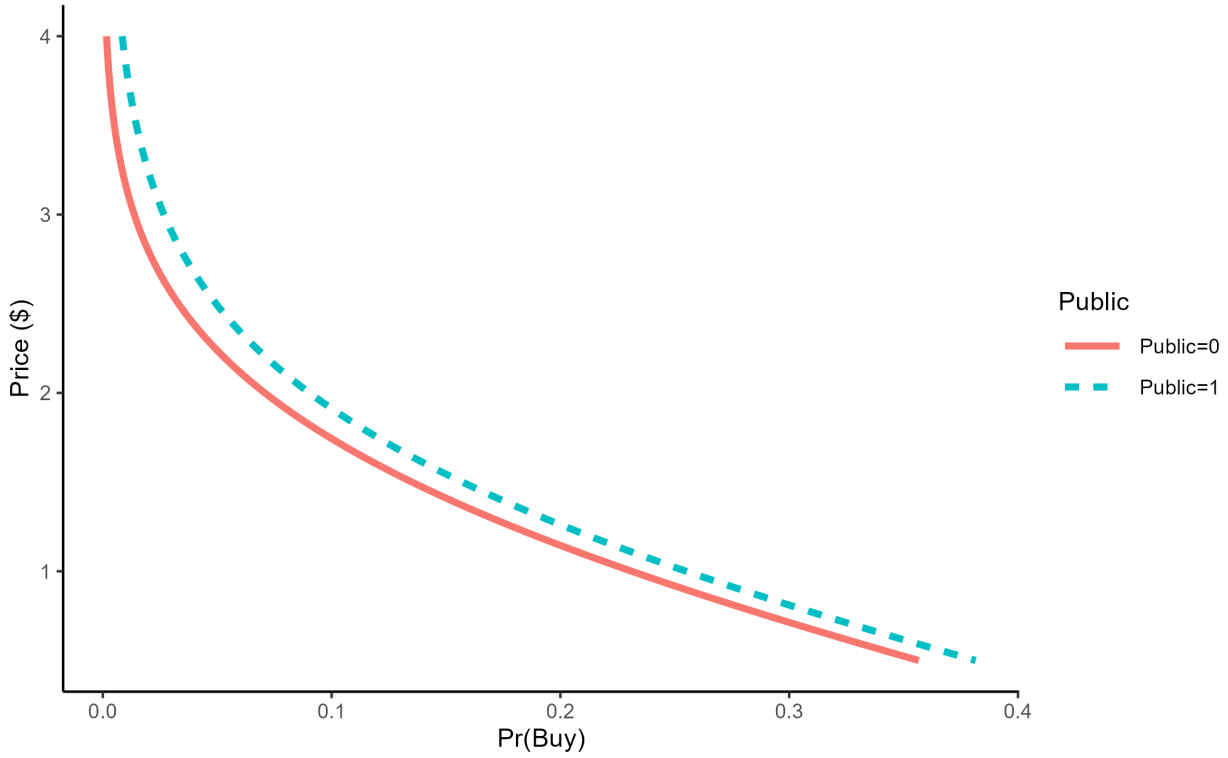


Figure A.17: Structurally Estimated Demand for Mitigating 145lbs. of Carbon

Notes: Figure depicts structural estimates of the demand for mitigating 145lbs. of carbon emissions in the invisible (solid red) and visible (dashed blue) settings. Second-order beliefs about uptake are instrumented for with a randomly generated signal. Second-order beliefs are assumed to coincide with actual uptake on average (e.g. a rational expectations assumption) to generate this figure.

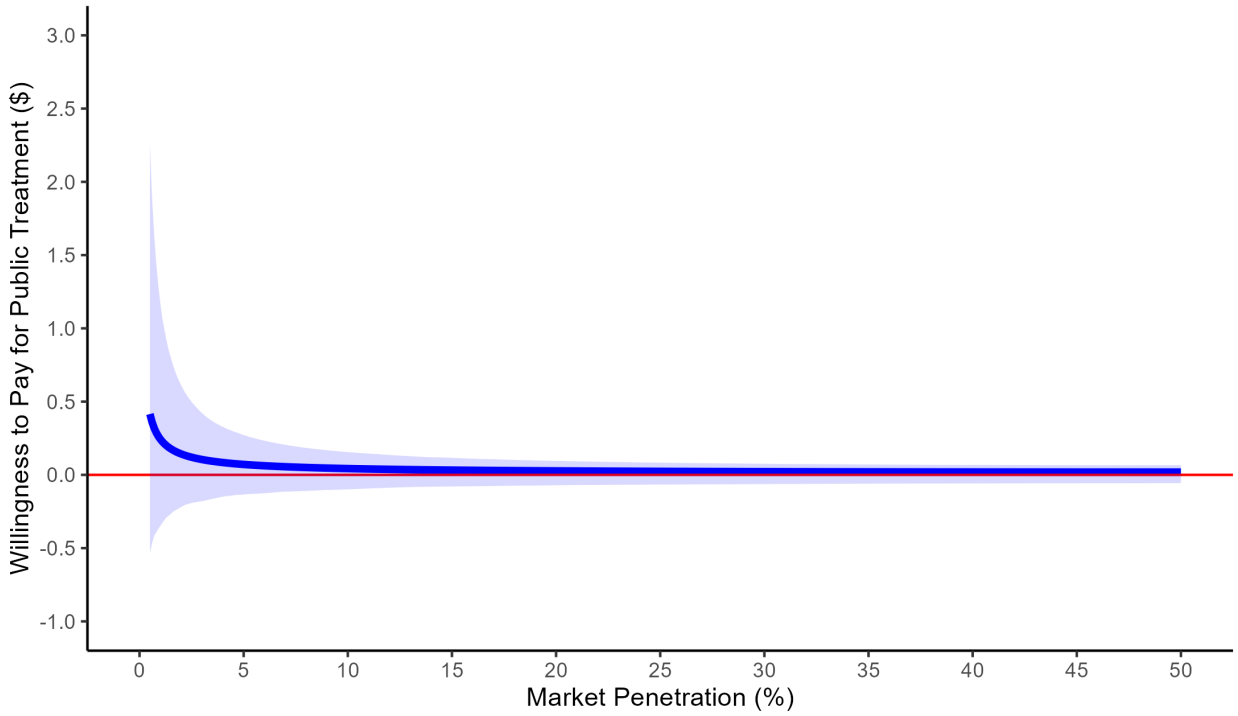


Figure A.18: Structurally Estimated WTP for Visibility by Market Penetration for Full Sample

Notes: Figure depicts structural estimates of the willingness-to-pay (WTP) for visibility from purchasing an offset by second-order belief of market penetration of carbon offsets, including for those who did not trust that the visibility treatment was visible. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

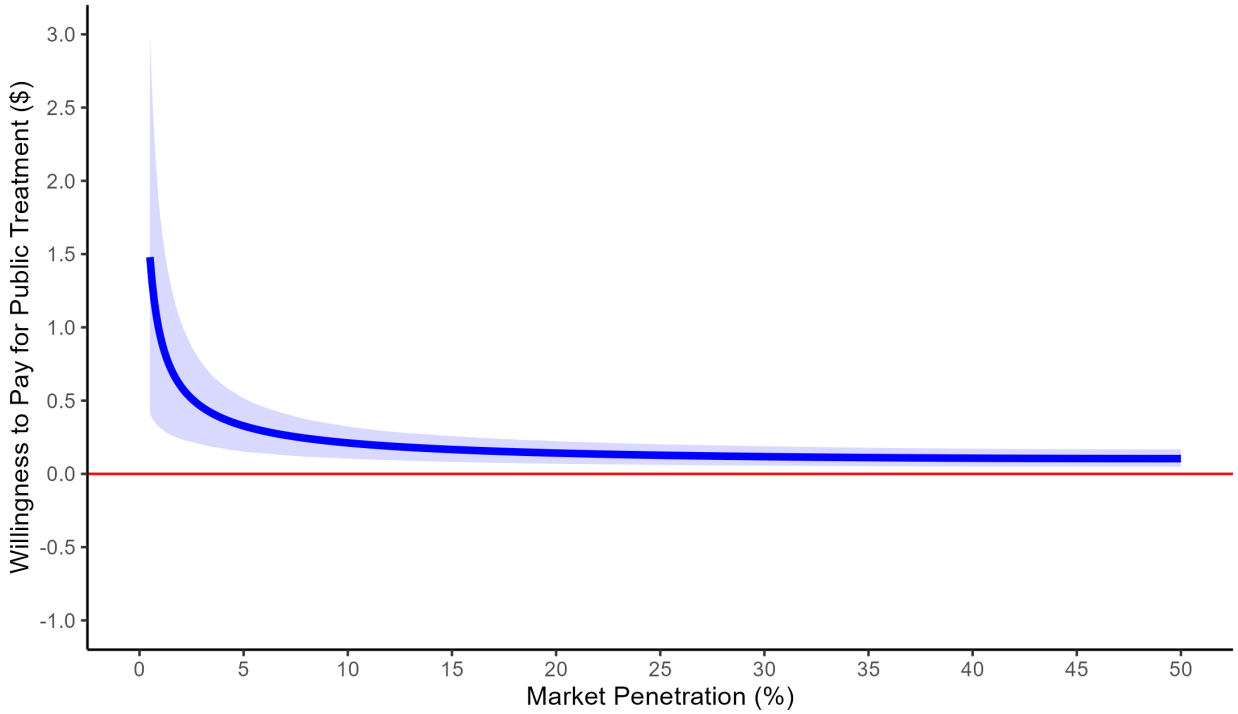


Figure A.19: Structurally Estimated WTP for Visibility by Market Penetration Omitting Conservatives

Notes: Figure depicts structural estimates of the willingness to pay for visibility from purchasing an offset by second-order beliefs of market penetration of carbon offsets, omitting all those who report being "Very Right" politically. Second-order beliefs are an individual's beliefs about the average of all other participants' beliefs on how common carbon offsets are. Second-order beliefs are instrumented for with a randomly generated signal. The blue shaded area depicts 90% confidence intervals and the solid blue line depicts point estimates. Standard errors are calculated via bootstrapping with 2,000 iterations.

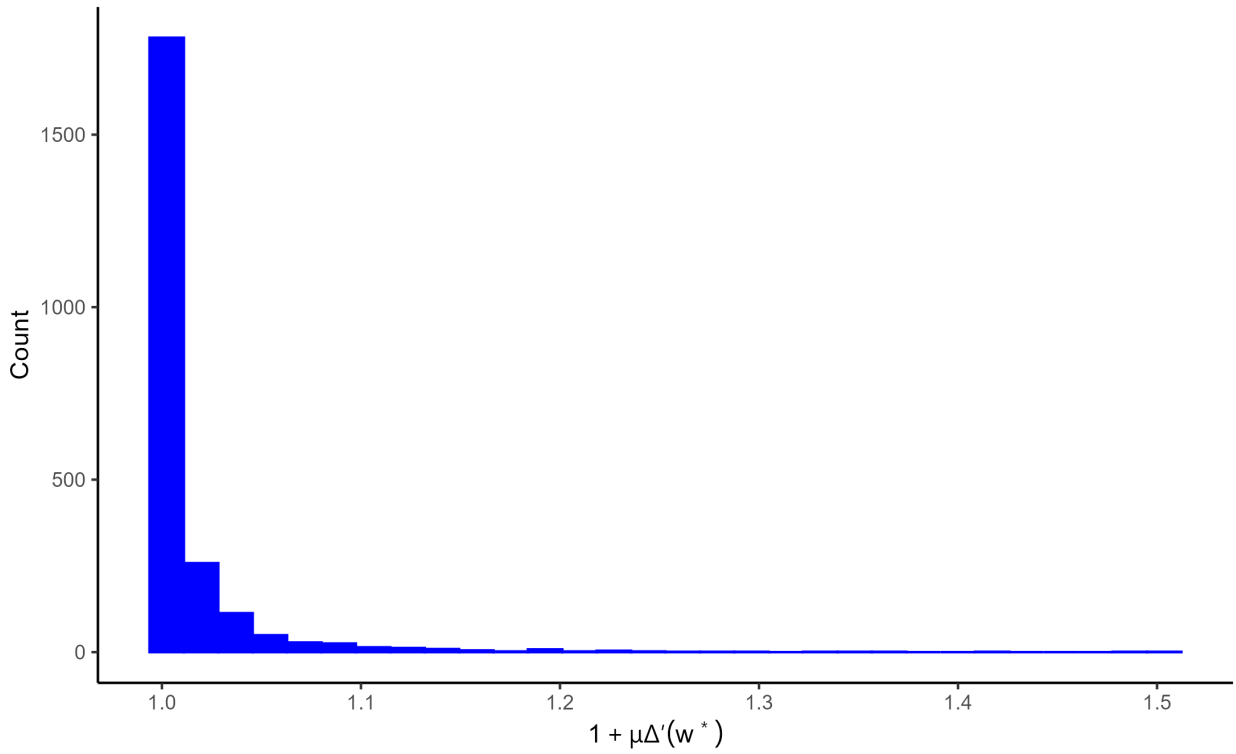


Figure A.20: Structural Estimates of the Model's Assumption

Notes: Figure depicts a histogram of structural estimates of the value of $1 + \mu\Delta'(w^*)$*, which must be positive in order for predictions 1-3 of the model to hold.*

APPENDIX B

APPENDIX FOR CHAPTER 2

B.1 Robustness Checks

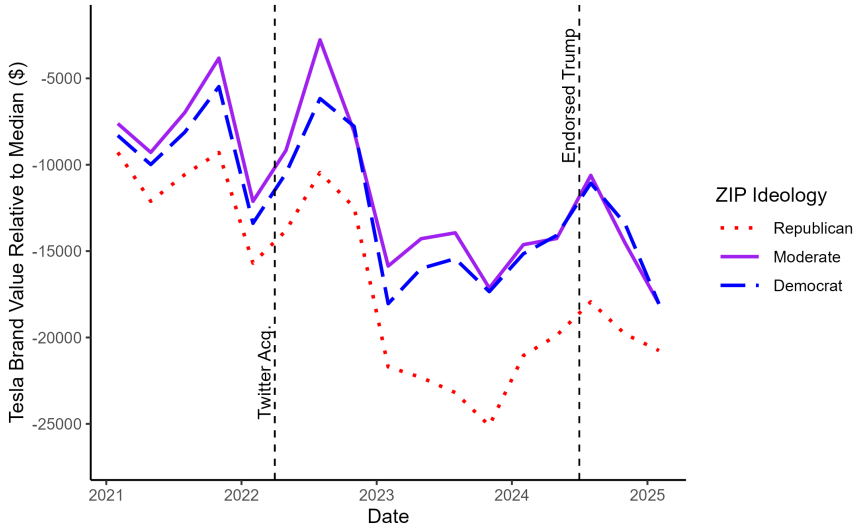
B.1.1 Varying Cutoffs of Political Classifications

The main body of the paper uses cutoffs of 40% and 60% vote share for Trump in the 2020 election to classify Democrat and Republican ZIPs. These cutoffs are arbitrary, but are chosen to balance a trade-off between exclusivity (i.e. “too many" moderate ZIP codes being counted as Democrat or Republican) and statistical power (too few ZIPs in each category to detect effects). However, it is important to understand the extent to which these cutoffs alter the conclusions of our paper.

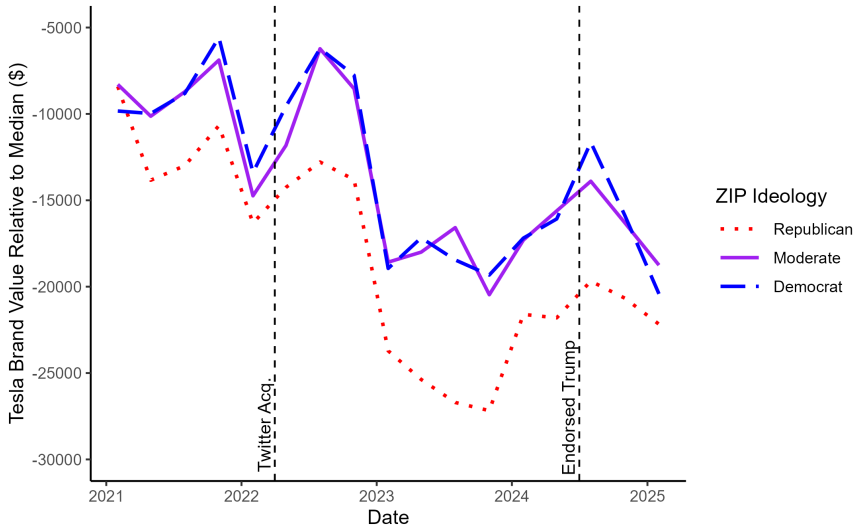
In Appendix Figure B.1, we show similar qualitative conclusions to the main paper when considering alternative classifications of Democrat, moderate, and Republican ZIP codes. In particular, all cutoffs show the same patterns as we find in our main specification. We also observe absolute decreases in moderate and Democrat WTP during the 2024 election cycle and little change in WTP among Republicans. Encouragingly, the magnitudes of estimates are similar when considering cutoffs of (45%, 55%) and (30%, 70%).

B.1.2 Using 2024 Vote Shares

In Appendix Figure B.2, we show that our results are robust to using 2024 vote shares to classify ZIP ideology instead of 2020 vote shares. We again find an absolute decrease in WTP for Teslas among Democrat and moderate ZIPs that is not present for Republican ZIPs. We note that the magnitude of the estimates appear to be larger for Democrat ZIPs and smaller for moderate ZIPs. This likely reflects that the 2024 election – where Trump received a greater share of votes than the 2020 election – led to only the most left-leaning of



Panel A: Cut-offs of 45% and 55%



Panel B: Cut-offs of 30% and 70%

Figure B.1: Alternative Cut-offs for Democrat and Republican ZIPs

Notes: Figure depicts estimates of consumer WTP for a Tesla vehicle relative to the median brand. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using varying cut-offs of voting for Trump in the 2020 election. The cutoffs are $\{45\%, 55\%\}$ *in Panel A and* $\{30\%, 70\%\}$ *in Panel B. WTP for the Tesla brand is calculated by taking Tesla's quarter-year specific fixed effects and comparing them to the median of quarter-year fixed effects of all other brands for each political leaning. This is then converted to WTP for a brand by normalizing the difference in fixed effects by an estimated price coefficient. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).*

ZIPs being in our Democrat classification.

