Essays in Behavioral Economics and Environmental Policy

by

Steven Eric Sexton

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

in

Agricultural and Resource Economics

in the

Graduate Division
of the
University of California, Berkeley

Committee in charge:

Professor David Zilberman, Chair
Professor Maximilian Auffhammer
Professor Enrico Moretti

Fall 2012

Essays in Behavioral Economics and Environmental Policy

Abstract

Essays in Behavioral Economics and Environmental Policy

by

Steven Eric Sexton
Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor David Zilberman, Chair

Social planners have long relied upon non-coercive interventions in order to achieve social welfare improvements that are not obtained by markets or direct policy. Such policies are perhaps nowhere more relevant and common than in environmental economics. Environmental goods and services are typically not traded in markets because of the difficulties of property rights assignment. And yet efforts to create markets or correct market failures by coercive policy are fraught with controversy. Thus, in addition to coercive mechanisms, social planners use information provision campaigns, appeals for cooperation, and "nudges" to improve the efficiency of environmental resource allocations. Non-coercive interventions have grown in popularity among social planners as behavioral economics has gained acceptance within the mainstream of the field. Indeed, such policies typically affect market outcomes and achieve environmental goals only insofar as they can exploit or correct decision making that deviates from standard theory.

In this dissertation, agent behavior is analyzed to assess the potential of non-coercive interventions to achieve socially preferred environmental outcomes. In a first essay, the concept of conspicuous conservation is introduced as a modern variant of conspicuous consumption that affords status for displays of austerity meant to signal environmental preferences rather than displays of ostentation meant to signal wealth. I identify conspicuous conservation in the automobile market and estimate a willingness to pay up to several thousand dollars for the "green" signal transmitted by ownership of the Toyota Prius.

In a second essay, I demonstrate how automatic bill payment programs can induce excessive consumption of goods and services by boundedly rational consumers who exhibit inattention to prices. As automatic payment programs have spread throughout industries characterized by recurring payments, from utility and telecommunication services to insurance and loan markets, this essay is the first to consider their implications for consumer demand and welfare. It is also the first to test empirically whether enrollment in such programs increases demand, as price salience theory suggests. It is shown that residential electricity consumption increases on average 2-4.5

A final essay examines the extent to which free transit fares and appeals for car-trip

avoidance reduce car pollution on smoggy days. With data on freeway traffic volumes and transit ridership, public appeals for cooperation are shown to have no significant effect on car trip demand. Free transit fares, however, do have a significant effect on car trip demand. But the effect is perverse in that it generates an increase in car trips and related pollution. Free fares also increase transit ridership. These results suggest that free transit rides do not induce motorists to substitute to transit, but instead subsidize regular transit rides and additional trips. Appeals for cooperation also have no affect on carpooling behavior.

Viewed in their totality, these essays communicate the importance of behavioral theories in formulating environmental policies and predicting agents' responses to such policies. Policies formulated without due regard for agents' bounded rationality and multifaceted motivations are doomed to unintended consequeces. However, recognition of these behavioral responses and their incorporation in policy design can result in improved environmental outcomes and efficient policies.

To my parents for their unwavering support.

Contents

List of Figures

List of Tables

Acknowledgments

I am indebted to my parents who instilled in me early on an appreciation for pursuit and who fully supported all my pursuits. My educational attainment is a consequence of their unwavering comittment to providing me with every conceivable opportunity for academic enrichment and to pushing me to my greatest potential. Their counsel, encouragement, and confidence proved particularly invaluable during my graduate studies. But above all, my parents inspired me.

It is by their devotion and by grace that I was blessed with many great teachers over the years with whom I share credit for this achievement. None has been more instrumental in my development as an economist than my advisor, David Zilberman, whose faith in me was unflagging. His ambition matched my own and inspired me to challenge ideas and the status quo. A more supportive and sage mentor I cannot imagine. My tenure at Berkeley was also greatly enriched by Max Auffhammer and Sofia Villas-Boas, who both guided me and buttressed my resolve and confidence. I am also honored to have Enrico Moretti as a member of my dissertation committee. His consequential research motivates me and galvanizes my appreciation for the power of economic analysis. Knowing that my work would be evaluated by a committee of such accomplished members consistently inspired my best effort.

And I am fortunate to share this milepost in life, as others, with my twin sister, Alison, who is my best friend. I am grateful for the strength she gave me, as well as her advice informed by shared experience. Finally, my time at Berkeley was made imminently more enjoyable by the friendship of classmate Sarah Dobson, whose patience I did not deserve but did appreciate.

Chapter 1

Introduction

Social planners have long relied upon non-coercive interventions in order to achieve social welfare improvements that are not achieved by markets or direct policy. Missing markets and demand-invariant and fixed prices can cause suboptimal allocations of scarce resources, while political considerations can forestall explicit rationing or coercive screening. Thus, during World War I, the U.S. government urged citizens to aid the war effort by planting "victory" gardens and conserving food. In World War II, citizens were asked to dress warmly to conserve fuel. More recently, governments have urged conservation of natural resources amid concern about environmental damage and sustainability. They also provide information to decision makers and alter the settings in which agents make decisions in order to affect choices and, thereby, achieve desired outcomes (e.g., Thaler and Sunstein 2008).

Non-coercive interventions are perhaps nowhere more relevant and common than in environmental economics. Environmental goods and services are typically not traded in markets because of the difficulties of property-rights assignment, and, yet, efforts to create markets or correct market failures by coercive policies are fraught with controversy. The absence of an international agreement to internalize the cost of greenhouse gas emissions in 2012 is, perhaps, the most prominent example. But others abound. For instance, robust water markets are absent in California and much of the water-starved American West because of concern about distributional impacts. Thus, in addition to coercive mechanisms, like standards, taxes, and quotas, social planners employ a range of non-coercive programs to improve the efficiency of environmental resource allocations. These include information provision campaigns like the Energy Star program and fuel economy disclosures for new cars, as well as appeals for conservation of energy, water, and clean air when they are in particularly scarce supply, and "nudges" that rely on social pressure or changes in the "choice architecture" to induce conservation.

Non-coercive interventions have grown in popularity among social planners as behavioral economics has gained acceptance within the mainstream of the field. Indeed, such policies typically affect market outcomes and achieve environmental goals only insofar as they can exploit or correct decision-making that deviates from standard theory. They elicit prosocial behavior by appealing to altruism and status-consciousness and induce technology adoption by correcting information asymmetries and drawing attention to otherwise insalient characteristics. They change the choice architecture that frames decision-making to increase the likelihood that socially preferred behaviors are provided by agents with bounded rationality or incomplete and transitive preferences. For non-coercive policies to achieve incremental improvements in outcomes of interest, it is necessary that some agents either exhibit irrationality or altruism, or are not fully informed about their decisions.

Among the many important contributions of behavioral economics, perhaps none are more relevant to environmental policy than the development of bounded rationality and the refinement of our understanding of altruism. The former proposes that individuals may lack the cognitive ability or the degree of attention required to make perfectly rational decisions all of the time (Kahneman 2003). Thus, they may rely on rules of thumb and heuristics and over-weight in decision making that information which is most salient, producing systematic biases in behavior. For instance, the energy paradox is attributed, in part, to the insalience of future cost savings from the adoption of energy-efficient technologies that typically require greater upfront costs than conventional alternatives.

Standard theory predicts the under-provision of public goods by private agents because of the classic free rider problem. The development of theory related to altruism has been iterative, with richer theory needed to explain increasingly puzzling phenomena. Altruism, and then impure altruism, were introduced to explain considerable private contributions to charity and other public goods that do not exhibit one-to-one crowd-out by other private or public contributions (Andreoni 1989, 1990). But they alone do not explain why prosocial behavior increases when it becomes more visible and decreases when it earns small pecuniary or non-pecuniary rewards. So Benabou and Tirole (2006a) defined reputational motivations that compel prosocial behavior for selfish, rather than altruistic, reasons. Thus, the more visible is the behavior and the less it can be confused with the pursuit of direct extrinsic rewards, the more common it will be, all else equal. A rich theory related to prosocial behavior is important to understand the success of green products and the relatively greater success of conspicuous green products like the Toyota Prius. As social planners seek to encourage resource conservation and other pro-environment behaviors, a full understanding of agents' motivations and the prevalence and strength of agent preferences is important.

A common critique of behavioral economics asserts that departures from behavior consistent with rational decision-making are sufficiently rare as to not impact market outcomes and that they should be fleeting as individuals learn from mistakes and markets enforce rationality. A growing body of empirical work, however, documents persistent deviations from standard theory in the field, and not merely in the laboratory. Learning need not occur, and markets need not enforce rationality (Thaler 1986). And if markets do not exist, as they typically do not for environmental goods and services, then systematic biases may persist (Shogren and Taylor 2008).

It is against this backdrop of recent advances in behavioral economics and their applications in the environmental field that my dissertation is set. In the following three chapters, agent behavior is analyzed to assess the potential of non-coercive interventions to achieve socially preferred environmental outcomes. First, I develop the concept of *conspicuous conservation* as a modern variant of conspicuous consumption, which the economic literature has long understood as an explanation for persistent price premia for luxury items that are functionally equivalent to less expensive alternatives. Amid growing concern about environmental protection, however, status may be attained by displays of austerity meant to signal environmental preferences rather than displays of ostentation meant to signal wealth. Consistent with Benabou and Tirole (2006a), I hypothesize that the private provision of environmental public goods can be achieved by reliance on status-seeking behavior rather than altruism. The hypothesis is tested in the market for passenger cars wherein the Toyota Prius exclusively signaled environmental virtues with a unique model design for many years. Drawing on vehicle registration data in two states, I identify a statistically significant conspicuous conservation effect and estimate a willingness to pay as much as several thousand dollars for the "green" signal transmitted by the Prius. Results suggest that conspicuous green products can gain market share relative to inconspicuous green products and that the private provision of public goods need not rely on altruism alone. The implications of conspicuous conservation for both coercive non-coercive environmental policy are considered.

In a second essay, I consider whether a modern feature of recurring payment programs combines with bounded rationality among consumers to induce greater consumption of goods and services than would occur by fully rational agents. Whereas bill payment historically required individuals to review their bills in order to process payments, the introduction of automatic payment mechanisms in about 2005 eliminated the need for individuals or firms to view their bills at all. If individuals and firms enrolled in automatic payment programs forego regular inspection of their recurring bills, then price salience is diminished. While diminished price salience does not impact behavior in standard models of perfectly rational economic agents, behavioral theories suggest that prices perceived by boundedly rational agents decline as price salience falls.

As automatic payment programs have spread throughout industries characterized by recurring payments, from utility and telecommunication services to insurance and loan markets, this essay is the first to consider their implications for consumer demand and welfare. It is also the first to test empirically whether enrollment in such programs increases demand, as price salience theory suggests. Based upon analysis of data for a South Carolina public utility, it is shown that residential electricity consumption increases on average 2-4.5% due to enrollment in automatic payment programs. Commercial electricity consumption is estimated to increase as much as 6%. Moreover, bill-smoothing programs that utilities offer to low-income households are shown to conceal the cost of deviations from mean consumption and induce an 8-9% increase in electricity consumption. Such billing programs, then, counteract the efforts of regulated utilities and social planners to reduce electricity consumption, partly by boosting the strength of the price signal.

Finally, in a third essay, I demonstrate that a well-intended program to induce conservation of clean air on smoggy days has the perverse effect of increasing car trips in the San Francisco Bay Area. Amid heightened concern about environmental damage and climate change, public appeals for cooperation and pecuniary incentives are frequently used to achieve resource conservation or induce other prosocial behavior. However, the relative effectiveness of the two instruments is poorly understood when pecuniary incentives are small. This essay examines the extent to which free transit fares and appeals for car trip avoidance reduce car pollution on smoggy days. With data on freeway traffic volumes and transit ridership, public appeals for cooperation are shown to have no significant effect on car trip demand. Free transit fares, however, do have a significant effect on car trip demand. But the effect is perverse in that it generates an increase in car trips and related pollution. Free fares also increase transit ridership. These results suggest that free transit rides do not induce motorists to substitute to transit, but instead subsidize regular transit rides and additional trips. Appeals for cooperation have no affect on carpooling behavior. Thus, the San Francisco program demonstrates the importance of fully considering the behavioral response of agents to various interventions and of a modern understanding of agents' motivations.

In total, these essays demonstrate that the limitations on rational decision-making and the varied extent of agent motivations not envisioned in standard theory are important to environmental policy. Deviations from the standard theory can be exploited by firms to increase revenues and gain market share in green markets, or they can be leveraged by social planners to expand private provision of environmental public goods. Failure to appreciate cognitive constraints and the breadth of agent motivations, including reputational motivations, however, can cause well-intentioned interventions to have perverse effects.

Chapter 2

Conspicuous Conservation: The Prius Halo and Willingness to Pay for Environmental Bona Fides

2.1 Introduction

Veblen explained in 1899 that "in order to gain and hold the esteem of man it is not sufficient merely to possess wealth or power. The wealth or power must be put in evidence, for esteem is awarded only on evidence" (p. 36). Since then, a considerable literature has explored the concept of conspicuous consumption and its implications in various settings, with particular focus on purchases that signal prestige, luxury and exclusivity.[1] Ownership of luxurious estates, automobiles, and fashion surely still afford a certain social status in the 21st Century. However, amid growing concern about environmental damage and global climate change, an evolution of social norms suggests status is increasingly conferred upon demonstration of austerity rather than ostentation–particularly austerity that minimizes the environmental impact of consumption. Consumers may, therefore, undertake costly actions in order to signal their preferences for the environment, a phenomenon we term *conspicuous conservation*.

The status conferred upon demonstration of one's preferences for environmental protection is sufficiently prized that some homeowners who purchase energy efficient home heating and cooling technologies display them prominently in their homes rather than relegate them to utility closets. Similarly, instead of boasting about the sizes of their homes, a growing number of homeowners are burying up to half of their homes underground in order to hide their magnitudes from view (WSJ 2011-06-17; WSJ 2010-06-15). And, as is demonstrated in the following sections, green signaling motivates automobile purchases.

Car ownership decisions are among the most visible consumption decisions households make. Consumers attach symbolic meanings to vehicles and use vehicle choices to communicate interests, beliefs, values, and status (e.g., Grubb and Hupp 1968; Sirgy 1985; Heffner et al. 2006). Since the U.S. introduction of the Toyota Prius in 2001, a growing number of vehicle models have been developed with features that reduce environmental impacts, particularly greenhouse gas emissions. They include small and light cars with conventional engines (like the DaimlerChrysler SmartCar), alternative fuel cars (like the Chevrolet flex-fuel fleet), and hybrid cars (like the Prius, the Honda Civic Hybrid, and others). Until the reintroduction of the Honda Insight in 2010, however, the Prius was the only model that at once provided standard features (e.g., climate control, four doors, luggage space, etc.), environmental amenities, and a design unique to the model.[2]

In 2010, the Prius was the clear leader among 24 different hybrid models available in the United States; 48% of the 290,271 hybrid cars sold in the U.S. in 2010 were Priuses. The success of the Prius is no doubt attributable, in part, to an aggressive and innova-

[1]See for instance: Leibenstein (1950), and more recently Frank (1985); Basu (1987); Braun and Wicklund (1989); Ireland (1998). More generally, other recent studies, including Akerlof (1980); Cole et al. (1992); ?); Fershtman and Weiss (1993); Bernheim (1994); Glazer and Konrad (1996) explore the impact of status consciousness on economic behavior.

[2]The Honda Insight was first introduced in 1999, two years before the Prius. However, it was a two-door subcompact car that sacrificed on amenities available in most passenger cars at the time. The Insight was re-introduced in 2010 as a four-door sedan, joining the class of four-door hybrids with unique model names and designs.

tive marketing effort by Toyota and to equity in the Toyota brand. However, national marketing effort and brand reputation do not explain why Prius ownership increases in green communities relative to other comparably green cars, a result demonstrated in this paper. The Toyota Camry Hybrid, for instance, loses market share relative to the Prius in greener communities even after controlling for the green attributes of the two models. Likewise, the Civic Hybrid achieves green ratings virtually identical to the Prius, yet it, too, is underrepresented in green locales relative to the Prius.

We hypothesize that these market-share differences are due to the unique design of the Prius, and, thus, its ability to signal environmental bona fides, the value of which can be expected to be greater in greener communities. Toyota executives reportedly instructed their engineers to develop a unique design, regardless of the quality of the styling. Prius design has been described as both utilitarian in order to maximize aerodynamics and "head-turning" with a shape unique to the model. The Civic Hybrid and other hybrids, in contrast, share body styling and model names with other trims in the model class that carry conventional drivetrains. The hybrid trims of these models typically carry only a small badge on the side or rear of the vehicles indicating their types. The Prius, therefore, has historically provided the most powerful signal of the owner's affinity for the environment of any vehicle in the U.S (e.g. The Washington Post 2004-08-23; Time Magazine 2007; The New York Times 2007-07-04).

We test empirically for the presence of a conspicuous conservation effect in vehicle purchase decisions. The empirical strategy follows from theory developed in Section 3 that the value of a green signal is increasing in the strength and prevalence of preferences for the environment within one's community. All else equal, then, a Prius is more valuable in communities with a strong green ethos like Boulder, Colorado than in communities with greater heterogeneity in attitudes toward the environment, like, for instance, Greeley, Colorado. Thus, while shares of all green car models are expected to be greater in green communities than "brown" communities due to the higher concentration of environmentally conscious consumers, shares of the uniquely designed Prius should be *disproportionately* greater in those communities if green purchases are motivated, at least in part, by green signaling.

Using observed variation in model ownership rates across communities in Colorado and Washington and relying on electoral outcomes as proxies for community greenness, we identify a statistically and economically significant conspicuous conservation effect in car purchases. This result has important implications for policies designed to improve the environment and for the private provision of public goods. Signaling motives can increase private provision of public goods, which, absent coercive policy, can be an important response to environmental challenges, like climate change. Signaling motives, however, cause the individual consumer's objective to deviate from that of the social planner. The planner seeks to maximize environmental benefit per conservation dollar where as the consumer may seek to maximize status gains per conservation dollar by undertaking highly visible conservation projects that may not yield the greatest environmental gain. Such distortion between the private and social objectives could generate an inefficient mix of conservation projects unless corrected by policy that treats preferentially inconspicuous conservation projects.

We are unaware of any prior research that has either developed a theory of green signaling or empirically tested for conspicuous conservation effects in any context. This paper thus contributes to the literature by introducing a modern variant of conspicuous consumption, developing a theory of green signaling, testing for a conspicuous conservation effect in automobile purchases, and estimating the willingness to pay for green status afforded by Prius ownership. This paper proceeds in Section 2 with a brief review of the theories related to conspicuous consumption and green markets in order to motivate the concept of conspicuous conservation. The self-interested motivations for private provision of public goods is also related to the vast literature on altruism. The theory of green signaling is developed in Section 3. Sections 4 and 5 present empirical methods and data, respectively, while Section 6 contains results. Section 7 estimates the willingness to pay for the green halo. Implications and conclusions are discussed in Section 8.

2.2 Status Seeking and Conspicuous Conservation

Economists since Veblen (1899) have relied on status-seeking motives to explain anomalies in consumption behavior like upward sloping individual demands and "non-additive" market demands (Leibenstein 1950; Frank 1985; Glazer and Konrad 1996; Ireland 1998, 2001; Barclay and Willer 2007). Much of this work has focused on ostentation as a signal of affluence and has provided a theoretical basis to understand consumer demand for luxury goods that are functionally equivalent to less costly alternatives. Ireland (1998) and Bernheim (1994), for instance, were concerned with "bizarre" premia for designer fashions and high expenditures on cars.

The pursuit of status is thought to motivate not only demonstrations of extravagance, but also displays of charity and other prosocial behavior. Glazer and Konrad (1996), for instance, theorized that the paradoxical abundance of charity and dearth of anonymous giving could be explained by efforts to signal wealth where consumption behavior is either unobservable or subject to imitation.

Other phenomena suggest that reputation is an important motivation for prosocial behavior and public goods provision. First, standard theory suggests small extrinsic rewards for prosocial behavior should increase their provision. Yet such rewards are documented to crowd-out intrinsic motivations like altruism and warm glow (e.g. Becker 1974; Andreoni 1989, 1990) in a number of contexts (see Frey and Oberholzer-Gee 1997; Frey and Jegen 2001 for surveys). For instance, schoolchildren collected less charity when they were given nominal performance bonuses (Gneezy and Rustichini 2000b), and parents became more delinquent in the timely retrieval of their children from childcare centers when small fines were imposed for late pick-ups (Gneezy and Rustichini 2000a). Second, the provision of prosocial behavior declines when such behavior becomes less observable and increases when it is made public. Funk (2010) showed that the introduction of mail-in balloting in Switzerland did not increase voter participation rates anywhere and reduced voter participation rates in small communities even though the time-inclusive costs of voting declined. Similarly, when one's election participation record is shared with neighbors, his participation increases (Gerber et al. 2008).

As Benabou and Tirole (2006a) explained, such phenomena are consistent with reputational motivations for prosocial behavior. Where prosocial behavior becomes less conspicuous or where it provides a less clear signal of altruism, its provision is expected to decline. Reputational motivations can induce behaviors consistent with pure and impure altruism that are fundamentally self-interested in the traditional sense. Relatedly, phychologists have defined a *competitive altruist* as one who contributes to the public good in order to attain status that can generate economic rewards (Hawkes et al. 1993; Roberts 1998; Hardy and Van Vugt 2006; Barclay and Willer 2007; Van Vugt et al. 2007).

As cultural norms change, it is natural to consider that the personal characteristics for which society confers status may change, and, so, too, the behaviors that individuals exhibit to their peers. Thus, as preferences for environmental protection grow stronger and more prevalent, a homeowner may wish to conceal the extravagance of his home or to draw attention to his fuel-efficient car. A number of studies have documented the importance of social norms in motivating conservation (e.g., Griskevicius et al. 2007; Griskevicius et al. 2010; Allcott 2009; Ayres et al. 2009; and Goldstein et al. 2008). While a small, energy efficient home or automobile certainly provides private benefits to the owner, it also jointly provides a public good by reducing the environmental impact of consumption relative to some benchmark (Kotchen 2006). Consumption of these "impure public goods," then, constitutes prosocial behavior that can generate status. As is true with conspicuous consumption, however, one's peers are likely only to award status if an inference about one's characteristics can be made on the basis of one's signaling. The following section considers the revelation problem and demonstrates that consumption of conspicuous green products can effectively signal one's green type.

2.3 Theoretical Foundations

Intrinsic motivation may explain positive willingness to pay for green product characteristics. But it does not explain why Prius market share increases disproportionately in greener communities. Much as the paucity of anonymous charity observed by Glazer and Konrad (1996) suggested the presence of status-seeking motives, so too does the relative success of highly visible green investments demand an explanation other than conventional altruism. We propose green signaling an alternative.

The success of green signaling hinges on two conditions. First is the observability of costly conservation effort, which may be reflected by willingness to pay premia for green product characteristics or by willingness to accept lower quality for products that generate less environmental damage in production or end-use than conventional products. Second is partial or full revelation of types through signaling. In wealth signaling models, consumption of luxury items permits revelation because declining rates of marginal substitution make high expenditures on ostentation (at the expense of other consumption) more tolerable for the affluent (Bernheim 1994). In the canonical Spence (1973) model of job market signaling, the benefits of signaling are homogeneous across types, but lower education costs to high ability types permit a separating equilibrium with full revelation of types.

Likewise, in a model of green signaling, tolerance of price premia or lower objective quality for green goods is increasing in the strength of preferences for the environment. One who values environmental protection gains utility from the aggregate provision of conservation or environmental protection (due to altruism) and likely also from *his* contribution to conservation (due to impure altruism). Thus, the net benefits of consuming a green product are greater for this type than for one who cares little about the environment, independent of any signalling benefit. Intuitively, one who derives utility from reductions in greenhouse gas emissions will sooner settle for the diminished quality of a green product relative to a comparably priced conventional product than one who is indifferent to environmental protection.

Moreover, the work of Akerlof and Kranton (2000, 2010) suggests that a "green" individual who undertakes green signaling will gain utility from identity conformance whereas a "brown" individual who signals green will suffer utility loss from identity deviance.[3] Distinct from wealth signaling and job-market signaling models, then, the cost of sending the green signal does not vary across types. The price premium for a green product is the same for greens and browns alike. However, the benefits of green signaling are greater for a green than for a brown because of intrinsic motivations. The wedge between net benefits of signaling to green and brown types can be sufficient to generate separation and full or partial revelation of types, as in Spence (1973).

To formalize this concept, define one receiver called Society (S) and two types of consumers (signalers), a green type (g) and a brown type (b). Among each type of signaler, there is a distribution of preferences for each of two brands of cars, T and H. Brand T produces two cars, one that is green, P, and one that is a standard car, B^T. Brand H likewise produces two cars, one green, C, and one brown, B^H. Assume that all vehicles provide the same direct private benefit called "driving services" and that the two green cars provide equal public benefits from reduced impact to the environment. Only P is conspicuously green, however. The green cars have the same price, which is greater than the price of standard cars, which also share a common price.[4] Let c denote the price premium for either green car.

