

Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments[†]

By HUNT ALLCOTT AND CHRISTOPHER KNITTEL*

It is often asserted that consumers are poorly informed about and inattentive to fuel economy, causing them to buy low-fuel economy vehicles despite their own best interest. This paper presents evidence on this assertion through two experiments providing fuel economy information to new vehicle shoppers. Results show zero statistical or economic effect on average fuel economy of vehicles purchased. In the context of a simple optimal policy model, the estimates suggest that current and proposed US fuel economy standards are significantly more stringent than needed to address the classes of imperfect information and inattention addressed by our interventions. (JEL C93, D12, D83, D91, L62, Q48)

Consumers constantly choose products under imperfect information. Most goods people buy have many attributes, and it is difficult to pay attention to and learn about all of them. This opens the door to the possibility that people might make mistakes: maybe they should have signed up for a better health insurance plan with a wider network and lower copays, and maybe they wouldn't have bought that coffee if they knew how many calories it has. Indeed, there is significant evidence that consumers can make systematic mistakes when evaluating products, either due to imperfect information about costs and benefits or by failing to pay attention to some attributes.¹

These issues are particularly important in the context of buying cars. Academics and policymakers have long argued that consumers are poorly informed and cognitively constrained when evaluating fuel economy. Turrentine and Kurani's (2007, 1213) structured interviews reveal that "when consumers buy a vehicle,

* Allcott: Department of Economics, New York University, 19 West 4th Street, 6th floor, New York, NY 10012, NBER, and E2e (email: hunt.allcott@nyu.edu); Knittel: Sloan School of Management and Center for Energy and Environmental Policy Research, MIT, 100 Main Street, Cambridge, MA 02142, NBER, and E2e (email: knittel@mit.edu). We are grateful to Will Tucker, Jamie Kimmel, and others at ideas42 for research management and to Skand Goel for research assistance. We thank Catherine Wolfram and seminar participants at the 2017 ASSA meetings, the MIT Center for Energy and Environmental Policy Research, and the University of California Energy Institute for comments. Funding was provided by the Ford-MIT Alliance, and we are grateful to Fords Emily Kolinski Morris for her collaboration and support of the experiments. Notwithstanding, Ford had no control over the data, analysis, interpretation, editorial content, or other aspects of this paper. This RCT was registered in the American Economic Association Registry for randomized control trials under trial number AEARCTR-0001421. Screen shots of the interventions and code to replicate the analysis are available from Hunt Allcott's website: <https://sites.google.com/site/allcott/research>.

[†] Go to <https://doi.org/10.1257/pol.20170019> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹ See, for example, Abaluck and Gruber (2011); Bollinger, Leslie, and Sorensen (2011); Barber, Odean, and Zheng (2005); Grubb (2009); Handel and Kolstad (2015); Hossain and Morgan (2006); Jensen (2010); Kling et al. (2012); and others.

they do not have the basic building blocks of knowledge assumed by the model of economically rational decision-making, and they make large errors estimating gasoline costs and savings over time.” Many have further argued that these errors *systematically* bias consumers against high fuel economy vehicles. For example, Kempton and Montgomery (1982, 826) describe “folk quantification of energy,” arguing that “[measurement inaccuracies] are systematically biased in ways that cause less energy conservation than would be expected by economically rational response to price.”² Such systematic consumer bias against energy conservation would exacerbate environmental externalities from energy use. As we discuss below, assertions of systematic bias have become one of the core motivations for Corporate Average Fuel Economy (CAFE) standards: the standards are justified largely on the grounds that inducing consumers to buy higher fuel economy vehicles will make them better off, independently of the additional externality reductions.

This important argument suggests a simple empirical test: does providing fuel economy information cause consumers to buy higher fuel economy vehicles? If consumers are indeed imperfectly informed about fuel costs or do not pay attention to fuel economy, then an informational intervention should cause people to buy higher fuel economy vehicles. If an informational intervention does not increase the average fuel economy of vehicles purchased, then the forms of imperfect information and inattention addressed by the intervention cannot be systematically relevant. Despite the importance of this debate and the CAFE regulation, such an experiment has not previously been carried out, perhaps because of the significant required scale and cost.