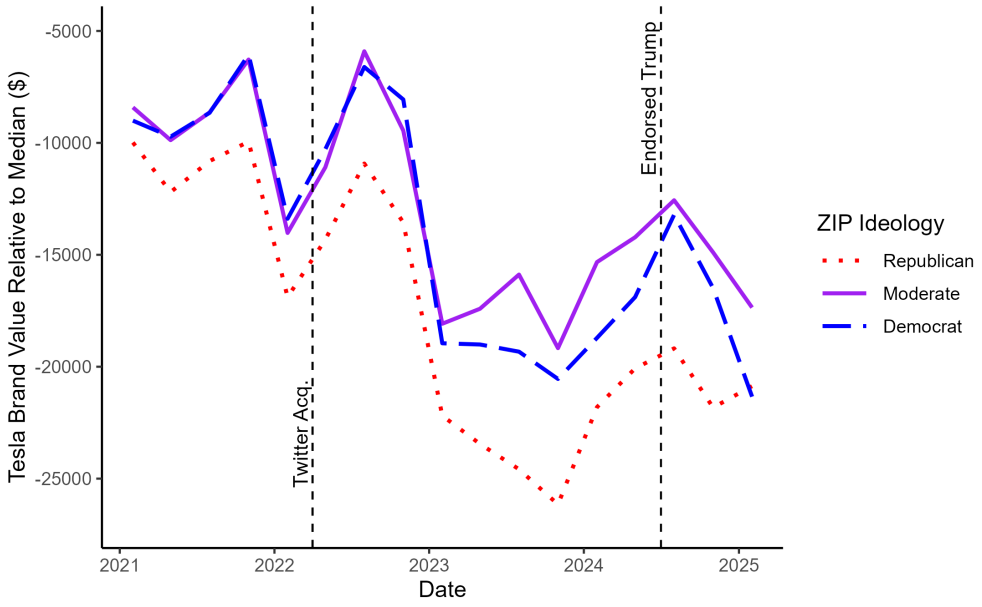


Figure B.2: Trends in WTP for Teslas Using 2024 Vote Shares

Notes: Figure depicts estimates of consumer WTP for a Tesla vehicle relative to the median brand. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2024 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2024 Trump vote share. WTP for the Tesla brand is calculated by taking Tesla's quarter-year specific fixed effects and comparing them to the median of quarter-year fixed effects of all other brands for each political leaning. This is then converted to WTP for a brand by normalizing the difference in fixed effects by an estimated price coefficient. Price is instrumented for by variation in eligibility for federal subsidies generated by the 2022 Inflation Reduction Act. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).

B.1.3 Instrument Measurement Error

We do not perfectly observe when subsidies are applied to vehicles, instead only observing subsidy eligibility. This is not an issue under point-of-sale transactions in which the subsidy is

immediately applied to the price.[1] However, subsidies prior to January 2024 were allocated as tax rebates. We have assumed that all vehicles prior to January 2024 receive the full amount of subsidy they are eligible for, but this likely creates measurement error in the 2SLS procedure. This measurement error raises concerns of exclusion restriction violations. We argue that the exclusion restriction is not violated, though potential measurement error likely increased noise in our estimates.

We control for both ZIP income and ZIP tax filing rates in Table 2.2, which alleviates concerns about measurement error.[2] The primary mechanisms through which individuals would not receive a subsidy they were eligible for prior to 2024 is that (i) their income is too low to have a high enough tax liability to receive the subsidy or (ii) they did not file taxes. Controlling for tax filing rates and income alleviates concerns that these mechanisms lead federal subsidy allocations to be correlated with drivers of demand aside from price. Conditional on these variables, we assume any additional measurement error is spurious, lacking correlation with the error term ξ_{jt}.

B.2 Sentiment Analysis

B.2.1 Reddit Comment Selection

In order to determine how sentiment for Elon Musk and Tesla has changed over time for Democrats, moderates, and Republicans, we collect data on Reddit comments. We source data from the Reddit Pushshift API, which contains data on all comments and all posts across the 40,000 largest subreddits up through December 31, 2024. For our analysis, we focus on the r/conservative, r/centrist, and r/liberal subreddits. These are popular subreddits which

1. Transaction prices are provided net of subsidies after January 2024. See box 38 here `https://www.txdmv.gov/sites/default/files/form_files/130-U.pdf`.

2. Income is collected from the 2023 ACS and tax filing rates from the IRS website here: `https://www.irs.gov/statistics/soi-tax-stats-individual-income-tax-statistics-zip-code-data-soi`.

provide a discussion platform for individuals espousing each of the corresponding ideologies. Within those subreddits, we examine the comments (as well as the content of the original post itself) on posts that include the key words “Musk," “Elon," “Tesla," or “SpaceX" in the title. This yields a total of 153,062 comments. Comments which are marked as deleted or removed by a Reddit moderator are omitted, resulting in 102,336 comments that we analyze: 1,506 comments on r/liberal, 16,879 comments on r/centrist, and 83,951 comments on r/conservative.

B.2.2 Analyzing Sentiment of Reddit Comments

For each Reddit comment, we generate two sentiment scores based on the sentiment expressed towards Musk and Tesla respectively. To conduct this sentiment analysis, we use ChatGPT’s GPT-4o model via OpenAI’s API. The usage of a large language model to measure sentiment is common practice in the literature [Alonso-Robisco and Carbó, 2023, Fischer et al., 2023, Jomni and Zakharov, 2024, Kwon et al., 2024, Rathje et al., 2024, Hu, 2025].

For each comment, we run the following prompt via the OpenAI API:

> You are a helpful assistant performing sentiment analysis on Reddit threads about Elon Musk and Tesla (the electric vehicle company). The comments are provided in the Comment column, and are grouped into threads by Identifier. Date specifies when a particular comment was posted. So you can identify previous comments in a thread by looking at entries with earlier Dates but with the same Identifier.
> You will analyze one comment at a time, taking into account the previous comments (context), and return two scores:
>
> 1. **ElonScore**
> - A score from -1 (very negative) to 1 (very positive), based on sentiment about Elon

Musk.

- If unrelated to Elon Musk, return 0.

2. **TeslaScore**

- A score from -1 (very negative) to 1 (very positive), based on sentiment about Tesla.

- If unrelated to Tesla, return 0.

If a comment contains [a quotation mark], what comes after [quotation mark] and until the next line break is a quote from a previous comment which the current comment is responding to.

So you should use the part between [quotation mark] to understand the context which the comment is responding to, but don't attribute the words there to the actual comment.

Repeat this analysis twice independently for each comment.

Respond in this format:

ElonScore1 (first trial): <score or 0>

ElonScore2 (second trial): <score or 0>

TeslaScore1 (first trial): <score or 0>

TeslaScore2 (second trial): <score or 0>

Do not provide an explanation for the scores, just output the scores.

Scoring notes:

- Use near-zero values (e.g., -0.1 or 0.1) if sentiment is vague or weak.

- Do not guess strong sentiment if it's not clear.

To ensure that ChatGPT did not "hallucinate" scores and to enhance reproducibility, we asked ChatGPT to score each comment twice independently. We find that the vast majority of scores do not differ meaningfully across the two iterations: the score difference between the two iterations never exceeds 0.1.

To exclude noise from weak, unclear, or uncertain scores, we define positive comments as comments where the average score across the two iterations is greater than or equal to 0.3. Negative comments are correspondingly defined as comments where the average score is less than or equal to -0.3. To produce the results shown in Figure 2.2, we add up the total number of negative and positive comments in each month between January 2021 and December 2024. Figures B.9 and B.10 show that the sentiment patterns are robust to changing these cutoffs, and look similar if we set the cutoff score for positive/negative comments at $>0/<0$ or $\geq 0.5/\leq$-0.5 respectively.

B.3 Additional Tables and Figures

	California	Texas	USA
Population	39,431,263	31,290,831	340,110,988
% White	32.57%	37.78%	56.27%
% Black	5.09%	11.98%	11.75%
% Hispanic	40.75%	40.27%	20.00%
% Asian	15.83%	6.04%	6.19%
Median household income	$100,149	$79,721	$81,604
Poverty rate	12.0%	13.7%	12.5%
Population density (per sq. mi.)	253.1	119.8	96.3
% Rural	5.8%	17.0%	20.4%
Total EV registrations	1,256,646	230,125	3,555,445
% of EVs among all registered vehicles	3.41%	0.89%	1.24%
% of new vehicle sales that are EVs	25.71%	6.20%	9.5%
EV Charging Ports	49,433	8,637	176,032
EV Charging Stations	16,381	3,104	64,315

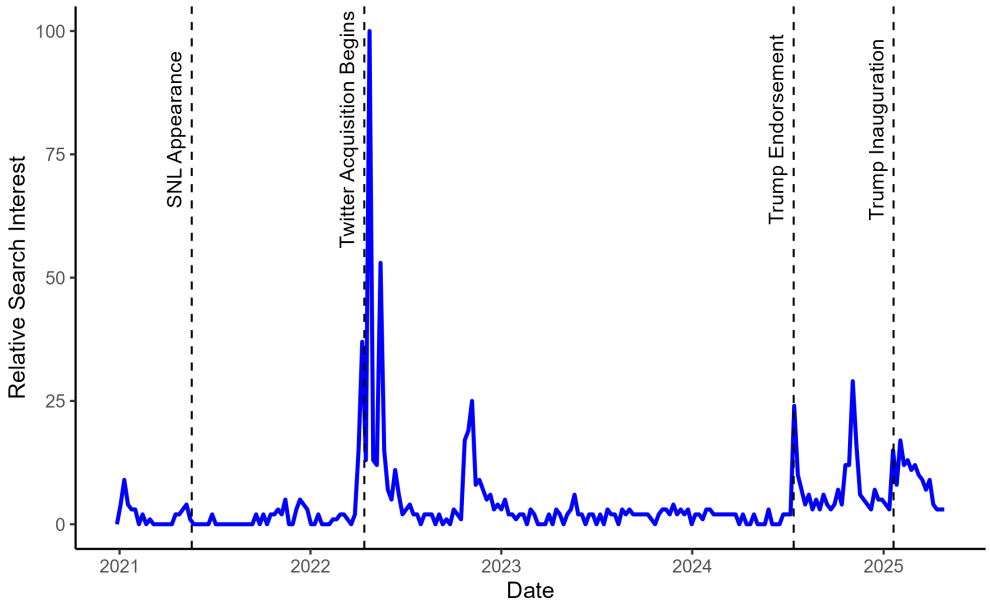
Table B.1: Descriptive Statistics for Texas, California, and the U.S.A.

Notes: This table contains a set of descriptive statistics for California, Texas, and the U.S. as a whole. Population estimates come from 2024 Census data. Population by race comes from 2024 American Community Survey (ACS) 1-year estimates, and includes only non-Hispanic people who identify as one race only as part of 'White', 'Black', and 'Asian' respectively. Median household income comes from the 2024 ACS 1-year estimates in 2024, in 2024 dollars. The percentage of a population estimated to live in a rural area and the poverty rate are based on ACS 1-year estimates for 2023. Land area for population density calculations comes from 2010 Census data. Data on electric vehicle registrations and charging infrastructure come from the U.S. Department of Energy data for 2023. Data on new vehicle sales come from the Alliance for Automotive Information Q4 2023 report, and reflects data for 2023 as a whole. New vehicle sales reflect Plug-in Hybrid EVs (PHEVs), Fuel Cell Electric Vehicles (FCEVs), and Battery Electric Vehicles (BEVs) as a proportion of the entire vehicle fleet, while EV registrations reflect BEVs only as a proportion of the entire vehicle fleet.

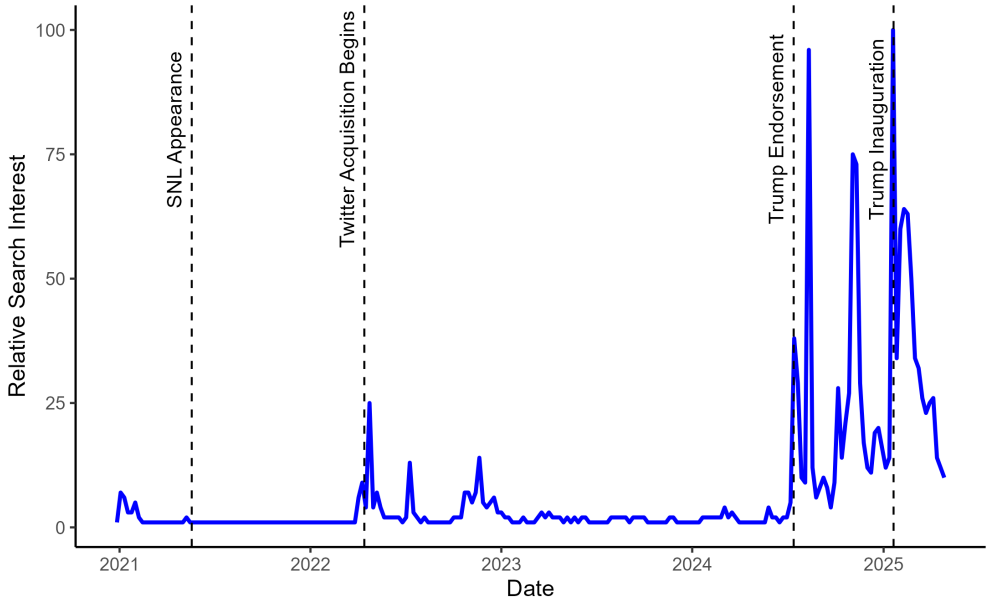
	(1)	(2)	(3)
Subsidy Amount	-0.621***	-0.663***	-0.622***
	(0.042)	(0.042)	(0.042)
MPG	-0.032***	-0.031***	-0.032***
	(0.001)	(0.001)	(0.001)
Horsepower	0.136***	0.136***	0.136***
	(0.000)	(0.000)	(0.000)
Class 3 Weight dummy	7.747***	7.806***	7.748***
	(0.038)	(0.038)	(0.038)
Class 1 Weight dummy	-6.875***	-6.820***	-6.871***
	(0.027)	(0.027)	(0.026)
Length (in.)	0.128***	0.127***	0.128***
	(0.002)	(0.002)	(0.002)
Observations	4,136,625	4,136,625	4,136,625
First-stage F-stat	220.56	249.61	220.85
Make x Quarter x Ideology FE	Yes	Yes	Yes
Make x Quarter x Income FE	Yes	Yes	Yes
County-Quarter FE	Yes	Yes	Yes
Charging Controls	No	Yes	Yes
Entry Controls	No	No	Yes

Table B.2: First Stage Results

Notes: Table depicts first stage estimates for price which is instrumented for with federal subsidy variation. Price and federal subsidy amount are both in thousands of USD. Column (1) depicts estimates allowing make-by-quarter year fixed effects to vary for ZIPs that voted less than 40% for Trump in 2020, more than 60% for Trump in 2020, and between 40-60% for Trump in 2020. Column (2) depicts estimates additionally including a month-year fixed effect. Finally, Column (3) additionally includes controls for ZIP-level income and tax filing rates. Tax filing rates data is taken from the IRS. Standard errors are clustered at the market level and depicted below estimates in parentheses. All specifications additionally control for vehicle body type which takes on one of ten categories and is thus omitted for readability.



Panel A: Searches for "Elon Musk Republican"



Panel B: Searches for "Elon Musk Trump"

Figure B.3: Google Trends Search Interest for Elon Musk's Politics

Notes: Figure depicts Google Trends data for the terms 'Elon Musk Republican' (Panel A) and "Elon Musk Trump" (Panel B). The four vertical dashed lines, from left to right, depict the dates Musk went on SNL, Musk initialized the Twitter acquisition, Musk endorsed Donald Trump for President, and Musk attended Trump's inauguration. The y-axis is normalized such that the maximum point is equal to 100 and all other points within the panel are relative to this maximum. This relative search frequency is determined separately for each panel.