In addition to a common utility from driving services, all consumers have utility, ID,

[3]Akerlof and Kranton (2000, 2010) developed a theory of individuals who self-select into social categories that encompass ideals of how one should behave. They defined the utility of individuals as increasing in their conformance to the norms of their chosen identities and decreasing in deviations from those norms. Identity, they argued, explains persistent gender biases in the workplace, like the over-representation of women in nursing and men in firefighting and other social phenomena.

[4]For each car, it is assumed there is one exogenous price.

over identity, gaining utility from identity conformance and losing utility from identity deviance (Akerlof and Kranton 2000, 2010). Specifically, assume $ID_{P,C}(g) = ID_B(b) > 0 > ID_B(g) = ID_{P,H}(b)$, where subscripts denote vehicle choice and parentheticals indicate the type of consumer. Thus, a brown who purchases a standard car derives the same identity utility as a green who purchases either green car, and a brown who consumes green incurs the same loss of utility as a green who consumes brown. Consumers of all types also derive utility from brand loyalty, such that consumers loyal to brand T (H) gain utility $l > 0$ for choosing a vehicle made by T (H) and zero utility from loyalty otherwise. Green consumers also derive utility w from their *individual* contributions to the public good, i.e. from choosing P or H.[5] Browns do not derive utility from w regardless of which vehicle they choose. That is, $w_{P,C}(g) > w_{P,C}(b) = w_B(g) = w_B(b) = 0$.

Finally, consumers have utility over status, which is conferred by S to green types. Specifically, assume that it is optimal for Society to confer to greens the status, $s^*_G(\hat{\theta})$, where $\hat{\theta}$ is a summary measure of the distribution of environmental preferences within a community and $\frac{\partial s^*_P}{\partial \hat{\theta}} > 0$, so that the status benefit is increasing in community greenness.[6] Society does not confer status to browns.

Perfect revelation of types requires that all browns choose B and that all greens choose P. A brown loyal to brand H will never masquerade as green by choosing P whenever a brown loyal to T prefers not to masquerade. Similarly, a brand T-loyal green will never masquerade as brown whenever a brand H-loyal green does not masquerade as brown. Thus, perfect revelation requires only that the incentive compatibility constraints of g^H, an H-loyal green, and b^T, a T-loyal brown, are satisfied. Moreover, a brown would never signal C because he incurs the cost of the green car price premium with no offsetting benefit. Therefore, the incentive compatibility constraints (ICC) of g^H and b^T are, respectively:

$$\begin{array}{c} s^*_P(\hat{\theta}) > l \\ s^*_P(\hat{\theta}) + ID_P(g) + w_G(g) - c_P > ID_B(g) + l \end{array}, \text{ and} \tag{2.1}$$

$$ID_B(b) > s^*_P(\hat{\theta}) + ID_P(b) - c_P \tag{2.2}$$

The first condition in (2.1) requires that a green loyal to H obtains greater utility from P than C. The second condition holds that he also obtains greater utility from P than B. Hence, he prefers P to the alternatives. If ICC(g^H) is satisfied, then ICC(g^T) is trivially satisfied as a green loyal to T gains loyalty utility from choosing P. The condition in (2.2) establishes that a brown loyal to T obtains greater utility from choosing B^T than from choosing P. Choosing B^H would make him unequivically worse off than choosing

[5]Green signalers may also derive utility from the aggregate provision of public goods. However, if they exclusively derive utility from the aggregate provision of the public good, then the standard crowd-out result may obtain (see Andreoni 1990).

[6]In the following empirical application, $\hat{\theta}$ is represented as a [0,1] index based upon voter party registration (Colorado) or party vote in the 2008 Presidential election (Washington).

B^T because of the loss of loyalty utility. Similarly, choosing C is dominated by choosing P because of the status benefit from P, in addition to the loyalty benefit.

If upon observing P, Society assumes the consumer is a green type with probability one and that any consumer that doesn't signal P is a brown with probability one (because C is indistinguishable from B), then a semi-separating perfect Bayesian equilibrium exists iff:

$$s_P^*(\hat{\theta}) < (ID_B(b) - ID_P(b)) + c \text{ and } s_P^*(\hat{\theta}) + w > c. \tag{2.3}$$

So long as the status benefit is less than the sum of the utility from identity conformance (net of identity loss from identity deviance) and the cost savings of choosing a standard car, a brown consumer does not masquerade. And if the status benefit and warm glow are greater than the green car premium, then a green always chooses a green car, though not every green neccesarily signals P.

In such an equilibrium, the standard car obtains in each community a market share equal to the share of browns in the community. The shares of P and C in each community are functions of the distributions of environmental and brand preferences as loyalty may induce brand-H-loyal greens to choose C rather than P. Because the status benefit is increasing in community greenness, the market share of P is non-decreasing in community greenness. Moreover, with constant utility from loyalty, the market share of P is increasing in community greenness if there exists a threshold level of community greenness, $\bar{\hat{\theta}}$, for which the status benefit causes a g^H consumer to switch from C to P. Then, at $\bar{\hat{\theta}}$, P market share increases discretely as both g^T and g^H types choose P.

The semi-separating equilibrium just described is not the only equilibrium that can obtain in this framework. For instance, a separating equilibrium would obtain if in addition to (2.3), $s_P^*(\hat{\theta}) > l$. For a sufficiently small status benefit or sufficiently high green car price premium, a pooling equilibrium could exist with all types choosing a brown car. Moreover, for the sake of parsimony, the foregoing model has assumed a two-point distribution for utility from each loyalty and status. If, instead, greater consumer heterogeneity were introduced by allowing a continuous distribution of brand loyalty and status consciousness among consumers, then multiple (infinitely many) semi-separating equilibria could obtain. With a continuous distribution of loyalty, for instance, P's market share could exhibit a continuous increase in community greenness as the g^H types increasingly switch from C to P, starting with the least loyal g^Hs. Likewise, with a continuous distribution of status preferences, the most status-conscious brown may masquerade as green while the least-status conscious green may choose the brown (standard) car.

If, instead, there were a continuum of types from brown to green, then in each community there would exist a threshold level of individual greenness, $\bar{\theta}_i$ at which an individual is just indifferent between signaling green or not. If, conditional on beliefs about the strategies of each type, Society's best response were to award status, $s_P^*(\hat{\theta})$, if $\bar{\theta}_i \geq \hat{\theta}$, then a semi-separating equilibrium would obtain in which those with individual greenness greater than $\hat{\theta}$ either choosing P to signal green or C to retain loyalty benefits. One could also imgaine Society's status offer to be a function of the expected greenness of those who signal green, in which case the status benefit would be increasing in community greenness at a given rate but also decreasing in community greenness as higher community greenness attracts less individually green consumers to signal green.

Finally, a more general model with a continuum of green types and a continuum of signaling actions by consumers can generate a separating equilibrium as Glazer and Konrad (1996) established in a similar but distinct setting. For any given community, a fully revealing equilibrium exists if (1) one's optimal green investment is a strictly monotonic function of his individual greenness, (2) beliefs about his type are correct, and (3) utility is increasing at a decreasing rate in status and warm glow net of green investment (see Glazer and Konrad (1996) for a proof). An implication of such an equilibrium is that greater levels of green investment are observed in greener communities.

While much of the signaling literature emphasizes the importance of the distribution of types and preferences among potential *signalers*, the foregoing model and extensions highlight the importance of the distribution of preferences among *receivers*. The model yields the proposition that the market share of a conspicuous green product is increasing with community greenness relative to comparably green inconspicuous products. An empirical test of that prediction is the subject of the following three sections.

2.4 Empirical Methods

In order to test empirically for the presence of status seeking in vehicle choice and to estimate willingness to pay for the "green halo" associated with hybrid vehicle ownership, we exploit spatial variation in vehicle model market share and in preferences for conservation and environmental protection in the states of Colorado and Washington.[7] Important to the empirical strategy of this paper is the insight that the value of the Prius signal, i.e., the green halo, is increasing in the greenness of the community in which the owner resides. The benefits of signaling one's green type should be greater the more one's peers are concerned about the environment. Kahn (2007) documented the clustering of Prius and Hummer ownership and showed that communities in California with more registered Green or Democrat party members are home to more Priuses. Communities with more Republicans have more Hummers.

Were there no status-seeking motivations for hybrid demand or were the Prius indistinct relative to other green cars, we would expect to see ownership patterns like those described by Khan, with hybrid cars enjoying greater market share in green communities. But a similar pattern should exist for all hybrid models with comparable greenness; their market shares should covary with measures of community environmentalism. If, instead, Prius owners also derive utility from the halo effect that is unique to a Prius, then, conditional on vehicle characteristics, the greater value of the halo in greener communities

[7]These states were chosen because of the availability and affordability of vehicle registration data and because of spatial variation in political preferences, which we use to measure community greenness.

should cause Prius ownership to increase disproportionately in those areas relative to other hybrids like the Civic Hybrid.

Following Kahn (2007) and Kahn and Vaughn (2009), we measure the relative greenness of communities using election data, particularly support for the Democrat party. As has been observed in a number of settings, political ideology is highly correlated with environmental ideology. Republican household energy consumption is less responsive to peer comparisons and may increase, whereas Democrat households decrease consumption on average (Costa and Kahn 2010); Households in communities with high Democratic and Green party registrations pay higher premia for homes with solar panels (Dastrop et al. 2010); per capita energy consumption has been trending upwards in majority Republican states but relatively flat in majority Democrat states; and public opinion surveys show Republicans are more than three times as likely as Democrats to think that the seriousness of global warming is exaggerated in the news media (Loewenstein 2009).

We focus on Democrat Party electoral data for this analysis. Green party affiliation could also be an important indicator of the strength and prevalence of preferences for environmental protection. Strategic voting, however, limits the Green party share of the electorate. Many environmentalists participate in Democrat politics to ensure their votes have the greatest impact on election and primary election outcomes.[8]

We define markets at the zip code level, the smallest geographical area for which car share data are available. We employ a reduced-form fixed-effects model that is effectively a regression-based difference in difference (DD) model with partial treatment. To motivate the full DD model, we first propose a two-by-two DD model in which we consider the market shares for the Prius and the Civic Hybrid in a green market and in a brown market. This 2x2 model assumes that the unique design of the Prius makes it a purchase that signals green status and that the Civic Hybrid is a perfect control for all attributes of the Prius except that it lacks a design that enables the owner to signal his green type. Its use further assumes green and brown markets are identical apart from preferences over the environment. Environmental preferences can be thought of as the policy parameter in treatment effects models with partial treatment. Then the DD estimate of the conspicuous conservation effect on market shares is given by:

$$\hat{\beta} = (S_G^P - S_B^P) - (S_G^C - S_B^C),$$

where S is market share and superscripts P and C denote Prius and Civic Hybrid, respectively, and subscripts G and B denote green and brown markets respectively.

Accepting the difficulty of identifying markets that are otherwise identical apart from greenness, and in order to exploit observations across a number of markets, we augment

[8]It is reasonable to assume that this measure of market greenness is exogenous to Prius market share. Indeed, while Kahn and Vaughn (2009) show that environmentally-minded individuals cluster around environmental amenities and mass transit access points, it is unlikely that individuals move to a community because they own a Prius or that Prius ownership induces individuals to align with the Democrat party.

the 2x2 model to consider a regression-based 2xN model, incorporating all zip codes (in the N-dimension), and using market fixed effects to condition on market characteristics other than the policy variable. We estimate:

$$S_{ij} = \xi V_i + \gamma D_j + \beta D_j * VOTE_i + \varepsilon_{ij} \tag{2.4}$$

where, for $j \in \{\text{Prius}, \text{Civic}\}$, the V_i are market fixed effects, D_j is a Prius indicator, $VOTE_i$ is a measure of the greenness of the market (i.e., the strength of the policy), and ε_{ij} is an idiosyncratic error. The coefficient of interest is β, which represents the change in Prius market share due to a one-unit change in $VOTE$.

Finally, aware that the Civic Hybrid, while similar to the Prius, is an imperfect control for attributes of the Prius apart from its unique design, we specify a full model that relies on the "common support" or "overlap" assumption on the distribution of covariates among treated and untreated groups. The empirical model incorporates many car models and controls for model heterogeneity with model fixed effects and for heterogeneous effects of green car characteristics according to market preferences for the environment by interacting a measure of a model's greenness, $GREEN_j$, with $VOTE_i$. This serves to control for the Prius attributes apart from the unique design that could cause its demand to increase disproportionately in green markets relative to other models. Specifically, we consider:

$$s_{ij} = \delta_j D_j + \xi_i V_i + \gamma GREEN_j * VOTE_i + \beta PRIUS_j * VOTE_i + \varepsilon_{ij}, \tag{2.5}$$

where D_j now represents a product fixed effect and interest again centers on the estimate of β.

Car manufacturer and dealer marketing effort may be correlated with the greenness of the region. In particular, one may be concerned that Toyota and Toyota dealers market Prius more heavily in green communities. Based on conversations with Toyota marketing executives, we believe these concerns are minimal. Toyota marketing is undertaken at the national, regional, and dealer levels. Colorado and Washington are each fully encompassed within their respective marketing regions, so regional marketing cannot confound. In addition, Toyota indicated that Prius success in specific markets, like Portland, Oregon, is largely independent of marketing effort. Data on model-specific marketing by dealers is unavailable. Nevertheless, in order to control for such effects, we defined dealer marketing areas by assigning each zip code to the nearest Toyota dealership using euclidean distance from the zip code centroid. We then included separate fixed effects for each product in each marketing area by interacting the product dummies with dealer dummies.[9]

We address concerns related to omitted variables bias arising from variation in the relative demand for different vehicle attributes by different demographic groups in two ways. First, because marketing data indicate that hybrid car ownership is positively

[9]This analysis includes all 19 Toyota dealerships in Colorado, i.e., 19 marketing areas. In Washington, we combined marketing areas for dealerships in the same cities or, in some instances, for proximal dealerships in nearby cities in order to improve the tractability of the econometric model. From the 30 dealerships in Washington, we created 18 marketing areas.

correlated with income and education, which are themselves highly correlated, and because both may be correlated with Democratic vote share, we allow for median household income to have a unique effect on the market share for each product. We do this by interacting the product dummies with median household income. In addition, while it is unclear whether the Toyota Prius should be in relatively higher demand in suburban areas or in cities, the high concentration of Democrat vote share in urban areas in our data suggests population density may also confound the conspicuous conservation effect. Therefore, we also allow population density to have a unique effect on the market share of each product by interacting product dummies with population density.

As to vehicle prices, our vehicle data are market shares based upon vehicle registrations, and, thus, reflect local vehicle prices over a period of years, none of which we observe. However, omission of price information is important to the analysis only if model prices exhibited considerable variation at the zip code level. Such variation is unlikely, however, because consumers can readily arbitrage such price differences by shopping outside their communities. They have clear incentive to arbitrage for a major purchase, such as a car. Nonetheless, car prices may be uniformly higher, say, in rural areas, reflecting nonzero arbitrage costs and some buyers' desire to buy local, but such an effect should be captured in the dealer area fixed effect and would not influence the distribution of shares among vehicles. The main issue for this analysis is if dealers were able to price discriminate, they would charge more for Prius in green communities. Although we think consumer arbitrage and dealer competition markedly limit such opportunities, it is worth noting that the effect, if it exists, would attenuate market shares of Prius in green communities, and bias the results from our model in the direction of finding no conspicuous conservation effect.

2.5 Data

Data on all registered vehicles in the states of Colorado and Washington were obtained from the states' respective vehicle licensing departments. For Colorado, 3.9 million vehicle identification number (VIN) records were matched to one of 511 5-digit zip codes. For Washington, 4.2 million VIN records were matched to one of 412 5-digit zip codes. A third-party, proprietary data set was used to decode the VINs in order to obtain make, model, and year of the car in each vehicle record, as well as the other characteristics used in this analysis, including the U.S. Environmental Protection Agency's fuel economy ratings. We defined products by iteration of make and model (i.e., model generation). In order to reduce dimensionality, we did not treat each model year as a distinct product but rather grouped models across years so long as model design was unchanged.[10]

We generated the average characteristics of each "product" and dropped products with

[10]For instance, the 2010 Toyota Camry is the sixth generation of Camry ever produced. The sixth generation was first introduced in 2007. We group Toyota Camry's from model years 2007-2010 as one product.

Manufacturer Suggested Retail Price (MSRP) greater than $100,000. In order to further reduce dimensionality, we restricted attention to all models manufactured by Acura, Cadillac, Chevrolet, Ford, General Motors, Honda, Lexus, Mercury, and Toyota since 2002–356 products in total and 27% and 25% of all cars registered in Colorado and Washington, respectively. These brands manufactured all but a few of the hybrid vehicle models available in the U.S. by 2010. Our measure of market greenness in Colorado is Democrat party registration data obtained from the Colorado Secretary of State. Washington state has an open primary system, so voters do not register with political parties. Therefore, our measure of market greenness in Washington is vote share for Democrat candidate Barak Obama in the 2008 Presidential election.

Green car ratings are used to condition for car characteristics that could have a heterogeneous effect on market share that varies with market greenness. For this rating, we used the American Council for an Energy Efficient Economy (ACEEE) "Green Book," which grades all models in the U.S. on a 100-point curve according to their environmental impacts, with tailpipe emissions ratings, fuel economy, and curb weight being the most important inputs into the grades.[11]

Summary statistics are reported in Table 2.1. Figure 1 shows Democrat party share of registered voters in Colorado by zip code along with Prius locations (Each green dot denotes five Priuses). Likewise, Figure 2 shows 2008 Presidential vote share for the Democrat party candidate by zip code in Washington along with Prius locations. Consistent with the findings of Kahn (2007), Priuses are clustered in the more Democrat areas.

2.6 Results

Results from the estimation of (2.4), the '2 x N' model, are reported in Table 2.2. We report results for Colorado using Democrat party registration as a measure of community greenness in column 1 and Green party registration as a measure of community greenness in column 2. We report both measures to demonstrate robustness of our results to alternative measures of community greenness. However, as described in Section 4, Democrat party measures are thought to be more reliable because of strategic voting concerns. The results from estimation of (2.5) are reported in Table 2.3 using two specifications. The top panel reports results from the model that includes only product-specific marketing area effects. The bottom panel reports our preferred estimates, which additionally include product-specific median-household-income and product-specific population-density effects in order to control for potential confounds. In all instances, we report the point estimate of the coefficient of interest, the standard error in parentheses, and the estimated elasticity of Prius market share with respect to Democrat (or Green) vote share in brackets.

The coefficient of interest on the interaction of the Prius indicator and the vote-share variable is positive and significant at the 99% level in each instance. These estimates

[11] For more information about ACEEE Greenbook ratings, see http://www.greenercars.org/greenbook_method.htm

Table 2.1: Summary Statistics - Vehicle Fleets in Colorado and Washington

	Colorado		Washington	
	Mean	Std. Dev.	Mean	Std. Dev.
Hwy MPG	22.27	7.61	23.54	8.11
City MPG	17.17	7.09	18.12	7.74
Green Score	31.05	7.45	33.61	7.45
MSRP	30,733	10,696	30,596	11,715
Dem. Share	0.30	0.11	0.54	0.14
Pop./sq. mi.	1,513	2,405	1,724	2,716
Med. Income ($)	48,506	14,538	45,303	13,339

Figure 2.1: Prius ownership and Democrat party share of registered voters in Colorado

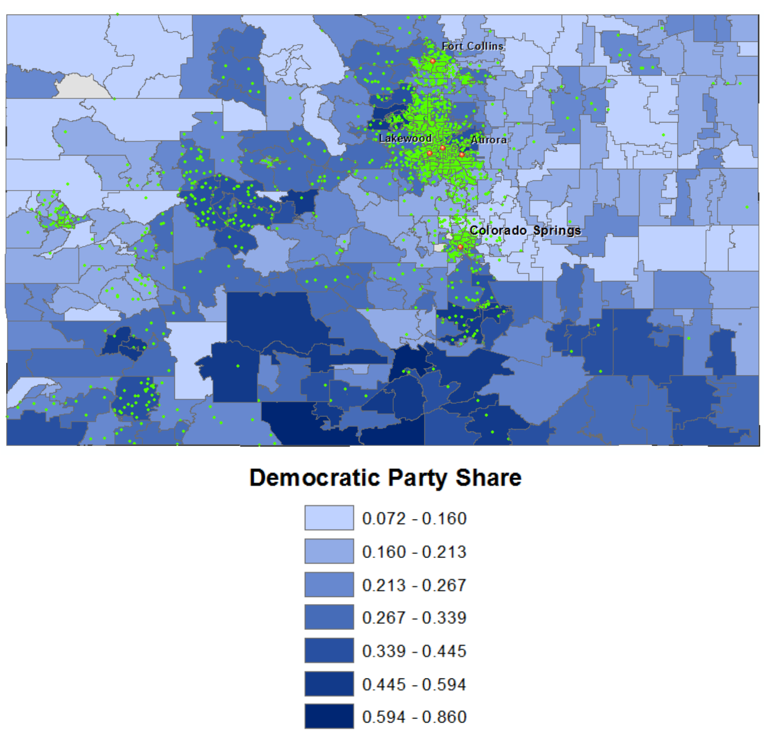

One green dot denotes 5 Priuses.

suggest economically significant conspicuous conservation effects on Prius market share. The elasticity of Prius market share with respect to Democrat share is 0.33 and 0.1 in Colorado and Washington, respectively, as reported in Table 2.3. The change in Prius market share induced by a one percentage point increase from the mean in Democrat share is 0.000052 and 0.000062 in Colorado and Washington, respectively.

A series of falsification tests were conducted by sequentially replacing the Prius indicator and vote share interactions in (2.5) with (1) a Civic Hybrid indicator and Civic Hybrid-vote-share interaction and (2) a Camry Hybrid indicator and Camry Hybrid-vote-share interaction. Because the Civic Hybrid and Camry Hybrid do not have unique designs, we expect the coefficient on theses interactions to be non-positive. If market share for these models were independent of Prius market share, we would expect the coefficients to be insignificant. As reported in Table 2.4, however, the estimated coefficients on the interaction variables are negative and statistically significant, indicating the absence of a conspicuous conservation effect, and, moreover, that Prius is likely being substituted in place of the Civic Hybrid and Camry Hybrid in these areas due to the value of the green halo signaled uniquely by the Prius.

An alternative explanation for the disproportionate increase in Prius market in green communities is a "social contagion" effect whereby an individual's adoption behavior is a function of the adoption behavior of others. However, the mechanisms of social contagion are either consistent with the theory presented in this paper or unlikely to have caused the observed variation in Prius ownership. Van den Bulte and Lilien (2001) defined four causal mechanisms of social contagion: normative pressures, competitive concerns, network effects, and information transfer. The first is consistent with the status seeking and identity motivations that we described as inducing demand for the Prius halo. The second mechanism, in the present context, is equivalent to competitive altruism, and therefore also subsumed within the theory developed in this paper. With respect to the third mechanism, there is no reason to expect network effects associated with Prius ownership just as there are not likely to be network effects associated with adoption of other conventional or hybrid car models.[12] Finally, it is unlikely that information transfer caused a social contagion in Prius ownership because the Prius was well marketed, discussed frequently in popular media, and has captured significant market share. Furthermore, specific information about Prius technology was shown to be unimportant to Prius buyers (Heffner et al. 2007). In addition, information transfer would confound this analysis only to the extent that information transfer is heterogeneous across hybrid cars, i.e. if information transfer is more important to adoption of Prius than other models, including other hybrids.

[12]There may be network effects associated with plug-in hybrids, but they were not introduced until the 2012 model year and, thus, are not reflected in this data.

Figure 2.2: Prius ownership and Democrat candidate vote share in Washington

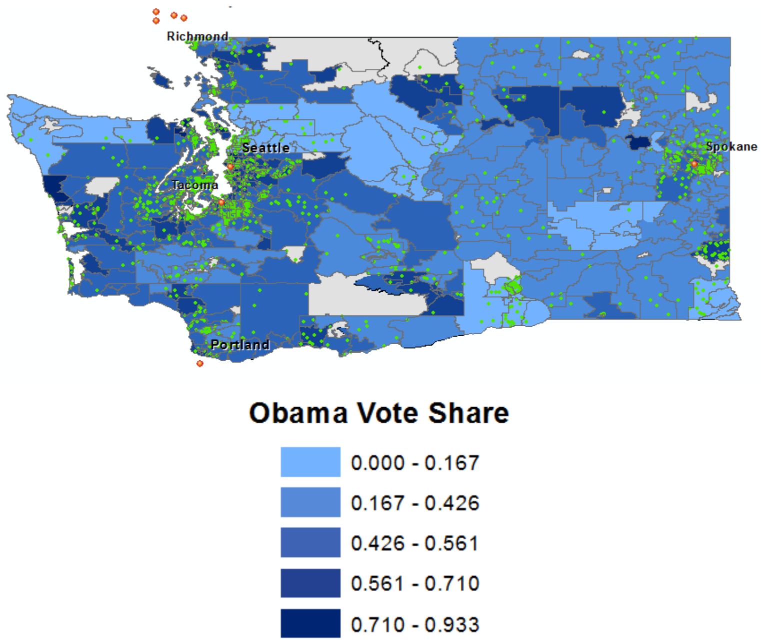

One green dot denotes 5 Priuses.