This paper presents the results of two experiments. The first provided fuel economy information to consumers via in-person intercepts at seven Ford dealerships nationwide. The second provided similar information to consumers in a nationwide online survey panel who reported that they were in the market to buy a new car. We later followed up with consumers to record what vehicles they bought. Our final samples for the dealership and online experiments comprise 375 and 1,489 vehicle buyers, respectively.

The core of the intervention was to provide individually tailored annual and lifetime fuel cost information for the several vehicles that the consumer was most closely considering, i.e., his or her “consideration set.” To make the cost information more salient, we also provided comparisons to common purchases: “that’s the same as it would cost for 182 gallons of milk” or for “8.7 tickets to Hawaii.” We designed the interventions to provide only hard information, minimizing demand effects and non-informational persuasion. We also took steps to ensure that the treatment group

² It is easy to find other examples of these arguments. For example, Greene et al. (2005, 758) write that “It could well be that the apparent undervaluing of fuel economy is a result of bounded rational behavior. Consumers may not find it worth the effort to fully investigate the costs and benefits of higher fuel economy.” Stern and Aronson (1984, 36) write that “The low economic cost and easy availability of energy made energy users relatively unaware of energy. As a result, energy was not a salient feature in family decisions about purchasing homes and automobiles... Energy has become invisible to consumers, so that even with some heightened awareness, they may be unable to take effective action.” Sanstad and Howarth (1994, 811) write that “problems of imperfect information and bounded rationality on the part of consumers, for example, may lead real world outcomes to deviate from the dictates of efficient resource allocation.”

understood and internalized the information provided, and recorded if they did not. In the dealership experiment, our field staff recorded that about 85 percent of the treatment group completed the intervention. In the online experiment, we ensured comprehension by requiring all respondents to correctly answer a quiz question about the information before advancing.

In the online experiment, we asked stated preference questions immediately after the intervention. Fuel cost information causes statistically significant but economically small shifts in stated preferences toward higher fuel economy vehicles in the consideration set, but interestingly, the information robustly causes consumers to *decrease* the general importance they report placing on fuel economy. In the follow-up surveys for both experiments, we find no statistically significant effect of information on average fuel economy of purchased vehicles. There are also no statistically significant fuel economy increases in subgroups that one might expect to be more influenced by information: consumers that were less certain about what vehicle they wanted, had spent less time researching, had more variation in fuel economy in their consideration set, or made their purchase sooner after receiving our intervention. The sample sizes deliver enough power to conclude that the treatment effects on fuel economy are also economically insignificant, in several senses. For example, we can reject with 90 percent confidence (in a two-sided test) that the interventions induced more than about 6 percent of consumers to change their purchases from the lower fuel economy vehicle to the higher fuel economy vehicle in their consideration sets.

Our results also help to evaluate part of the motivation for Corporate Average Fuel Economy standards, which are a cornerstone of energy and environmental regulation in the United States, Japan, Europe, China, and other countries. As we discuss in Section V, both regulators and academics have long argued that along with reducing carbon emissions and other externalities, an important possible motivation for CAFE standards is that they help to offset consumer mistakes such as imperfect information and inattention. In Section V, we formalize this argument in a simple optimal policy model. We then show formally that if an intervention that corrects misperceptions increases fuel economy by Q miles per gallon (MPG), but it's not practical to implement that intervention at scale, then the second-best optimal fuel economy standard to address misperceptions also increases fuel economy by Q MPG.

Our 90 percent confidence intervals rule out that the interventions increased fuel economy by more than 1.08 and 0.29 MPG, respectively, in the dealership and online experiments. Estimates are naturally less precise when reweighting the samples to match the nationwide population of new car buyers on observables, but the confidence intervals still rule out increases of more than 3.14 and 0.62 MPG. By contrast, CAFE standards are expected to require increases of 5.7 and 16.2 MPG by 2016 and 2025, respectively, relative to 2005 levels, after accounting for various alternative compliance strategies. Thus, in our samples, the CAFE regulation is significantly more stringent than can be justified by the classes of imperfect information and inattention addressed by our interventions.