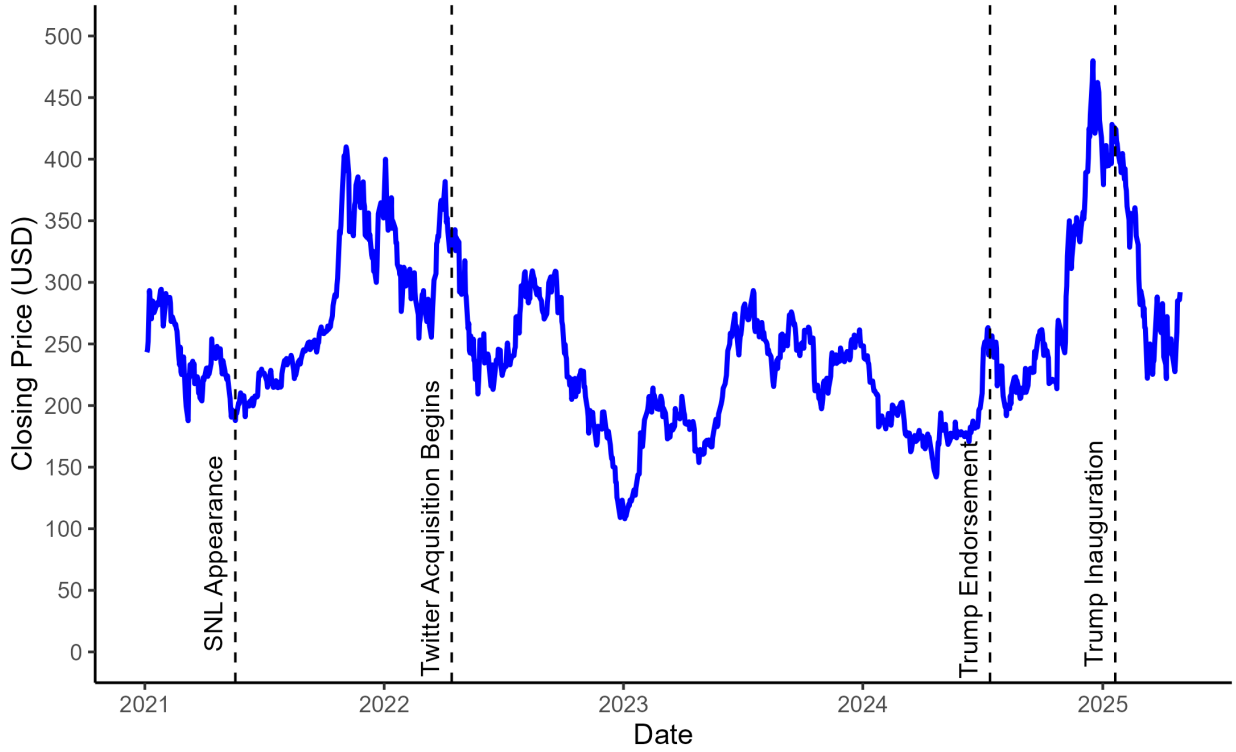


Figure B.4: Trends in Tesla Stock Price

Notes: Figure depicts Tesla's closing stock price each day from January 2021 through April 2025.

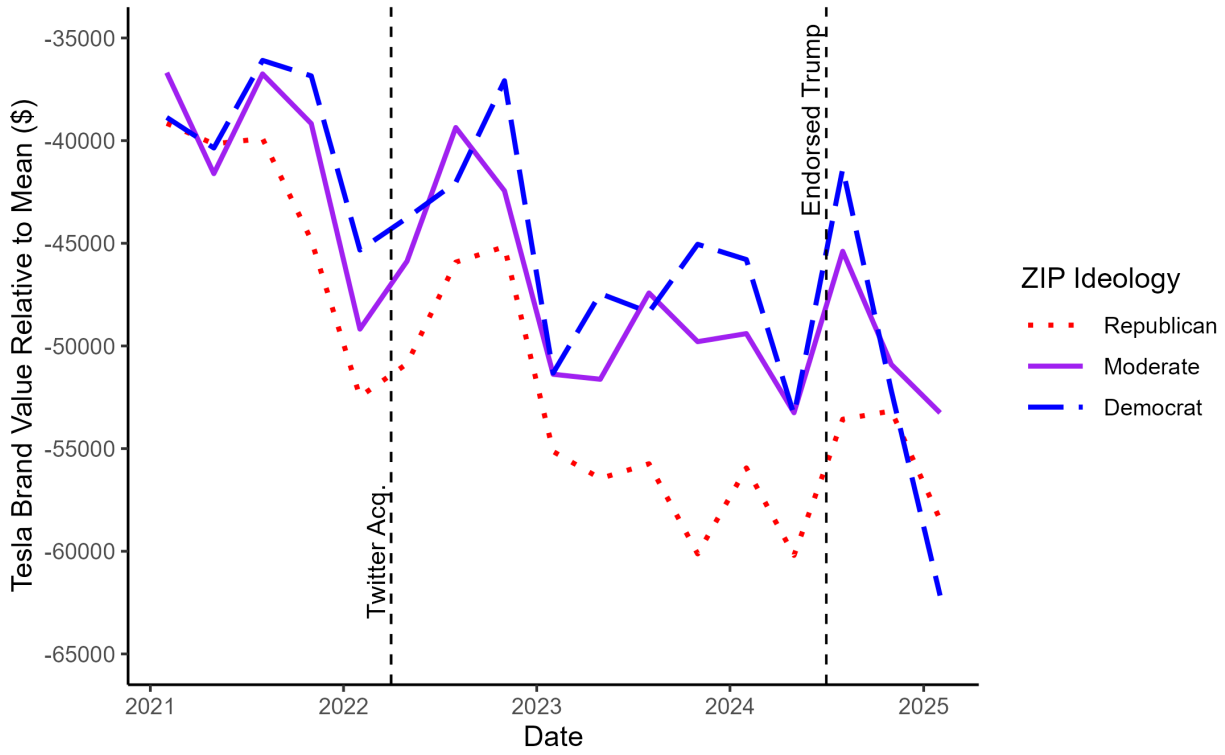


Figure B.5: Trends in WTP for Teslas Relative to the Average Brand

Notes: Figure depicts estimates of consumer WTP for a Tesla vehicle relative to the average brand. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Tesla's quarter-year specific fixed effects are then compared to the average of quarter-year fixed effects of all other brands for each political leaning. This is then converted to WTP for a brand by normalizing the difference in fixed effects by an estimated price coefficient. Price is instrumented for by variation in eligibility for federal subsidies generated by the 2022 Inflation Reduction Act. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).

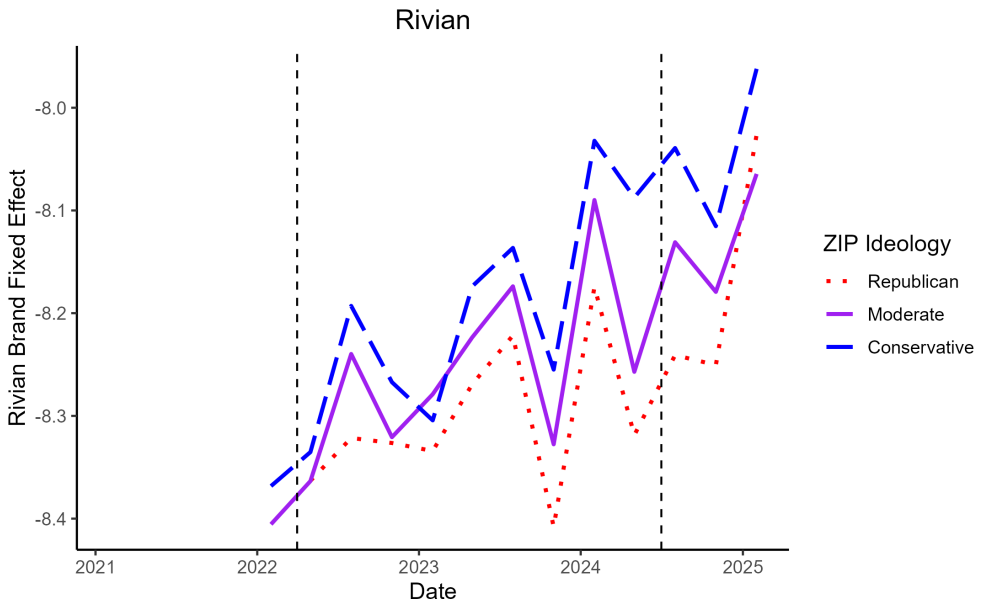


Figure B.6: Brand Fixed Effect for Rivian

Notes: Figure depicts estimates of fixed effects for the Rivian, the second most prominent manufacturer of solely EVs. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Fixed effects are all relative to the excluded group, which is the outside option of not purchasing a new vehicle. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).

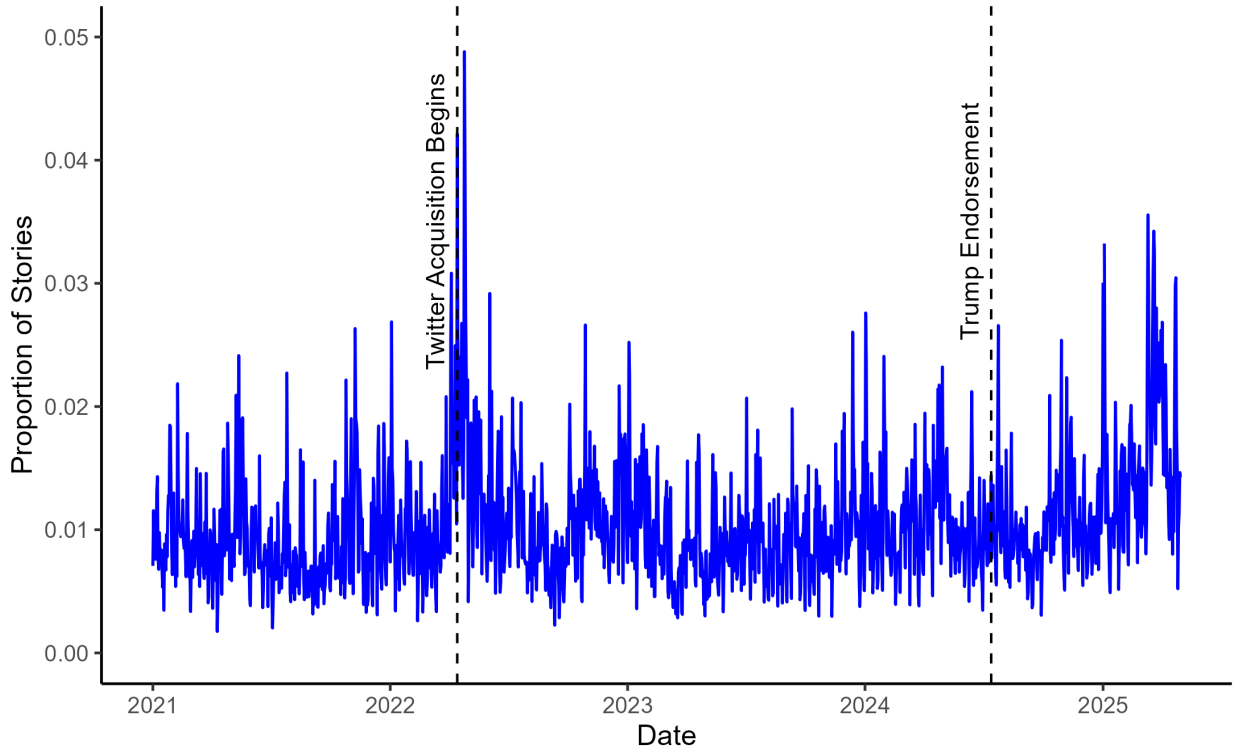
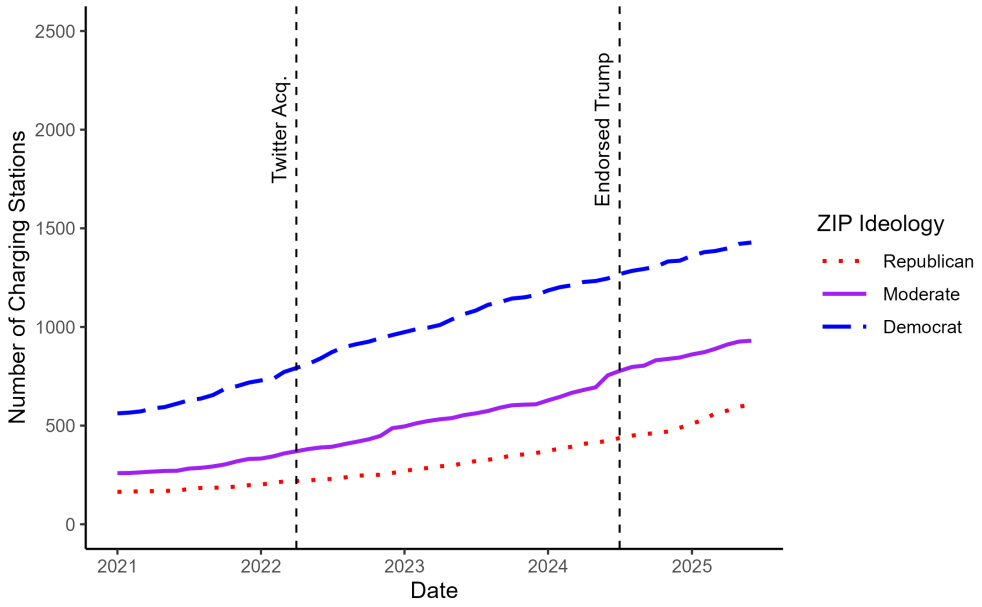
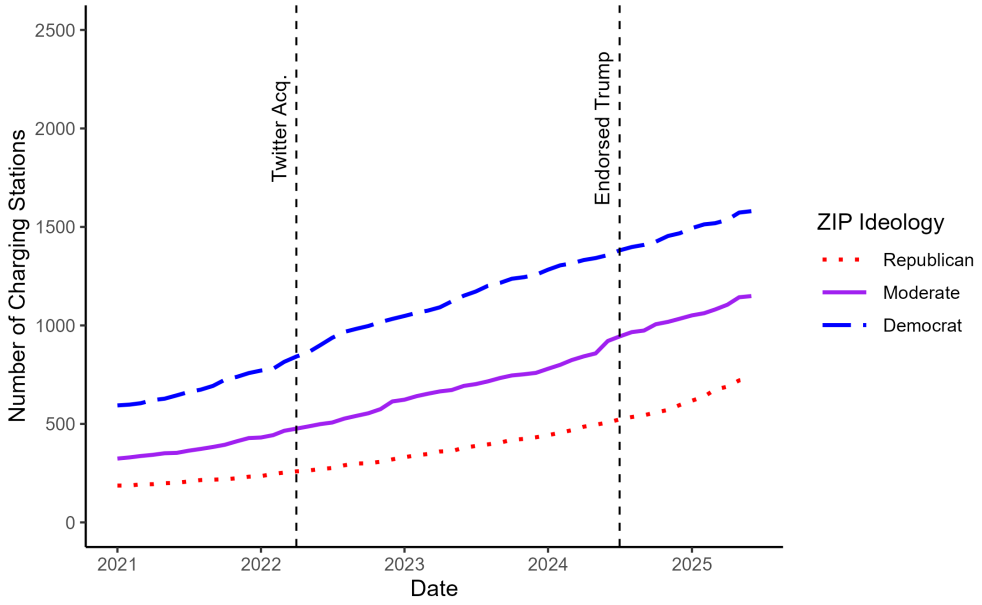


Figure B.7: Trends in News Stories Mentioning Tesla

Notes: Figure depicts Media Cloud data on the proportion of online news stories that mention the keyword "Tesla" over time. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).



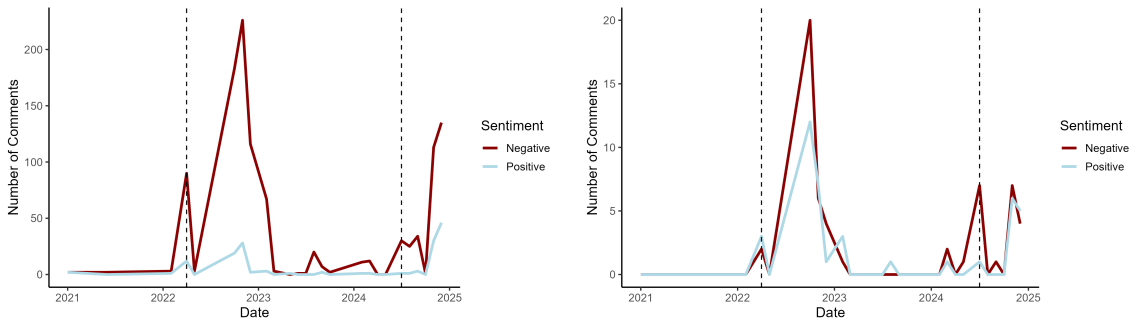
Panel A: Stations Able to Charge Tesla EVs



Panel B: Stations Able to Charge non-Tesla EVs

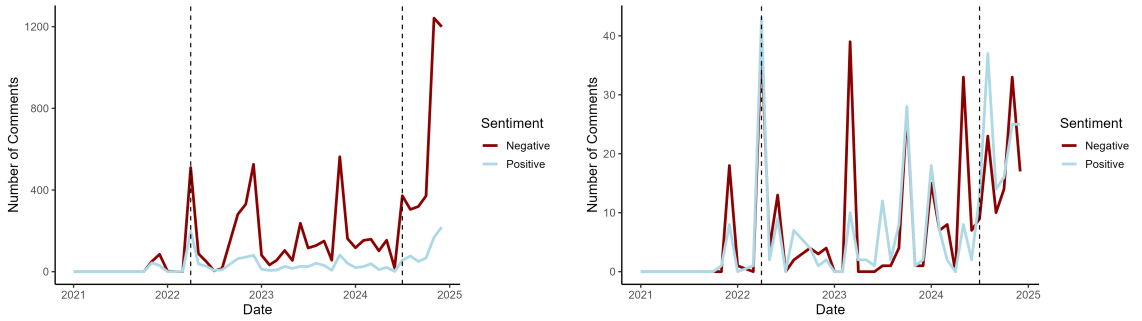
Figure B.8: Trends in Charging Station Availability

Notes: Figure depicts monthly total stock of charging stations which are compatible with Tesla EVs (Panel A) and with non-Tesla EVs (Panel B). Stations are presented separately by ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).