Table 2.2: Conspicuous Conserv. Effect on Prius Market Share: '2 x N' Model for Colorado

Dependent Variable: Product Market Share		
	(1) Democrat	(2) Green
PRIUS*VOTE	0.0094*** (0.0007) [0.59]	1.0139*** (0.1163) [0.42]

Robust standard errors in parentheses
Elasticity of Prius share with respect to Democrat share in brackets
*** p<0.01, ** p<0.05, * p<0.1

Table 2.3: Conspicuous Conserv. Effect on Prius Market Share: Full Model

Dependent Variable: Product Market Share		
	(1) Colorado	(2) Washington
Product-specific Marketing Area Effects		
PRIUS*VOTE	0.0052***	0.0113***
	(0.0024)	(0.0023)
	[0.33]	[0.18]
Product-specific Marketing, Income, and Population Density Effects		
PRIUS*VOTE	0.0052***	0.0062***
	(0.0014)	(0.0026)
	[0.33]	[0.10]

Robust standard errors in parentheses
Elasticity of Prius share with respect to Democrat share in brackets
*** p<0.01, ** p<0.05, * p<0.1

Table 2.4: Conspicuous Conserv. Effect on Civic Hybrid and Camry Hybrid Market Shares

Dependent Variable: Product Market Share

	(1) Colorado	(2) Washington
	Honda Civic Hybrid	
CIVIC_HYB*VOTE	-0.0046***	-0.0047***
	(0.0009)	(0.0013)
	[-87.3]	[-90.4]
	Toyota Camry Hybrid	
CAMRY_HYB*VOTE	-0.0036***	0.0028*
	(0.0012)	(0.0014)
	[-45.5]	[-44.4]

Robust standard errors in parentheses
Mean conspicuous consumption effect as percent of share in brackets
*** p<0.01, ** p<0.05, * p<0.1

2.7 Estimating benefits from the Green Halo

Absent sales price data, econometrics alone cannot price the Prius halo. A parametric model informed by the foregoing econometrics, however, can provide estimates of consumer willingness to pay for the status signal afforded by Prius ownership. A locally linear Prius demand is assumed and the conspicuous conservation effect is treated as a (parallel) demand shifter.[13] For any community in the data, it is first determined what magnitude of Prius demand would, for given price, generate an equilibrium market share equal to the model estimate of actual market share. Then, the hypothetical market share without the green halo is estimated by subtracting the estimated conspicuous conservation effect from the observed share, generating a hypothetical shift back in demand for a Prius without the green halo.

This approach is illustrated in figure 2.3. D_1 represents Prius demand in a given community measured in market share assuming the Prius has all of its existing features except the halo-generating unique design. D_2^B represents actual Prius demand (including the green halo) if the community is "brown". A green-halo effect is present, but it is less valuable than if the community is "green", as represented by a demand shift to D_2^G. As previously argued, Toyota must price the Prius nationally, and that price is represented in the figure as P^*. Given P^*, the market share for Prius without the green halo is s_1. Because of the conspicuous conservation effect induced by the signaling capacity of Prius, the Prius market share in the brown community is s_2^B; it is s_2^G in the green community. The incremental willingness to pay in the brown and green communities is measured by the vertical shift in demand, the amounts $W^B - P^*$ and $W^G - P^*$, respectively.

To quantify the approach set forth in figure 2.3, we use our econometric estimates of the magnitude of the conspicuous conservation demand shifter from the previous section and fit locally linear demand curves using estimated price elasticities of demand for individual vehicle models from the literature. To our knowledge, there are no price elasticity estimates in the literature for the Prius or for individual hybrid models. Therefore, we relied on estimated elasticities for similar models. Specifically, Mannering and Hani (1985) estimated a Toyota Corolla elasticity of 1.59, while Mannering and Winston (1985) estimated a Corolla price elasticity of 1.7. Honda Accord elasticities were estimated to be 2.0 and 4.8 by Mannering and Hani (1985) and Berry et al. (1995), respectively. Because of the uniqueness of the Prius, its price elasticity likely falls in the low end of this range. Thus, although we present willingness to pay estimates using Berry's elasticity estimate of 4.8 to form a lower bound on willingness to pay, we believe the cluster of estimates in the 1.6 - 2.0 range likely generate the more realistic set of values for willingness to pay.

With these elasticity estimates and the estimated magnitude of the conspicuous conservation effect, the willingness to pay for the Prius halo is estimated. Using preferred specifications from the bottom panel of Table 2.3, the willingness to pay for the green

[13] Assumption of a parallel shift is of no great consequence for the analysis here. Our data do not allow estimation of slope changes in demand because we are unable to construct interaction terms with a price variable.

halo in a Colorado community with mean Democrat party registration share (0.30) is estimated to be between \$1,402.84 and \$4,208.53, depending upon the demand elasticity estimate chosen to parameterize the model. In Washington willingness to pay for the green halo in a community with mean Democrat vote share (0.53) is estimated to fall in the range \$430.45 and \$1,291.34. These results are reported in Table 2.5.

Spatial variation in green intensity implies that the willingness to pay for the Prius halo is greater in some communities than in others, as figure 2.3 illustrates. Table 2.6 reports for each of the three price elasticities the difference from state-mean willingness to pay for two green cities with high Democrat registration shares in Colorado (Denver and Boulder) and two brown cities with low Democrat registration shares in Colorado (Longmont and Loveland). Similar estimates are reported for two green cities with high Obama vote share in Washington (Seattle and Spokane) and two brown cities with low Obama vote share (Yakima and Richland). The table also reports for each city the difference from statewide mean willingness to pay as a percent of statewide mean willingness to pay. In Denver, for instance, the value of the green signal is estimated to be more than twice the state mean. The signal is worth as much as \$5,000 more in Denver than the average Colorado community. Likewise, it is worth as much as \$3,400 more in Boulder. In less green communities, like Longmont and Loveland, however, the value of the green signal is worth as much as \$300 and \$700 less than the state mean, respectively. In Seattle, the signal is worth 54% more than the Washington state mean (as much as \$698 more), whereas in Richland, it is worth 27% less than the state mean (as much as \$350 less).

The benefit to Toyota from designing the Prius to convey the green halo is not easy to estimate because Toyota's pricing decisions likely involve fleet considerations and constraints imposed by CAFE standards. For example, given the high fuel economy of the Prius, Toyota could rationally choose to sacrifice profits from the Prius by setting price below a monopoly optimum in order to expand market share and improve its CAFE rating. However, back-of-the envelope calculations can provide insight into how much of a price premium Toyota is likely to capture on the Prius due to the green halo. Such a calculation depends upon an estimate of the price of a Prius without signaling value. Because Toyota does not manufacture such a counterfactual Prius, one is created by adding the estimated cost of Toyota's hybrid system to the price of the most comparable conventional car in the Toyota fleet. According to Consumer Reports, the Toyota Corolla is "the closest available" non-hybrid alternative to the Prius, considering "all factors, including performance, safety, and features" (Consumer Reports 2011). Thus, the price of a hypothetical Prius without signaling value is calculated as the price of a hypothetical hybrid version of the Toyota Corolla.[14] The price of the most comparable Corolla, the Corolla LE, was \$15,615 in 2008. In order to estimate the price of a Corolla with a hybrid system, we infer the cost of Toyota's hybrid system by comparing the price of the Camry Hybrid with a comparably equipped non-hybrid version of the Camry. Apart from the hybrid system, a 2008 Camry Hybrid priced at \$25,200, was comparable to the Camry LE, priced at \$20,025, and to

[14]Toyota has not offered a hybrid version of the Corolla.

Figure 2.3: Prius demand with and without the Conspicuous Conservation Effect

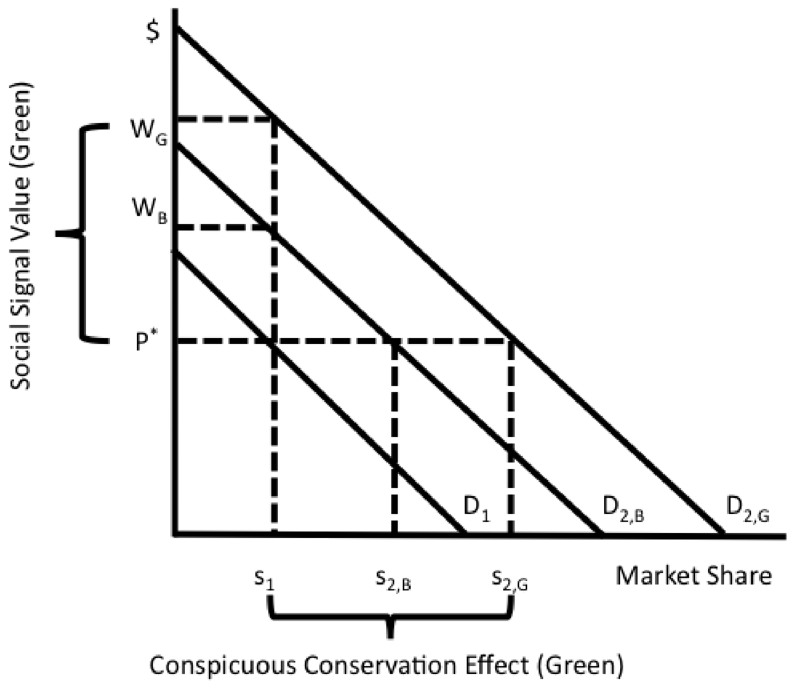

Table 2.5: Estimated Mean Willingness to Pay for the Prius Halo (in dollars)

	Price Elasticity		
Percent Change in Share	-1.6	-2.0	-4.8
10.1 (WA)	1,291.34 [229.11, 2,353.57]	1,033.07 [183.29, 1,882.85]	430.45 [76.37, 784.52]
32.9 (CO)	4,208.53 [2,023.33, 6,393.73]	3,366.83 [1,618.67, 5,114.99]	1,402.84 [674.44, 2,131,24]

95% confidence interval is reported in brackets holding elasticity constant

Table 2.6: Willingness to pay for the Prius Halo (in dollars) Relative to State Mean

City	% Diff. in WTP relative to state mean	Price Elasticity -1.6	-2.0	-4.8
		Colorado		
Denver	120.75%	5,082.13	4,065.71	1,694.04
Boulder	81.22%	3,418.13	2,734.51	1,138.38
Longmont	-7.74%	-325.87	-260.69	-108.62
Loveland	-17.63%	-741.86	-593.49	-247.29
		Washington		
Seattle	54.10%	698.66	558.92	232.89
Spokane	22.54%	291.07	232.85	97.02
Yakima	-7.17%	-92.55	-74.03	-30.85
Richland	-27.59%	-356.28	-285.02	-118.76

the Camry SE, priced at \$21,240. Assuming Toyota's margin is the same across Camry trims, then the 2008 cost of the hybrid system was between \$3,960 and \$5,175. The price of a hypothetical Corolla hybrid, then, would be \$19,575 to \$20,790, assuming a constant absolute markup. A rough estimate of the markup extracted by Toyota due to the Prius' green halo is the difference in the price of a 2008 Prius (\$21,500) and the hybrid-inclusive price of the Corolla, which is between \$710 and \$1,925.

Cognizant of the likely complexity of Toyota's pricing decision for the Prius, this estimated Prius halo markup is consistent with estimates of consumer willingness to pay earlier in this section. A monopoly seller and consumers share the benefits of a demand shift according to the functional forms of demand and the seller's marginal cost. In the benchmark case of linear demand and constant marginal cost, the seller captures in higher prices exactly half of the incremental willingness to pay.

2.8 Conclusion

This paper introduces conspicuous conservation as a modern variant of wealth signaling that is intended to garner status. In a theoretical model of green signaling, intrinsic motivations make purchases of impure environmental public goods more valuable to those with preferences for environmental protection than those without, allowing for inference about individual characteristics based on consumption behavior. Status, then, is attained by consumption of conspicuous green products and its value is increasing in the strength of environmental preferences one's peers. This phenomenon is identified in the demand for the Toyota Prius using zip-code-level data on vehicle ownership in Colorado and Washington. Empirical results suggest that consumers in Colorado and Washington are willing to pay up to several thousand dollars to signal their environmental bona fides through their car choices. These results, and the theoretical model of green signaling, suggest that status-seeking by environmentalists can increase the private provision of environmental public goods via purchases of impure public goods in the green market. The presence of green signaling motivations also suggests opportunities for firms to price discriminate, differentiate, and gain market share by developing conspicuous green products.

While much of the literature on conspicuous consumption emphasized the wastefulness of spending to signal wealth, *conspicuous conservation* may improve social welfare by moving towards optimal provision of environmental protection, particularly in the presence of market failures that under-value environmental amenities and in the absence of first-best policies. In wealth signaling models characterized by utility over one's consumption in relation to a peer, conspicuous consumption leads to the "Keeping Up with the Joneses" arms race that induces wasteful and costly signaling that does not change the relative position of competitors. All are made worse off (Hopkins and Kornienko 2004). In a similar model of green signaling, however, the competitors may be made worse off while others, and society writ large, may be better off because the competition is over provision of public goods, not wastefulness. Thus, there may exist a role for policy to enhance the conspicuousness of private behaviors that bear on the stock and quality of environmental public goods. Household energy billing records, for instance, are public record in some jurisdictions in the United States but protected from disclosure by law in others. If such information were made public and easily accessible, a once private behavior would become conspicuous and perhaps induce conservation.

The social welfare implications of conspicuous conservation, however, depend upon substitution effects with respect to conservation effort. The social signaling motive can distort private incentives from those of the social planner, generating conservation investment that is individually rational but not social welfare maximizing. For instance, homeowners may over-invest in residential solar power because of its conspicuousness and under-invest in home insulation improvements, energy efficient heating and cooling systems, and window sealing treatments because of the relative inconspicuousness of those investments.

Policy makers, then, should re-evaluate subsidies for conspicuous green products such as green cars and solar panels in order to better align private incentives with behaviors that are in the public interest. Subsidies should be targeted toward inconspicuous conservation in order to achieve a more efficient mix of conservation effort. Moreover, policy makers should be mindful of the potential to crowd out intrinsic motivation with extrinsic rewards like tax-breaks and subsidies. Because conspicuous-conservation goods enable their purchasers to signal their willingness to sacrifice to enhance the environment, the public subsidy of such goods diminishes the value of such goods as social signals. In the context of the green signaling model of this paper, public incentives for green product purchases reduce the likelihood that greens and browns separate, and, therefore they weaken the inferences society can make on observing green signals. Public support for conspicuous green products may, therefore, have the perverse effect of reducing their demand.

Chapter 3

Automatic Bill Payment and Salience Effects: Evidence from Electricity Consumption

3.1 Introduction

In standard theory, the consumption decisions of rational agents are invariant to the salience of product attributes. And yet, if cognitive ability is constrained or if attention is, itself, a scarce resource, then agents may, in the course of decision-making, overweight information that is easily recalled or prominent (Simon 1955, Tversky and Kahneman 1974). Motivated by such insights from psychology, economists have demonstrated that individuals are less responsive to shipping fees than to auction prices on Ebay (Brown et al. 2010; Hossain and Morgan 2006), to rebates for car purchases than to car purchase price (Busse et al. 2006), to taxes excluded from product prices than to taxes included in prices (Chetty and Saez 2005; Chetty et al. 2009), to income tax incentives than to sales tax incentives (Gallagher and Muehlegger 2007), and to earnings statements issued early in the week than to earnings statements issued just before the weekend (DellaVigna and Pollet 2009).

Consumer inattention to less prominent prices explains the persistence of "shrouded" add-on prices (e.g. parking and internet fees associated with hotel room reservations, and transaction and minimum balance charges associated with checking accounts). The proliferation of add-on prices for checked baggage, seat assignments, telephone bookings, and other perquisites of air travel garnered the attention of regulators in 2011. In 2012, they required fare advertisements to include all mandatory taxes and fees and weighed additional rules on the disclosure of optional fees. In short, there is growing evidence that inattention to the less-salient components of product prices generates systematic biases in individuals' consumption decisions that firms may exploit.

Given that the manner in which prices are displayed is potentially important in consumer decision making, this paper considers whether enrollment in automatic bill payment programs affects consumption of goods and services regularly procured by recurring payments, like telecommunications and gas, electric and water utility services. Voluntary automatic bill payment (ABP) programs permit the timely payment of recurring bills via automatic credit and debit card transactions or deposit account withdrawals without requiring individuals to ever view their bills. Historically, such transactions were initiated by individuals who were obligated to view their bills in order to transmit appropriate funds to service providers and avoid penalties for delinquency. Absent certain penalties for inattention to recurring payments, the incentive for consumers to attend to billing records is diminished. Consequently, inattention to accounts serviced by ABP may increase, reducing the price salience of products and services financed by those accounts, and potentially inducing consumption above levels that would be chosen with full attention to price.

The use of ABP grew quickly at the turn of the century, following a broader trend in developed countries of growing reliance on electronic payments, like credit and debit card transactions, and diminishing use of paper-based payments, especially checks. The growing popularity of ABP mechanisms can be attributed to perceived benefits to retailers and service providers in the form of reduced billing transactions costs and greater certainty of payment and to consumers in the form of convenience and avoided postage and penalties for delinquency (Mastercard 2006; Visa 2006). It is estimated that among internet-connected households in the U.S., 41% of all recurring bills are paid automatically (Fiserv 2010) and that two-thirds of consumers with recurring bills use automatic payment mechanisms. By 2005, 53% of credit or debit cardholders in the U.S. paid recurring telephone bills automatically, while 44% and 37% respectively paid recurring cable television and utility bills automatically. Survey results suggest roughly $23 trillion worth of bill payments in the U.S. are transacted automatically. In the UK, three in four consumers made at least one recurring payment automatically in 2010 (Payment Council 2010).

To my knowledge, this paper is the first to estimate the causal effect of ABP enrollment on firm and household consumption. The customer-level data necessary for such analysis across the many industries that employ ABP mechanisms are typically proprietary and closely guarded. The regulated status of utility companies in some states, however, requires their compliance with state public record laws that mandate, in some circumstances, the disclosure of annonymized customer billing data. Thus, this paper employs monthly observations on electricity consumption among residential and commercial customers in South Carolina in order to test the prediction that the diminished price salience caused by ABP enrollment induces increased electricity consumption.

Estimates suggest that ABP enrollment causes a 2.4-5% increase in electricity use among residential customers and as much as a 6% increase in consumption by relatively new commercial accounts. This paper also provides the first credible estimates of the electricity consumption impact of budget billing programs that smooth bills of participating, low-income customers over a 12-month period. Consistent with theory that the price salience of deviations from mean consumption is diminished with enrollment in budget billing, it is estimated that enrollment in such programs increases residential electricity consumption 8-9%.

Results presented here suggest that utility billing programs like ABP and budget billing interfere with efforts by regulated utilities and policymakers to boost energy price salience and reduce energy demand amid concern about anthropogenic climate change and the health and environmental damages caused by energy production and consumption. Such efforts include billions of dollars of investment in electricity infrastructure upgrades to deliver real time use and price information to customers. The results of this analysis may also partly explain the "energy paradox," a phenomenon in which individuals fail to make investments in energy efficiency that offer positive returns from energy cost savings (Allcott et al. 2010, 2011; Allcott and Greenstone 2012; Gillingham et al. 2009; Sanstad and Howarth 1994). The persistence of the energy efficiency gap is typically understood to be a consequence of the diminished salience of and inattention to future cost savings relative to upfront expenditures for energy efficient technologies. If the price for energy consumption is, itself, insalient, then future cost savings from efficiency improvements will be undervalued.

The salience effect induced by ABP mechanisms is perhaps most important in energy markets because of policy interest in reducing energy demand and because of their prevalence. Each of the ten largest electric utilities in the U.S. employ ABP. However, the salience effect from ABP enrollment estimated here in the context of electricity demand is likely also present in markets for other goods and services characterized by recurring payments, including markets for other natural resources like water, natural gas, and heating oils, as well as telecommunications, insurance, and home maintenance. ABP enrollment can induce overconsumption not just of marginal units when a consumer faces volumetric charges. If the consumer faces multipart tariffs and increasing block rate pricing, overconsumption can also occur when changes in individual demand, technology and prices would induce a fully attentive individual to select a less costly product bundle or forego consumption of the product entirely. Thus, even for products priced predominantly by fixed fees, like some mobile telephone or cable television packages, diminished price salience may still result in increased consumption if consumers fail to select into less costly bundles as prices change or as their product demands evolve or become known with greater certainty.

This paper proceeds in section 2 by developing a theoretical model of price salience effects on consumption. In section 3, the data and empirical model are introduced. Results and discussion are provided in sections 4 and 5, respectively. And section 6 concludes.

3.2 Bounded Rationality and Salience Effects

Bounded rationality holds that agents seek to maximize utility subject to informational and cognitive constraints that may lead to mistakes in decision making (Simon 1955; Tversky and Kahneman 1974). Such decision errors are observed in laboratories and in the field, and their consequences are the subject of a growing literature (see DellaVigna 2009 for a survey). While traditional models assume agents are perfectly informed, behavioral models acknowledge that attention may be a scarce resource. In this context, consumer demand can be altered as a consequence of changes in the salience of various product characteristics. Diminished price salience, for instance, is shown in a number of contexts to increase demand as "perceived" prices fall. Theory related to budget billing impacts on energy consumption was considered by McDermott et al. (1980) and Beard et al. (1998), so attention here is focused on developing an understanding of how more prevalent ABP programs impact consumption.

In order to flexibly model diminished price salience associated with ABP enrollment across heterogeneous pricing regimes, let the exogenous price of a good X, $P(x)$, be comprised of a fixed fee, a, and a non-linear tariff whereby the volumetric charge is p_1 for quantities less than k and $p_2 = p_1 + n$ for quantities above k. Then the price of X is given by:

$$P(x) = \begin{cases} (x-k)p_2 + kp_1 + a & \text{if } x > k \\ xp_1 + a & \text{if } 0 < x \leq k \\ 0 & \text{if } x = 0 \end{cases} .$$

Within this framework, a typical electric bill would be characterized by a small $a > 0$, $p_1 > 0$ and $n > 0$, reflecting increasing block rate pricing. A typical gym membership is represented by $a > 0$, $p_1 = 0$ and $n = 0$. Mobile telephone service is represented by $a > 0$, $p_1 = 0$, and $n > 0$, where k is the limit on free calling or data. Throughout, we will consider non-decreasing block rate structures, i.e. $n \geq 0$, which characterize a growing share of utility pricing regimes.

Assume individuals derive utility from consuming a numeraire, L, and M, which is produced from X according to $M = \alpha x$, where α is a technology parameter.[1] Though the framework is general, one may think of X as electricity that is used to produce household activities, M.[2] The individual utility function is of the quasi-linear form $U(L, M) = L + \theta V(M)$, where θ is a taste parameter that functions as a demand shifter, and $V(\cdot)$ is a strictly concave and twice continuously differentiable function. Individuals maximize utility subject to the budget constraint $I = L + P(x|\alpha)$.

Following DellaVigna (2009), I incorporate inattention associated with enrollment in ABP by assuming that individuals enrolled in ABP may perceive only a portion of the price of X and of the technical efficiency with which M is produced from X. Thus, define the perceived entrance fee of X, $\tilde{a} = (1-\beta)a$, the perceived volumetric price of X, $\tilde{p_j} = (1-\beta)p_j$ for $j \in \{1, 2\}$, and the perceived technical efficiency of production of M, $\tilde{\alpha} = \bar{\alpha} + (1-\beta)\ddot{\alpha}$, where $\beta \in [0, 1)$ is an inattention parameter with $\beta = 0$ denoting full attention and $\beta = 1$ denoting full inattention, and where $\bar{\alpha}$ denotes a baseline value of the technology parameter and $\ddot{\alpha}$ denotes deviations from the baseline (e.g., following replacement of an old air conditioner). This framework allows that an individual enrolled in ABP may be inattentive to technology-induced changes in consumption levels, in addition to changes in prices. A constant inattention parameter is assumed for simplicity. For notational convenience, define $\delta = 1 - \beta$.[3]

The individual's perceived objective is:

[1] This is a simplification of the household production function (Lancaster 1966).

[2] The technology parameter, α, that governs the household production function is introduced in order to identify a dimension of technological change in household production to which individuals may be inattentive amid diminished price salience. For $\alpha = 1$, X can be considered a good that yields utility directly.