The interpretation of the above empirical and theoretical results hinges on the following question: how broad are “the classes of imperfect information and

inattention addressed by our interventions?” On one extreme, one might argue that our interventions did provide the exact individually tailored fuel cost information that consumers would need, and the interventions did literally “draw attention” to fuel economy for at least a few minutes. On the other hand, there are many models of imperfect information and inattention, including models where cognitive costs prevent consumers from taking into account all information that they have been given; memory models in which consumers might forget information if it is not provided at the right time; and models where the presentation or trust of information matters, not just the fact that it was presented. Our interventions might not address the informational and attentional distortions in these models, so such distortions, if they exist, could still systematically affect fuel economy. This question is especially difficult to resolve if one believes that nuances of how the interventions were implemented could significantly impact the results. At a minimum, these results may move priors at least slightly toward the idea that imperfect information and inattention do not have large systematic effects on fuel economy, although it is crucial to acknowledge the possibility that the interventions could have been ineffective for various reasons.

The paper’s main contribution is to provide the first experimental evidence on the effects of fuel economy information on vehicle purchases, and to draw out the potential implications for optimal policy. Our work draws on several literatures. First, it is broadly related to randomized evaluations of information provision in a variety of contexts, including Choi, Laibson, and Madrian (2010) and Duflo and Saez (2003) on financial decisions; Bhargava and Manoli (2015) on takeup of social programs; Jin and Sorensen (2006), Kling et al. (2012), and Scanlon et al. (2002) on health insurance plans; Bollinger, Leslie, and Sorensen (2011) on calorie labels; Dupas (2011) on HIV risk; Hastings and Weinstein (2008) on school choice; Jensen (2010) on the returns to education; Ferraro and Price (2013) on water use; and many others; see Dranove and Jin (2010) for a review. There are several large-sample randomized experiments measuring the effects of energy cost information for durable goods other than cars, including Allcott and Sweeney (2017), Allcott and Taubinsky (2015), Davis and Metcalfe (2016), and Newell and Siikamäki (2014), as well as total household energy use, including Allcott (2011b), Dolan and Metcalfe (2013), and Jessoe and Rapson (2015).

Second, one might think of energy costs as a potentially “shrouded” product attribute in the sense of Gabaix and Laibson (2006), and information and inattention as one reason why “shrouding” arises. There is thus a connection to the empirical literatures on other types of potentially shrouded attributes, including out-of-pocket health costs (Abaluck and Gruber 2011), mutual fund fees (Barber, Odean, and Zheng 2005), sales taxes (Chetty, Looney, and Kroft 2009), and shipping and handling fees (Hossain and Morgan 2006). An earlier literature on energy efficiency, including Dubin and McFadden (1984) and Hausman (1979), studied similar issues using the framework of “implied discount rates.”

Third, our simple model of optimal taxation to address behavioral biases builds on work by Farhi and Gabaix (2015); Gruber and Köszegi (2004); Allcott, Lockwood, and Taubinsky (2018); Mullainathan, Schwartzstein, and Congdon (2012); and O’Donoghue and Rabin (2006). Energy efficiency

policy evaluation has been an active subfield of this literature, including work by Allcott, Mullainathan, and Taubinsky (2014); Allcott and Taubinsky (2015); Heutel (2015); and Tsvetanov and Segerson (2013).

Finally, we are closely connected to the papers estimating behavioral bias in automobile purchases. There is significant disagreement in this literature. A 2010 literature review found 25 studies, of which 12 found that consumers “undervalue” fuel economy, 5 found that consumers overvalue fuel economy, and 8 found no systematic bias (Oak Ridge National Laboratory 2010). The recent literature in economics journals includes Allcott (2013); Allcott and Wozny (2014); Busse, Knittel, and Zettelmeyer (2013); Goldberg (1998); Grigolon, Reynaert, and Verboven (2017); and Sallee, West, and Fan (2016). These recent papers use different identification strategies in different samples, and some conclude that there is no systematic consumer bias, while others find mild bias against higher fuel economy vehicles. Our work complements this literature by using experimental designs instead of observational data, by focusing primarily on new car sales instead of used car markets, and slightly strengthening the case that informational and behavioral distortions may not have large systematic effects on fuel economy.

Sections I–VI present the experimental design, data, baseline beliefs about fuel costs, treatment effects, theoretical model of optimal policy, and conclusion, respectively.