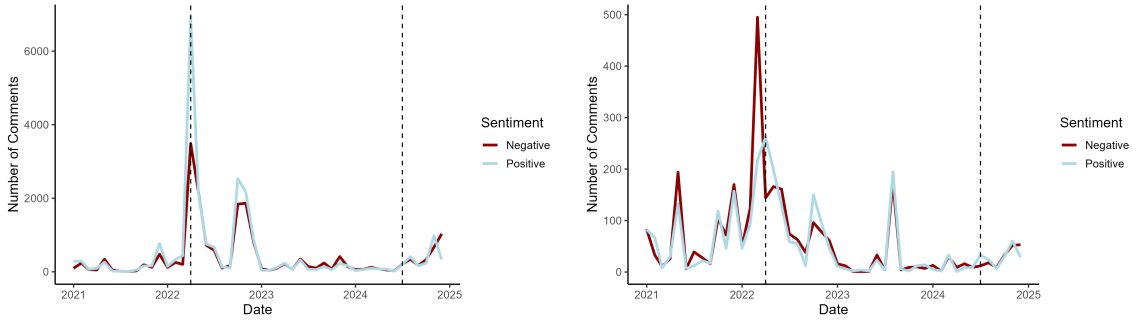
r/liberal: Elon Musk Sentiment

r/liberal: Tesla Sentiment



r/centrist: Elon Musk Sentiment

r/centrist: Tesla Sentiment

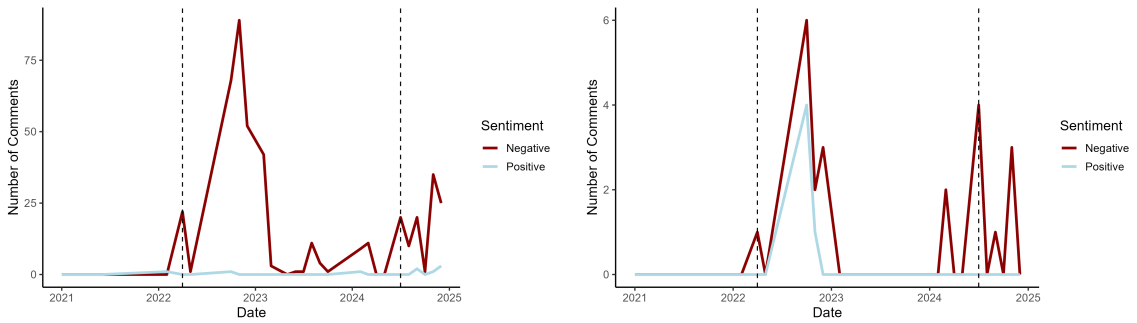


r/conservative: Elon Musk Sentiment

r/conservative: Tesla Sentiment

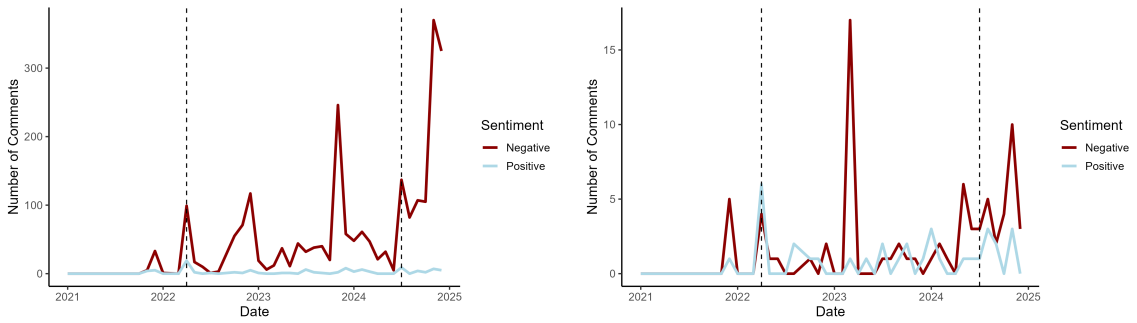
Figure B.9: Sentiment Analysis - Alternative Definition of Positive and Negative Comments

Notes: This figure depicts the sentiment analysis results from the subreddits r/liberal, r/centrist, and r/conservative, using an alternative cutoff to define positive and negative comments. Positive comments are defined as comments with a score above 0, and negative comments are defined as comments with a score below 0. The figures on the left depict sentiment about Elon Musk, while the figures on the right depict sentiment about Tesla. In each subfigure, the first dashed vertical line corresponds to the month immediately prior to the month of Musk's Twitter acquisition, and the second dashed vertical line corresponds to the month immediately prior to the month of Musk's endorsement of Trump.



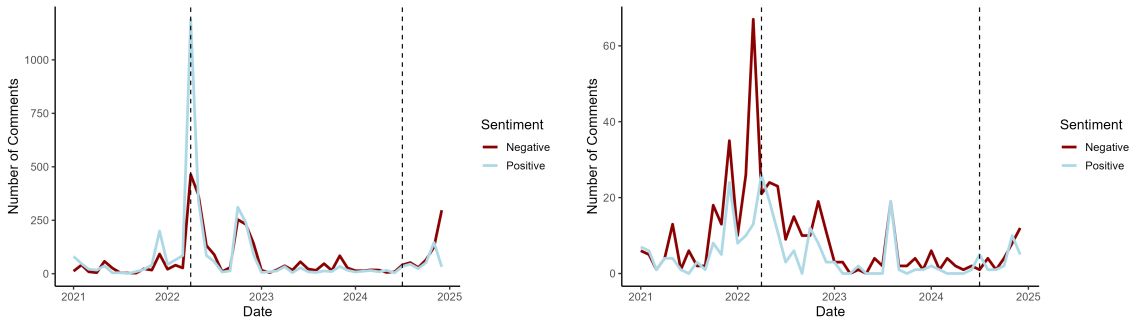
r/liberal: Elon Musk Sentiment

r/liberal: Tesla Sentiment



r/centrist: Elon Musk Sentiment

r/centrist: Tesla Sentiment

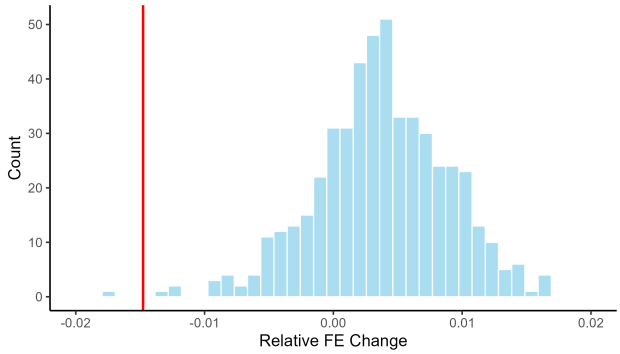


r/conservative: Elon Musk Sentiment

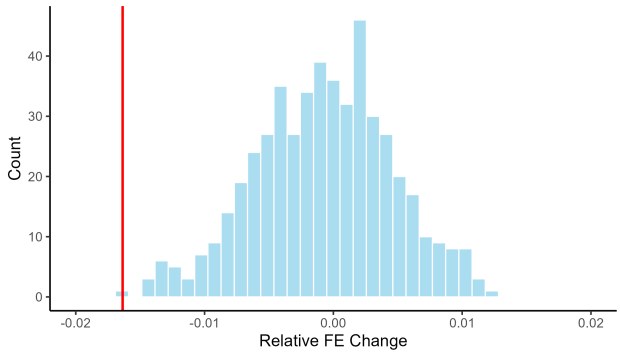
r/conservative: Tesla Sentiment

Figure B.10: Sentiment Analysis - Alternative Definition of Positive and Negative Comments

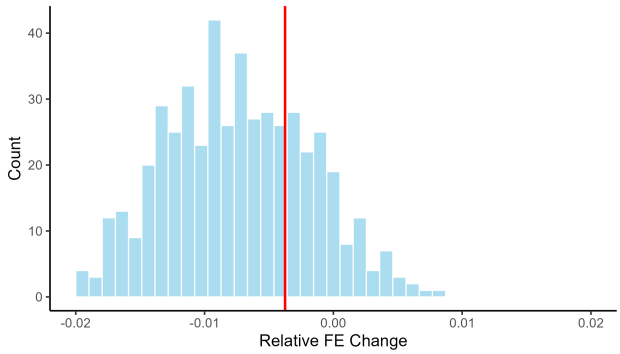
Notes: This figure depicts the sentiment analysis results from the subreddits r/liberal, r/centrist, and r/conservative, using an alternative cutoff to define positive and negative comments. Positive comments are defined as comments with a score greater than or equal to 0.5, and negative comments are defined as comments with a score less than or equal to -0.5. The figures on the left depict sentiment about Elon Musk, while the figures on the right depict sentiment about Tesla. In each subfigure, the first dashed vertical line corresponds to the month immediately prior to the month of Musk's Twitter acquisition, and the second dashed vertical line corresponds to the month immediately prior to the month of Musk's endorsement of Trump.



Panel A: Fall in Democrat Fixed Effects



Panel B: Fall in Moderate Fixed Effects



Panel C: Fall in Republican Fixed Effects

Figure B.11: Permutation Test on Estimated Tesla WTP Changes

Notes: Figure depicts permuted estimates of the 2024 falls in Democrat fixed effects (Panel A), moderate fixed effects (Panel B), and Republican fixed effects (Panel C). Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. Tesla's quarter-year specific fixed effects are then compared to the median quarter-year fixed effects of all other brands for each political leaning. This regression is run separately for permutations that permute logit market shares within each ZIP-month-year and recalculate the resulting within-nest shares.

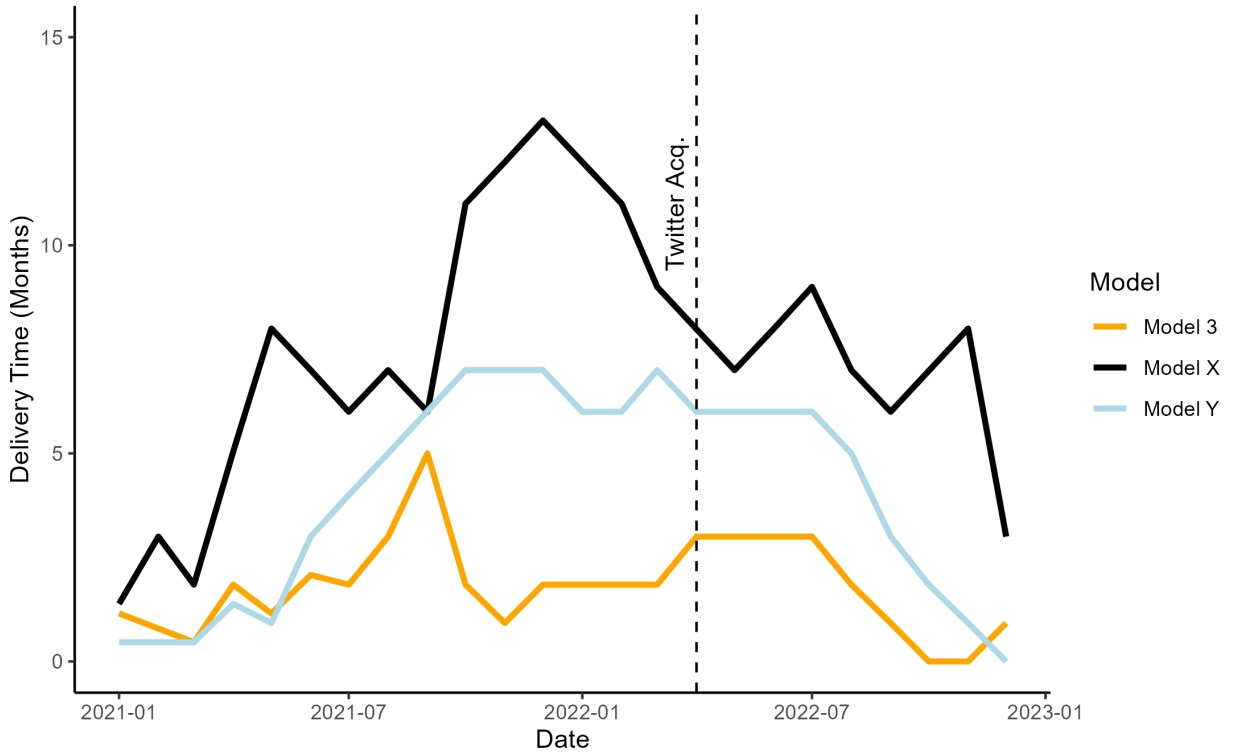


Figure B.12: Tesla Delivery Times

Notes: Figure depicts estimated delivery times for three Tesla models from January 2021 to December 2022. Data is taken from Wayback Machine snapshots of Tesla.com. Data is not consistently available on delivery times from January 2023 on.

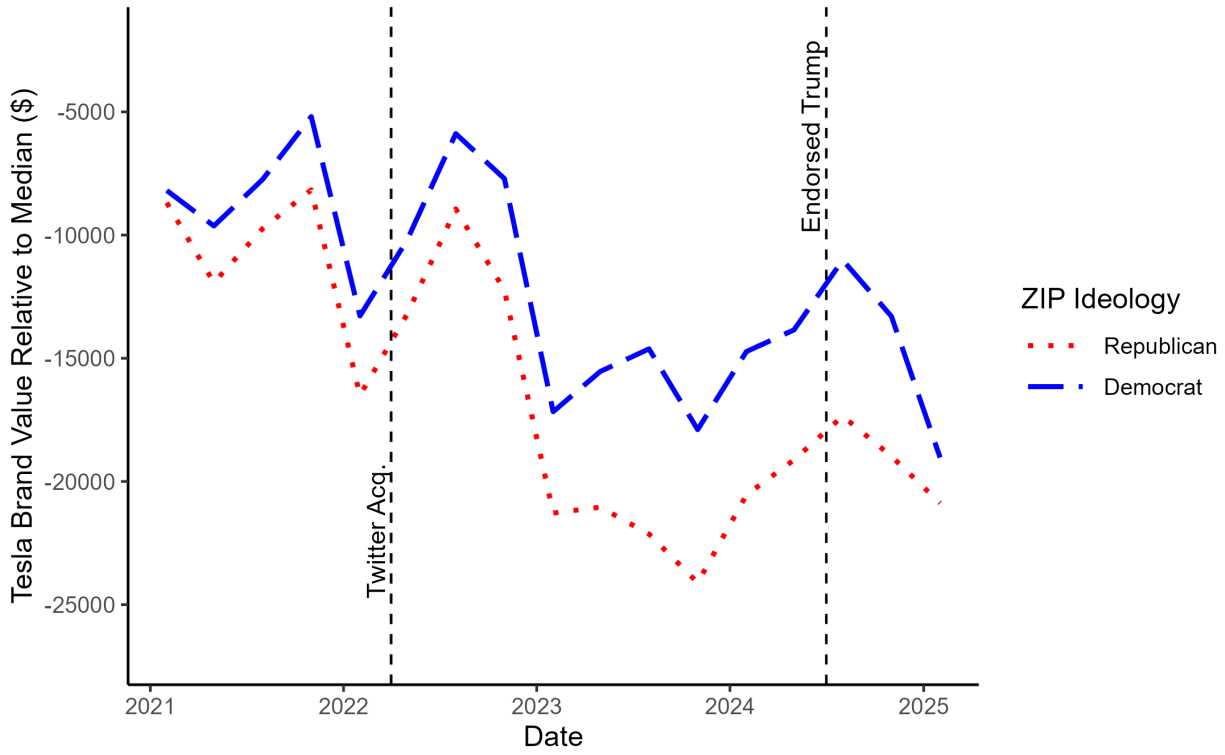
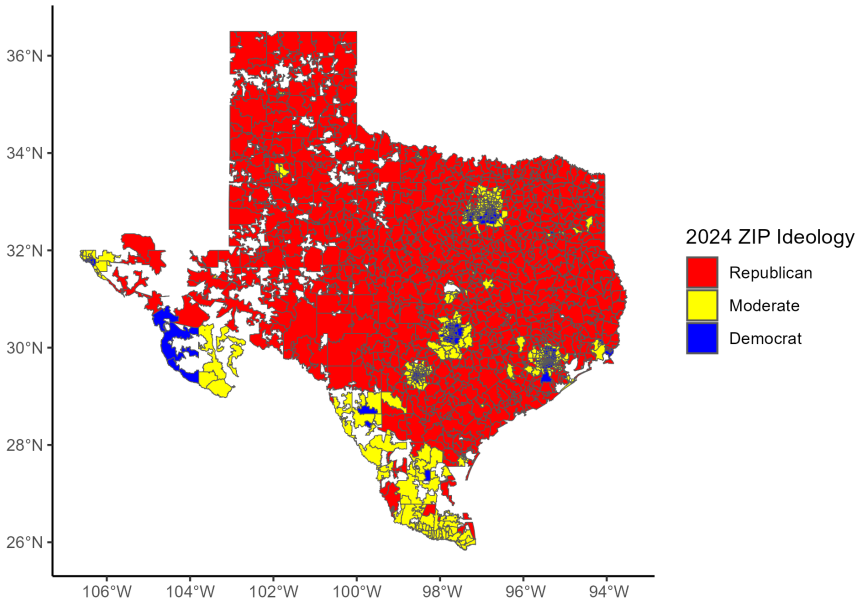
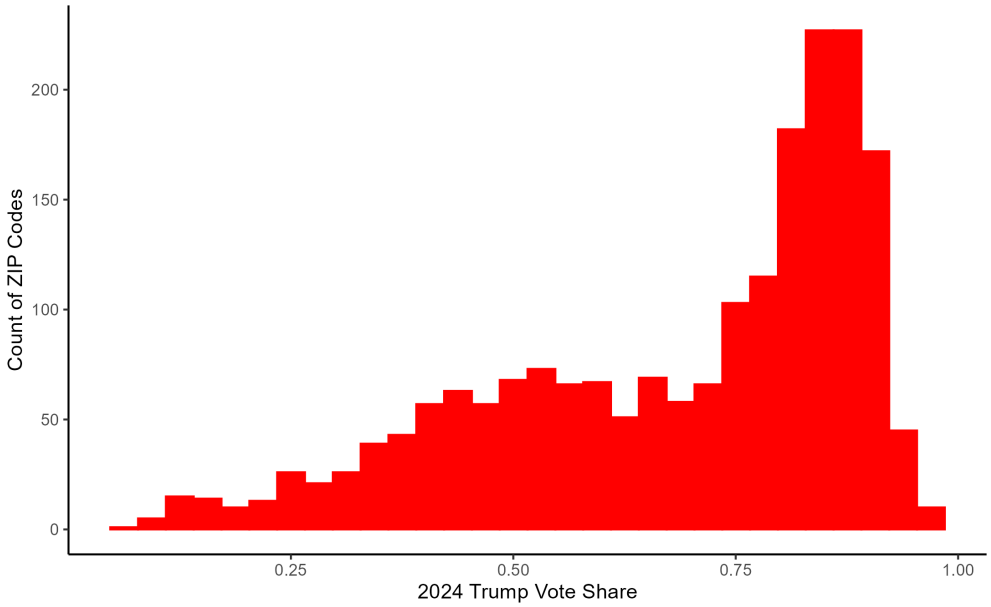


Figure B.13: Trends in WTP for Teslas Relative to the Median Brand with Two Political Classifications

Notes: Figure depicts estimates of consumer WTP for a Tesla vehicle relative to the median brand. Estimates are generated by a logit demand estimation exercise which includes brand-quarter-year fixed effects interacted with ZIP political leanings. ZIP political leanings are calculated using cutoffs of below 50% (Democrat) and above 50% (Republican) voting for Trump in the 2020 election. Tesla's quarter-year specific fixed effects are then compared to the median of quarter-year fixed effects of all other brands for each political leaning. This is then converted to WTP for a brand by normalizing the difference in fixed effects by an estimated price coefficient. Price is instrumented for by variation in eligibility for federal subsidies generated by the 2022 Inflation Reduction Act. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).