[3] Moreover, the assumption of a single inattention parameter, β, is for analytical convenience. A priori, there is no reason to expect the degree of inattention to be the same across various price components and technology.

$$\max_{x} \quad U(L,x) = L + \theta V(\alpha x)$$
$$s.t. \quad I = \begin{cases} L + \left((\bar{\alpha} + \delta\ddot{\alpha})x - k \right) \delta n + (\bar{\alpha} + \delta\ddot{\alpha})x\delta p_1 + \delta a & \text{if } x > k \\ L + (\bar{\alpha} + \delta\ddot{\alpha})x\delta p_1 + \delta a & \text{if } 0 < x \leq k \\ L & \text{if } x = 0 \end{cases}. \tag{3.1}$$

Given the concave utility function and concave budget set, a unique solution to the consumer's problem exists at either a corner solution, the kink in the piecewise linear budget set that occurs at $x = k$, or at a point of tangency along one of the two linear segments of the budget set.

Consider an individual whose solution to (3.1) is $x^* \in (0, k)$ defined by:

$$V'(M) = \frac{(\bar{\alpha} + \delta\ddot{\alpha})\delta p_1}{\theta}. \tag{3.2}$$

Greater inattention reduces the RHS of (3.2), causing greater consumption of X in equilibrium because $V'(M) \geq 0$ and $V''(M) < 0$. Thus, *ceteris paribus,* an inattentive individual consumes more X than a fully attentive individual. The magnitude of "excess" consumption is increasing in the degree of inattention. If the solution to the consumer's problem lies in the region $x > k$, then an analogous result obtains.

This result and associated welfare losses are depicted in figure 3.1, in which an inattentive consumer perceives the budget constraint defined by the line segments ABC and optimizes by intending to consume (l^*, x^*). This point, however, lies outside the feasible set defined by the origin and the line segments EFG. Consequently, in order to consume x^*, the inattentive individual must forgo more consumption of L than he perceives when he commits to consuming x^*. His consumption bundle lies on a lower indifference curve associated with the bundle $(\hat{l}, x^*)$ that lies on the true budget constraint. Given that the slopes of the true and perceived budget constraints differ at x^*, the intersection of the lower indifference curve and the budget constraint at x^* is necessarily not a point of tangency, which implies there exists a bundle (l^{**}, x^{**}) that is feasible and yields higher utility. Because the slope of the true budget constraint is greater than that of the perceived budget constraint for $\beta \in (0, 1)$, such a point must satisfy $x^{**} < x^*$. The difference $x^* - x^{**}$ is defined as the overconsumption or "excess" consumption of X due to consumer inattention. It results in a welfare loss equal to m' in figure 3.1, which represents the amount of income that could be taken from the consumer while still enabling him to achieve the utility associated with the bundle $(\hat{l}, x^*)$ if the budget were allocated optimally between goods L and X, i.e., bundle (l^{**}, x^{**}) in figure 3.1.[4]

Inattention to price can result in other types of consumption errors in the presence

[4]It is acknowledged that attention may, itself, be scarce, in which case a reduction of overconsumption of X would come at a cost in terms of foregone consumption of L.The interpretation of $x^* - x^{**}$ as overconsumption assumes it is not a consequence of rational inattention.

of corner solutions and multi-part tariffs, which characterize the pricing regimes for several goods subject to ABP. Figure 3.2 depicts a case where the fully attentive consumer achieves an optimum at the kink of the budget constraint $(l^{**}, x^{**} = k)$, but the inattentive consumer seeks to consume the combination (l^*, x^*), which is infeasible. The welfare loss, m', is defined as before. Intuitively, a fully attentive individual will only consume a unit of x beyond k units if the marginal benefit of the $(k+1)$th unit exceeds the higher marginal cost imposed by the increasing block rate. But an inattentive individual perceives only a fraction δ of the $\frac{n}{p_1} \times 100$ percent increase in marginal price above $x = k$ units and thus will sooner consume the $(k+1)$th unit than a fully attentive individual.

Finally, overconsumption and welfare loss can result from the inattentive consumer undertaking positive levels of consumption when the full attention optimum occurs at $(l^{**} = I, x^{**} = 0)$, as depicted in figure 3.3. Again optimizing according to the perceived budget constraint, the inattentive individual chooses x^*. But consumption of x^* units of X does not permit l^* units of consumption of L as the inattentive individual infers. Rather, the indifference curve that satisfies $x = x^*$ and the true budget constraint is lower than the indifference curve that maximizes utility subject to the true budget constraint. If the optimal solution for the attentive individual is $x^{**} = 0$, then the vertical distance between I and the vertical intercept of the indifference curve that satisfies the true budget constraint and the condition $x = x^*$ is the welfare loss in terms of units of L. Intuitively, the inattentive individual perceives a lower marginal price *and* a lower fixed "entrance" fee than the attentive individual who is dissuaded by the full entrance fee from consuming any X at all.

Individuals who enroll in ABP programs may be fully attentive to the cost of some baseline level of consumption of M, such as the cost of consumption at the point of enrollment in ABP. Nevertheless, they may exhibit inattention to changes in preferences, θ, technology, α, and price, $P(x)$, which should induce changes in optimal consumption of X and therefore cause the cost of recurring bills to change. Let us assume $0 < x^* < k$ or $x^* > k$ so that the solution to the individual's problem is on the interior but not at a kink in the budget constraint. First consider the impact of changes in preferences on the optimal consumption of the attentive and inattentive consumer. An increase in utility from M increases θ so that for given price and technology, the optimal consumption of x must increase from x^* to $x^{**} > x^*$. Because the RHS of (3.2) is decreasing in β, then the smaller is β the greater must be the increase in consumption of x to satisfy (3.2). Likewise, for a given decrease in utility from M, the magnitude of the optimal decrease in x is declining in β. Intuitively, the greater is inattention, the less the individual perceives the costs of consuming additional M or realizes the benefits of conserving M.

The optimality condition also implies by the implicit function theorem that:

$$\left.\frac{\partial \frac{dm}{dp}}{\partial \beta}\right|_{\ddot{\alpha}=0} = -\frac{\bar{\alpha} - 2(1-\beta)\ddot{\alpha}}{V''(M)} > 0 \tag{3.3}$$

and

Figure 3.1: Overconsumption from inattention to marginal price

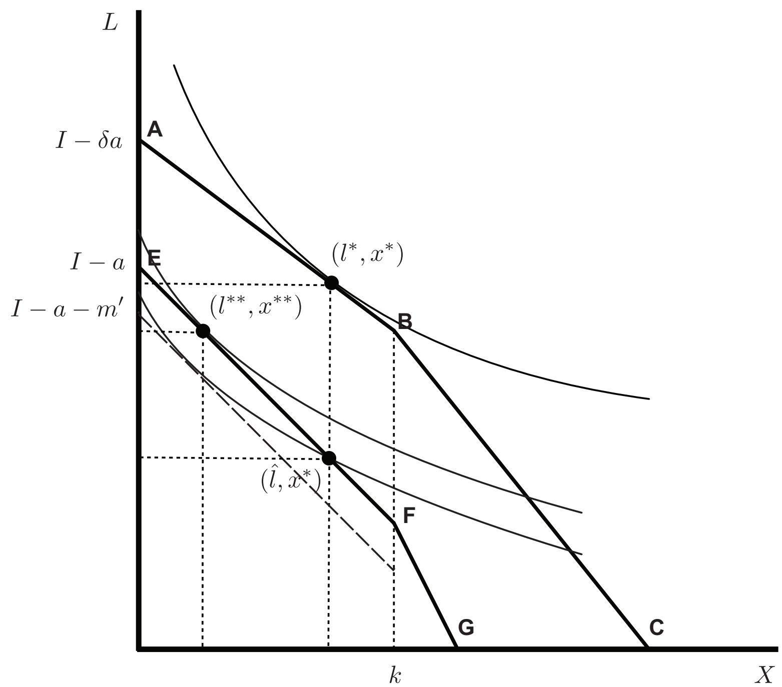

An inattentive individual perceives the budget constraint identified by ABC (ignoring the corner solution, $L = I$, which is considered in Figure 3.3) and intends to optimize according to (l^*, x^*), which is, in fact, infeasible. The feasible set is defined by the origin and the budget constraint identified by EFG. Having committed to consume x^{**} units of X, the individual can only consume $\hat{l} < l^{**}$ units of L, yielding less utility than the optimizing bundle (l^{**}, x^{**}) that is chosen by a fully attentive individual.

Figure 3.2: Overconsumption from inattention to an increasing block rate

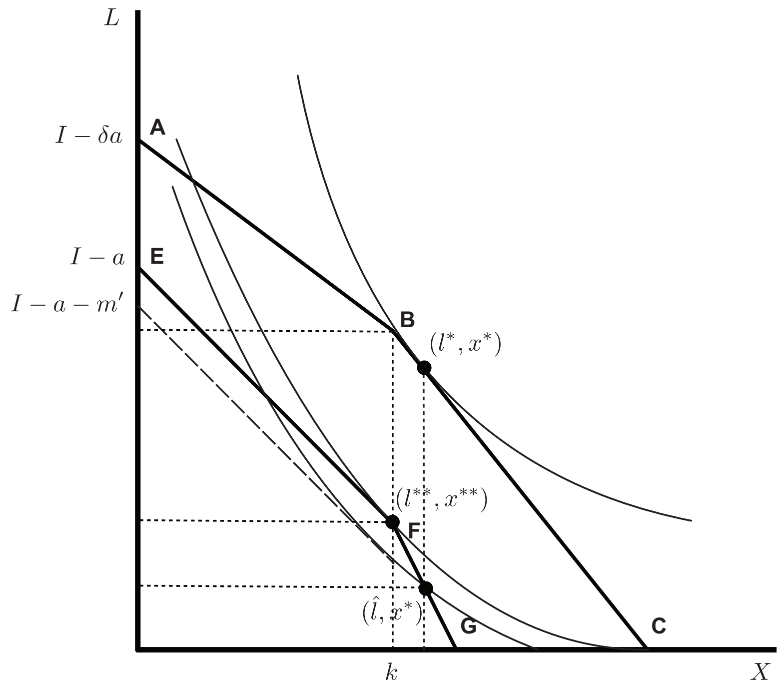

An inattentive individual intends to consume (l^*, x^*), which is infeasible, whereas the attentive individual optimizes at (l^{**}, x^{**}), which occurs at the kink of the budget constraint. The inattentive individual perceives less of the increase in marginal price for the $(k + 1)$th unit of X than the attentive individual.

Figure 3.3: Overconsumption from inattention to a fixed entrance fee

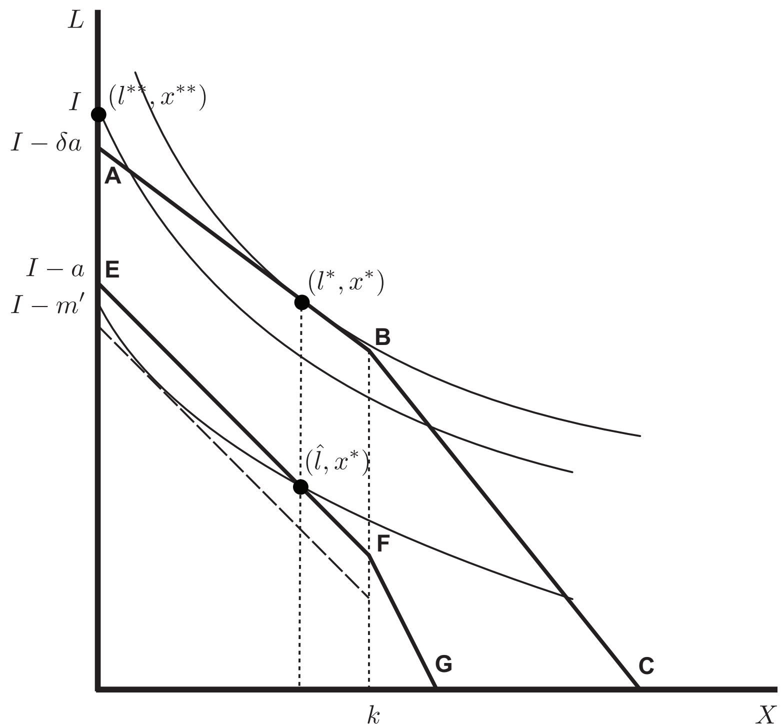

An inattentive individual intends to optimize according to (l^*, x^*), which is infeasible. Having committed to consume x^{**} units of X, the individual can only consume $\hat{l} < l^{**}$ units of L, yielding less utility than the optimizing bundle (l^{**}, x^{**}) chosen by a fully attentive individual. The attentive individual perceives the full entrance fee and avoids it by foregoing consumption of X, whereas the inattentive individual perceives only a portion of the entrance fee.

$$\frac{\partial \frac{dm}{d\alpha}}{\partial \beta} = \frac{2(1-\beta)p}{V''(M)} < 0. \tag{3.4}$$

Thus, given $\frac{dm}{dp} < 0$ and for constant technology, demand is less responsive to changes in price as inattention increases, as shown in (3.3). An inattentive individual increases consumption less in response to price declines and reduces consumption less in response to price increases than the fully attentive individual. Likewise, given $\frac{dm}{d\alpha} > 0$, inattention reduces the responsiveness of demand to changes in technical efficiency, as seen in (3.4).

Finally, individuals may be inattentive to changes in the technical efficiency with which X is converted to M. The greater is this inattention, the lesser is the increase in consumption of X and M due to a positive shock in technical efficiency, i.e., $\ddot{\alpha} > 0$. Inattention to technical efficiency diminishes the Jevon's paradox or "rebound effect" that is understood to diminish input conservation from efficiency improvements because of price and income effects. Thus, while inattention is often blamed for suboptimally high levels of energy consumption, inattention to efficiency improvements reduces energy consumption. Inattention to negative shocks to technical efficiency, however, results in suboptimally small reductions in consumption of M as a consequence of the greater input costs in production of M.

This stylized model demonstrates that inattention to the price of a good results in excess consumption of the good relative to the full attention optimum. Inattentive individuals overreact to a relative increase in preference for the good and under-react to a relative decline in preference for the good. They under-respond to changes in technical efficiency, too. And, importantly, an increase in price induces too small of a reduction in consumption from an inattentive consumer, while a decrease in price induces too small of an increase in consumption. As retail electricity prices exhibit a persistent upward trend, inattention to price is likely to induce overconsumption.

In a dynamic formulation of this two-good model, one may expect the individual to learn from his optimization error upon realizing that consumption of l^* units of L is infeasible given x^* units of consumption of X. But in the context of a consumption bundle that contains hundreds of goods and services, it would be unlikely for an inattentive individual to identify the source of the optimization error upon realizing the infeasibility of the preferred bundle. As Thaler has noted (e.g., 1986) learning requires immediate feedback, which markets may not always provide. Thus, the optimization error may persist with limited learning.

Regardless, a critique that repeated decision making and associated learning undermines the potential for persistent decision errors can be equally levied against much of the work on bounded rationality and salience effects, including that of Chetty et al. (2009), DellaVigna (2009), and DellaVigna and Pollet (2009), among others. Indeed, the potential for learning and for markets to enforce discipline in consumption decisions is central to critiques of behavioral economics (e.g. Friedman 1953) and to debate over the external validity of laboratory experiments that showed deviations from standard theory (e.g., Mullainathan and Thaler 2001,List 2003, and Levitt and List 2007). Such critiques not withstanding, the growing body of behavioral work, including accumulating evidence from the field in support of behavioral theories, demonstrates that decision errors can persist and impact market equilibria, that competition does not always save decision makers from their mistakes, and that arbitrage opportunities, where they exist, do not always enforce rationality (Akerlof and Yellen 1985; Thaler 1986; Conlisk 1996; Mullainathan and Thaler 2000; Gabaix and Laibson 2006).

3.3 Data and Methods

In order to identify the average effect of enrollment in budget billing or ABP programs on electricity consumption, I employ an extensive panel of monthly household-level electricity consumption in a fixed-effects framework. Such a framework controls for selection on time-invariant unobservables and is robust to heterogeneous treatment effects.

Define y_{it} as the monthly electricity consumption (in kilowatt hours) of household i in period t, such that

$$y_{it} = \lambda_t + c_i + z_i + w_{it}^A \gamma_i^A + w_{it}^B \gamma_i^B + u_{it}, \tag{3.5}$$

where λ_t are a full set of time effects, c_i are observed and unobserved individual-specific time-invariant heterogeneity, z_i are zip-code fixed effects intended to improve precision, w_{it}^A is a treatment indicator equal to one if household i is enrolled in ABP in period t and zero otherwise, w_{it}^B is a similarly defined treatment indicator for enrollment in budget billing, and u_{it} is an idiosyncratic error.[5,6] Equation (3.5) allows heterogeneous treatment effects, which permits the behavioral response to enrollment in the billing programs to vary across households. Some households may enroll in the billing programs but remain perfectly informed about and attentive to monthly electricity costs. In such instances, enrollment should have no impact on electricity consumption. In other cases, electricity costs may become less salient and induce an increase in consumption as the perceived price of electricity falls.

For a large cross-section and a small time-series dimension, estimation of the individual-specific treatment effects, γ_i^j for $j = \{A, B\}$ reflecting ABP and budget billing treatments, respectively, is problematic. However, a valid estimator of the average treatment effect can be obtained if

$$E(\gamma_i^j | \ddot{w}_{it}) = E(\gamma_i^j) = \gamma^j \ \forall t, \tag{3.6}$$

[5]Estimation is carried out with standard errors clustered by zip code to allow within zip-code error correlation.

[6]In a separate specification, equation (3.5) is augmented by an interaction between the ABP treatment indicator w_{it}^A and, sequentially, a zip code-level measure of median household income, median home value, and the share of the population less than 25 years of age. This tests whether the treatment effect is correlated with these demographic characteristics.

where $\ddot{w}_{it} = w_{it} - \bar{w}_i$. That is, individual-specific treatment effects can be correlated with the average propensity to receive treatment, $\bar{w}_i$, but not with the deviations in any time period (Wooldridge 2002). In the context of a voluntary treatment regime, condition (3.6) constitutes a self-selection constraint that ensures program participation decisions are not systematically related to expected treatment effects. Here, the treatment impacts potential outcomes through inattention or salience effects. Thus, the treatment effect is essentially unintended and the self-selection constraint is likely satisfied. Intuitively, it is unlikely an individual selects into ABP or BB because he expects to consume more electricity than he otherwise would. Conditional on (3.6), a good estimator of γ^j is $\hat{\gamma}^j = N^{-1}\sum_{i=1}^{N} \hat{\gamma}_i^j$.

In addition to (3.6) and the standard overlap assumption, consistent estimation of α^j, the population-averaged treatment effect (PATE), also relies upon strict exogeneity of treatment (Wooldridge 2002; Imbens and Wooldridge 2009; Rosenbaum and Rubin 1983; Rubin 1990). That is, the propensity to receive treatment and potential outcomes must be independent conditional on covariates. Though treatment assignment is non-random in this context, if (3.6) is satisfied, then strict exogeneity can only be violated by correlation between deviations in unobservable time-varying household characteristics and deviations in treatment status. But because nearly all households in these data that enroll in ABP and budget billing programs persist in the programs until their accounts are terminated, strict exogeneity is likely satisfied.[7] Indeed, as Wooldridge and Imbens (2007) noted, if $w_{it} = 1$ whenever $w_{ir} = 1$ for $r < t$, then strict exogeneity is a reasonable assumption. Intuitively, if enrollment in period r were induced by a stochastic positive shock to an underlying characteristic, like preference for convenience, then a stochastic negative shock to the same characteristic at a time $t > r$ should induce program withdrawal in period t. Moreover, as Imbens (2004) noted, even when individuals are presumed to choose treatment optimally, unconfoundedness will hold if the (potentially unobserved) characteristics determining treatment are unrelated to the outcome of interest. This is particularly true if the objective of the decision-maker is distinct from the outcome of interest. In this case, an individual may optimally select into ABP because of a preference for convenience. But it is unlikely the individual selects into ABP because he seeks to influence consumption. The outcome of interest is a consequence of inattention and suboptimal decision-making. If desire for convenience or other characteristics determining ABP enrollment are unrelated to inattention or cognitive constraints, as seems plausible, then even correlated changes in unobservable household characteristics and treatment status do not jeopardize unconfoundedness.

While strict exogeneity is necessary to interpret $\hat{\gamma}^A$ and $\hat{\gamma}^B$ as estimates of PATEs, their interpretation as population-averaged treatment effects on the treated (PATTs) does not depend on independence of treatment status and potential outcomes. As it is unlikely that enrollment in ABP or BB would ever be compulsory, particularly given the results of

[7] Only 0.33% of the time does enrollment status change more than once per household in these data.

this analysis, policy interest likely centers on the PATT rather than the PATE. The effect of treatment on those who will never receive treatment is of no particular policy interest (Heckman and Hotz 1989; Heckman et al. 1997).

The principal set of data employed in this paper is a large panel of monthly observations on residential and commercial electricity consumption from 1994 to 2010. These data were obtained pursuant to a public records request from Santee Cooper, a publicly owned electric utility in South Carolina that provides electricity to 164,680 customers located along the Atlantic coast between Charleston and Myrtle Beach. Due to privacy limitations in state public record law, personally identifiable information is exempt from disclosure. Consequently, each customer account is identified only by an account number. It is matched only to zip-code, limiting information about the account-holders and their homes. Aggregate demographic information from the 2000 Census is used to identify heterogeneous effects according to zipcode-level measures of household income, home value, and age distribution. Separate random samples of these data are drawn for various rate classifications in order to demonstrate robustness across rate codes for residential and commercial accounts. ABP and BB accounts are oversampled. Summary statistics for each sample are reported in Table 1 for residential accounts. Summary statistics for commercial accounts are reported in Table 2.

3.4 Results

This analysis yields evidence consistent with the hypothesis that diminished price salience associated with enrollment in automatic bill payment programs induces additional electricity consumption. In particular, enrollment induces a 2% increase in electricity consumption on average among all residential customers, as reported in column 1 of Table 3.3. This estimate is significant at the 5% level. Restricting attention to newer accounts, particularly those initiated in the year 2000 or later ("Recent"), the salience effect increases to a highly significant 4.5% average increase in consumption, as reported in column 2 of Table 3.3. These estimates are fairly consistent across the utility's rate classes, as reported in columns 3-8 of Table 3.3 for all accounts by rate class and the subset of accounts initiated since 2000. In particular, the estimated effect of ABP enrollment on electricity consumption across all accounts in a rate class ranges from 2.3% to 2.7%. And it is significant at the 1% level in two out of three rate classes. The estimated effect on recent accounts is roughly 100% larger than the estimated effect across all accounts and it is significant at the 1% level for all rate classes.

The larger magnitude of the salience effect on recent accounts suggests that inattention to electricity bills is greater among new account holders, who tend to be younger. This could reflect a generational change in attention to financial accounts or relatively lower financial literacy of younger account holders. That is, more senior account holders may

Table 3.1: Summary Statistics - Residential Accounts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Resid.		R1-R4[a]		R5, R6, RS, RG[b]		RE[c]	
	All	Recent	All	Recent	All	Recent	All	Recent
Observations	847,068	862,774	2,211,233	914,449	1,273,753	421,455	1,515,196	443,898
No. of accounts	24,586	29,114	58,082	32,792	40,004	19,105	53,259	22,670
No. of ABP Accounts	12,358	29,114	20,074	9,581	30,519	13,905	34,735	13,232
No. of BB Accounts	7,544	4,884	2,345	1,085	3,098	1,307	3,432	1,024
Mean length of ABP enrollment (months)	57.86	36.56	61.50	45.97	54.88	36.51	66.37	50.90
	(0.56)	(0.21)	(0.43)	(0.49)	(0.33)	(0.30)	(0.33)	(0.38)
Mean length of BB enrollment (months)	60.61	40.32	69.78	58.24	63.97	41.36	69.86	53.73
	(0.78)	(0.57)	(1.45)	(1.75)	(1.29)	(1.10)	(1.18)	(1.51)
Mean date of account initiation	475.66	532.61	492.57	535.73	480.10	532.99	474.29	533.40
	(0.36)	(0.15)	(0.23)	(0.17)	(0.28)	(0.22)	(0.25)	(0.20)
Mean date of ABP initiation	492.84	545.71	500.98	547.42	501.39	545.36	472.44	540.78
	(0.99)	(0.34)	(0.78)	(0.63)	(0.58)	(0.48)	(0.63)	(0.66)
Mean no. months before ABP	27.48	13.73	25.08	14.18	25.10	13.33	22.81	15.29
	(0.69)	(0.24)	(0.51)	(0.45)	(0.40)	(0.34)	(0.40)	(0.48)
Mean log consumption (kWh)	6.70	6.60	6.85	6.73	6.53	6.56	6.85	6.82
	(0.0008)	(0.0008)	(0.0004)	(0.0009)	(0.0007)	(0.001)	(0.0006)	(0.001)

Standard errors in parentheses.

[a] Rates for customers living in homes meeting energy efficiency standards dictated by the utility.

[b] Standard residential service rates.