I. Experimental Design

Both the dealership and online experiments were managed by ideas42, a behavioral economics think tank and consultancy. While the two interventions differed slightly, they both had the same two key goals. The first was to deliver hard information about fuel costs to the treatment group, without attempting to persuade them in any particular direction, and also without affecting the control group. The second was to make sure that people understood the interventions, so that null effects could be interpreted as “information didn’t matter” instead of “people didn’t understand the information” or “the intervention was delivered poorly.”

The two experiments had the same structure. Each began with a baseline survey, then the treatment group received fuel economy information. Some months later, we delivered a follow-up survey asking what vehicle consumers had bought.

A. Dealership Experiment

We implemented the dealership experiment at seven Ford dealerships across the United States: in Baltimore, MD; Broomfield, CO; Chattanooga, TN; Naperville, IL (near Chicago); North Hills, California (near Los Angeles); Old Bridge Township, NJ (near New York City); and Pittsburgh, PA. In each case, Ford’s corporate office made initial introductions, then ideas42 met with dealership management and recruited them to participate. We approached nine dealerships in different areas of the country chosen for geographic and cultural diversity, and

these were the seven that agreed to participate.³ This high success rate reduces the likelihood of site selection bias (Allcott 2015). Online Appendix Figure A1 presents a map of the seven dealership locations.

In each dealership, ideas42 hired between one and three research assistants (RAs) to implement the intervention. Ideas42 recruited the RAs through Craigslist and university career services offices. Of the 14 RAs, ten were male and four were female. The median age was 25, with a range from 19 to 60. Nine of the 14 (64 percent) were White, and the remainder were Indian, Hispanic, and African American.

Ideas42 trained the RAs using standardized training materials, which included instructions on what to wear and how to engage with customers. Importantly, the RAs were told that their job was to provide information, not to persuade people to buy higher (or lower) fuel economy vehicles. For example, the RA training manual stated that “our explicit goal is not to influence consumers to pursue fuel-efficient vehicles. Rather, we are exploring the ways in which the presentation of information affects ultimate purchasing behavior.”

The RAs would approach customers in the dealerships and ask them if they were interested in a gift card in exchange for participating in a “survey.”⁴ If they refused, the RA would record the refusal. The RAs recorded visually observable demographic information (gender, approximate age, and race) for all people they approached.

For customers who agreed to participate, the RAs would engage them with a tablet computer app that asked baseline survey questions, randomized them into treatment and control, and delivered the intervention. The tablet app was designed by a private developer hired by ideas42. The baseline survey asked people the make, model, submodel, and model year of their current car and at least two vehicles they were considering purchasing; we refer to these vehicles individually as “first-choice” and “second-choice,” and collectively as the “consideration set.” The tablet also asked additional questions, including two questions measuring how far along they are in the purchase process (“how many hours would you say you’ve spent so far researching what car to buy?” and “how sure are you about what car you will purchase?”) and three questions allowing us to calculate annual and “lifetime” fuel costs (“if you purchase a car, how many years do you plan to own it?,” “how many miles do you expect that your vehicle will be driven each year?,” and “what percent of your miles are City versus Highway?”) The baseline survey concluded by asking for contact information.

The tablet computer randomly assigned half of participants to treatment versus control groups. For the control group, the intervention ended after the baseline survey. The treatment group first received several additional questions to cue them to start thinking about fuel economy, including asking what they thought the price of gas will be and how much money it will cost to buy gas for each vehicle in the consideration set. We use these fuel cost beliefs in Section IIIB, along with similar fuel cost belief questions from the follow-up survey.

³ We failed to engage one dealership in Massachusetts that was under construction, and our Colorado location was a replacement for another Colorado dealership that declined to participate.

⁴ For the first few weeks, we did not offer any incentive, and refusals were higher than we wanted. We then experimented with \$10 and \$25 Amazon or Target gift cards and found that both amounts reduced refusals by a similar amount, so we used \$10 gift cards for the rest of the experiment.

The treatment group then received three informational screens. The first was about MPG Illusion (Larrick and Soll 2008), describing how a two-MPG increase in fuel economy is more valuable when moving from 12 to 14 MPG than when moving from 22 to 24 MPG. The second provided individually tailored annual and lifetime fuel costs for the consumer's current vehicle and each vehicle in the consideration set, given the participant's self-reported years of ownership, driving patterns, expected gas price. To make these costs salient, the program compared them to other purchases. For example, "A Ford Fiesta will save you \$8,689 over its lifetime compared to a Ford Crown Victoria. That's the same as it would cost for 8.7 tickets to Hawaii." Figure 1 presents a picture of this screen. The third screen pointed out that "fuel costs can vary a lot within models," and presented individually tailored comparisons of annual and lifetime fuel costs for each submodel of each vehicle in the consideration set. After the intervention, we emailed a summary of the information to the participant's e-mail address.