Panel A: Spatial Distribution of Trump 2024 Vote Shares



Panel B: Histogram of ZIP Code Vote Shares

Figure B.14: Texan Vote Shares for Donald Trump in 2024

Notes: Figure depicts a map of vote shares for Donald Trump in 2024 (Panel A) and a corresponding ZIP-code level histogram of vote shares (Panel B). Missing values on vote shares are omitted in the map (depicted as empty white space).

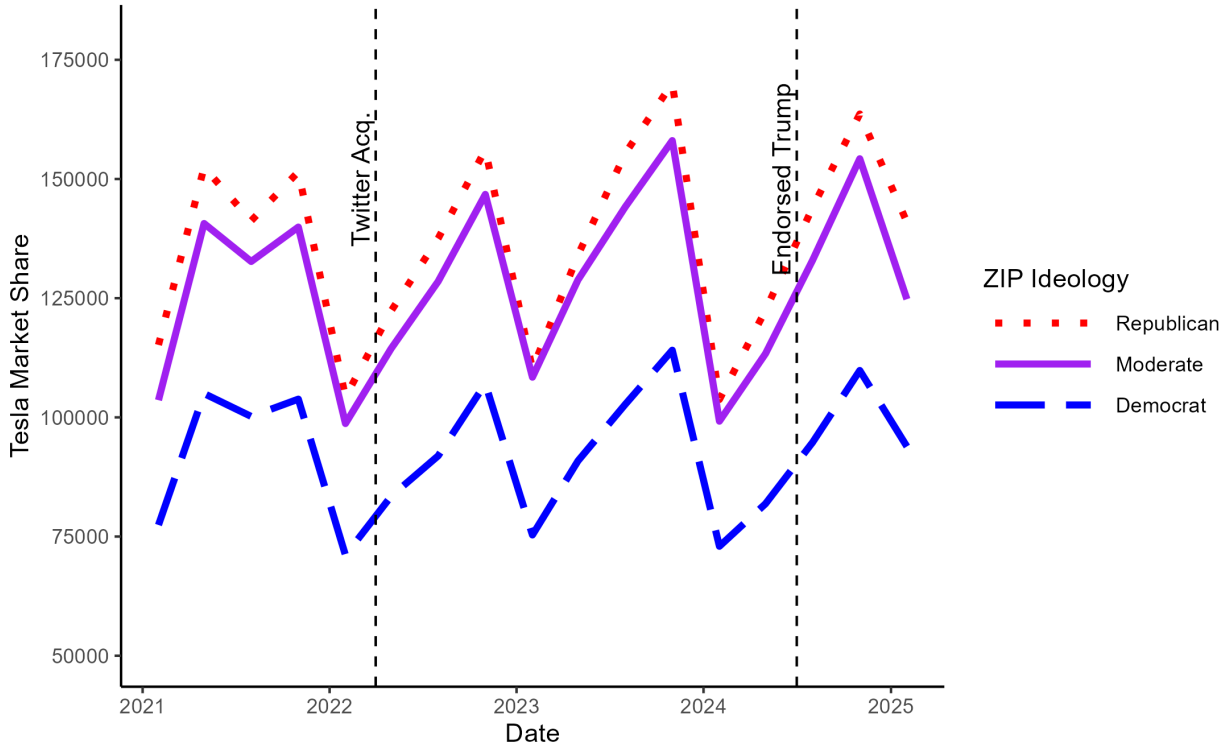
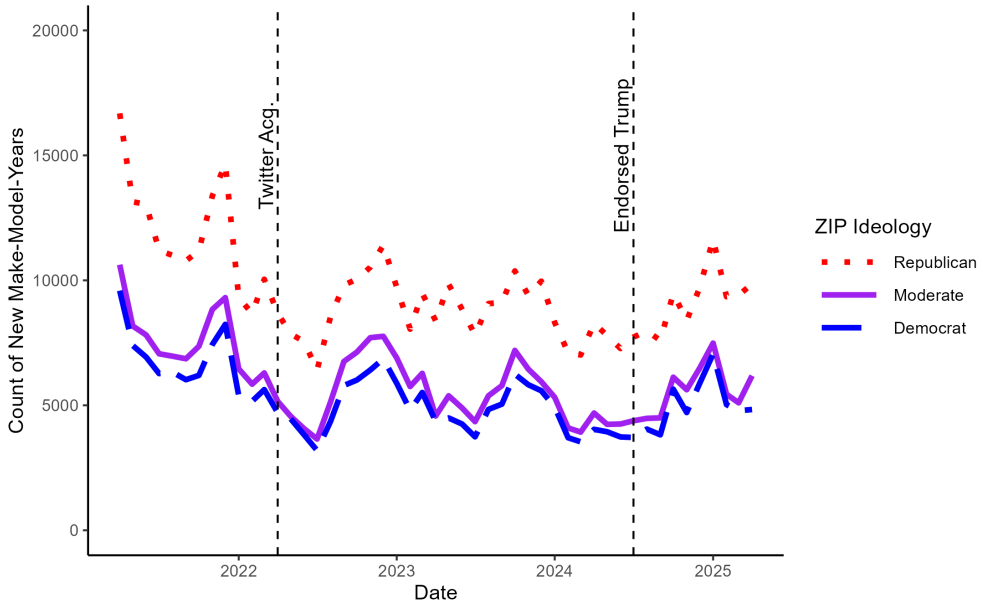
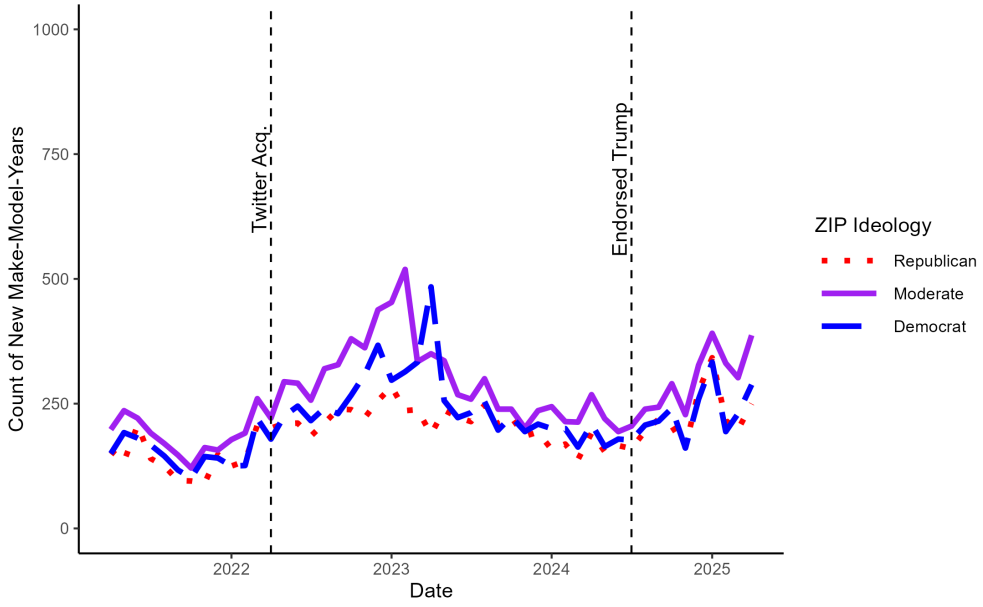


Figure B.15: Trends in Total Vehicles

Notes: Figure depicts trends in total vehicles being registered by ZIP Ideology. ZIP ideology is calculated using cutoffs of below 40% (Democrat) and above 60% (Republican) voting for Trump in the 2020 election. Moderate ZIPs are all ZIPs between these two cutoffs in terms of their 2020 Trump vote share. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right).



Panel A: All New Vehicles Entering Markets



Panel B: New EVs Entering Markets

Figure B.16: Trends in Competitor Entry Into Vehicle Markets

Notes: Figure depicts monthly total counts of vehicle make-model-years that have previously never been registered in a ZIP code. These counts are calculated at the ZIP code level and then aggregated based on ZIP political leanings. Black vertical lines depict Elon Musk initializing the Twitter acquisition (left) and initiating involvement in the 2024 Presidential Election cycle (right). Panel A depicts entries for all vehicles. Panel B depicts entries only for EVs.

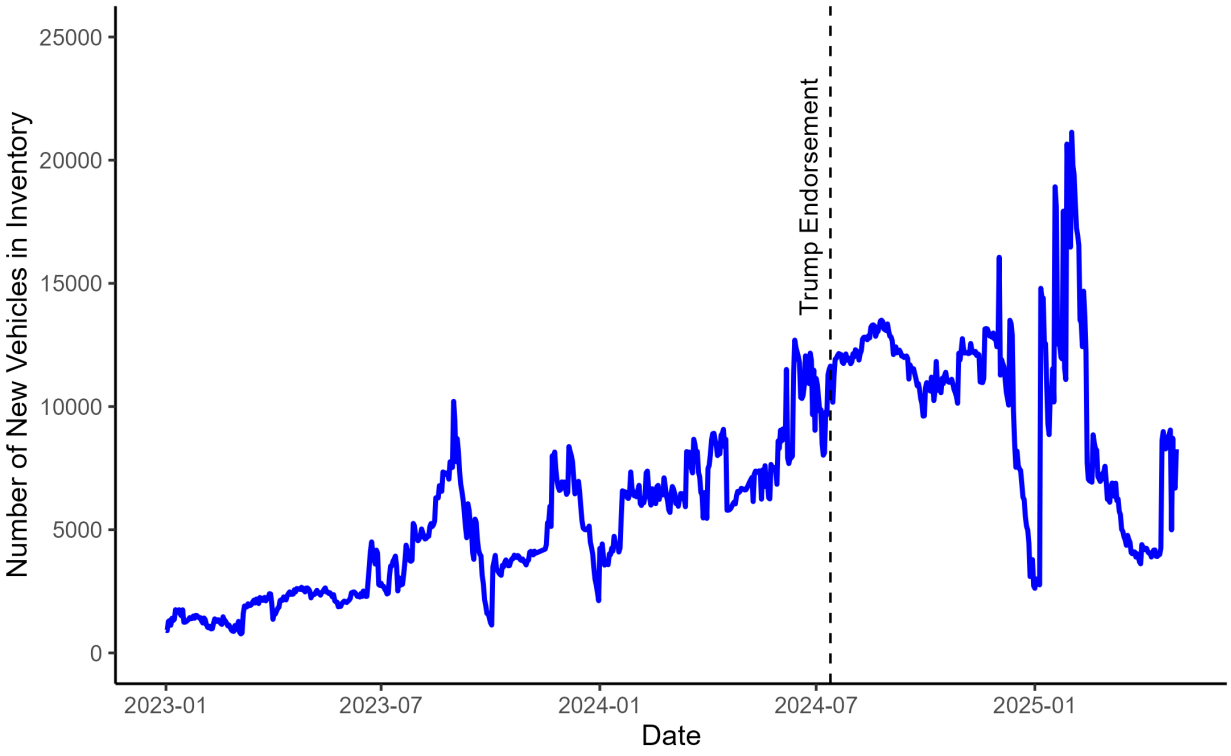


Figure B.17: Supply Chain Constraints in Tesla Deliveries Over Time

Notes: Figure depicts estimates of new vehicles in inventory for Tesla over time. Figure aggregates across the Tesla Model Y, X, S, and 3. Estimates are of inventories covering the entire U.S.A. Estimates are generated by Tesla-Info.com [2025] who scrape listed vehicles from Tesla's website daily and post data up to one year in the past. Historical data is collected using Internet Archive [2025].

Nest	Price	MPG	Horsepower	Length (in.)	Diesel	EV	Class 1 Wt.	Class 3 Wt.
1	73248	14.59	391.28	235.62	0.46	0.00	0.02	0.40
2	65373	18.07	351.02	215.19	0.11	0.00	0.08	0.04
3	63468	19.34	349.40	218.58	0.03	0.01	0.21	0.01
4	54348	20.78	343.40	207.55	0.01	0.01	0.39	0.00
5	49471	22.21	296.63	197.60	0.01	0.01	0.69	0.00
6	43192	24.40	256.32	187.71	0.00	0.01	0.79	0.00
7	35712	27.57	206.12	182.13	0.00	0.01	0.96	0.00
8	32416	30.32	196.68	183.60	0.00	0.01	0.97	0.00
9	31187	33.31	188.94	183.22	0.00	0.01	0.97	0.00
10	41718	63.27	281.67	185.43	0.00	0.34	0.90	0.00

Table B.3: Summary Statistics of Nests

Notes: Table depicts separately by logit nest the summary statistics of newly registered passenger vehicles in Texas from January 2021 to April 2025. Tesla vehicles are included in nest 10.

APPENDIX C

APPENDIX FOR CHAPTER 3

C.1 Discussion of Missing Data Imputation

Because a large share of listing characteristics (e.g. number of bathrooms, number of balconies, building height, etc.) are optional for listing on Housing.com, we observe a non-trivial number of missing data. Fortunately, we observe a large number of characteristics for every single observation. In particular, we can observe presence in a gated community, furnishing level, number of bedrooms, listing type (e.g. “flat" or “independent house"), listing date, and listing ward for every single listing. As such, we can use these characteristics to impute the missing characteristics of each row.

We do so using predictive mean matching (PMM) as in [Rubin, 1986, Little, 1988]. Consider a variable x which contains missing values in some observations and a set of variables z which are observed in each observation. The PMM algorithm works as follows [Allison, Paul, 2015]:

1. With observations that are not missing x, run a regression of x on z, producing estimated coefficients $\hat{\beta}$
2. Make some random draw of coefficients $\hat{\beta}'$ based off the estimated distribution of $\hat{\beta}$
3. Use $\hat{\beta}'$ to generate predicted values of x for all observations, including those that have values of x
4. For each observation that has missing x, identify observations with observed x whose predicted values of x are closest to the predicted value of x for the observation. Define these as ‘close cases.’
5. Randomly assign one observed value of x from the close cases as the missing value of x

6. Repeat 2-5 for all completed data

Notably, the choice of how many close cases are considered in the above algorithm is important for generating valid imputation. We choose to consider the 5 closest cases in our main results. In Appendix Figure C.1, we depict results that instead use 1 case and 10 cases. As can be seen, the results are largely unchanged from our main result, suggesting that our findings are robust to imputing under different specifications for PMM.

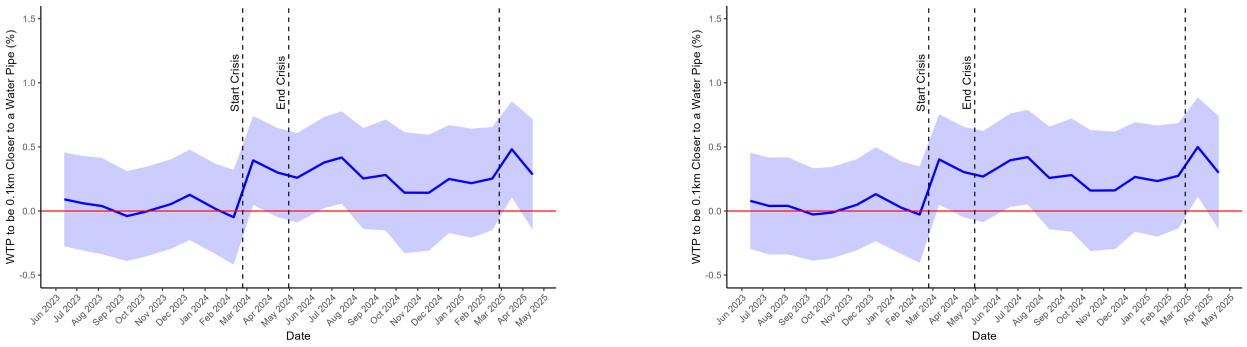


Figure C.1: Hedonic Estimates with $m = 1$ and $m = 10$ PMM Imputation

Notes: Figure depicts estimated rent premium for rental listings being 0.1km closer to the nearest pipe in each month of our data. The left panel does so relying on data imputation from predictive mean matching (PMM) with $m = 1$ closest candidates. The right panel does so relying on data imputation from PMM with $m = 10$ closest candidates. Estimates are generated by running regressions of the log of monthly rent prices on the distance (km) from the nearest water pipe, interacted by month. Regressions control for a rich set of listing-level characteristics and include ward and year-month fixed effects. Standard errors are clustered at the ward and year-month levels. The shaded blue area represents a 95% confidence interval, whereas the solid blue line represents point estimates. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

C.2 Additional Survey Findings

C.2.1 Information Frictions on Water Sources

We find that residents in Bangalore are unaware in many cases where their water is sourced from. We measure in our survey that 69% of people were at most only somewhat aware of where a building sources its water from before the crisis, a percentage which is largely unchanged at 70% post-crisis. This mirrors the fact that many renters do not directly interact with BWSSB, nor with individuals who manage their supplying borewells. It further mirrors that listing sites such as Housing.com do not explicitly advertise the sources of water for each listing.