[c] Rates for customers who restrict peak season (July-September) demand as a function of off-peak season demand

Table 3.2: Summary Statistics - Commercial Accounts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)		
	Comm.		GN[a]		GS[b]		GL[c]		GC[d]	
	All	Recent	All	Recent	All	Recent	All	Recent	All	Recent
Observations	701,920	182,818	641,119	178,057	628,995	161,573	7,484	1,844	63,933	10,320
No. of accounts	18,453	7,572	16,339	6,959	12,409	4,728	122	37	1,051	332
No. of ABP Accounts	13,087	5,402	11,771	5,054	2,125	927	7	6	362	102
Mean length of ABP enrollment (months)	47.56	32.13	43.77	29.83	42.86	35.40	15	16	75.46	42.16
	(0.50)	(0.44)	(0.72)	(0.45)	(1.05)	(1.11)	(2.86)	(3.46)	(3.46)	(3.87)
Mean date of account initiation	471.98	530.63	474.49	531.76	467.19	529.18	454.42	516.76	459.23	536.70
	(0.41)	(0.36)	(0.49)	(0.38)	(0.49)	(0.45)	(4.28)	(4.30)	(1.77)	(1.74)
Mean date of ABP initiation	505.69	543.96	511.50	549.13	501.48	535.83	579.50	579.50	491.19	550.48
	(0.85)	(0.73)	(0.91)	(0.79)	(2.14)	(1.97)	(0.50)	(0.50)	(5.50)	(5.90)
Mean no. months before ABP	35.25	15.23	37.43	18.73	37.43	11.78	73	73	37.43	11.82
	(0.71)	(0.56)	(0.76)	(0.65)	(0.76)	(1.22)	(23)	(23)	(4.60)	(3.38)
Mean log consumption (kWh)	6.94	7.07	6.55	6.72	9.39	9.36	12.14	12.13	7.25	7.07
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)	(0.10)	(0.24)	(0.05)	(0.08)

Standard errors in parentheses.

[a] Small commercial service (<7,500 kWh/month)

[b] Medium commercial service.

[c] Large commercial service (>90,000 kWh/month for at least one month per year).

[d] Service for places of worship (i.e., churches).

be more likely to review electricity bills even after enrolling in ABP, thereby mitigating the diminution of price salience. It is important to note that these results constitute an estimate of a population-averaged effect. The treatment effect is expected to be heterogeneous by customer according to bill viewing habits before and after program enrollment. For those households that continue to regularly review bills, the effect of enrollment should be negligible, producing zeros in the estimation of individual-specific effects. But for those program customers who stop regularly viewing bills absent the incentives imposed by account delinquency under manual payments, the magnitude of the treatment effect could be considerably larger than the population averaged effects reported here. The magnitude of the salience effect is invariant to zip code-level measures of median home value, household income, and age distribution.[8]

Table 3.3 also provides new evidence of the impact of budget billing programs on electricity consumption. Though the literature includes previous estimates of the impact these programs have on electricity consumption (McDermott et al. 1980; Williams et al. 1990; Beard et al. 1998; Ha et al. 1993), none afford the internal validity of the fixed-effects estimation in this paper. On average, budget billing induces a 9.1% increase in consumption across all residential accounts, as shown in column 1 of Table 3.3. This result is significant at the 1% level. Across rate classes, the effect is always significant at the 5% level or better, and ranges from 5.2% to 10.5%. These results are similar to those of Ha et al. (1993), who estimated that households enrolled in budget-smoothing programs consumed 10% more electricity than those not enrolled in the programs. They did not control for selection. Beard et al. (1998) did not report readily interpretable results. Distinct from the estimated impacts of ABP enrollment, the magnitude of consumption increases due to budget billing do not vary consistently by the recency of account initiation. This undoubtedly reflects the fact the budget billing conceals from customers the marginal cost of deviations from mean consumption irrespective of whether they review their bills or not. This result, thus, lends credence to the interpretation of the estimated consumption increases due to budget billing and automatic payment as a consequence of diminished price salience.

An additional set of results examines whether diminished price salience due to ABP enrollment also induces over-consumption by commercial account holders. As reported in Table 3.4, there is some evidence that it does, though mostly only among recent account holders. For instance, there is no statistically significant effect of ABP enrollment on consumption among a random sample of all permanent (not temporary) commercial accounts, as reported in column 1 of Table 3.4. But among recent accounts, there is a statistically significant increase in consumption due to ABP enrollment of 6.3% on average (see column 2). Among small commercial accounts (GN) initiated after 2000, there is an estimated 10.3% increase in electricity consumption that is statistically significant at the 1% level (column 4). There is no significant effect among all small commercial accounts (column 3). Among all medium commercial accounts (GS), however, there is an estimated 2.9%

[8]Results are omitted for brevity, but are available from the author upon request.

increase in consumption (column 5). Among recent medium-sized commercial accounts, the estimated effect is of a similar magnitude (2.7%), but not statistically significant (column 6). The point estimates of the ABP effect on electricity consumption among large commercial accounts (GL) are of approximately the same magnitudes as estimates for all residential accounts, but they are not significant because of relatively large standard errors (columns 7 and 8). Finally, among recent accounts held by places of worship (GC), it is estimated that ABP enrollment induces a large and statistically significant 22.1% increase in electricity consumption, though the significance and magnitude of the effect disappears among all church accounts.

Thus, evidence on over-consumption among commercial accounts that enroll in ABP is mixed. One might have expected there to be no impact among commercial account holders who may employ professional financial services or more closely monitor budgets than households. On the other hand, small commercial account holders are likely to be small business owners who may not have any more financial literacy than homeowners and who may not employ anyone who does. These accounts are likely to have more elastic electricity demand, too. Large commercial account holders, on the other hand, are likely to have less elastic demand, reflecting the invariance of operations to electricity prices. They are also more likely to employ individuals who closely monitor operating expenses. Despite the lack of significance of the ABP effect among large firms, the results are difficult to reconcile with any theory about the relative elasticity of demand among commercial account holders or their relative financial literacy and attention to operating expenses. Results nevertheless suggest that the diminished price salience induced by ABP enrollment may lead to overconsumption among some commercial customers.

3.5 Discussion

Since at least the 1970s, energy conservation has been a priority of federal energy policy. By 2012, it took on greater urgency amid heightened concern about global climate change and other environmental and human health consequences of energy production and consumption. But as electric utilities and regulators undertook considerable effort to reduce residential electricity consumption by non-coercive means, especially by strengthening price signals to consumers, many utilities also employed one or more billing programs that this analysis suggests lead to excessive electricity consumption due to consumer inattention. While the magnitude of the estimated consumption impact of enrollment in budget billing is greater than the estimated impact of ABP enrollment, the latter is offered by more utilities, including each of the ten largest utilities in the U.S., and is taken up by customers at considerably higher rates. Approximately one in four customers of Santee Cooper were enrolled in ABP in 2010. Thus ABP is likely to have a significant aggregate impact on electricity consumption across the country. In fact, the estimated 2% increase

Table 3.3: Results- Residential Accounts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Resid.		R1-4		R5-6, RG, RS		RE	
	All	Recent	All	Recent	All	Recent	All	Recent
ABP	0.0204**	0.0447***	0.0234***	0.0584***	0.0278***	0.0509***	-0.00211	0.0405***
	(0.0104)	(0.00668)	(0.00664)	(0.0115)	(0.00690)	(0.00999)	(0.00594)	(0.00962)
BB	0.0838***	0.0913***	0.0523***	0.0548**	0.105***	0.0888***	0.0818***	0.0788***
	(0.00869)	(0.0102)	(0.0125)	(0.0239)	(0.0137)	(0.0157)	(0.0131)	(0.0239)
Observations	739,539	670,467	1,897,882	704,311	1,101,064	326,614	1,438,610	409,755
No. of accounts	14,146	18,240	40,154	22,882	18,637	9,173	33,182	15,703

Standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3.4: Results- Commercial Accounts

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Comm.		GN		GS		GL		GC	
	All	Recent	All	Recent	All	Recent	All	Recent	All	Recent
ABP	0.011	0.063**	0.015	0.103***	0.029**	0.027	0.018	0.020	0.067	0.221**
	(0.021)	(0.028)	(0.021)	(0.032)	(0.014)	(0.019)	(0.094)	(0.056)	(0.072)	(0.109)
Observations	630,553	153,490	544,261	135,856	590,772	143,476	6,943	1,483	60,101	9,069
No. of accounts	9,703	3,964	8,327	3,470	10,645	3,938	107	24	782	251

Standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

in electricity consumption due to ABP enrollment is expected to cause the average Santee Cooper customer enrolled in ABP to consume an additional 194.98 kWh per year at a cost of $20. Throughout the Santee Cooper service area, consumption increases an aggregate 7.1 million kWh per year due to ABP, the same amount of electricity consumed by 730 average households. ABP-induced overconsumption cost the utility's residential customers approximately $745,500 in 2010. To the extent ABP induces overconsumption by commercial account holders, too, the cost to utility customers is greater.

These results suggest that elimination of ABP programs could contribute to energy and environmental policy goals by reducing residential electricity consumption. Notably, the magnitude of the estimated increase in electricity consumption due to ABP enrollment is large relative to other factors impacting electricity demand that have attracted the attention of economists and policymakers. For instance, empirical evidence suggests the decades-old practice of shifting clocks forward in the spring and backward in the fall in order to reduce energy consumption has an effect on residential consumption that is dwarfed by the ABP effect found here, undermining the principal rationale for Daylight Saving Time. Kotchen and Grant (2011) estimated that Daylight Saving Time *increases* electricity consumption 1% on average, while Kellogg and Wolff (2008) found that an extension of Daylight Saving Time in Australia during the 2000 Summer Olympic Games did not significantly impact electricity consumption. In the U.S., at least 47 utilities in 21 states, including each of the ten largest utilities, rely on social norms to achieve demand reductions by issuing home energy reports that compare a household's electricity consumption to that of its neighbors. The reports are estimated to reduce electricity consumption by 2% (Allcott 2011), though the impact is heterogeneous. Relatedly, a review of pilot programs that introduce energy monitoring devices to households in order to boost price signals suggests such programs can induce 7% demand reductions among active users of the devices (Faruqui et al. 2010).

The elimination of ABP programs could achieve residential and commercial electricity demand reductions at virtually no pecuniary cost to consumers or utilities. If ABP-induced consumption increases are not a consequence of rational inattention, then consumer welfare would be improved. And to the extent that the objective functions of utilities include electricity demand reductions (as they typically do given regulations), then the elimination of ABP programs could constitute a Pareto improvement, achieving demand reductions at virtually no pecuniary costs and affording potential welfare gains to both utilities and consumers. In contrast, home energy reports cost approximately $1 per household per report and deliver demand reductions at a cost of $0.033 per kWh. Moreover, consumer welfare is diminished by foregone utility from the energy consumption that is sacrificed to achieve conservation. The magnitude of the welfare costs have not been estimated. More generally, Gillingham et al. (2004) estimated that utility-based demand-side management programs and voluntary and information programs, like the Energy State voluntary label program, achieve electricity conservation at a cost of $0.034 to $0.038 per kWh. Arimura et al. (2011) estimated the cost-effectiveness of demand-side management by utilities to be $0.06 per kWh. While many of the demand-side management policies pursued by utilities and regulators target the residential sector alone, efforts to mitigate the diminution of price salience associated with ABP programs could generate reduced electricity consumption in the commercial sector as well as the residential sector, again, at virtually no cost.

Given that utilities are said to benefit from the payment certainty of ABP and that consumers benefit from convenience, ABP programs could be modified to retain most of the benefits but minimize the price salience declines. For instance, utilities could require customers to open an email containing an electronic bill or access the bill online before the automatic payment is transacted. If the ABP-induced increase in consumption is due, in part, to rational inattention, then one-click payment processing could avert overconsumption without requiring much attention or the loss of much convenience.

Finally, while the implications of ABP-induced overconsumption are important in the setting of this empirical analysis, namely the electricity market, theory suggests these results should not be specific to energy consumption. As the theory developed in Section 2 suggests, overconsumption from price insalience can occur across a wide range of pricing regimes, including those employed by water and gas utilities, as well as home and mobile telecommunications services providers. Because consumer-level consumption data is proprietary in many of these other industries, it is difficult to estimate the degree to which ABP enrollment also induces overconsumption of other goods. However, because ABP programs are implemented across a range of goods and services providers and because the programs are increasingly popular among consumers, it is likely consumer overconsumption due to ABP-induced inattention is considerable. As the results in the preceding section indicate that the magnitude of overconsumption is greater among holders of newer accounts, perhaps reflecting generational changes in budget management, the effect may grow over time.

3.6 Conclusions

Behavioral economists have theorized that agents are less responsive to information that is less salient. These theoretical predictions have been confirmed empirically in a variety of contexts. They motivate the hypothesis of this paper that enrollment in budget billing and automatic bill payment programs reduces price salience, causing residential electricity consumption to increase. After first developing a theory of consumption increases due to ABP that is robust to flexible, non-linear tariffs, like those employed by utilities and telecommunications service providers, the theoretical prediction of salience induced overconsumption is empirically tested using a panel of individual-level commercial and residential electricity consumption. Results suggest that ABP enrollment causes a 2-4.5% increase in residential electricity consumption, with the larger effects exhibited among accounts initiated after 2000. The magnitude of the salience effect is invariant to zip code-level measures of median home value, household income, and age distribution. The greater effect of ABP among recent account holders suggests a generational effect that could cause overconsumption from diminished price salience to increase in coming years. Budget billing is shown to induce an 8-9% increase in electricity consumption among residential accounts. The results are largely invariant to the recency of accounts, consistent with the equal diminution of price salience for deviations from mean consumption among households that regularly review bills and those that do not.

There is also some evidence that ABP increases commercial electricity consumption among newer accounts, though results are less consistent than they are among residential customers. This analysis suggests that the elimination of ABP and budget billing programs or changes to the programs that mitigate the decline in price salience could contribute to energy conservation objectives of utilities and regulators, and do so cost effectively relative to other programs that have attracted the attention of researchers and policymakers alike. It also contributes to a growing literature that documents the importance of information salience to the decision-making of agents characterized by bounded rationality.

Chapter 4

Paying for Pollution? How General Equilibrium Effects Undermine the "Spare the Air" Program

4.1 Introduction

Appeals for cooperation and pecuniary incentives are used in a variety of settings to reduce depletion of scarce resources that are not properly rationed by prices, either due to missing markets or demand-invariant and fixed prices. During World War I, the U.S. government urged citizens to aid the war effort by planting "victory" gardens and conserving food. In World War II, citizens were asked to dress warmly to conserve fuel and to collect aluminum cans to be converted into ammunition. More recently, conservation has been urged amid shortages of energy, water, and medical supplies. These mechanisms have become increasingly common amid heightened concern about environmental damage and increasing scarcity of traditional water and energy resources. They are also used in the context of climate change mitigation, wherein the quantity of environmentally acceptable greenhouse gas emissions is constrained by the greenhouse effect. Whereas explicit rationing and coercive screening can be controversial and politically unsustainable responses to shortages, voluntary conservation programs are generally more acceptable. Whether they are effective is another matter, and the subject of a growing body of research that has yielded mixed results.

Voluntary conservation and pecuniary incentives are used in a number of localities across the United States to reduce atmospheric ozone concentrations on days when air pollution is expected to reach unsafe levels. While air quality has generally improved in recent years throughout the U.S., nearly half of all regions that monitor ground-level ozone periodically exceed federal standards for safe ozone concentrations. Because ozone levels peak on hot, calm, sunny days, ozone exceedances are seasonal and predictable. Therefore, they can, in principle, be avoided with episodic pollution-control regimes rather than permanent pollution-reduction programs, which can be more costly.

This paper studies the effectiveness of non-coercive policy mechanisms in the context of the Spare the Air (STA) program in the San Francisco Bay Area, which, like programs in other cities across the country, aims to reduce ozone pollution on bad-air days by motivating motorists to reduce car trips. It specifically urges motorists to substitute to public transit trips, carpool, or eliminate trips. While the STA initially entailed only public appeals for car trip reduction, it later also provided free transit rides on smoggy days. If the program were to elicit cooperation, then it is expected that car trips would fall on STA days, carpooling would increase, and public transit demand would climb. I test these hypotheses using data on freeway traffic volumes and transit ridership. To preview the results, I find that public appeals for cooperation have no significant effect on prosocial behaviors like carpooling and car trip reductions, though they do increase transit demand. The marginal effect of free transit fares, however, is significant, but significantly positive, meaning car trips *increase* on days when transit is free. Thus, paradoxically, the pecuniary incentives cause additional pollution.

Previous studies of the effectiveness of these voluntary episodic pollution control programs produced conflicting results (Cummings and Walker 2000; Lu et al. 2004; Schreffler 2003; Welch et al. 2005). The work of Cutter and Neidell (2009), henceforth C-N, found a statistically significant 2.5-3.5% decline in car trips as a result of the San Francisco STA program.

This paper advances this literature in several ways. First, whereas C-N analyzed data from 2001 to 2004, this paper draws on data from 2002 through 2009, which not only permits a more current estimate of program effectiveness but also exploitation of a change in the program in 2004 that made free transit rides available on some STA days. This change in program parameters enables separate estimation of the effectiveness of pecuniary incentives and public appeals. Understanding the relative effectiveness of these two components of the STA is important for motivating behavior that is in the interest of society but not individually rational in the traditional sense. The separate estimation of the impact of these two incentives for cooperation supports behavioral theories that extrinsic motivations like rewards and punishments crowd-out intrinsic motivations like altruism. Results demonstrate that expenditures to fund free transit rides are counter-productive relative to program objectives.

Second, beyond estimating the average treatment effect of STA on car trips, analysis unique to this paper provides evidence of whether motorists respond to the appeals for cooperation by making short-run substitutions toward carpooling, as intended by air quality regulators.

Third, this paper considers the validity of regression discontinuity estimates of STA impacts, like those offered by C-N, and finds evidence that inference based on those estimates may be invalid.

Finally, this paper investigates whether responsiveness to STA declines on consecutive STA days. Policy fatigue and the rising cost of intertemporal trip substitutions suggests that it could. However, program salience is likely heightened across consecutive days, which could generate greater responsiveness on STA days preceeded by STA days.

This paper proceeds in Section 2 with a review of the existing literature related to non-coercive rationing programs generally and episodic pollution control specifically. Section 3 provides an overview of the theory of car-trip reduction. Section 4 describes the data and empirical approach of this paper. Section 5 presents empirical results. Section 6 concludes.

4.2 Literature Review

Appeals for cooperation and pecuniary incentives for prosocial behavior can allocate resources in socially beneficial ways by inducing voluntary conservation among resource-consuming populations. Such mechanisms are common in the context of environmental protection, where missing markets are a frequent problem and policies to correct market failures are rare. Appeals for conservation are also made following natural disasters like storms, fires, and earthquakes when safe water, food, and shelter may be in short supply. While voluntary conservation is more politically viable than coercive screening and more efficient than non-pecuniary cost rationing and quotas, it achieves social objectives only to the extent that it induces behavior that is not individually rational in the neoclassical sense. (Olson 1971; Hardin 1968). The costs of cooperation are borne by the cooperators, but benefits accrue to cooperators and shirkers alike, making cooperation subject to the free-rider and under-provision problems typically associated with public goods.

The effectiveness of non-coercive rationing depends upon the inducement of prosocial behavior among a subset of the target population. Prosocial behavior may be generated by intrinsic motivation, such as other-regarding preferences (pure altruism), utility from the "warm glow" of giving, (impure altruism), and benefits derived from social signaling (competitive altruism). It may also be induced by extrinsic motivations, such as rewards for pro-social behavior or punishments for anti-social behavior (Aaron and Schwartz 1990; Andreoni 1989; Batson and Powell 2003; Becker 1965; Benabou and Tirole 2006b). Given that cooperation is expected to be under-provided, public appeals for conservation can have the perverse effect of increasing demand for scarce resources by increasing salience and inducing hoarding (Adelman 2004; De Janvry et al. 2008; Folger 1992; Lynn 1992). Similarly, extrinsic rewards, like pecuniary incentives, can have the perverse effect of reducing cooperation by crowding out intrinsic motivation (Benabou and Tirole 2006b; Ariely et al. 2009; Frey and Jegen 2001; Janssen and Mendys-Kamphorst 2004).

With theory providing inconclusive predictions about the impact of non-coercive rationing, the determination of its effectiveness is the domain of empiricists. Attempts have been made in a variety of settings to determine whether appeals for conservation reduce demand from a cooperative public, or whether they induce greater demand from consumers afraid of resource depletion. Much of the early work in this regard surrounded energy conservation efforts during the energy crisis induced by the Arab oil embargo of 1974. Peck and Doering III (1976) found no reduction in home energy demand due to appeals for conservation. Mayer (1978), however, attributed a ten-percent decrease in home energy use in New Jersey in 1974 to cooperation and changing social norms. Taylor et al. (1984) found that 29 states consumed less energy during the crisis than expected on the basis of demand forecasts and realized prices; only two states consumed less energy than expected one year before the crisis. More recently, Reiss and White (2008) reported that a conservation campaign during the California energy crisis of 2000-01 generated a prolonged 6% drop in demand.

De Janvry et al. (2008) investigated the response to a public appeal for cooperation at the University of California, Berkeley during the 2004 U.S. flu vaccine shortage. Members of the campus community were asked to refrain from consuming scarce flu shots in order to save them for at risk groups like the elderly and infirm. They found that provision of information about scarcity had the effect of increasing demand due to increased salience of the scarce commodity and decreased procrastination in getting shots. The increase in demand was only partly offset by self-restraint exhibited by some individuals.

A number of regions across the country, including two air quality districts in California, operate voluntary episodic pollution control programs that appeal to motorists to reduce car trips on smoggy days. When pollution damages are seasonal and predictable, episodic controls may reduce health and environmental damages from exposure to abnormally low air quality more cost-effectively than permanent controls. Because of the nature of ozone formation, ozone pollution is well-suited for episodic control. The extant literature includes several analyses of these programs. Schreffler (2003) used telephone surveys on daily travel activities to infer a 4.8% reduction in car trips due to the San Francisco STA. Cummings and Walker (2000) found no significant effect of a program in Atlanta, GA, while Welch et al. (2005) observed that a program in Chicago increased traffic volumes during peak demand and reduced traffic volumes during off-peak hours.

C-N improved upon the self-reports of Schreffler and, through reliance upon the regression discontinuity (RD) framework, hoped to obtain more robust estimates of treatment effects than those reported in Cummings and Walker (2000) and Welch et al. (2005). C-N estimated a 2.5-3.5% reduction in traffic volumes due to STA. Their results also indicated some substitution toward public transit and supported their theory that discretionary trips are more likely to be avoided than work-related commutes. Analysis of data comparable to that employed by C-N suggests, however, that regression discontinuity (RD) estimates may be invalid. This paper considers the appropriateness of the C-N methodology and presents evidence that the RD estimates cannot be interpreted as the causal effect of treatment. I compare estimates from the RD design to results from parametric estimation like that used by Cummings and Walker (2000) and Welch et al. (2005).

4.3 Theory of STA trip reduction

STA succeeds in reducing vehicle pollution on smoggy days to the extent that it induces motorists to either carpool, use transit, or take fewer trips. Like other appeals for cooperation, STA relies upon altruistic preferences to motivate car trip reductions. If the sole effect of STA were to elicit prosocial behavior among altruists, then the program should reduce car trips and increase carpooling and transit demand.

Indirect effects associated with STA alerts, however, may lead to increased demand for car trips and less demand for transit. For instance, STA alerts may function as warnings about air pollution, triggering transit passengers to substitute to car trips in order to minimize exposure to bad air. Would-be transit riders may also opt for car trips on STA days in order to avoid expected congestion on mass transit caused by the public appeals. Furthermore, while neoclassical theory holds that adding pecuniary incentives like free fares to the intrinsic incentives of STA alerts should increase cooperation, evidence in the behavioral literature suggests it may decrease cooperation by clouding the social signal from cooperation and lessening the sense of civic duty (Gneezy and Rustichini 2000a,b; Akerlof and Dickens 1982; Frey and Palacios-Huerta 1997; Frey and Jegen 2001; Benabou and Tirole 2006b).

C-N developed a model of utility maximizing travelers who can choose to drive, take transit, or not travel. Utility is defined over consumption of a composite good, health, "environmental altruism", and travel time, where health is decreasing in exposure to bad air and exposure is greater for trips made by transit than by car. STA serves as a signal of poor air quality, which is assumed to be otherwise unobserved.[1] Travel time by car is assumed to be less on STA days due to less congestion, while travel time by transit is assumed to be unchanged. Given this framework, utility from not traveling increases on STA days because of "warm glow" and diminished exposure to ozone. Utility from transit rides may either increase or decrease depending on whether utility from warm glow exceeds disutility from exposure to smog. Expectations of reduced travel time for car trips (due to reduced congestion) *increases* utility from car trips on STA days in the C-N model.