Figure 2 presents a Consort diagram of the dealership experiment and sample sizes.⁵ The dealership intercepts happened from December 2012 to April 2014. The follow-up surveys were conducted via phone from August 2013 to September 2014. Of the 3,981 people who were initially approached, 1,740 refused, and 252 accepted but had already purchased a vehicle. Of the remaining 1,989 people, 958 were allocated to treatment and 1,031 to control. Of those allocated to treatment or control, 1,820 people (92 percent) completed the baseline survey.

A subcontractor called QCSS conducted the follow-up survey by phone in three batches: August 2013, January–April 2014, and August–September 2014. There was significant attrition between the baseline and follow-up surveys—some people gave incorrect phone numbers, and many others did not answer the phone. Of those who completed the baseline survey, 399 people (22 percent) completed the follow-up survey. While high, this attrition rate was not unexpected, and 22 percent is a relatively high completion rate for a phone survey. Twenty-four people had not purchased a new vehicle, leaving a final sample of 375 for our treatment effect estimates.

Especially given that we will find a null effect, it is crucial to establish the extent to which the treatment group engaged with and understood the informational intervention. We designed the tablet app to measure completion of the treatment in two ways. First, the participants had to click a "Completed" button at the bottom of the Fuel Economy Calculator screen (the top of which is pictured in Figure 1) in order to advance to the final informational screen. Second, after the intercept was over, the tablet app asked the RA, "Did they complete the information intervention?" Of the treatment group consumers who also completed the follow-up survey and thus enter our treatment effect estimates, 87 percent clicked "completed," and the RAs reported that 85 percent completed the information.

RA comments recorded in the tablet apps suggest that for the 13 to 15 percent of the treatment group that did not complete the intervention, there were two main

⁵The CONSolidated Standards Of Reporting Trials (Consort) diagram is a standardized way of displaying experimental designs and sample sizes. See <http://www.consort-statement.org/consort-statement/flow-diagram> for more information.