We show that individuals who live far from the piped water network incorrectly believe that they have a piped water connection. In Figure C.2, we plot locally estimated regression lines which link the probability that a respondent reported having a piped water connection to their distance from the pipe network. Panel A does so including the entire sample, whereas Panel B omits individuals who are less aware of their pipe connection status. We estimate that 40% of individuals living more than 0.5kms away from the nearest pipe report having a piped water connection, falling to 0% when omitting those who are unsure.

Because individuals incorrectly believe that areas far from the pipe network have pipe connections, our estimates of the change to WTP for pipes following the crisis are likely attenuated relative to what they would be under full information. This motivates future analysis which considers removing these informational barriers to understand how consumers respond under better information sets.

C.2.2 Risk Perceptions Pre- and Post-Crisis

We find that residents in Bangalore had perceptions of drought risk which increased as a result of the water crisis. We present the results of belief elicitation using Likert scales

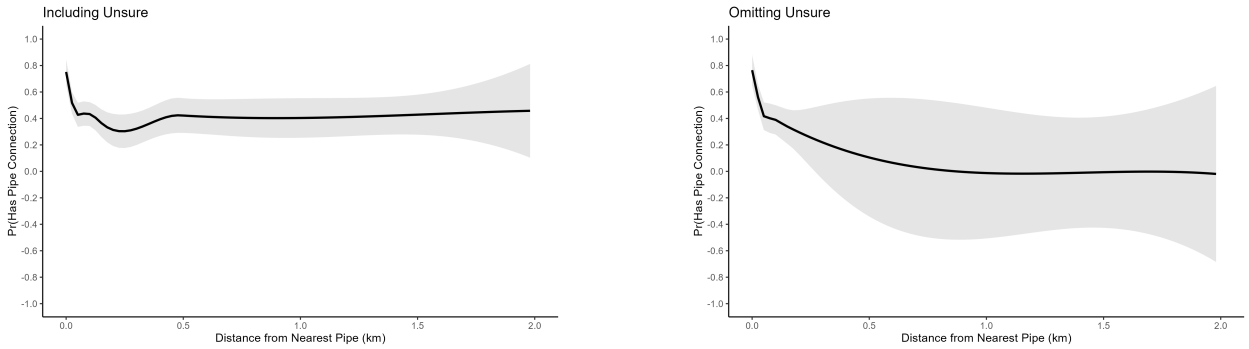


Figure C.2: Self-reported Pipe Access by Distance to Nearest Pipe

Notes: Figure depicts locally estimated regression lines of probability that a respondent reports having pipes by their distance from the nearest pipe. The left figure does so with the whole sample, whereas the right figure does so omitting those who report not being sure about pipe access in their area.

in Figure C.3, with 5 indicating high perceived risk and 1 indicating low perceived risk. Despite likely recall bias in measuring pre-crisis risk perceptions, we observe that perceived risks are higher for a drought to occur within 5 years of the survey date. This is particularly pronounced at the tails – fewer individuals believe a crisis is very unlikely and more believe that a crisis is very likely. These results suggest that long-term perceived drought risk has increased in Bangalore, which we have argued is tied to our estimates of WTP for water pipes.

C.2.3 Micro-level Migration Patterns

We observe in our survey data that individuals who migrated within Bangalore after the 2024 water crisis tended to move closer to water pipes. We present pre- and post-crisis distances from pipes in Figure C.4. A mass of individuals migrate from being 0.1-0.5km from the nearest pipe to being less than 0.1km. This has resulted in the first quartile of distance from nearest pipe decreasing by roughly 40% and the median distance falling by 22%. This lends support to our estimates that WTP to be nearer to the piped water network increased

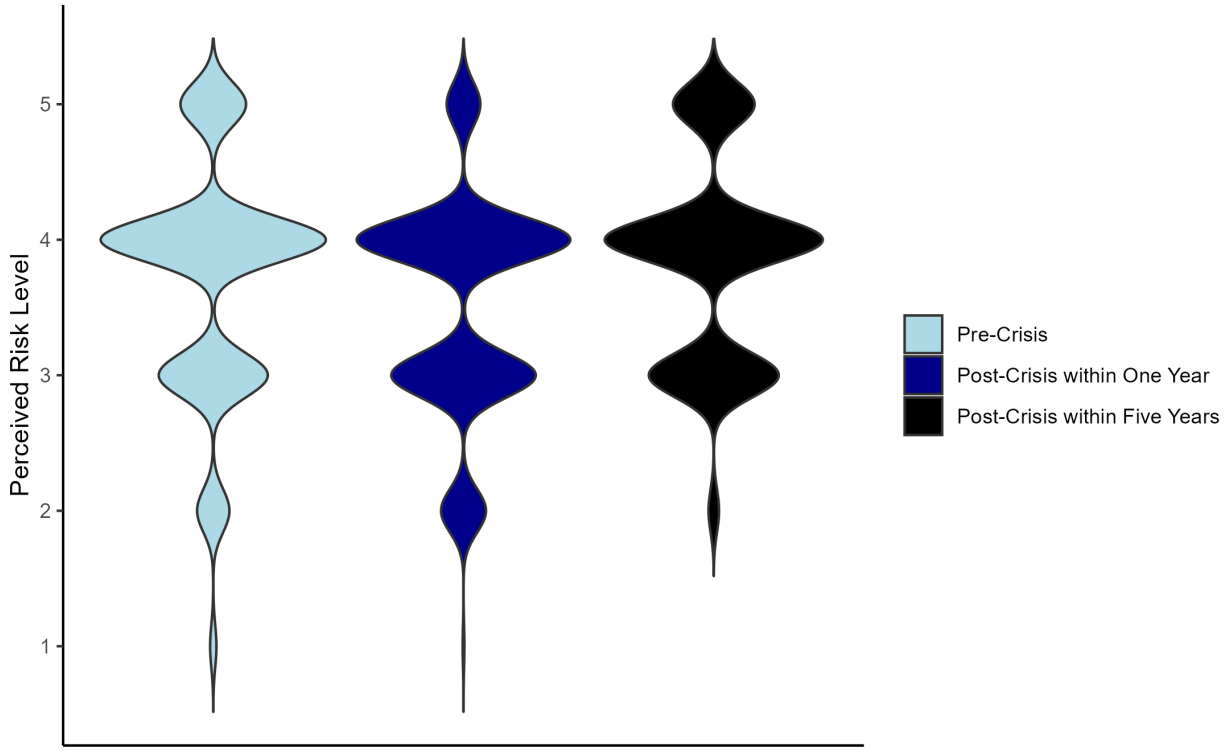


Figure C.3: Perceptions of Drought Risk

Notes: Figure depicts violin plots of Likert scales measuring perceived drought risks, where 5 indicates high perceived risk and 1 indicates low perceived risk. The leftmost violin asks individuals to recall their perceptions of risk prior to the crisis. The middle violin asks individuals about their current perceptions of a crisis happening within one year. The rightmost violin asks individuals about their current perceptions of a crisis happening within five years.

following the crisis.

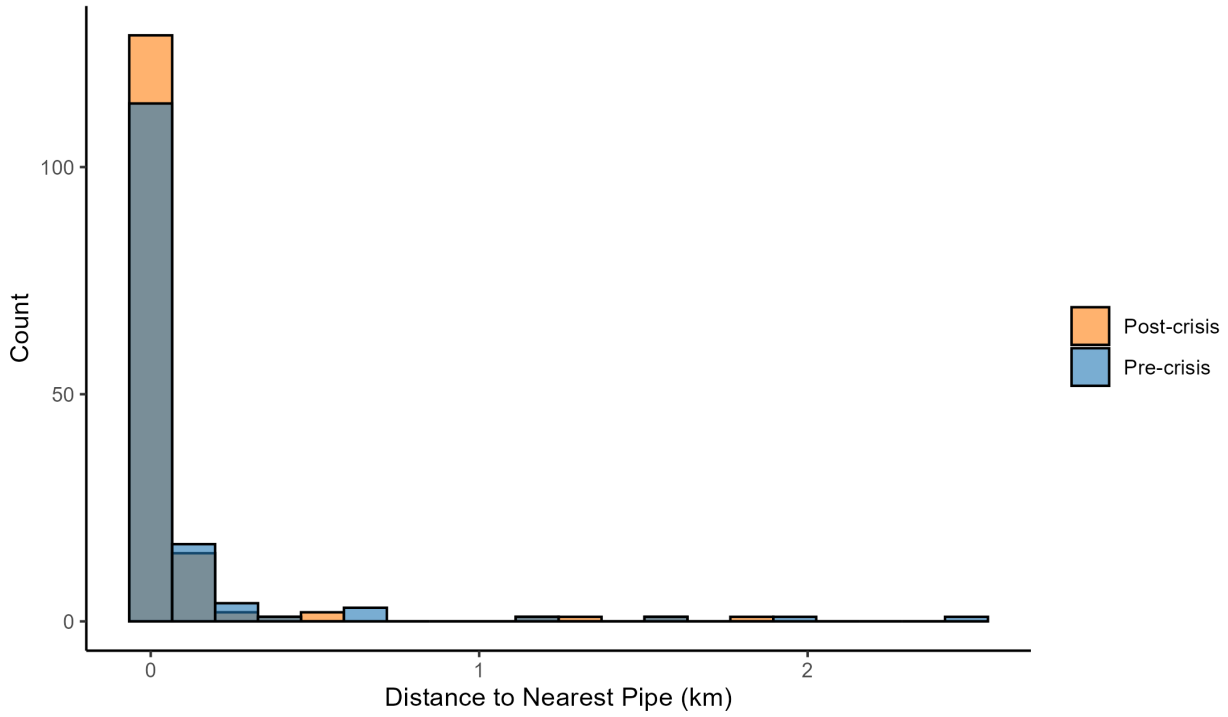


Figure C.4: Surveyed Distance from Nearest Pipe

Notes: Figure depicts distance from the nearest pipe for individuals who migrated within Bangalore between January 2024 and August 2025. Figure does so for initial location (blue) and for destination location (orange). Overlap in the two distributions is depicted in gray. Distance to nearest pipe is calculated by capturing coordinates of locations either via Google Maps pins or Google Maps geolocating of provided addresses.

C.3 Additional Tables and Figures

Panel A: Utility Piped Water Impacted		
City	**Dates**	**Piped Water Access (%)**
Barcelona, Spain	Mid 2008, Mid 2023	>95%
Cape Town, South Africa	Mid 2017, Mid 2018	80-85%
Chennai, India	Mid 2019	40-50%
Karachi, Pakistan	Early 2022, Early 2025	90-95%
São Paulo, Brazil	Late 2014, Late 2015	85%
San Antonio, Texas	Early 2025	>95%
Sana'a, Yemen	Early 2024, Early 2025	50-55%
Tehran, Iran	Early 2025	80-90%
Panel B: Only Private Water Supply Impacted		
City	**Dates**	**Piped Water Access (%)**
Bangalore, India	Early 2024	50-60%
Jakarta, Indonesia	Late 2023	60-65%
Lima, Peru	Early 2025	90-95%
Mumbai, India	Early 2025	40-45%
Nairobi, Kenya	Late 2019	50-60%
New Delhi, India	Mid 2019	90%

Table C.1: Summary of Other Water Crises

Notes: Table depicts details on an array of water crises which have occurred across the world. Details are collected from a variety of sources, including Wikipedia pages, utility websites, NGO documents, and academic articles.

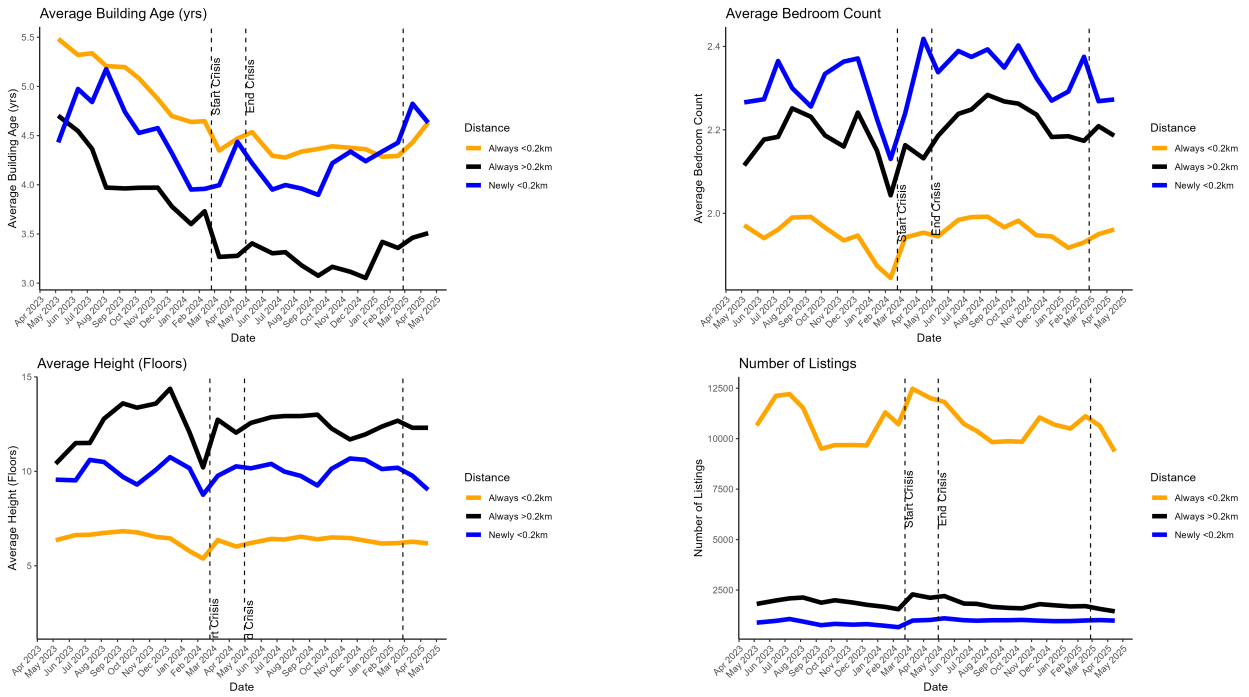


Figure C.5: Trends in Listing Characteristics and Count by Distance

Notes: This figure depicts the average building age (top left), average bedroom count (top right), average building floor count (bottom left), and total number of listings (bottom right) by month. Red lines are for listings less than 0.2km from the nearest water valve and blue lines are for listings at least 0.2km away. The orange trend line is for listings less than 0.2km away from the nearest water pipe at all times before and after the September 2024 pipe expansion. The black trend line is for listings always at least 0.2km away from the nearest pipe both before and after the 2024 expansion. The blue line is for listings that were originally more than 0.2km away from the nearest pipe but moved within 0.2km of the nearest pipe due to the 2024 expansion. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

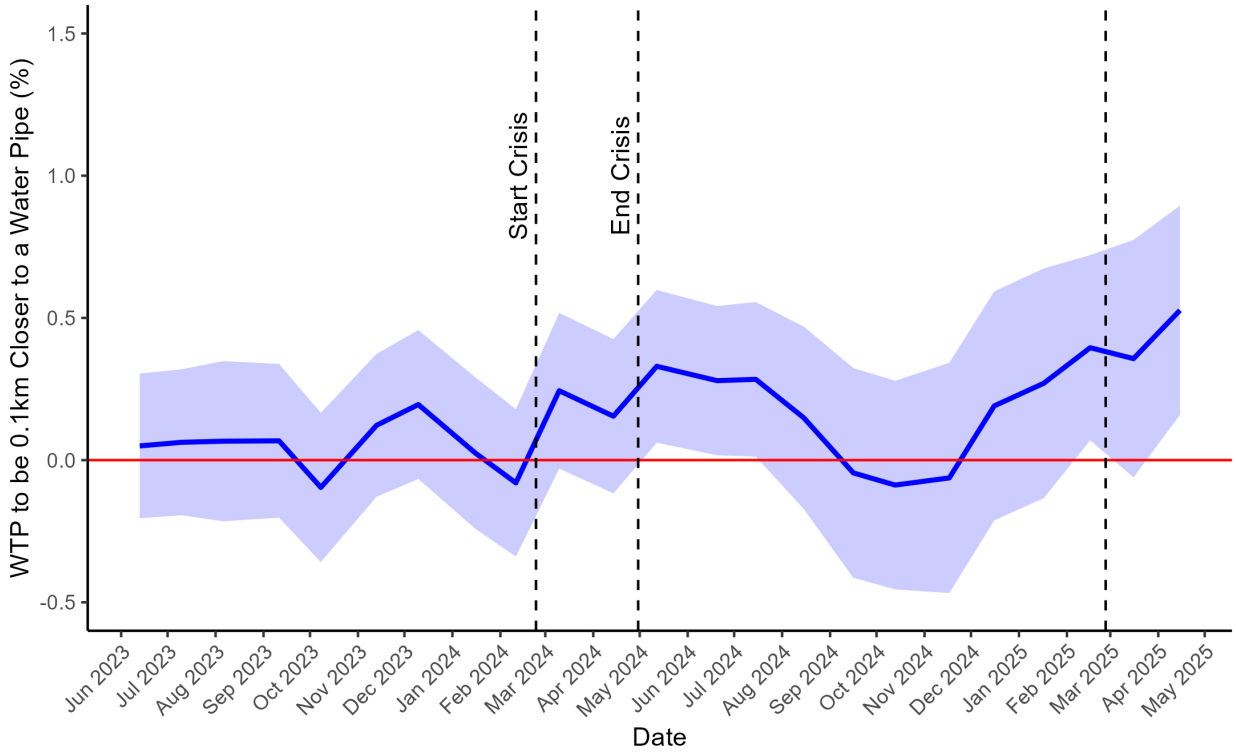


Figure C.6: The Impact of an Urban Water Crisis on Housing Rents: Building-level Panel

Notes: Figure depicts estimated rent premium for rental listings being 0.1km closer to the nearest pipe in each month of our data, where data is aggregated to a building by month panel. Estimates are generated by running regressions of the log of monthly rent prices on the distance (km) from the nearest water pipe, interacted by month. Regressions control for a rich set of listing-level characteristics and include ward and year-month fixed effects. All characteristics average across a building's listings within that month. Standard errors are clustered at the ward and year-month levels. The shaded blue area represents a 95% confidence interval, whereas the solid blue line represents point estimates. Vertical dashed lines indicate roughly the beginning and end of the water crisis.