This paper departs from the C-N framework by considering several additional effects that may influence the overall impact of a STA alert, including the possibility that travelers respond by using carpools. Travelers' consideration of congestion effects on both highways and public transit may be an important factor influencing trip decisions and mode choice. These congestion effects are market-level effects, whereas the C-N model is focused solely on the utility of individual travelers. It is reasonable to assume that travel times for both transit and car trips are increasing in the number of total trips made via the respective modes, as is inconvenience associated with travel by public transit. If STA is expected to induce additional demand for transit, then expected travel time on transit increases, causing some transit riders to substitute to car trips or no trips. Thus, trip decisions and mode choice based on expectations about these market effects, can lead to unintended outcomes, like increased car trips.

To incorporate expectations of market impacts and the choice of carpooling in response to a STA declaration, consider the following adaptation of the C-N model. Assume that each individual chooses from among car (C), carpool (CP), public transit (PT), and no trip (NT) for each travel opportunity. The individual chooses the mode that yields the maximum utility among the available options. Thus, for individual i:

$$U_i = \max\left\{U_i^C, U_i^{CP}, U_i^{PT}, U_i^{NT}\right\},$$

where $U_i^C, U_i^{CP}, U_i^{PT}$ and U_i^{NT} denote individual i's indirect utility from each trip mode. Let $T_i^j = 1$ if $U_i^j = \max\left\{U_i^C, U_i^{CP}, U_i^{PT}, U_i^{NT}\right\}$ and $T_i^j = 0$ otherwise, for $j \epsilon \{C, CP, PT, NT\}$. Then define market demands as:

$$Q^j = \sum_i T_i^j$$

for j defined as above. Each trip mode j generates intrinsic utility for individual i,

[1] The federal Clean Air Act requires metropolitan statistical areas with populations greater than 350,000 to report air quality information and provide detailed reports to local media. While air quality forecasts are typically provided with weather forecasts in major metropolitan newspapers, such information may lack the salience of STA alerts, which are typically reported on local newscasts and radio news reports. Sexton (2012) shows that individuals adjust behavior in response to air quality alerts by reducing outdoor activity 18%, but they do not respond to forecasted or realized air quality, itself. Likewise Graff Zivin and Neidell (2009) find a 5-15% decline in attendance at outdoor attractions when air quality alerts are issued.

V_i^j. A constant utility U^0 is also derived from each unit of consumption of a composite commodity with price equal to 1.0.

Utility from each trip mode net of the impact on consumption of the composite commodity is given by:

$$U_i^C = V_i^C - f\left[t^c(Q^C, Q^{CP})\right] - P^C U_i^0 \tag{4.1}$$

$$U_i^{CP} = V_i^{CP} - f\left[t^{CP}(Q^C, Q^{CP})\right] - P^{CP} U_i^0 + s_i^{CP}(STA) \tag{4.2}$$

$$U_i^{PT} = V_i^{PT} - f\left[t^{PT}(Q^{PT}), G(Q^{PT})\right] - P^{PT}(STA)U_i^0 + s_i^{PT}(STA, P^{PT}) - h_i\left[t^{PT}(Q^{PT}), STA\right], \tag{4.3}$$

$$\text{and } U_i^{NT} = V_i^{NT} + s_i^{NT}(STA), \tag{4.4}$$

where P^j denotes the inclusive pecuniary cost of transit mode j for $j \epsilon \{C, CP, PT, NT\}$. Equations 4.1-4.3 incorporate that car trips, carpools and public transit involve disutility from time spent in transit, t^j, which is a function of the total demand for travel by that trip mode.[2] The presence of carpool lanes means that car and carpool travel times may differ. Public transit trips cause disutility from transit time and from congestion, G, which is also a function of total demand, Q^{PT}. Carpools, public transit and trip avoidance (NT) are all associated with utility from the "warm glow" effect, s_i^j, on STA days. The effect may vary by transit mode. In particular, free fares may diminish the warm glow from public transit trips. PT also exposes users to an expected adverse health effect, $h_i\left[t^{PT}(Q^{PT}), STA\right]$, which is a function of exposure to outside air, as determined by transit time and STA status, which serves as a proxy for air quality.

STA impacts behavior only to the extent it motivates a mode switch among individuals. As equations 4.1-4.4 demonstrate, the impacts of STA on trip mode choice are complex. Notably carpooling, a response option not previously investigated, unambiguously yields higher utility on a STA day due to both the presence of warm glow and a likely reduction of congestion on roadways. Travel by private car, of course, yields no warm glow, but the absolute utility from car travel rises if road congestion declines. Equation 4.3 demonstrates that the impact of STA declaration on transit demand is ambiguous as expected trip delays and congestion, as well as potential adverse health impacts, offset utility from warm glow. Free fares increase the individual's net-utility of choosing transit, but may reduce the warm glow consumers receive from altruistic behavior.

Equations 4.1-4.4 also make clear that trip time and congestion impacts are determined at the market level, based upon the aggregate impacts of individuals' actions. Because individuals must commit to their travel mode choice before actually observing

[2] C-N assume away transit time costs associated with elevated demand for public transit on STA days because transit is intended to operate on fixed schedules. However, elevated demand increases both the frequency and duration of stops to board and disembark passengers. If trains and buses are filled to capacity, then transit riders may be delayed by waiting for subsequent buses and trains.

these aggregate effects, their decisions must be based upon expectations formed by past experiences with STA days and media coverage. Given the countervailing effects associated with STA that are described in equations 4.1-4.4, determination of the net-effect of STA on traffic volumes is an empirical question. The next section presents a strategy to estimate the effect of STA and free fares on car trips, carpooling, and transit ridership.

4.4 Econometric Specification and Data

STA days are not exogenously assigned. Rather, they are determined by forecasts of the Air Quality Index (AQI), a measure of air quality defined by the U.S. Environmental Protection Agency. AQI is correlated with contemporaneous and forecasted weather characteristics that influence trip demand and transit-mode decisions. Moreover, air quality may directly impact trip demand. Therefore, estimates of the impact of STA alerts on trip behavior must account for the non-random assignment of STA days. Because STA days are the consequence of an administrative decision that declares a STA day whenever a threshold AQI is reached, the regression discontinuity (RD) framework can theoretically be used to recover estimates of the STA effect with strong internal validity. This is the empirical framework of C-N, which estimated a 3% decrease in car trips due to STA alerts. As popularly understood, identification in the RD design requires only that factors determining car trip demand evolve smoothly with respect to AQI about the treatment threshold and that AQI, the forcing variable, is not subject to endogenous sorting. These assumptions, perhaps trivially satisfied in the context of STA, typically allow that STA days are "as good as randomly assigned" in a neighborhood around the treatment threshold (Lee and Lemieux 2010).

A primary objective of this paper is to determine the effectiveness of small pecuniary incentives for car trip avoidance and to estimate changes in carpooling as a consequence of STA alerts. In estimating these effects, however, this paper also provides additional information about changes in car-trip demand in response to appeals for cooperation on STA days. In particular, this paper presents estimates of the STA effect based on recent data. It also reconsiders the appropriateness of the RD design in this context. As is demonstrated in the subsequent section, the foundational assumption for valid inference in the RD design is likely to be violated in the data used in this analysis and the analysis of C-N. Therefore, and because free fare days can be considered conditionally randomly assigned, the effect of free fares is estimated using a second econometric model that relies on overlap of treatment and control group covariate distributions rather than the RD assumptions. Likewise, estimates of carpooling response are valid irrespective of the RD design in a differencing framework.

4.4.1 STA Effect on Car Trips

The effect of STA alerts on car trip demand is considered first as a baseline specification. Individual changes in trip choices on STA days are not observed. Instead, daily aggregate volumes of car traffic, carpool traffic, and trips on the Bay Area Rapid Transit (BART) commuter rail service are observed. Consequently, an aggregate trip demand equation is specified according to:

$$q_{it} = \beta_1 STA_t + \psi_1 g(o3_t) + \varphi q_{it-1} + STA_{t-1} + \boldsymbol{\delta}\mathbf{Z_t} + \boldsymbol{\xi}\mathbf{D_t} + \theta_i + \gamma_t + \mu_{it}, \tag{4.5}$$

where subscripts i and t denote traffic-monitoring station and day, respectively and q_{it} denotes traffic volume. STA_t is a treatment indicator equal to one if day t is an STA day and zero otherwise, and $g(o3_t)$ is a function relating ozone (in ppm) to traffic volume. The vector $\mathbf{Z}$ represents daily contemporaneous, lagged, and forecasted weather characteristics, including continuous measures of high and low temperatures and precipitation and indicators for forecasted sky conditions (clear and cloudy). The vector $\mathbf{D}$ denotes indicators for day of the week and the interaction of month and year dummies. θ_i is an unobserved monitor-specific effect that accounts for potential correlation in the residuals of monitoring stations across time, and γ_t is an unobserved time effect that accounts for potential correlation in the residuals of days across stations. The idiosyncratic error is μ_{it}. For compactness, let $\varepsilon_{it} = \theta_i + \gamma_t + \mu_{it}$ denote the composite error. Lagged traffic volume, q_{it-1}, and lagged STA day, STA_{t-1} are included in order to control for potential serial correlation and to improve the precision of coefficient estimates (e.g., Lee and Lemieux 2010).

Equation 4.5 is similar to equation 5 in C-N, which modeled the station and time effects as "random effects" in preferred specifications. The random station effect is intended to control for unobserved station heterogeneity, while the random time effect is intended to control for the aggregate nature of the treatment (e.g. Cameron and Miller 2010, Bertrand et al. 2004, Kezdi 2004 and Moulton 1986). Their parametric approach to controlling for two-dimensional error correlation only yields unbiased standard errors if the monitor and time effects are fixed (i.e., they exhibit no decay over time and space, respectively). Because the precise nature of the error correlation is unknown, it is preferable to cluster on two dimensions (Cameron and Miller 2010, Petersen 2009, and Cameron et al. 2006). Therefore, and in order to relate the preferred error specification of this paper to C-N, (4.5) is estimated both with random effects and with two-way clustering. The preferred specification for the RD design in this paper augments (4.5) to include the interaction of $g(o3_t)$ and the treatment indicator, STA, allowing for the slope of the function relating ozone to traffic volume to differ on each side of the treatment threshold (Lee and Lemieux 2010 and Imbens and Lemieux 2008). This yields:

$$q_{it} = \beta_1 STA_t + \psi_1 g(o3_t) + \psi_2 g(o3_t) * STA_t + \varphi_1 q_{it-1} + \varphi_2 STA_{t-1} + \boldsymbol{\delta}\mathbf{Z_t} + \boldsymbol{\xi}\mathbf{D_t} + \varepsilon_{it} \tag{4.6}$$

The coefficient of interest in these models is β_1, which represents the mean change in traffic volume due to STA alerts.

4.4.2 STA and Free Fare Effects on Car Trips

In order to separately determine how travelers respond to public appeals and free transit fares, this paper exploits variation in the use of free fares during the recent history of the STA program. The use of free fares in conjunction with public appeals is conditional on funding to reimburse regional transit agencies for foregone fare revenue. The constraint of funding availability generates exogenous variation in the use of free fares both across years and within years. Thus, the population-averaged marginal effect of free fares on car trip demand is estimated by the coefficient on $FREE$, β_2, where $FREE$ is added to (4.6) and denotes an indicator equal to one if free fares are offered in conjunction with STA and zero otherwise. That is:

$$q_{it} = \beta_1 STA_t + \beta_2 FREE_t + \psi_1 g(o3_t) + \psi_2 g(o3_t) * STA_t + \varphi_1 q_{it-1} + \varphi_2 STA_{t-1} + \varphi_3 FREE_{t-1} + \boldsymbol{\delta}\mathbf{Z_t} + \boldsymbol{\xi}\mathbf{D_t} + \varepsilon_{it}. \quad (4.7)$$

The RD design is estimated using a variety of bandwidths and orders of the ozone concentration polynomial as recommended by Lee and Lemieux (2010) and detailed in the next section. The separate estimation of the effect of free fare days permits interpretation of the coefficient on STA as a "cooperation effect" and the coefficient on $FREE$ as a "price effect."

Because the assumptions for valid inference in the RD design likely do not hold in this application, it is important to observe that identification of the free fare effect, which is of principal interest in this paper, does not depend on the smoothness of the conditional distribution of traffic volume. Rather, an unbiased estimator of the population-averaged effect of free fares on traffic volume is achieved by standard linear methods if, conditional on STA days, free fare days are orthogonal to other determinants of traffic volume, which seems reasonable. Therefore, the marginal effect of free fare days is also estimated by:

$$q_{it} = \beta_1 STA_t + \beta_2 FREE_t + \psi_1 g(o3_t) + \varphi_1 q_{it-1} + +\varphi_2 STA_{t-1} + \varphi_3 FREE_{t-1} + \boldsymbol{\delta}\mathbf{Z_t} + \boldsymbol{\xi}\mathbf{D_t} + \varepsilon_{it}, \quad (4.8)$$

where (4.8) differs from (4.7) only in that the interaction between the STA indicator and $g(\cdot)$ is omitted.

It is perhaps important to note that the "parametric" approach represented by (4.8) and the "locally linear" approach represented by (4.7) are not substantially different, particularly if the STA and ozone interaction is omitted from (4.7) as in some RD specifications. As Lee and Lemieux (2010) noted, the parametric approach in (4.8) can be considered a local linear regression with a very large bandwidth. Likewise, (4.7) can be considered a parametric regression that eliminates the influence of data points in the tails of the forcing variable. Regardless, it is important to note that in estimating (4.8), the smoothness assumptions on the conditional distribution of traffic volume are not maintained, so that the estimate of β_1 is biased if there are unobserved determinants of traffic volume that are correlated with air quality and if the functional form imposed on the relationship between observables and traffic volume is wrong. The linear relationship between traffic volume and ozone is relaxed by considering second and third-order polynomials in ozone concentration.

4.4.3 STA and Free Fare Effects on Public Transit Demand

The effect of public appeals and free transit fares on public transit demand is estimated using data on trips made on BART in order to understand the nature of behavioral responses to public appeals and free fares. While both treatments are expected to reduce traffic volume to the extent that individuals exhibit altruism or respond to pecuniary incentives, they are expected to increase transit demand for the same reasons. Because the data are generated from BART station entrance counts, they do not exhibit the same error correlation as traffic volume data. Nevertheless, (4.7) and (4.8) are estimated for transit trips using two-way clustered errors.

4.4.4 Lagged Effects and Evidence on STA Fatigue

This paper also investigates the dynamics of trip demand surrounding STA and free fare days. First, STA days and free fare days are intended to induce short-run substitutions away from car trips. To the extent that they are successful, they either induce individuals to abandon trips altogether, to substitute to transit or carpools, or to postponse car trips until another day. Car trip postponement could lead to increased demand for car trips the day after STA and free fare days. Free fare days, however, could induce individuals to move up planned transit trips in order to take advantage of the free rides. This could lead to a reduction in transit ridership on days following free fare days. In addition, it is expected that the cost of making short-run substitutions is increasing for each consecutive STA day, such that the response to a consecutive STA day may be diminished. This effect is termed policy fatigue.

On the other hand, a period of bad air quality may generate media attention and, consequently, heighten the salience of the STA program, which could induce greater car trip avoidance on a consecutive STA day. In order to investigate these effects, (4.8) is separately estimated with a suite of lagged and consecutive-day interactions on both the traffic volume data and BART ridership data. Specifically, $L.STA_t$ is a one-day lag of STA and is equal to one if $STA_{t-1} = 1$ and zero otherwise. $L.FREE_t$ is a similarly defined one-day lag of $FREE$. Then $L.STA \times STA$ and $L.FREE \times FREE$ constitute indicators for consecutive STA and free fare treatments, respectively, and $L.FREE \times STA$ is an indicator equal to one if $STA_t = 1$, $FREE_t = 0$, and $FREE_{t-1} = 1$. While in theory, a free fare day could be declared absent the declaration of a STA day, in practice this is observed only once in the data, and that day is omitted from the analysis. Thus, $FREE_t = 1 \Rightarrow STA_t = 1$ and there does not exist a day in the analysed data that is a free fare day and not a STA day. Consecutive STA days occur only 2% of the time during the STA season. Conditional on a STA day occurring, however, the probability that it will be followed by at least one consecutive STA day is 0.44. Likewise, consecutive $FREE$ days occur only 0.6% of the time during the STA season, but conditional on a$FREE$ day occurring, the probability that it is succeeded by a $FREE$ day is 0.67.

4.4.5 Carpooling Response

Finally, in order to test the magnitude of STA and $FREE$ on carpooling, the subset of traffic monitors located along carpool, or High Occupancy Vehicle (HOV), segments was identified. For a traffic monitoring station located within an HOV segment, it is known that lane 1 is the HOV lane and that other lanes within the Bay Area are "mainline" or unregulated lanes. Within the Bay Area, HOV lanes are actuated during morning and evening commute hours when traffic congestion is a problem.[3] During these periods of actuation, only those vehicles that carry the minimum number of passengers (typically two in the Bay Area, but three along some segments) may lawfully travel in the HOV lane. Vehicles that travel in the HOV lane with fewer passengers are subject to fines in excess of $300. Restricted access to HOV lanes during actuation typically provides travel time savings to vehicles traveling in the HOV lanes. Identification of the carpool response relies upon these characteristics of HOV lanes to generate separating equilibria in which carpools self-select into carpool lanes to reduce travel times and mainline traffic selects out of HOV lanes in order to avoid fines.

The analysis of carpool response follows a generalization of the triple differencing framework that leverages arbitrary treatment patterns and repeated observations on each station. The change in HOV-lane traffic volume during actuation is compared to changes in mainline traffic volume during actuation and traffic volume in HOV lanes outside actuation periods, controlling for time, lane, and station fixed effects. The triple difference is estimated in a regression context in order to improve precision by conditioning on observed covariates. Specifically, the mean effect of STA days on carpooling is estimated by:

$$q_{itj} = \alpha_0 + \alpha_1 HOV_i + \alpha_2 STA_j + \alpha_3 ACT_t + \alpha_4 HOV_i * STA_j + \alpha_5 HOV_i * ACT_t + \alpha_6 STA_j * ACT_t + \alpha_7 HOV_i * ACT_t * STA_j + \varepsilon_{itj}, \tag{4.9}$$

where i indexes station, t indexes time of day, and j indexes date. HOV is an indicator equal to 1 if the monitor is in lane 1 of an HOV segment and 0 otherwise, ACT is an indicator equal to 1 if the time of day is during a period of carpool lane actuation and 0 otherwise, and STA is as previously defined. Equation (4.9) is estimated by OLS with two-way clustered standard errors and with the full complement of observed covariates to improve precision. Interest is in the coefficient on the triple interaction, α_7. In order to reduce measurement error from inertia in lane movement at the beginning and end of carpool actuation periods and carpool lane segments, one-hour traffic volumes on each

[3] Exact times vary by HOV segment, but are typically between 6-10AM and 3:30-7PM.

side of the morning and afternoon actuation periods are dropped and attention is restricted exclusively to monitors along HOV segments for the baseline specification. In order to further boost the power of the signal, attention is also restricted in some specifications to HOV monitors and hours when HOV segments experience congestion. This increases the likelihood that the separating equilibria obtain and that all carpools are reflected in traffic volumes for HOV monitors. It is worth noting that identification does not depend on perfect sorting. Rather, it is sufficient that sorting behavior not vary across STA days, which seems reasonable given the relatively small "market-level" effects from STA days reported in the subsequent section. Finally, the restriction to analysing data from HOV segments omits a fourth dimension of differencing: the difference in lane 1 across HOV and non-HOV segments. As a final specification, a quadruple difference is specified, such that interest rests in the coefficient on the quadruple interaction between an HOV segment indicator (equal to one if the station is located within an HOV segment and zero otherwise), a lane indicator (equal to one if the lane is lane 1 and zero otherwise), and the actuation and STA day indicators as previously defined.

4.4.6 Data

Data on traffic volumes are generated by a network of vehicle detectors embedded in the roadways. The network is managed by the California Department of Transportation, in conjunction with the UC Berkeley Department of Electrical Engineering and Computer Science. It includes 11,716 stations throughout the state and 1,275 stations in the Bay Area that offer multiple observations on traffic volume across all lanes of traffic along each major highway segment. This analysis relies on observations of ten randomly selected stations along each of 40 freeway segments in the Bay Area, except where there are fewer than ten stations, in which case all stations along these segments are utilized. This yields a sample of 316 stations throughout the Bay Area.

The network generates reports with traffic counts for each station every five minutes. This analysis relies on aggregated hourly and daily reports on traffic volumes from 2002-2009.[4] Because stations were installed at different times and are subject to malfunction, the number of observations varies by station, yielding an unbalanced panel.[5] It is assumed that these station errors are uncorrelated with STA and therefore do not bias the estimation. These same reports provide traffic counts separately for each lane. Observations on HOV lane volumes are derived from these same data. Data on transit ridership were obtained from station entry data provided by BART pursuant to a public records request and cover the period 2002-2008.

Data on STA days and ozone forecasts were obtained from the Bay Area Air Quality Management District (BAAQMD). Sixty STA days and 12 STA free fare days are observed from 2002 to 2009. Following the BAAQMD's administrative decision rule for issuing STA

[4]Data were not available before September 2001.

[5]For a thorough description of these data and the Freeway Performance Management System, see Chen (2003).

Table 4.1: Summary Statistics: Number of STA days by year

			Bandwidth		
			All	Wide	Narrow
Year	STA=1	FREE=1	STA=0	STA=0	STA=0
2002	7	0	172	35	6
2003	8	0	172	83	21
2004	4	2	176	43	8
2005	2	1	176	66	12
2006	11	6	169	77	13
2007	2	2	178	42	12
2008	14	0	165	103	26
2009	12	0	168	63	15
Total	60	11	1376	512	113

days, the maximum forecasted ozone concentration and AQI across BAAQMD regions is used as a covariate. The STA season initially ran from June through October, though it was later expanded to include May. Consequently, this analysis considers observations from May to October for the period 2002-2009. Data on contemporaneous weather-related variables (high and low temperatures and precipitation) were obtained from the Surface Summary of the Day provided by the National Climatic Data Center. GIS software was used to assign observations from eight Bay Area weather stations to traffic stations according to proximity between weather and traffic stations. Forecasted high and low temperatures and sky conditions were obtained from the NCDC's coded city forecasts. Forecasts from four Bay Area weather stations that issue such reports were assigned to traffic stations in the same manner as contemporaneous weather variables.

Summary statistics are reported in tables 4.1-4.2. Table 4.2 demonstrates that the covariates are relatively well balanced across RD bandwidth specifications.

Table 4.2: Summary Statistics: Means of dependent variables and covariates

	All Observations	RD Wide	RD Narrow
Traffic Volume	58,237.16 [25,531.88]	58,290.53 [25,561.37]	57,661.46 [25,611.52]
BART Volume	6490.61 [6565.77]	6602.18 [6667.63]	6607.69 [6750.40]
Precipitation	0.01 [0.98]	0.00 [0.001]	0.00 [0.001]
Max. Temp.	74.05 [9.89]	78.70 [11.08]	84.39 [11.72]
Min. Temp.	54.73 [3.84]	55.53 [3.40]	57.54 [4.03]
Forecast High	77.91 [9.89]	83.36 [10.31]	89.38 [10.16]
Forecast Low	54.47 4.79]	56.18 [4.94]	58.61 [4.48]
Forecast Day Clear	0.56 [0.50]	0.72 [0.45]]	0.79 [0.41]
Forecast Day Rainy	0.03 [0.17]	0.005 [0.07]	0.00 [0.00]
Forecast Night Clear	0.57 [0.50]	0.68 [0.47]	0.75 [0.43]
Forecast Night Rainy	0.02 [0.14]	0.004 [0.07]	0.00 [0.00]
Holiday	0.02 [0.15]	0.02 [0.14]	.0.02 [0.14]

Std. Dev in brackets.

4.5 Results

In order to provide a baseline comparison to C-N, an aggregate STA effect was first estimated by the RD specifications in (4.5) and (4.6) for data from 2002-2004, roughly the period of study in C-N. In recognition of the tradeoff in bandwidth selection between the precision of larger bandwidths and the bias from functional form misspecification, and in order to demonstrate stability in coefficient estimates, the two bandwidth specifications of C-N are adopted (in addition to a bandwidth that covers the full support of the forcing variable). The widest bandwidth ("wide") includes observations within 20 parts per billion (ppb) of the ozone concentration that triggers STA status. A "narrow" bandwidth uses observations within 10ppb.[6] Throughout, (4.6) is estimated with cross-sectional clustering on assigned weather station, the highest level of cross-sectional aggregation among covariates (Cameron et al. 2006).

The results of the baseline estimation are reported in Table 4.3. The top panel of the figure presents the coefficients and standard errors reported by C-N. The middle panel presents estimates of the STA effect using the C-N specification on data from 2002-2004.[7] Like C-N, it shows a statistically significant negative effect of STA days on traffic volumes across bandwidths. The point estimates suggest a magnitude of effect comparable to that estimated by C-N (1.7-2.8% decrease versus 1.7-3.5% decrease in C-N). The bottom panel of Table 4.3 reports estimates based on the two-way clustering specification in (4.6) that flexibly allows for differing slopes on either side of the treatment threshold. A significant effect is estimated in only the narrowest bandwidth. In each panel, the point estimates exhibit a lack of stability across bandwidths, which raises doubts about the appropriateness of the RD design.