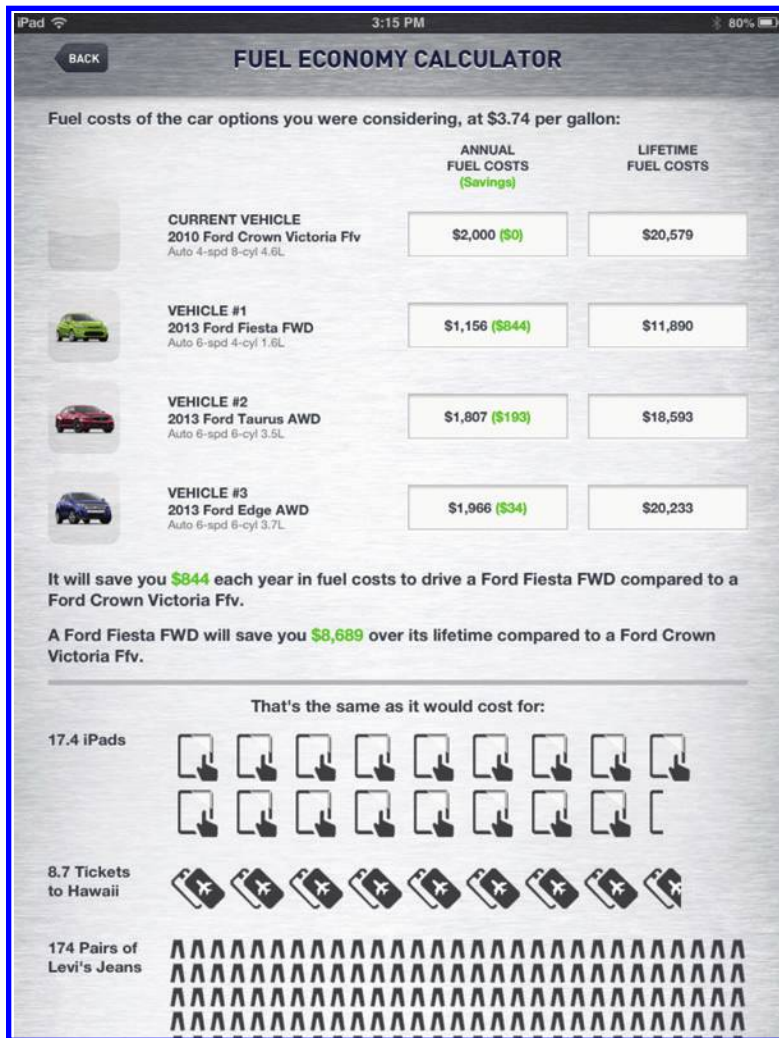


FIGURE 1. DEALERSHIP TREATMENT SCREEN

Notes: This is a screen capture from part of the dealership informational intervention, which was delivered via tablet computer. Vehicles #1, #2, and #3 were those that the participant had said he/she was considering purchasing, and fuel costs were based on self-reported driving patterns and expected gas prices.

reasons: distraction (example: “we’re in a hurry to leave the dealership”) and indifference (example: “was not very concerned with fuel efficiency, was looking to purchase a new Mustang for enjoyment”). If non-completion is driven by distraction, we should think of our treatment effects estimates as intent-to-treat, and the local average treatment effect would be 1/0.85 to 1/0.87 times larger. On the other hand, if non-completion is because people are already well-informed or know that their purchases will be unaffected by information, our estimates would reflect average treatment effects.

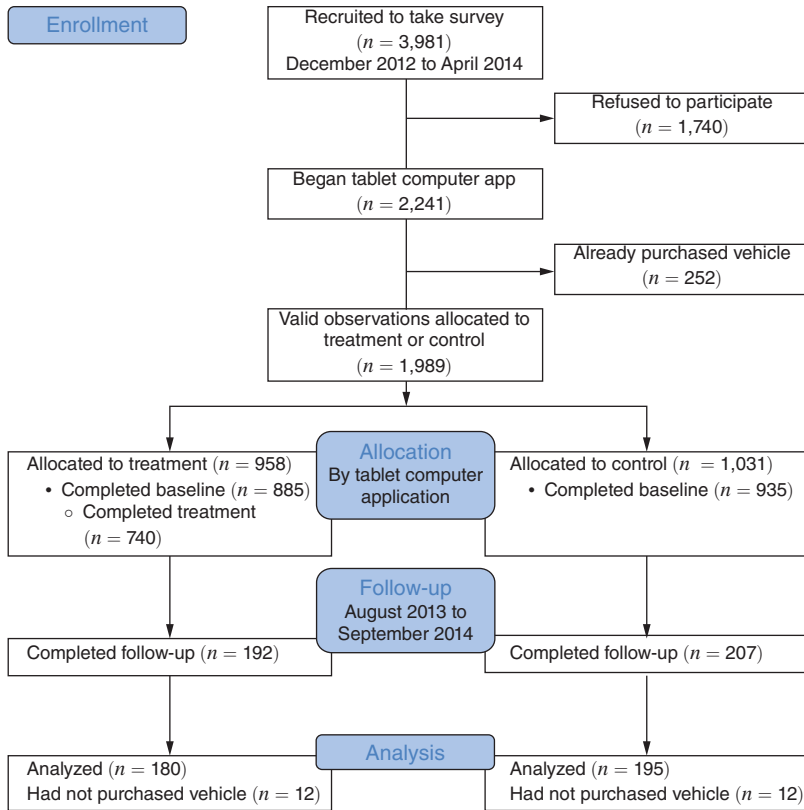


FIGURE 2. DEALERSHIP EXPERIMENT CONSORT DIAGRAM

In the follow-up survey, we also asked, “did you receive information from our researchers about the gasoline costs for different vehicles you were considering?” We would not expect the full treatment group to say “yes,” both because they might have forgotten in the months since the dealership interaction and because someone else in the household could have spoken with the RA. We also might expect some people in the control group to incorrectly recall the interaction. We find that 48 percent of the treatment group recalls receiving information many months later, against 16 percent of the control group.