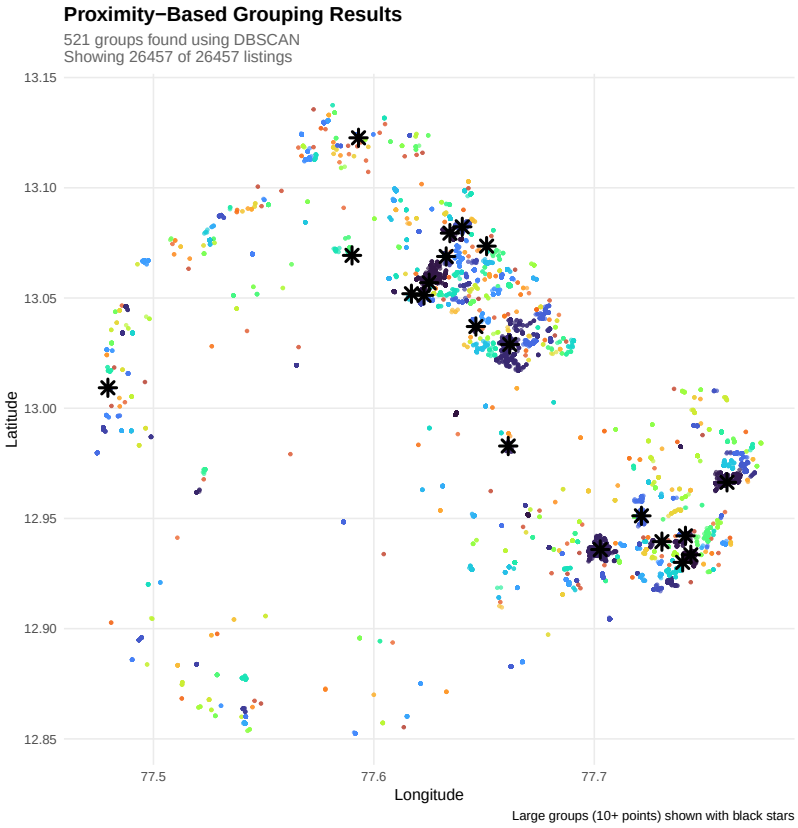


Figure C.7: The Grouping of Listings More Than 0.2 km Away from Piped Water

Notes: Figure depicts the distribution of listings in Bangalore post-crisis that are more than 0.2 km away from the piped water network. We use the density-based spatial clustering of applications with noise (DBSCAN) algorithm to cluster listings together that are within 0.2km of each other. The figure classifies listings in the same cluster by using different colors; clusters that contain more than 10 listings are represented as a large star.

	100m cutoff	200m cutoff	300m cutoff
Prop. Within 100m × Post-Drought	0.691**		
	(0.304)		
Prop. Within 200m × Post-Drought		0.595*	
		(0.304)	
Prop. Within 300m × Post-Drought			0.540*
			(0.319)
Observations	1,620	1,620	1,620
Mean of build	0.12	0.12	0.12
Fixed effects	Subsection, Month	Subsection, Month	Subsection, Month

Table C.2: Proportional Coverage Estimates on New Residential Construction

Notes: Table depicts estimates of the change in construction frequency for areas with piped water access versus those without. Column (1) does using the proportion of an area that is within 100m of a pipe. Column (2) does the same except with a 200m cutoff. Column (3) does the same except with a 300m cutoff. Estimates are interpreted as percent increases after being scaled by 100. For instance, column (3) finds an estimate of 0.540, which indicates a 54.0% larger increase in construction in areas that are 100% within 300m of a pipe compared to areas that are 0% within 300m of a pipe. Standard errors are clustered at the subsection and date levels and reported in parentheses below the pertinent estimate.



Figure C.8: Estimated Price Gradient Before and After Water Crisis

Notes: Figure depicts the estimated price gradient along distance to the piped water network using the method described in Bishop and Timmins [2019] and Robinson [1988]. We estimate a locally qudratric polynomial: We let f_t be locally quadratic and estimate: $P(Z_{it}, H_{it}, \varepsilon_{it}; \beta_t); \beta_t) - H'_{it}\hat{\beta}_t^H = \beta_{0,i^*,t} + \beta_{1,i^*,t} Z_{it} + \beta_{2,i^*,t} Z_{it}^2 + \varepsilon_{it}$.

REFERENCES

Openstreetmap. https://www.openstreetmap.org, 2025. Accessed June 24, 2025.

Alexander C Abajian, Cassandra Cole, Kelsey Jack, Kyle C Meng, and Martine Visser. Dodging day zero: Drought, adaptation, and inequality in cape town. Technical report, National Bureau of Economic Research, 2025.

Gabriel M Ahlfeldt, Nathaniel Baum-Snow, and Remi Jedwab. The skyscraper revolution: Global economic development and land savings. 2023.

Prottoy Akbar, Victor Couture, Gilles Duranton, and Adam Storeygard. Mobility and congestion in urban india. *American Economic Review*, 113(4):1083–1111, 2023a.

Prottoy A Akbar, Victor Couture, Gilles Duranton, and Adam Storeygard. The fast, the slow, and the congested: Urban transportation in rich and poor countries. Technical report, National Bureau of Economic Research, 2023b.

Jesper Akesson, Robert W. Hahn, Robert D. Metcalfe, and Itzhak Rasooly. Race and redistribution in the united states: An experimental analysis. *NBER Working Paper W30426*, 2022.

Hunt Allcott and Judd B. Kessler. The welfare effects of nudges. *American Economic Journal: Applied Economics*, 11(1):236–276, 2019a.

Hunt Allcott and Judd B. Kessler. The welfare effects of nudges. *American Economic Journal: Applied Economics*, 11(1):236–276, 2019b.

Hunt Allcott and Todd Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104 (10):3003–3037, 2014a.

Hunt Allcott and Todd Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104 (10):3003–3037, 2014b.

Hunt Allcott, Reigner Kane, Maximilian S. Maydanchik, Joseph S. Shapiro, and Felix Tintelnot. The effects of 'buy american': Electric vehicles and the inflation reduction act. *NBER Working Paper W33032*, 2024a.

Hunt Allcott, Reigner Kane, Maximilian S. Maydanchik, Joseph S. Shapiro, and Felix Tintelnot. The effects of 'buy american': Electric vehicles and the inflation reduction act. *NBER Working Paper W33032*, 2024b.

Allison, Paul. Imputation by predictive mean matching: Promise & peril. https://statisticalhorizons.com/predictive-mean-matching/, 2015. Accessed June 25, 2025.

Andres Alonso-Robisco and José Manuel Carbó. Analysis of cbdc narrative by central banks using large language models. *Finance Research Letters*, 58:Article 104643, 2023.

Marcella Alsan and Claudia Goldin. Watersheds in child mortality: The role of effective water and sewerage infrastructure, 1880–1920. *Journal of Political Economy*, 127(2):586–638, 2019.

Alternative Fuels Data Center. Alternative fueling station locator, 2025. URL `https://afdc.energy.gov/stations#/analyze?country=US&fuel=ELEC&ev_levels=all&access=public&access=private`. Accessed June 19, 2025.

Shifrah Aron-Dine. Rebuild or relocate? 2024.

Nava Ashraf, Oriana Bandiera, and B. Kelsey Jack. No margin, no mission? a field experiment on incentives for public service delivery. *Journal of Public Economics*, 120:1–17, 2014a.

Nava Ashraf, Oriana Bandiera, and Scott S. Lee. Awards unbundled: Evidence from a natural field experiment. *Journal of Economic Behavior & Organization*, 100:44–63, 2014b.

Karl M. Aspelund and Anna Russo. Additionality and asymmetric information in environmental markets: Evidence from conservation auctions. *Working paper*, 2024.

Rüdiger Bachmann, Gabriel Ehrlich, Ying Fan, Dimitrije Ruzic, and Benjamin Leard. Firms and collective reputation: a study of the volkswagen emissions scandal. *Journal of the European Economic Association*, 21(2):484–525, 08 2022. ISSN 1542-4766. doi:10.1093/jeea/jvac046. URL `https://doi.org/10.1093/jeea/jvac046`.

Clare Balboni, Gharad Bryan, Melanie Morten, and Bilal Siddiqi. Transportation, gentrification, and urban mobility: The inequality effects of place-based policies. *Preliminary Draft*, 3, 2020.

Bangalore Mirror Bureau. Bm property: Bengaluru's rising urban population drives housing growth. `https://bangaloremirror.indiatimes.com/bangalore/others/review-of-freedom-fighter-pension/articleshow/116019621.cms`, December 2024. Updated6Dec202406:00IST; Accessed13August2025.

H Spencer Banzhaf. Difference-in-differences hedonics. *Journal of Political Economy*, 129 (8):2385–2414, 2021a.

H. Spencer Banzhaf. Difference-in-Differences Hedonics. *Journal of Political Economy*, 129 (8):2385–2414, August 2021b. ISSN 0022-3808. doi:10.1086/714442. URL `https://www.journals.uchicago.edu/doi/full/10.1086/714442`. Publisher: The University of Chicago Press.

Roland Bénabou and Jean Tirole. Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678, 2006.

Roland Bénabou and Jean Tirole. Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678, 2006.

Roland Bénabou and Jean Tirole. Laws and norms. *NBER Working Paper W17579*, 2011.

Franziska Bendisch, Gretchen Larsen, and Myfanwy Trueman. Fame and fortune: A conceptual model of ceo brands. *European Journal of Marketing*, 47(3-4):596–614, 2013.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Steven T. Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242–262, 1994.

Bill Bishop. *The Big Sort: Why the Clustering of Like-Minded America Is Tearing Us Apart*. Houghton Mifflin Harcourt, New York, 2009.

Kelly C Bishop and Christopher Timmins. Estimating the marginal willingness to pay function without instrumental variables. *Journal of Urban Economics*, 109:66–83, 2019.

Nicholas Bloom, Mark Schankerman, and John van Reenen. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393, 2013.

Bryan Bollinger and Kenneth T. Gillingham. Peer effects in the diffusion of solar photovoltaic panels. *Marketing Science*, 31(6):900–912, 2012.

Bryan Bollinger, Jesse Burkhardt, and Kenneth T. Gillingham. Peer effects in residential water conservation. *American Economic Journal: Economic Policy*, 12(3):107–133, 2020a.

Bryan Bollinger, Jesse Burkhardt, and Kenneth T. Gillingham. Peer effects in residential water conservation. *American Economic Journal: Economic Policy*, 12(3):107–133, 2020b.

Bryan Bollinger, Kenneth T. Gillingham, A. Justin Kirkpatrick, and Steven Sexton. Visibility and peer influence in durable good adoption. *Marketing Science*, 41(3):453–476, 2022.

Levi Boxell, Matthew Gentzkow, and Jesse M. Shapiro. Cross-country trends in affective polarization. *Review of Economics and Statistics*, 106(2):557–565, 2024.

Bart J. Bronnenberg and Jean-Pierre Dubè. The formation of consumer brand preferences. *Annual Review of Economics*, 9:353–382, 2017.

Bart J. Bronnenberg, Sanjay K. Dhar, and Jean-Pierre Dubè. Brand history, geography, and the persistence of brand shares. *Journal of Political Economy*, 117:87–115, 2009.

Bart J. Bronnenberg, Jean-Pierre Dubè, and Matthew Gentzkow. The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, 102(6): 2472–2508, 2012.

Bart J. Bronnenberg, Jean-Pierre Dubè, Matthew Gentzkow, and Jesse M. Shapiro. Do pharmacists buy bayer? informed shoppers and the brand premium. *Quantitative Journal of Economics*, 130(4):1669–1726, 2015.

Fiona Burlig, James Bushnell, David Rapson, and Catherine Wolfram. Low energy: Estimating electric vehicle electricity use. *AEA Papers and Proceedings*, 111:430–435, 2021.

Fiona Burlig, Amir Jina, and Anant Sudarshan. The value of clean water: Experimental evidence from rural india. Technical report, National Bureau of Economic Research, 2025.

Leonardo Bursztyn and Robert Jensen. Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure. *Annual Review of Economics*, 9: 131–153, 2017.

Raphael Calel, Jonathan Colmer, Antoine Dechezleprêtre, and Matthieu Glachant. Do carbon offsets offset carbon? *American Economic Journal: Applied Economics*, 17(1):1–40, 2025.

Car and Driver. Various pages, 2025. Accessed June 1, 2025.

Stefano Carattini, Kenneth T. Gillingham, Xiangyu Meng, and Erez Yoeli. Peer-to-peer solar and social rewards: evidence from a field experiment. *Centre for Climate Change Economics and Policy Working Paper 408*, 2022.

Vanesa Castán Broto, HS Sudhira, and Hita Unnikrishnan. Walk the pipeline: Urban infrastructure landscapes in bengaluru's long twentieth century. *International Journal of Urban and Regional Research*, 45(4):696–715, 2021.

Damien Cave. India's silicon valley faces a water crisis that software cannot solve. `https://www.nytimes.com/2024/03/31/world/asia/india-bangalore-water.html`, March 2024. Accessed 13 August 2025.

Eugene Y. Chan and Jasmina Ilicic. Political ideology and brand attachment. *International Journal of Research in Marketing*, 36(4):630–646, 2019.

Kenneth Y Chay and Michael Greenstone. Does air quality matter? evidence from the housing market. *Journal of Political Economy*, 113(2):376–424, 2005.

M. Keith Chen and Ryne Rohla. The effect of partisanship and political advertising on close family ties. *Science*, 360(6392):1020–1024, 2018.

C. Robert Clark. Advertising restrictions and competition in the children's breakfast cereal industry. *Journal of Law & Economics*, 50(4):757–780, 2007.

Ronald J. Cohen. Brand personification: Introduction and overview. *Psychology & Marketing*, 31(1):1–30, 2013.

Christopher Conlon and Jeff Gortmaker. Best practices for differentiated products demand estimation with pyblp. *RAND Journal of Economics*, 51(4):1108–1161, 2020.

Jacob Conway and Levi Boxell. Consuming values, 2024. Working Paper.

Thomas R. Covert and Richard L. Sweeney. Winds of change: Estimating learning by doing without cost or input data. *Working Paper*, 2024.

José-Luis Cruz and Esteban Rossi-Hansberg. The economic geography of global warming. *Review of Economic Studies*, 91(2):899–939, 2024.

Lucas Davis, Jing Li, and Katalin Springel. Political ideology and u.s. electric vehicle adoption. *NBER Working Paper W33591*, 2025.

Stefano DellaVigna, John A. List, and Ulrike Malmendier. Testing for altruism and social pressure in charitable giving. *The Quarterly Journal of Economics*, 127(1):1–56, 2012.

Stefano DellaVigna, John A. List, Ulrike Malmendier, and Gautam Rao. Voting to tell others. *The Review of Economic Studies*, 84(1):143–181, 2017.

Klaus Desmet and Esteban Rossi-Hansberg. On the spatial economic impact of global warming. *Journal of Urban Economics*, 88:16–37, 2015.

Klaus Desmet, Robert E Kopp, Scott A Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H Strauss. Evaluating the economic cost of coastal flooding. Technical report, National Bureau of Economic Research, 2018.

Florencia Devoto, Esther Duflo, Pascaline Dupas, William Parienté, and Vincent Pons. Happiness on tap: Piped water adoption in urban morocco. *American Economic Journal: Economic Policy*, 4(4):68–99, 2012.

DNHS. Bengaluru water crisis: Karnataka govt to take over all private tankers, says dks. `https://www.deccanherald.com/india/karnataka/bengaluru/karnataka-government-to-takeover-private-water-tankers-in-bengaluru-amid-severe-water-shortage-deputy-cm-2919285`, March 2024. Accessed 13, August 2025.

Ulrich Doraszelski and Sarit Markovich. Advertising dynamics and competitive advantage. *RAND Journal of Economics*, 38(3):557–592, 2007.

Carolina Tojal R dos Santos and Bruna Morais Guidetti. Expansion of piped water and sewer networks: The effects of regulation. Technical report, Working Paper, 2024.

Benjamin D. Douglas, Patrick J. Ewell, and Markus Brauer. Data quality in online human-subjects research: Comparisons between mturk, prolific, cloudresearch, qualtrics, and sona. *PLoS ONE*, 18(3):e0279720, 2023.

Pascaline Dupas, Vivian Hoffmann, Michael Kremer, and Alix Peterson Zwane. Targeting health subsidies through a nonprice mechanism: A randomized controlled trial in kenya. *Science*, 353(6302):889–895, 2016.

Deniz Dutz, Ingrid Huitfeldt, Santiago Lacouture, Magne Mogstad, Alexander Torgovitsky, and Winnie van Dijk. Selection in surveys: Using randomized incentives to detect and account for nonresponse bias. *NBER Working Paper W29549*, 2025.

Edmunds. Various pages, 2025. Accessed June 1, 2025.

Ivar Ekeland, James J Heckman, and Lars Nesheim. Identification and estimation of hedonic models. *Journal of political economy*, 112(S1):S60–S109, 2004.

Marc Fetscherin. *CEO Branding: Theory and Practice*. Routledge, 1 edition, 2015.

F&I Tools USA. New car sales by state, 2025. URL `https://www.factorywarrantylist.com/car-sales-by-state.html`. Accessed June 5, 2025.