4.5.1 STA and Free Fare Effects on Car Trips in the RD Framework

As interest in this paper centers on estimating the changes in transportation mode and trip demand induced by free fares, the RD design is also employed to separately estimate the effect of STA on traffic volumes and the marginal effect on traffic volumes of $FREE$. Results of this estimation, based on (4.7), are reported in Table 4.4. Results are reported for the two bandwidths previously defined and a middle bandwidth (± 0.015 ppm. of the treatment threshold), as well as for each of three orders of the forcing variable polynomial (first order, second order, and third order), in order to demonstrate whether estimates

[6]The U.S. EPA changed the correspondence between AQI and ozone concentration in 2008. Hence, the wide bandwidth corresponds to AQI between 51 and 151 before 2008 and between 47 and 151 after 2008. The narrow bandwidth corresponds to AQI of 77-127 before 2008 and 71-127 after 2008.

[7]This paper uses the same data as C-N. However, data from the year 2001 were not available at the time of this analysis. Data administrators at the California Department of Transportation reported that 2001 data were never available from the department's repository. Because 2001 data were not available and because C-N sampled the data as this paper does, it is impossible to provide a precise replication.

Table 4.3: Effect of STA Day on All Day Traffic: Pre-2005

	(1) All observations	(2) ±0.02 of threshold	(3) ±0.01 of threshold
		C-N (2009)	
Monitor and station random effects	-1105.97	-2332.26**	-2009.98*
	(823.08)	(857.49)	(1010.08)
	-[1.7]	-[3.5]	-[3.0]
No. obs.	70,805	24,073	8,768
No. of days	536	179	67
No. of monitors	142	142	142
		Random Effects 2002-2004	
Monitor and date random effects	-996.93***	-1118.34***	-1634.26***
	(127.53)	(171.89)	(271.51)
	-[1.7]	-[1.9]	-[2.8]
No. obs.	161,077	59,931	15,215
Mean No. of days	509.7	167.5	48.1
No. of monitors	316	316	316
		Two-way Clustered Standard Errors 2002-2004	
Weather station and date clustered errors	-2516.47	-2605.83	-5037.52***
	(1849.85)	(1969.78)	(1793.14)
	-[4.3]	-[4.5]	-[8.7]
No. Obs.	161,077	59,931	15,215
Mean No. of days	509.7	167.5	48.1
No. of Monitors	316	316	316

$*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Std. errors in parentheses.

Change as a percent of mean in brackets.

are robust to various functional form assumptions (Lee and Lemieux 2010). Results are also presented for the C-N specification (equation 4.5 with the inclusion of $FREE$) and the preferred specification (equation 4.7).

As Table 4.4 reveals, however, the RD estimates do not appear robust to the inclusion of higher order polynomials or to varying bandwidth specifications. The point estimates for STA are particularly unstable. They range from significantly negative to significantly positive, suggesting in some instances statistically significant effects of such a large magnitude that they exceed even the maximum recorded traffic volumes along any monitor in the Bay Area. Such estimates are likely a consequence of over-fitting the data with higher order polynomials, particularly within small bandwidths (Lee and Lemieux 2010). Nevertheless, the lack of robustness is cause for concern. In contrast, and except at the smallest bandwidth, the FREE effects exhibit greater stability and are generally statistically significant at the 10% level, at least, and robust to higher order polynomials and varying bandwidths. Outside the narrowest bandwidth, the statistically significant point estimates on FREE range from 843.78 to 1480.44, suggesting an *increase* in traffic volumes due to free fares in the range of 1.4-2.5%. It is also notable that columns 1 and 2 provide complementary estimates to those reported in Table 4.3 for the period 2002-2009.

The instability exhibited in estimates of STA in Table 4.4 and in Table 4.3 suggest the RD estimates are of questionable validity. As suggested in the previous section, the assumptions of the RD design as commonly understood are seemingly trivially satisfied. That is, it is assumed there is no endogenous sorting about the treatment threshold as motorists and other commuters are unable to influence AQI forecasts and air quality regulators are expected to have little incentive to manipulate the assignment of STA days. The "as good as random assignment" result, therefore is likely to hold. However, though perhaps frequently taken for granted, inference in the RD design requires either that the conditional distribution of the outcome and such features of that distribution as the conditional expectation evolve smoothly in the forcing variable everywhere except at the treatment threshold or that any discontinuities in the conditional distribution with respect to the forcing variable be explained on substantive grounds. Otherwise, the interpretation of a discontinuity at the threshold as a causal effect of the treatment is dubious (Imbens and Lemieux 2008).

Often, simple graphical analysis is sufficient to confirm that this condition is satisfied (e.g. Lee 2008 and Lee et al. 2004), and it is therefore recommended that RD analyses include a few select graphical presentations. In the present context, however, graphical analysis of the conditional mean of traffic volume raises serious concerns about the validity of the RD design. Figure 4.1a plots mean traffic volume in each of 39 equal-sized AQI bins using substantially similar data to that employed by C-N. In contrast to the conditional expectation functions exhibited in Lee (2008) and Lee et al. (2004), the conditional expectation of traffic volumes does not evolve smoothly near the treatment threshold or away from it. Figure 4.1b plots, according to AQI bins, the mean residual from a regression of traffic volume on year, month, and day-of-the-week indicators, lagged traffic volume, and a host or contemporaneous and lagged weather outcomes and forecasts. Even conditioning

Table 4.4: STA and Free Effects in the Regression Discontinuity Design

Bandwidth		(1)	(2)	(3)	(4)	(5)	(6)
All Observations	STA	338.17	-516.51***	-8535.61**	-5627.14***	-125,812.9	-34,553.35***
		(1876.61)	(84.55)	(3684.71)	(1612.46)	(155,857.3)	(7819.96)
No. of obs.: 379,815							
No. of monitors: 316	Free	1224.67**	866.61***	1244.03**	887.78***	1277.18**	884.80***
Mean No. of days: 1201.9		(554.41)	(156.11)	(548.21)	(156.52)	(646.27)	(159.07)
±0.02 of threshold	STA	3369.90	-729.89***	5887.31	14,272.01**	-125,812.9	-117,092***
		(3667.84)	(62.12)	(29,998.51)	(6387.07)	(155,857.3)	(48,762.82)
No. of obs.: 146,785							
No. of monitors: 316	Free	1436.58**	1219.3***	1480.44**	902.13***	1277.18**	1093.47***
Mean No. of days: 464.5		(664.49)	(83.97)	(686.94)	(156.28)	(646.27)	(171.75)
±0.15 of threshold	STA	1919.92	-562.45***	91,138.56	18,906.94***	680,518.6	-5763.50
		(3795.05)	(66.04)	(76,685.8)	(7129.51)	(1,235,961)	(65,563.37)
No. of obs.: 108,468							
No. of monitors: 316	Free	971.2264	1179.09***	1277.32*	843.78***	1295.26*	226.33
Mean No. of days: 343.5		(689.02)	(84.48)	(775.06)	(156.09)	(768.39)	(145.96)
±0.01 of threshold	STA	2058.66	-674.52***	224,864.6**	112,909***	680,518.6	-3310840***
		(3831.17)	(108.94)	(100,195.6)	(24,958.71)	(1235961)	(621048.5)
No. of obs.: 40,952							
No. of monitors: 316	Free	971.23	-395.83***	5.45	-214.21	1295.26*	221.81
Mean No. of days: 129.6		(689.02)	(107.65)	(1011.83)	(207.91)	(767.63)	(236.34)
S.E. Correlation Correction		Cluster	RE	Cluster	RE	Cluster	RE
Order of O3 ppm polynomial		1	1	2	2	3	3

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Std. errors in parentheses.

on a suite of observables thought to influence car trip demand, the relationship between traffic volumes and AQI exhibits no smooth relationship nor even a visible discontinuity at the treatment threshold, as is seen in Lee et al. (2004) and Lee (2008), for example. Rather, the relationship remains idiosyncratic. Restricting attention to observations from 2002 to 2004 yields nearly identical figures. Such figures were not included in C-N. Far from supporting the RD approach and providing a visual representation of the treatment effect, this analysis raises further doubts about inference based on RD estimates.

More formal tests of RD validity draw from parallel tests of pseudo outcomes in the treatment effects literature based on unconfoundedness, e.g., Imbens (2004). Such tests include replacement of the dependent variable with covariates in the RD specification. In general, any sign of a discontinuity at the treatment threshold would violate the RD assumptions and suggest that an observed discontinuity could be caused by factors other than the treatment. In particular, lagged covariates provide a compelling test of RD validity if they are realized before the forcing variable (Lee and Lemieux 2010, Lee et al. 2004). In the present context, lagged traffic volumes can be substituted for contemporaneous traffic volumes in (4.7) to provide a test of RD validity. Given that STA days are declared one day in advance, STA is likely to have no effect on a one-day lag of traffic volume and should certainly have no effect on a two-day lag.[8] However, across several specifications, the null hypothesis of no effect of STA day is consistently rejected at the 5% and 10% significance level in favor of the alternative of a decrease in traffic volume. Similarly, based on tests regressing contemporaneous traffic volume on the treatment and the full complement of covariates, the null hypothesis of smoothness away from the treatment threshold is rejected in various specifications and at arbitrary bandwidths on each side of the threshold (Imbens and Lemieux 2008). The sum of this evidence, then, suggests that the RD framework is inappropriate in estimating the average treatment effect of STA days using traffic monitor data. The interpretation of discontinuities as the causal effect of STA alerts is dubious. Unbiased estimation of the effect of $FREE$, however, does not rely on the RD assumptions. And, as is shown below, the point estimates of the marginal effect of free fares are relatively consistent across the parametric and locally linear RD specifications.

4.5.2 "Parametric" Estimation of Free Fare Impacts on Traffic Volumes and BART Ridership

Because assignment of free fare days is exogenous conditional on STA days, estimation of the marginal effect of free fares on car trip demand is unencumbered by the potential pitfalls of the RD approach previously enumerated. Instead, estimation can proceed based on "parametric" regression using (4.8). These results are reported in columns 1-

[8] To the extent that individuals respond at time t to an announcement at time t of an STA day at time $t+1$, then one would expect to find an increase in traffic volumes at time t as individuals make intertemporal substitutions away from car trips at time $t+1$.

Figure 4.1: Conditional Distribution of Traffic Volume

(a) Unrestricted Mean Traffic Volume by AQI Bin

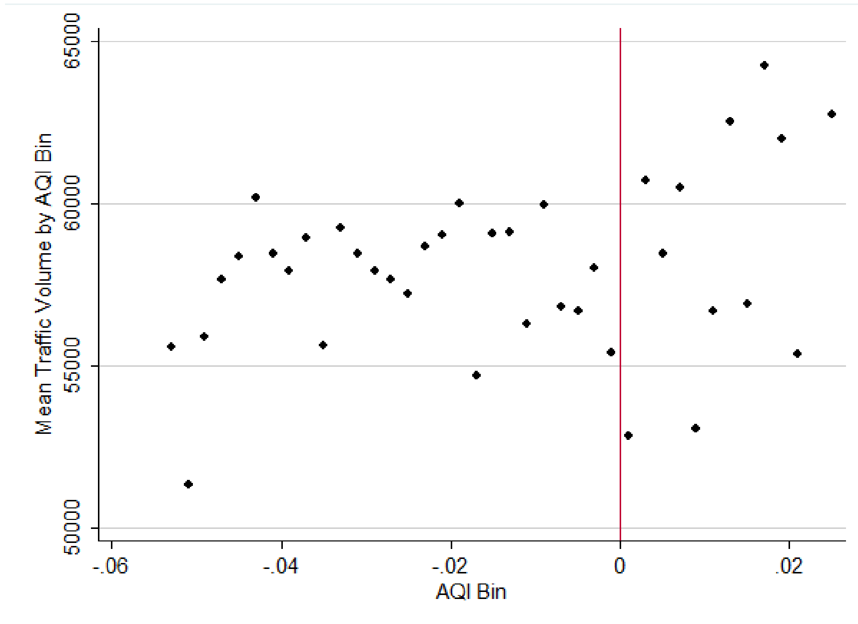

(b) Mean Residuals by AQI Bin (from regressing traffic volume in full complement of covariates except STA and $FREE$)

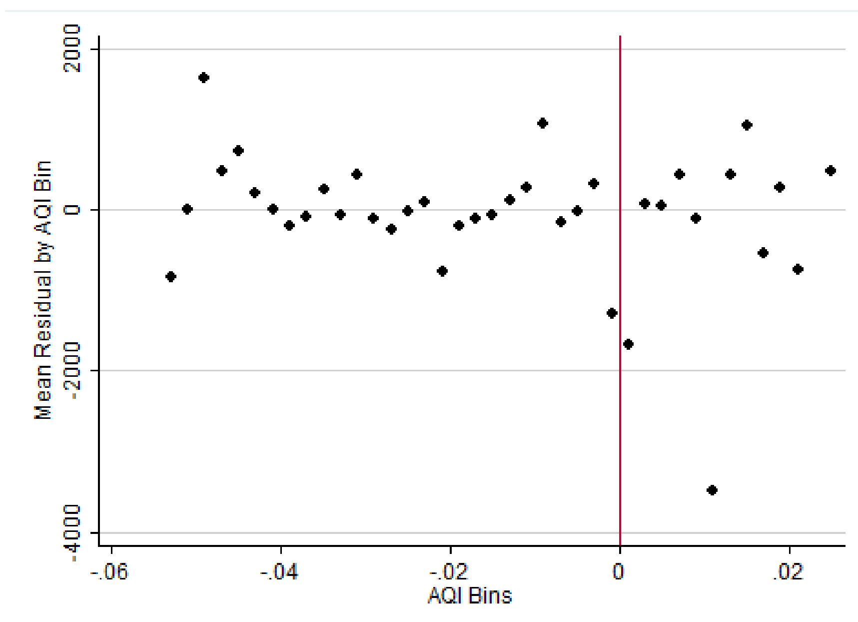

3 of Table 4.5 with two-way clustered standard errors for each of three orders of the ozone polynomial. The linearity assumption is maintained for other covariates. In each specification, free fares are estimated to cause a statistically significant *increase* in traffic volumes, on the order of 1.8% of daily traffic volume. This confirms estimates of the causal effect of free fares from the RD design and suggests a perverse consequence of free transit rides, which are intended to reduce car trips on STA days. An interpretation of this result is provided in the following section. Table 4.5 also reports estimates of the effect of STA, which is a neccesary control in this context. While the estimates are of the expected sign, none are statistically significant, which may reflect a downward bias from the imposed linearity assumption, although figures 4.1a and 4.1b suggest a higher order polynomial would provide little improvement in fit.

The parallel analysis is conducted for changes in transit ridership due to free fare days using data on BART station entrances. These results are reported in columns 4-6 of Table 4.5.[9] Across the polynomial specifications, free fare days are shown to have a statistically and economically significant marginal effect on transit demand. The estimates are highly stable, significant at the 1% level and of the expected sign. Free fares are estimated to cause a 3.6% increase in trips on BART. In addition, the point estimates of the coefficient on STA are highly significant suggesting BART trips increase 1.5% in response to STA alerts. The point estimates for STA should, however, be interpreted with caution as they may suffer from omitted variables bias. Nevertheless, the direction of the bias should be the same as the direction of the bias in estimation of STA effects on traffic volume. That is if the insignificance of STA in determining traffic volumes is explained as a consequence of attenuation bias, then it is expected the true effect of STA on BART ridership is greater than 1.5%. If it is assumed that the estimates of STA effects are biased down, then the true aggregate effect of a free fare day is greater than a 5.1% increase in BART demand (given by the sum of the estimated STA and FREE effects).

4.5.3 Lagged Effects, Salience, and Policy Fatigue

Additional information about aggregate responses to STA alerts and pecuniary incentives can be inferred from the lagged treatment variables and their interactions with contemporaneous treatment, as discussed in section 4.4.4. The results from estimating a version of (4.8) augmented to include these variables on traffic and BART data are reported in Table 4.6.[10] They show that, while there is no statistically significant effect from

[9]STA day and free fare day effects on BART demand were also estimated by the RD specification. The marginal effect of free fares on STA days was estimated to be statistically significant (at the 1% level) and positive in lower order polynomial specifications with “wide” and “medium” bandwidths. These point estimates range from 228.83-262.38. Estimates of the STA day effect vary widely across specifications and are statistically significant at only higher order polynomials in the “middle” bandwidth. The complete table of results is omitted for compactness, but is available from the author.

[10]Previously reported estimates in this paper conditioned on lagged treatment to control for serial correlation, though the point estimates for those lags were not reported or discussed. Here they are expressly considered in conjunction with the lagged and contemporaneous treatment interactions that

Table 4.5: Contemporaneous effects of Free fares on traffic volumes and BART ridership

	(1)	(2)	(3)	(4)	(5)	(6)
	Traffic Volume			Transit Ridership		
STA	-366.99	-476.95	-304.57	99.03***	98.41***	113.82***
	(429.21)	(453.97)	(450.43)	(27.47)	(27.35)	(29.02)
	–	–	–	[1.5]	[1.5]	[1.8]
Free	1025.93**	1008.18*	1054.29**	234.31***	234.32***	234.19***
	(519.66)	(527.99)	(519.86)	(42.52)	(43.18)	(47.69)
	[1.8]	[1.7]	[1.8]	[3.6]	[3.6]	[3.6]
Order of O3 ppm polynomial	1	2	3	1	2	3
No. obs.	379,815	379,815	379,815	42,714	42,714	42,714
No. of days	1202	1202	1202	1017	1017	1017
No. of monitors	316	316	316	42	42	42

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Std. errors in parentheses.

Change as a percent of mean in brackets where statistically significant.

a free fare day one day ago on traffic and BART volumes today (see $L.FREE$), netting out the lagged effects and interactions generates larger point estimates for the marginal effect of $FREE$. The magnitude of the effect of free fare days on traffic volumes increases roughly 50% to 2.7-2.9% of traffic volume (see columns 1-3). The magnitude of the free fare effect on BART demand almost doubles to 6.3-6.4% (see columns 4-6). However, the statistically significant effect of STA on BART ridership observed in Table 4.5 is lost. In both the traffic and transit data sets, STA and lagged STA are insignificant.

To the extent that STA alerts and free fares induce short-run trip avoidance or mode substitution, it is expected that responses will be muted for consecutive treatments as the willingness to repeatedly make such substitutions declines, i.e. policy fatigue sets in. Policy fatigue suggests that the marginal effect of the consecutive treatment, as represented by the coefficient on the interaction of lagged and contemporaneous STA and $FREE$ should be of the opposite sign of the main effect. That is, the marginal effect of consecutive STA and free fare days should be positive for traffic volumes and negative for BART ridership. However, if the salience of STA alerts and free fares is increasing over consecutive exposures, then the marginal effects may be of the same sign as the main effects. As shown in Table 4.6, there is conflicting evidence as to whether the salience effect or policy fatigue effect dominates. On BART, consecutive STA days have a highly significant and positive effect on demand, suggesting the importance of salience effects. On the other hand, consecutive free fare days have a highly significant negative effect on ridership that exceeds the main effect. This result for consecutive free fare days is suggestive of policy fatigue. Moreover, it suggests that individuals may make intertemporal substitutions to transit trips on free fare days in order to capitalize on the free transit rides. If planned, non-commute trips are moved up to free fare days, then overall transit demand may fall on subsequent days. Finally, in the BART context, the highly statistically significant and negative point estimates for the coefficient on "L.Free x STA" suggests that a STA day preceded by a free fare day induces no additional transit demand, and, given that the main effect of STA days on transit demand is indistinguishable from zero, the net effect is negative. There is evidence of a similar effect on car trip demand.

4.5.4 Impacts on carpooling

Though this paper finds no consistent and compelling evidence of an aggregate effect of STA alerts on car-trip demand, it is possible that the aggregate effects conceal responses in carpooling behavior. The triple differencing framework afforded by the robustness of traffic volume data allows that STA is as good as randomly assigned. However, estimation yields no evidence of a significant effect of STA alerts on carpooling. As described in section 4.4.5, identification of the carpooling response to STA days depends upon obtaining a separating equilibrium whereby carpools self-select into HOV lanes for travel-time savings and non-carpool traffic self-selects out of HOV lanes to avoid penalties. Traffic counts

are unique to this specification.

Table 4.6: Lagged Effects of STA, Free and Treatment Dynamics on Traffic and Transit Demands

	(1)	(2)	(3)	(4)	(5)	(6)
	Traffic Volumes			Bart Ridership		
STA	-630.33	-700.78	-568.75	-15.05	-4.88	14.98
	(567.77)	(570.79)	(555.36)	(23.94)	(23.79)	(28.54)
	–	–	–	–	–	–
Free	1596.71***	1573***	1681.52***	409.40***	409.76***	415.33***
	(639.76)	(657.08)	(634.06)	(57.68)	(58.55)	(62.69)
	[2.74]	[2.7]	[2.9]	[6.3]	[6.3]	[6.4]
L.STA	221.41	209.32	187.14	17.88	19.66	15.08
	(395.44)	(397.32)	(395.99)	(37.54)	(38.10)	(38.17)
	–	–	–	–	–	–
L.Free	467.65	461.38	433.48	-96.82	-96.16	-98.98
	(375.85)	(374.33)	(377.46)	(116.54)	(119.35)	(121.17)
	–	–	–	–	–	–
L.STA x STA	894.01	869.91	1144.06*	348.01***	351.62***	395.59***
	(684.26)	(689.39)	(395.99)	(69.05)	(62.88)	(57.29)
	–	–	[2.0]	[5.4]	[5.4]	[6.1]
L.Free x Free	-1060.51	-988.57	-1521.87	-453.07***	-453.17***	-461.19***
	(951.51)	(992.55)	(1012.67)	(71.38)	(74.66)	(87.52)
	–	–	–	[7.0]	[7.0]	[7.1]
L.Free x STA	-1389.76**	-1432.54**	-995.53	-233.52***	-236.10***	-241.10***
	(613.12)	(636.01)	(691.01)	(51.27)	(48.06)	(44.27)
	[2.3]	[2.5]	–	[3.6]	[3.6]	[3.7]
Order of O3 ppm polynomial	1	2	3	1	2	3
No. obs.	379,815	379,815	379,815	42,714	42,714	42,714
No. of days	1202	1202	1202	1017	1017	1017
No. of monitors	316	316	316	42	42	42

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

Std. errors in parentheses.

Change as a percent of mean in brackets where statistically significant.

are observed separately by lane. In the first several specifications, attention is restricted to HOV segments in a triple differencing framework to avoid measurement error due to inertia in the response of motorists to changes in the HOV status of highway segments.

I estimate the triple differences estimator in regressions that (1) include all complete hours of actuation and omit all hours surrounding actuation (to reduce measurement error from driver inattention to temporal changes in HOV status); (2) include only actuation hours during peak morning and evening commutes and exclude all other actuation hours and surrounding hours; and (3) include only observations from HOV segment-hours that experience congestion. These iterations of the triple differences regression were all intended to boost the signal and diminish noise. None of these regressions generated estimated treatment effects that were significant at even the 10% level. Finally, I also exploited a fourth dimension of variation in a quadruple differencing framework that included non-HOV segments in order to exploit variation in lane 1 HOV status. Following the same iterative process as before, this procedure did not generate any significant treatment effects. A table of these regression results is omitted for brevity.

4.6 Discussion and conclusions

Previous studies have estimated the responsiveness of motorists to episodic appeals for car trip reductions (Cutter and Neidell 2009; Welch et al. 2005; Lu et al. 2004; Schreffler 2003; Cummings and Walker 2000). However, this is the first analysis to separately estimate responses to public appeals for cooperation and pecuniary incentives, which form a dual treatment in some programs, including the STA program in San Francisco. While the assignment of STA days according to forecasted exceedances of threshold air quality levels can theoretically yield estimated treatment effects with strong internal validity, this paper presents evidence that raises doubts about the validity of RD estimates from aggregate traffic monitor and BART ridership data, like those offered by C-N. Absent a compelling RD framework, this analysis finds no evidence of car trip reductions in response to public appeals for cooperation on STA days. "Parametric" estimates may, however, suffer from omited variables bias and are not robust to flexible functional form assumptions. There is, however, evidence that STA alerts boost transit ridership. Estimates of the effect of STA days on BART demand suggest public appeals for cooperation generate a 1.5-1.8% increase in BART demand on STA days. The result seems to be driven by consecutive STA days, which suggests that heightened salience of STA alerts may be important to generating intended responses.