Pierre Filion and Roger Keil. Contested infrastructures: Tension, inequity and innovation in the global suburb. *Urban Policy and Research*, 1:7–19, 2016.

Eric Fischer, Rebecca McCaughrin, Saketh Prazad, and Mark Vandergon. Fed transparency and policy expectation errors: A text analysis approach. Technical report, Federal Reserve Board of New York Staff Report No. 1081, 2023.

Sebastian Galiani, Paul Gertler, and Ernesto Schargrodsky. Water for life: The impact of the privatization of water services on child mortality. *Journal of political economy*, 113(1): 83–120, 2005.

Shanti Gamper-Rabindran, Shakeeb Khan, and Christopher Timmins. The impact of piped water provision on infant mortality in brazil: A quantile panel data approach. *Journal of Development Economics*, 92(2):188–200, 2010.

Amit Gandhi and Jean-François Houde. Measuring substitution patterns in differentiated-products industries. *NBER Working Paper w26375*, 2019.

Anita George, Rashad-Rudolf Kaldany, and Joseph Losavio. The world is facing a $15 trillion infrastructure gap by 2040. here's how to bridge it, April 2019. URL `https://www.weforum.org/stories/2019/04/infrastructure-gap-heres-how-to-solve-it/`.

Alan S. Gerber, Donald P. Green, and Christopher W. Larimer. Social pressure and voter turnout: Evidence from a large-scale field experiment. *American Political Science Review*, 102(1):33–48, 2008.

Karen Gift and Thomas Gift. Does politics influence hiring? evidence from a randomized experiment. *Political Behavior*, 37(3):653–675, 2015.

Kenneth T. Gillingham and Bryan Bollinger. Social learning and solar photovoltaic adoption. *Management Science*, 67(11):7091–7112, 2021a.

Kenneth T. Gillingham and Bryan Bollinger. Social learning and solar photovoltaic adoption. *Management Science*, 67(11):7091–7112, 2021b.

Kenneth T. Gillingham, Matthew Kotchen, James A. Levinsohn, and Barry J. Nalebuff. The musk partisan effect on tesla sales. *NBER Working Paper w34413*, 2025.

James G. Gimpel and Iris S. Hui. Seeking politically compatible neighbors? the role of neighborhood partisan composition in residential sorting. *Political Geography*, 48:130–142, 2015.

Google. Google trends, 2025. URL `https://trends.google.com/`. Accessed June 10, 2025.

Greer Gosnell, Stefano Carattini, and Alessandro Tavoni. Observing the unobservable: A field experiment on early adopters of a climate-friendly behavior. *Centre for Climate Change Economics and Policy Working Paper 393*, 2021.

Michael Greenstone. The continuing impact of sherwin rosen's "hedonic prices and implicit markets: product differentiation in pure competition". *Journal of Political Economy*, 125 (6):1891–1902, 2017.

Michael Greenstone and Justin Gallagher. Does hazardous waste matter? evidence from the housing market and the superfund program. *The Quarterly Journal of Economics*, 123(3): 951–1003, 2008.

Paul L E Grieco, Charles Murry, and Ali Yurukoglu. The evolution of market power in the u.s. automobile industry*. *The Quarterly Journal of Economics*, 139(2):1201–1253, 09 2023. ISSN 0033-5533. doi:10.1093/qje/qjad047. URL `https://doi.org/10.1093/qje/qjad047`.

Vladas Griskevicius, Joshua M. Tybur, and Bram Van den Bergh. Going green to be seen: Status, reputation, and conspicuous conservation. *Journal of Personality and Social Psychology*, 98(3):392–404, 2010.

Kevin Haninger, Lala Ma, and Christopher Timmins. The value of brownfield remediation. *Journal of the Association of Environmental and Resource Economists*, 4(1):197–241, 2017.

Mariaflavia Harari. Cities in bad shape: Urban geometry in india. *American Economic Review*, 110(8):2377–2421, 2020.

Jonathan Hartley. The macroeconomic effects of oil supply chain shocks: Evidence from us extreme weather events affecting refineries. *Available at SSRN 5165200*, 2024.

Johannes Haushofer, Michael Kremer, Ricardo Maertens, and Brandon Joel Tan. Water treatment and child mortality: Evidence from kenya. Technical report, National Bureau of Economic Research, 2021.

Gea Hoogendoorn, Bernadette Sütterlin, and Michael Siegrist. When good intentions go bad: The biased perception of the environmental impact of a behavior due to reliance on an actor's behavioral intention. *Journal of Environmental Psychology*, 64:65–77, 2019.

Richard Hornbeck. The enduring impact of the american dust bowl: Short-and long-run adjustments to environmental catastrophe. *American Economic Review*, 102(4):1477–1507, 2012.

Richard Hornbeck and Suresh Naidu. When the levee breaks: black migration and economic development in the american south. *American Economic Review*, 104(3):963–990, 2014.

Housing.com. Housing.com: Real estate listings and resources. `https://housing.com`, 2025. Various webpages accessed May 2023 through April 2025.

Ye Hu. Diversity matters: How film critic ratings vary with critic and movie cast racial profiles. *Journal of Marketing*, 89(5):111–129, 2025.

Gregory A. Huber and Neil Malhotra. Political homophily in social relationships: Evidence from online dating behavior. *Journal of Politics*, 79(1):269–283, 2017.

IEA. Global ev outlook 2024, 2024. URL `https://www.iea.org/reports/global-ev-outlook-2024/trends-in-electric-vehicle-charging?utm_source=web&utm_medium=article&utm_campaign=did_you_know`. Accessed June 19, 2025.

Internet Archive. Wayback machine. `https://web.archive.org/`, 2025. Accessed: 2025–6–1.

Shanto Iyengar, Gaurav Sood, and Yphtach Lelkes. Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, 76(3):405–431, 2012.

Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. The origins and consequences of affective polarization in the united states. *Annual Review of Political Science*, 22:129–146, 2019.

Japan International Cooperation Agency. Unveiling of the jica-funded 5th stage of cauvery water supply scheme. `https://www.jica.go.jp/english/overseas/india/information/press/2024/1553733_53431.html`, 2024. Accessed August 8, 2025.

Edward Jee, Anne Karing, and Karim Naguib. Optimal policy with social image concerns: Experimental evidence from deworming. *Working Paper*, 2024.

Hana Jomni and Nikita Zakharov. Do terrorist attacks polarize politicians? evidence from the european parliamentary speeches on migration. *University of Freiburg, Department of International Economic Policy Discussion Paper Series No. 50*, 2024.

Emir Kamenica, Robert Naclerio, and Anup Malani. Advertisements impact the physiological efficacy of a branded drug. *Proceedings of the National Academy of Sciences*, 110(32):12931–12935, 2013.

Anne Karing. Social signaling and childhood immunization: A field experiment in sierra leone. *The Quarterly Journal of Economics*, 139(4):2083–2133, 2024.

Kelley Blue Book. Various pages, 2025. Accessed May 30, 2025.

Romana Khan, Kanishka Misra, and Vishal Singh. Ideology and brand consumption. *Psychological Science*, 24(3):326–333, 2013.

Michael Kremer, Jessica Leino, Edward Miguel, and Alix Peterson Zwane. Spring cleaning: Rural water impacts, valuation, and property rights institutions. *The Quarterly Journal of Economics*, 126(1):145–205, 2011.

Randi Kronthal-Sacco and Tensie Whelan. Sustainable market share index. Technical report, NYU Stern Center for Sustainable Business Report, 2024.

R. Kulranjan, S. Palur, and M. Nesi. How water flows through bengaluru: Urban water balance report. Water, Environment, Land and Livelihoods (WELL) Labs at Institute for Financial Management and Research, 2023. Accessed 13 August 2025.

O. Hwang Kwon, Katie Vu, Naman Bhargava, Mohammed I. Radaideh, Jacob Cooper, Veda Joynt, and Majdi I. Radaideh. Sentiment analysis of the united states public support of nuclear power on social media using large language models. *Renewable and Sustainable Energy Reviews*, 200:Article 114570, 2024.

Ashley Langer and Derek Lemoine. Designing dynamic subsidies to spur adoption of new technologies. *Journal of the Association of Environmental and Resource Economists*, 9(6): 1197–1234, 2022.

Benjamin Leard and Yidi Wu. New passenger vehicle demand elasticities: Estimates and policy implications. *Resources for the Future Working Paper 23-33*, 01 2023.

Benjamin Leard, Joshua Linn, and Yichen Christy Zhou. How much do consumers value fuel economy and performance? evidence from technology adoption. *The Review of Economics and Statistics*, 105(1):158–174, 01 2023. ISSN 0034-6535. doi:10.1162/rest_a_01045. URL `https://doi.org/10.1162/rest_a_01045`.

Chuhui Li, Donald S. Poskitt, Frank Windmeijer, and Xueyan Zhao. Binary outcomes, ols, 2sls, and iv probit. *Econometric Reviews*, 41(8):859–876, 2022.

Roderick J. A. Little. Missing-data adjustments in large surveys. *Journal of Business & Economic Statistics*, 6:287–296, 1988.

Jiarui Liu. Who should be subsidized for electric vehicles? demand estimation and policy design under network effects, 2022. Working Paper.

Jill Luoto, Nusrat Najnin, Minhaj Mahmud, Jeff Albert, M Sirajul Islam, Stephen Luby, Leanne Unicomb, and David I Levine. What point-of-use water treatment products do consumers use? evidence from a randomized controlled trial among the urban poor in bangladesh. *PloS one*, 6(10):e26132, 2011.

Charlie Mahtesian and Calder McHugh. Tesla's elon musk problem, 2025. URL https://www.politico.com/newsletters/politico-nightly/2025/02/26/teslas-elon-musk-problem-00206353. Politico. Accessed June 19, 2025.

Christopher McConnell, Yotam Margalit, Neil Malhotra, and Matthew Levendusky. The economic consequences of partisanship in a polarized era. *American Journal of Political Science*, 62(1):5–18, 2018.

Daniel McFadden. Econometric models of probabilistic choice. In Charles F. Manski and Daniel McFadden, editors, *Structural Analysis of Discrete Data with Econometric Applications*, pages 198–272. MIT Press, Cambridge, MA, 1981.

Media Cloud. Media cloud news database, 2025. URL https://search.mediacloud.org/search. Accessed June 18, 2025.

MIT Election Data and Science Lab. U.s. elections project data repository, 2024. URL https://electionlab.mit.edu. Accessed January 15, 2024.

Anahit Mkrtchyan, Jason Sandvik, and Vivi Z. Zhu. Ceo activism and firm value. *Management Science*, 70(10):6843–7343, 2023.

Rhea Mogul, Meenketan Jha, and Sania Farooqui. India's 'silicon valley' is running dry as residents urged to take fewer showers and use disposable cutlery. https://edition.cnn.com/2024/03/14/india/india-bangalore-water-crisis-impact-intl-hnk-dst, March 2024. Accessed 13, August 2025.

Jonathan Mummolo and Clayton Nall. Why partisans do not sort: The constraints on political segregation. *Journal of Politics*, 79(1):45–59, 2017.

Kevin M. Murphy and Robert H. Topel. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*, 3(4):370–379, 1985.

NBC. Hart research associates/public opinion strategies,". https://pos.org/wp-content/uploads/2025/03/NBC-March-2025-Poll-3-16-2025-Release.pdf, 2025. Accessed: 2025–9-20.

NORC. Climate attitudes, 2024. URL https://apnorc.org/wp-content/uploads/2024/06/EPIC-2024-Combined-Fact-Sheets.pdf. Accessed 2024.

Clair Null, Christine P Stewart, Amy J Pickering, Holly N Dentz, Benjamin F Arnold, Charles D Arnold, Jade Benjamin-Chung, Thomas Clasen, Kathryn G Dewey, Lia CH Fernald, et al. Effects of water quality, sanitation, handwashing, and nutritional interventions on diarrhoea and child growth in rural kenya: a cluster-randomised controlled trial. *The Lancet Global Health*, 6(3):e316–e329, 2018.

Daniel O'Callaghan and Sean Prior. Central technical appraisal parameters: Discount rate, time horizon, shadow price of public funds and shadow price of labour. Technical report, IGEES Technical Paper, 2018.

Ashton Pallottini and Sofia Shchukina. Changing gears: Elon musk and the role of ideology in demand for environmental goods. *Working Paper*, 2025.

Costas Panagopoulos, Donald P. Green, Jonathan Krasno, Michael Schwam-Baird, and Kyle Endres. Partisan consumerism: Experimental tests of consumer reactions to corporate political activity. *Journal of Politics*, 82(3):996–1007, 2020.

Ricardo Perez-Truglia and Guillermo Cruces. Partisan interactions: Evidence from a field experiment in the united states. *Journal of Political Economy*, 125(4):1208–1243, 2017.

Andhika Putra Pratama, Muhammad Halley Yudhistira, and Eric Koomen. Highway expansion and urban sprawl in the jakarta metropolitan area. *Land use policy*, 112:105856, 2022.

Steve Rathje, Dan-Mircea Mirea, Ilia Sucholutsky, Raja Marjieh, Claire E. Robertson, and Jay J. Van Bavel. Gpt is an effective tool for multilingual psychological text analysis. *PNAS*, 121(34):e2308950121, 2024.

Real Estate Regulatory Authority – Karnataka. Official portal of the real estate regulatory authority – karnataka (k-rera). `https://rera.karnataka.gov.in`, 2025. Accessed June 24, 2025.

Jennifer Roback. Wages, rents, and the quality of life. *Journal of political Economy*, 90(6): 1257–1278, 1982.

Peter M Robinson. Root-n-consistent semiparametric regression. *Econometrica: journal of the Econometric Society*, pages 931–954, 1988.

Matthias Rodemeier. Willingness to pay for carbon mitigation: Field evidence from the market for carbon offsets. *IZA Working Paper 15939*, 2024.

Sherwin Rosen. Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, 82(1):34–55, 1974.

Jordan Daniel Rosenthal-Kay. Urban costs around the world. 2025.

Donald B. Rubin. Statistical matching using file concatenation with adjusted weights and multiple imputations. *Journal of Business & Economic Statistics*, 4:87–94, 1986.

Christian Rudeloff and Julius Damms. Building brand loyalty on twitter: The case of elon musk and tesla. *Journal of Digital & Social Media Marketing*, 10(3):238–248, 2023.

Karri S.D. Sarkar, A. Water quality index to ascertain the tanker water quality in a bangalore rural area: a case study. *Int J Energ Water Res*, 9:609–635, 2025.

Richard Schmalansee. Advertising and entry deterrence: An exploratory model. *Journal of Political Economy*, 91(4):636–653, 1983.

Steven E. Sexton and Alison L. Sexton. Conspicuous conservation: The prius halo and willingness to pay for environmental bona fides. *Journal of Environmental Economics and Management*, 67(3):303–317, 2014a.

Steven E. Sexton and Alison L. Sexton. Conspicuous conservation: The prius halo and willingness to pay for environmental bona fides. *Journal of Environmental Economics and Management*, 67(3):303–317, 2014b.

Sangwoo Shin, Sanjog Misra, and Dan Horsky. Disentangling preferences and learning in brand choice models. *Marketing Science*, 31(1):115–137, 2012.

Jeanne Sorin Le Guével. Essays in development, urban, and environmental economics. 2025.

Michael Spence. Cost reduction, competition, and industry performance. *Econometrica*, 52 (1):101–122, 1984.

Tristian R Stolte, Hans de Moel, Elco E Koks, Marthe LK Wens, Felix van Veldhoven, Snigdha Garg, Neuni Farhad, and Philip J Ward. Global drought risk in cities: Present and future urban hotspots. *Environmental Research Communications*, 5(11):115008, 2023.

Andrea Szabo. The value of free water: analyzing south africa's free basic water policy. *Econometrica*, 83(5):1913–1961, 2015.

Jenny Tang, Eleanor Birrell, and Ada Lerner. Replication: How well do my results generalize now? the external validity of online privacy and security surveys. In *Proceedings of the Eighteenth Symposium on Usable Privacy and Security*, 2022.

Tesla-Info.com. Tesla info — inventory, release notes and buyers guides. `https://tesla-info.com/`, 2025. Accessed: 2025–6–1.

Texas Department of Motor Vehicles. Vehicle registration data, 2025. Accessed via public records request May 10, 2025.

Richard Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, 1(1):39–60, 1980.

U.S. Census Bureau. American community survey 2023 (acs 1-year estimates), 2024. URL `https://www.census.gov/programs-surveys/acs`. Accessed January 15, 2024.

U.S. EPA. Epa report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances. Technical report, Technical Report, 2023.

U.S. EPA. fueleconomy.gov dataset, 2025. URL `https://www.fueleconomy.gov/feg/download.shtml`. Accessed January 15, 2024.

U.S. White House. Fact sheet: President biden sets 2035 climate target aimed at creating good-paying union jobs, reducing costs for all americans, and securing u.s. leadership in the clean energy economy of the future. Official Statement, 2024.

R. Vaidyanathan. Bwssb water supply lines maps of bengaluru. Dataset, OpenCity (data.opencity.in), 2022. URL https://data.opencity.in/dataset/bwssb-water-supply-lines-map-of-bengaluru. Accessed June 24, 2025; Public Domain.

Watchful1. Subreddit comments/submissions 2005-06 to 2024-12, 2025. URL https://www.reddit.com/r/pushshift/comments/1itme1k/separate_dump_files_for_the_top_40k_subreddits/. Accessed June 13, 2025.

Jeffrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*. The MIT Press, 2010.

World Resources Institute. The realities of current urbanization in the global south. World resources report working paper, World Resources Institute, 2021. Accessed July 2, 2025.

Jianwei Xing, Benjamin Leard, and Shanjun Li. What does an electric vehicle replace? *Journal of Environmental Economics and Management*, 107:102432, 2021. ISSN 0095-0696. doi:https://doi.org/10.1016/j.jeem.2021.102432. URL https://www.sciencedirect.com/science/article/pii/S0095069621000152.