The estimated increase in BART rides due to STA alerts contrasts with C-N, who found no statistically significant effect over the period 2001-2004. Because the STA effects net out the impact of pecuniary incentives on free fare days, increased transit demand on STA days can be interpreted as reflecting cooperation motivated by altruism or other instrinsic considerations. It is unlikely the significant STA effects reflect added or additional trips. Rather, they should reflect substitutions from car trips. This result is difficult to reconcile with the lack of a statistically significant reduction in car trips due to STA alerts. However, it may suggest that estimates of the STA effect on traffic volumes are biased downward by omitted variables or measurement error. On the other hand, the BART estimates should suffer a bias in the same direction.

While appeals for car trip reduction on STA days may go unheeded, the offer of pecuniary incentives in the form of free transit fares has the perverse effect of increasing car trips in addition to increasing transit trips. BART demand is estimated to increase 3.6% on free fare days, or roughly 12,600 trips. Traffic volumes are estimated to increase by 1.7-1.8% or 300,000 car trips. The fact that car trips and BART ridership both increase on free fare days suggests that free fare days generate additional transit trips, i.e., trips that otherwise would not have been undertaken, *and* substitutions away from BART towards car trips for regular commutes. Interpretation of these results as evidence of intertemporal substitution of discretionary trips (as opposed to work commute trips) is supported by newspaper accounts of BART trains inundated on FREE days by passengers destined for shopping districts and other leisure destinations. They also suggest that substitution to car trips from transit trips is motivated by fears of crowding on trains and even violence among those making discretionary trips. These results are also consistent with theories that intrinsic motivations, like those that could induce car trip avoidance, are crowded out by extrinsic rewards, like pecuniary incentives for choosing alternative transportation modes. The substantial increase in car trips likely represents not only substitutions away from BART, but also substitutions away from buses, bicycling, or walking that may be imotivated by health concerns induced by STA alerts.[11] A caveat to these results is that they rely on traffic counts along highways and major roadways and ridership data for commuter rail. To the extent that car trip demand is systematically different across surface streets and highways, these results will not be representative of the system-wide car-trip demand response to the STA program. Likewise, high-speed commuter rail demand may systematically differ from demand for bus trips, for instance, limiting the appropriateness of extrapolation to system-wide impacts.

Regardless, this analaysis points out that failure to consider general equilibrium effects of policies like the STA program can lead to wildly inaccurate predictions of behavioral responses. Far from a pollution control program, STA in the San Francisco Bay Area has been a pay-for-pollution program. The expenditure of pollution control funds, estimated at $2.5 million per free fare STA day, generates increased car trips and subsidizes added trips on public transit. To the extent STA induces additional transit trips or inter-temporal trip substitutions to capitalize on free rides, then STA may not only result in air quality regulators paying for pollution, but also in the inducement of low-income individuals, who are perhaps most likely to take up free fares, to travel on those days when it is least healthy to do so.

This paper underscores the importance for decision makers to consider the general

[11] Evidence suggests that air quality directly induces exposure avoidance among some segments of the population (Sexton 2012; Graff Zivin and Neidell 2009). Regardless, the free fare effect is robust to a number of specifications of the function relating air quality to trip demand.

equilibrium effects associated with policies intended to elicit cooperation. While free rides on public transit seem relatively benign, this analysis has shown that not only may the subsidized transit trips not generate net substitution towards transit, they may generate net substitution away from transit. As decision makers come to increasingly rely upon non-coercive or voluntary programs like the STA to achieve environmental objectives, it is important to understand the extent to which private provision of public goods can be motivated. This paper suggests that elicitation of prosocial behavior may require more than appeals for cooperation and small pecuniary incentives.

Bibliography

Aaron, H. and WB Schwartz, "Rationing health care: the choice before us," *Science*, 1990, *247* (4941), 418.

Adelman, M.A., "The real oil problem," *Regulation-Washington*, 2004, *27* (1), 16–21.

Akerlof, G.A., "A theory of social custom, of which unemployment may be one consequence," *Quarterly Journal of Economics*, 1980, *94* (4), 749–775.

_ **and J.L. Yellen**, "Can small deviations from rationality make significant differences to economic equilibria?," *The American Economic Review*, 1985, *75* (4), 708–720.

_ **and R.E. Kranton**, "Economics and identity," *Quarterly Journal of Economics*, 2000, *115* (3), 715–753.

_ **and** _ , *Identity Economics: How Our Identities Shape Our Work, Wages, and Wellbeing*, Princeton Univ Pr, 2010.

_ **and W.T. Dickens**, "The economic consequences of cognitive dissonance," *The American Economic Review*, 1982, *72* (3), 307–319.

Allcott, H., "Social norms and energy conservation," *MIT Center for Energy and Environmental Policy Working Paper*, 2009, pp. 2009–014.

_ , "Social norms and energy conservation," *Journal of Public Economics*, 2011.

_ **and M. Greenstone**, "Is There an Energy Efficiency Gap?," Technical Report, National Bureau of Economic Research 2012.

_ **, N.N. Wozny, Massachusetts Institute of Technology. Center for Energy, and Environmental Policy Research**, *Gasoline prices, fuel economy, and the energy paradox*, MIT Center for Energy and Environmental Policy Research, 2010.

_ **, S. Mullainathan, and D. Taubinsky**, "Externalizing the Internality," *New York University Working Paper*, 2011.

Andreoni, J., "Giving with impure altruism: Applications to charity and Ricardian equivalence," *Journal of Political Economy*, 1989, *97* (6), 1447–1458.

_ , "Impure altruism and donations to public goods: a theory of warm-glow giving," *Economic Journal*, 1990, *100* (401), 464–477.

Ariely, D., A. Bracha, and S. Meier, "Doing good or doing well? Image motivation and monetary incentives in behaving prosocially," *American Economic Review*, 2009, *99* (1), 544–555.

Arimura, T.H., S. Li, R.G. Newell, and K. Palmer, "Cost-effectiveness of electricity energy efficiency programs," Technical Report, National Bureau of Economic Research 2011.

Ayres, I., S. Raseman, and A. Shih, "Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage," Working Paper 15386, National Bureau of Economic Research, Cambridge, Mass., USA 2009.

Barclay, P. and R. Willer, "Partner choice creates competitive altruism in humans," *Proceedings of the Royal Society B: Biological Sciences*, 2007, *274* (1610), 749.

Basu, K., "Monopoly, quality uncertainty and status goods," *International Journal of Industrial Organization*, 1987, *5* (4), 435–446.

Batson, C.D. and A.A. Powell, *Altruism and prosocial behavior*, Vol. 5, Hoboken, NJ: Wiley,

Beard, T.R., D.M Gropper, and J.E. Raymond, "Bill Averaging Programs and Consumer Behavior: Theory and Evidence," *Journal of Regulatory Economics*, 1998, *13* (1), 19–35.

Becker, G.S., "A Theory of the Allocation of Time," *The Economic Journal*, 1965, *75* (299), 493–517.

_ , "A theory of social interactions," *Journal of Political Economy*, 1974, *82* (6), 1063–1093.

Benabou, R. and J. Tirole, "Incentives and prosocial behavior," *American Economic Review*, 2006, *96* (5), 1652–1678.

_ **and** _ , "Incentives and Prosocial Behavior," *American Economic Review*, 2006, *96* (5), 1652–1678.

Bernheim, B.D., "A theory of conformity," *Journal of Political Economy*, 1994, *102* (5), 841–877.

Berry, S., J. Levinsohn, and A. Pakes, "Automobile prices in market equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Bertrand, M., E. Duflo, and S. Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Braun, O.L. and R.A. Wicklund, "Psychological antecedents of conspicuous consumption," *Journal of Economic Psychology*, 1989, *10* (2), 161–187.

Brown, J., T. Hossain, and J. Morgan, "Shrouded attributes and information suppression: Evidence from the field," *The Quarterly Journal of Economics*, 2010, *125* (2), 859–876.

Busse, M., J. Silva-Risso, and F. Zettelmeyer, "1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *The American Economic Review*, 2006, *96* (4), 1253–1270.

Cameron, A.C. and D.L. Miller, "Robust inference with clustered data," *Handbook of Empirical Economics and Finance*, 2010, pp. 1–28.

_ , J.B. Gelbach, and D.L. Miller, "Robust inference with multi-way clustering," *NBER working paper*, 2006.

Chen, C., "Freeway performance measurement system (PeMS)," PATH Research Report UCB-ITS-PRR-2003-22, University of California 2003.

Chetty, R., A. Looney, and K. Kroft, "Salience and Taxation: Theory and evidence," *The American Economic Review*, 2009, *99* (4), 1145–1177.

_ and E. Saez, "Dividend Taxes and Corporate Behavior: Evidence From the 2003 Dividend Tax Cut*," *The Quarterly Journal of Economics*, 2005, *120* (3), 791–833.

Cole, H.L., G.J. Mailath, and A. Postlewaite, "Social norms, savings behavior, and growth," *Journal of Political Economy*, 1992, *100* (6), 1092–1125.

Conlisk, J., "Why bounded rationality?," *Journal of economic literature*, 1996, *34* (2), 669–700.

Consumer Reports, "Hybrids and Diesel's: Do They Save Money?," May 2011.

Costa, D.L. and M.E. Kahn, "Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment," *NBER Working Paper*, 2010.

Cummings, R.G. and M.B. Walker, "Measuring the effectiveness of voluntary emission reduction programmes," *Applied Economics*, 2000, *32* (13), 1719–1726.

Cutter, W.B. and M. Neidell, "Voluntary information programs and environmental regulation: Evidence from 'Spare the Air'," *Journal of Environmental Economics and Management*, 2009, *58*, 253–265.

Dastrop, S., J.G. Zivin, D.L. Costa, and M.E. Kahn, "Understanding the solar home price premium: electricity generation & green social status," Working Paper 17200, National Bureau of Economic Research, Cambridge, Mass., USA 2010.

DellaVigna, S., "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 2009, *47* (2), 315–372.

_ **and J.M. Pollet**, "Investor inattention and Friday earnings announcements," *The Journal of Finance*, 2009, *64* (2), 709–749.

den Bulte, C. Van and G.L. Lilien, "Medical innovation revisited: Social contagion versus marketing effort," *ajs*, 2001, *106* (5), 1409–35.

Faruqui, A., S. Sergici, and A. Sharif, "The impact of informational feedback on energy consumption–A survey of the experimental evidence," *Energy*, 2010, *35* (4), 1598–1608.

Fershtman, C. and Y. Weiss, "Social status, culture and economic performance," *Economic Journal*, 1993, *103* (419), 946–959.

Fiserv, "2010 Billing Household Survey: Consumer Survey of Offline and Online Billing and Payment Practices," Research Paper, Fiserv, Inc., Brookfield, WI, USA 2010.

Folger, R., "On wanting what we do not have," *Basic and Applied Social Psychology*, 1992, *13* (1), 123–133.

Frank, R.H., "The demand for unobservable and other nonpositional goods," *American Economic Review*, 1985, *75* (1), 101–116.

Frey, B.S. and F. Oberholzer-Gee, "The cost of price incentives: An empirical analysis of motivation crowding-out," *American Economic Review*, 1997, *87* (4), 746–755.

_ **and I. Palacios-Huerta**, *Not just for the money: An economic theory of personal motivation*, Edward Elgar Cheltenham, 1997.

_ **and R. Jegen**, "Motivation crowding theory," *Journal of Economic Surveys*, 2001, *15* (5), 589–611.

Friedman, M., *Essays in positive economics*, University of Chicago Press, 1953.

Funk, P., "Social incentives and voter turnout: evidence from the Swiss mail ballot system," *Journal of the European Economic Association*, 2010, *8* (5), 1077–1103.

Gabaix, X. and D. Laibson, "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Gallagher, K. and E. Muehlegger, "Giving green to get green? The effect of incentives and ideology on hybrid vehicle adoption," *Journal of Environmental Economics and Management, forthcoming. doi*, 2007, *10.*

Gerber, A.S., D.P. Green, and C.W. Larimer, "Social pressure and voter turnout: Evidence from a large-scale field experiment," *American Political Science Review*, 2008, *102* (01), 33–48.

Gillingham, K., R. Newell, and K. Palmer, "Retrospective examination of demand-side energy efficiency policies," *Washington, DC: Resources for the Future. Original Edition, June*, 2004.

_ , R.G. Newell, and K. Palmer, "Energy efficiency economics and policy," Technical Report, National Bureau of Economic Research 2009.

Glazer, A. and K.A. Konrad, "A signaling explanation for charity," *American Economic Review*, 1996, *86* (4), 1019–1028.

Gneezy, U. and A. Rustichini, "A fine is a price," *Journal of Legal Studies*, 2000, *29* (1), 1–17.

_ and _ , "Pay enough or don't pay at all," *Quarterly Journal of Economics*, 2000, *115* (3), 791–810.

Goldstein, N.J., R.B. Cialdini, and V. Griskevicius, "A room with a viewpoint: Using social norms to motivate environmental conservation in hotels," *Journal of Consumer Research*, 2008, *35* (3), 472–482.

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

_ , _ , J.M. Sundie, R.B. Cialdini, G.F. Miller, and D.T. Kenrick, "Blatant benevolence and conspicuous consumption: When romantic motives elicit strategic costly signals," *Journal of Personality and Social Psychology*, 2007, *93* (1), 85.

Grubb, E.L. and G. Hupp, "Perception of self, generalized stereotypes and brand selection.," *Journal of Marketing Research*, 1968.

Ha, M., S.E. Williams, and M.J. Weber, "Utility Averaging Policies Impact on Consumer's Energy Usage," *Journal of Consumer Affairs*, 1993, *27* (2), 284–299.

Hardin, G., "The tragedy of the commons," *Science*, 1968, *162* (3859), 1243–1248.

Hardy, C.L. and M. Van Vugt, "Nice guys finish first: The competitive altruism hypothesis," *Personality and Social Psychology Bulletin*, 2006, *32* (10), 1402.

Hawkes, K., J. Altman, S. Beckerman, R.R. Grinker, H. Harpending, R.J. Jeske, N. Peterson, E.A. Smith, G.W. Wenzel, and J.E. Yellen, "Why hunter-gatherers work: an ancient version of the problem of public goods," *Current Anthropology*, 1993, pp. 341–361.

Heckman, J., H. Ichimura, and P. Todd, "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Program," *Review of Economic Studies*, 1997, *64*, 605–654.

Heckman, J.J. and V.J. Hotz, "Choosing Among Alternative Nonexperimental Methods for Estimating the Impact of Social Programs: The Case of Manpower Training," *Journal of the American Statistical Association*, 1989, *84* (408), 862–874.

Heffner, R.R., K.S. Kurani, and T.S. Turrentine, "Symbolism in California's early market for hybrid electric vehicles," *Transportation Research Part D: Transport and Environment*, 2007, *12* (6), 396–413.

_ **, T.S. Turrentine, and K.S. Kurani**, "A primer on automobile semiotics," Research Report UCD-ITS-RR-06-01, UC Davis Institute of Transportation Studies, Davis, Calif 2006.

Hopkins, E. and T. Kornienko, "Running to keep in the same place: consumer choice as a game of status," *American Economic Review*, 2004, pp. 1085–1107.

Hossain, T. and J. Morgan, "Plus Shipping and Handling," *Advances in Economic Analysis and Policy*, 2006, *6* (2).

Imbens, G.W., "Nonparametric estimation of average treatment effects under exogeneity: A review," *Review of Economics and Statistics*, 2004, *86* (1), 4–29.

_ **and J.M. Wooldridge**, "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 2009, *47* (1), 5–86.

_ **and T. Lemieux**, "Regression discontinuity designs: A guide to practice," *Journal of Econometrics*, 2008, *142* (2), 615–635.

Ireland, N.J., "Status-seeking, income taxation and efficiency," *Journal of Public Economics*, 1998, *70* (1), 99–113.

_ , "Optimal income tax in the presence of status effects," *Journal of Public Economics*, 2001, *81* (2), 193–212.

Janssen, M.C.W. and E. Mendys-Kamphorst, "The price of a price: on the crowding out and in of social norms," *Journal of Economic Behavior & Organization*, 2004, *55* (3), 377–395.

Janvry, A. De, E. Sadoulet, and S. Villas-Boas, "Can Appeals to Cooperation be Effective in Managing the Scarcity of a Vital Good? Responses to the 2004 Flu Vaccine Shortage," Working Paper Series 1013, Department of Agricultural and Resource Economics, UC Berkeley 2008.

Kahn, M.E., "Do greens drive Hummers or hybrids? Environmental ideology as a determinant of consumer choice," *Journal of Environmental Economics and Management*, 2007, *54* (2), 129–145.

_ **and R.K. Vaughn**, "Green market geography: the spatial clustering of hybrid vehicles and LEED registered buildings," *The BE Journal of Economic Analysis & Policy*, 2009, *9* (2), 2.

Kahneman, D., "Maps of Bounded Rationality: Psychology for Behavioral Economics," *American Economic Review*, 2003, *93* (5), 1449–1475.

Kellogg, R. and H. Wolff, "Daylight time and energy: Evidence from an Australian experiment," *Journal of Environmental Economics and Management*, 2008, *56* (3), 207–220.

Kezdi, G., "Robust standard error estimation in fixed-effects panel models," *Hungarian Statistical Review*, 2004, *9*, 95–116.

Kotchen, M.J., "Green markets and private provision of public goods," *Journal of Political Economy*, 2006, *114* (4), 816–834.

_ **and L.E. Grant**, "Does daylight saving time save energy? Evidence from a natural experiment in Indiana," *Review of Economics and Statistics*, 2011, *93* (4), 1172–1185.

Lancaster, K.J., "A new approach to consumer theory," *The Journal of Political Economy*, 1966, *74* (2).

Lee, D.S., "Randomized experiments from non-random selection in US House elections," *Journal of Econometrics*, 2008, *142* (2), 675–697.

_ **and T. Lemieux**, "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 2010, *48* (2), 281–355.

_ **, E. Moretti, and M.J. Butler**, "Do voters affect or elect policies? Evidence from the US House," *The Quarterly Journal of Economics*, 2004, *119* (3), 807.

Leibenstein, H., "Bandwagon, snob, and Veblen effects in the theory of consumers' demand," *Quarterly Journal of Economics*, 1950, *64* (2), 183–207.

Levitt, S.D. and J.A. List, "What do laboratory experiments measuring social preferences reveal about the real world?," *The Journal of Economic Perspectives*, 2007, *21* (2), 153–174.

List, J.A., "Does market experience eliminate market anomalies?," *The Quarterly Journal of Economics*, 2003, *118* (1), 41–71.

Loewenstein, G., *Using Decision Errors To Combat Climate Change*, Washington, DC: Behavior, Energy, and Climate Change Conference, 2009.

Lu, J., J. Weir, E.N. Schreffler, T. McLarney, R. Sarles, T. Higgins, and W. Johnson, "Development of Quantification Method for Measuring Travel and Emissions Impacts of Episodic Ozone Alert Programs," *Transportation Research Record: Journal of the Transportation Research Board*, 2004, *1864*, 153–159.

Lynn, M., "The psychology of unavailability: Explaining scarcity and cost effects on value," *Basic and Applied Social Psychology*, 1992, *13* (1), 3–7.

Mannering, F. and C. Winston, "A dynamic empirical analysis of household vehicle ownership and utilization," *RAND Journal of Economics*, 1985, *16* (2), 215–236.

_ **and M. Hani**, "Consumer valuation of foreign and domestic vehicle attributes: econometric analysis and implications for auto demand," *Transportation Research Part A: General*, 1985, *19* (3), 243–251.

Mastercard, "Revealing Attitudes on Recurring Payments," Technical Report, Mastercard Worldwide, Purchase, NY, USA 2006.

Mayer, L.S., "Estimating the effects of the onset of the energy crisis on residential energy demand," *Resources and Energy*, 1978, *1* (1), 57–92.

McDermott, K.A., J.M. Guldmann, K.A. Pfister, and K.V. Kumari, "Budget Billing Plans for Electric and Gas Utilities: An Analysis and Some Recommendations for Change," Technical Report, The National Regulatory Research Institute, Columbus, OH 1980.

Moulton, B.R., "Random group effects and the precision of regression estimates," *Journal of Econometrics*, 1986, *32* (3), 385–397.

Mullainathan, S. and R. Thaler, *Behavioral Economics* International Encyclopedia of the Social and Behavioral Sciences, Amsterdam: Elsevier, 2001.

_ **and R.H. Thaler**, "Behavioral economics," Technical Report, National Bureau of Economic Research 2000.

Olson, M., *The logic of collective action: Public goods and the theory of groups*, Cambridge, MA: Harvard University Press, 1971.

Payment Council, "The Way We Pay 2010," Technical Report, Payments Council, London, UK April 2010.

Peck, AE and OC Doering III, "Voluntarism and price response: consumer reaction to the energy shortage," *The Bell Journal of Economics*, 1976, *7* (1), 287–292.

Petersen, M.A., "Estimating standard errors in finance panel data sets: Comparing approaches," *Review of Financial Studies*, 2009, *22* (1), 435–480.

Reiss, P.C. and M.W. White, "What changes energy consumption? Prices and public pressures," *RAND Journal of Economics*, 2008, *39* (2), 636–663.

Roberts, G., "Competitive altruism: from reciprocity to the handicap principle," *Proceedings of the Royal Society B: Biological Sciences*, 1998, *265* (1394), 427.

Rosenbaum, P.R. and D.B. Rubin, "The central role of the propensity score in observational studies for causal effects," *Biometrika*, 1983, *70* (1), 41–55.

Rubin, D.B., "Formal models of statistical inference for causal effects," *Journal of Statistical Planning and Inference*, 1990, *25* (3), 279–292.

Sanstad, A.H. and R.B. Howarth, "Consumer rationality and energy efficiency," in "Proceedings of the ACEEE" 1994.

Schreffler, E., "Quantification methods for identifying emission reductions resulting from seasonal and episodic public education programs," *Air Resource Board Final Research Report, ARB contract*, 2003.

Sexton, Alison L., "Responses to Air Quality Alerts: Do Americans Spend Less Time Outdoors?" Ph.d. dissertation, University of Minnesota 2012.

Shogren, J.F. and L.O. Taylor, "On Behavioral-Environmental Economics," *Review of Environmental Economics and Policy*, 2008, *2* (1), 26.

Simon, H.A., "A Behavioral Model of Rational Choice," *The Quarterly Journal of Economics*, 1955, *69* (1), 99–118.

Sirgy, M.J., "Using self-congruity and ideal congruity to predict purchase motivation," *Journal of Business Research*, 1985, *13* (3), 195–206.

Spence, M., "Job market signaling," *Quarterly Journal of Economics*, 1973, *87* (3), 355–374.

Taylor, Lester D., Gail R. Blattenberger, and Robert K. Rennhack, *Econometric Models of the Demand for Energy*, Greenwich, Conn: JAI Press, Inc.,

Thaler, R.H., "The psychology and economics conference handbook: Comments on Simon, on Einhorn and Hogarth, and on Tversky and Kahneman," *The Journal of Business*, 1986, *59* (4), 279–284.

_ **and C.R. Sunstein**, *Nudge: Improving decisions about health, wealth, and happiness*, Yale Univ Pr, 2008.

The New York Times, "Say 'Hybrid' and Many People will Hear 'Prius'," *The New York Times*, 2007-07-04, p. via nytimes.com.

The Washington Post, "Toyota's Prius Proving to Be the Hotter Ride in Hybrids," 2004-08-23.

Time Magazine, "Heroes of the Environment: Toyota Prius Design Team," 2007.

Tversky, A. and D. Kahneman, "Judgment Under Uncertainty: Heuristics and Biases," *Science*, 1974, *185* (4157), 1121–1134.

Veblen, T., *The Theory of the Leisure Class* New York, McMillan, 1899.

Visa, "Visa Payment Panel Study: 2006 Payment Trends Summary," Technical Report, Visa USA Research Services, San Francisco, USA 2006.

Vugt, M. Van, G. Roberts, and C. Hardy, "Competitive altruism: Development of reputation-based cooperation in groups," *Handbook of Evolutionary Psychology*, 2007, pp. 531–540.

Welch, E., X. Gu, and L. Kramer, "The effects of ozone action day public advisories on train ridership in Chicago," *Transportation Research Part D: Transport and Environment*, 2005, *10* (6), 445–458.

Williams, S., M. Weber, M. Ha, S. Routh, J. Girard, and L. Worthington, "Utility Averaging Programs: A Contributing Factor to Household Energy Consumption," Technical Report, Oklahoma State University Center for Energy Research, Stillwater, OK 1990.

Wooldridge, J.M., *Econometric Analysis of Cross Section and Panel Data*, The MIT press, 2002.

_ **and G.M. Imbens**, "What's New in Econometrics," in "in" National Bureau of Economic Research Cambridge, MA July 2007.

WSJ, The Wall Street Journal, "In the New Dream Home, Majestic Boilers and Designer Pipes," 2010-06-15.

_ , "The Half There House," 2011-06-17.

Zivin, J. Graff and M. Neidell, "Days of haze: Environmental information disclosure and intertemporal avoidance behavior," *Journal of Environmental Economics and Management*, 2009, *58* (2), 119–128.