B. Online Experiment

For the online experiment, we recruited subjects using the ResearchNow market research panel. The ResearchNow panel includes approximately 6 million members worldwide, who have been recruited by email, online marketing, and customer loyalty programs. Each panelist provides basic demographics upon enrollment, then takes up to six surveys per year. They receive incentives of approximately \$1 to \$5 per survey, plus prizes. We began with a subsample that were US residents at least 18 years old who reported that they are intending to purchase a car within the next six months.

The online experiment paralleled the dealership experiment, with similar baseline survey, informational interventions, and a later follow-up survey. As in the dealership experiment, we elicited beliefs about annual fuel costs for each vehicle in the consideration set, in both the baseline and follow-up surveys. However, the online experiment offered us the opportunity to ask additional questions that were not feasible in the more time-constrained dealership environment. In the initial survey, before and after the informational interventions, we asked participants the probability that they would buy their first- versus second-choice vehicles if they had to choose between only those two vehicles, using a slider from 0 to 100 percent. Also immediately after the informational interventions and on the follow-up survey, we asked participants to rate the importance of five attributes on a scale of one to ten, as well as how much participants would be willing to pay for four additional features. These questions allow us to construct stated preference measures of the intervention's immediate and long-term effects.

The ResearchNow computers assigned 60 percent of people to treatment and 40 percent to control using an algorithm that we discuss below. The base treatment was to provide information similar to the dealership experiment tablet app, including annual and "lifetime" (over the expected years of ownership) costs for the first-choice and second-choice vehicles, as well as for the highest-MPG vehicle in the same class as the first choice. Figure 3 presents a picture of the key information treatment screen. As in the dealership experiment, we compared these fuel costs to other tangible purchases: "that's the same as it would cost for 182 gallons of milk" or for "16 weeks of lunch."

Because we had fully computerized experimental control instead of delivering the treatment through RAs, we decided to implement four information treatment arms instead of just one. The "Base Only" treatment included only the above information, while the other three treatments included additional information. The "Base + Relative" treatment used the self-reported average weekly mileage to compare fuel savings to those that would be obtained at the national average mileage of about 12,000 miles per year. The "Base + Climate" treatment compared the social damages from carbon emissions (monetized at the social cost of carbon) for the same three vehicles as in the Base sub-treatment. The "Full" treatment included all of the Base, Relative, and Climate treatments. There were also four control groups, each of which paralleled one of the treatment arms in length, graphics, and text, but contained placebo information that was unrelated to fuel economy and would not plausibly affect purchases.⁶

To ensure that people engaged with and understood the information, participants were given a four-part multiple choice question after each of the treatment and control screens. For example, after the base treatment screen in Figure 3, participants were asked, "What is the difference in total fuel costs over [self-reported ownership period] years between the best-in-class MPG model and your first choice

⁶ The Base control group was informed about worldwide sales of cars and commercial vehicles in 2007, 2010, and 2013. The second control group received the Base information plus information on average vehicle-miles traveled in 2010 versus 1980. The third control group received the Base information plus data on the number of cars, trucks, and buses on the road in the United States in 1970, 1990, and 2010. The fourth control group received all control information.



FIGURE 3. ONLINE TREATMENT SCREEN

Notes: This is a screen capture from part of the online informational intervention. Choices #1 and #2 were the participant's first-choice and second-choice vehicles, and fuel costs were based on self-reported driving patterns and expected gas prices.

vehicle?" Four different answers were presented, only one of which matched the information on the previous screen. Sixty-nine, 79, and 79 percent of the treatment group answered the Base, Relative, and Environment quiz questions correctly on the first try. Seventy-seven, 66, and 84 percent of the control group answered the three control group quiz questions correctly on the first try. Every participant was required to answer the questions correctly before advancing.

Figure 4 presents a Consort diagram for the online experiment. The baseline survey and intervention were delivered in March 2015. We conducted the follow-up surveys in two rounds, the first from July to November 2015 and the second in August and September 2016. Here, 6,316 people planned to purchase vehicles and agreed to participate in the survey, of whom 5,014 finished the baseline survey and treatment or control intervention. There is natural attrition over time in the ResearchNow panel, and 3,867 people began the follow-up survey when it was fielded. Of those who began the follow-up survey, 2,378 had not bought a new vehicle or had incomplete data, leaving a final sample of 1,489 people for our treatment effect estimates.

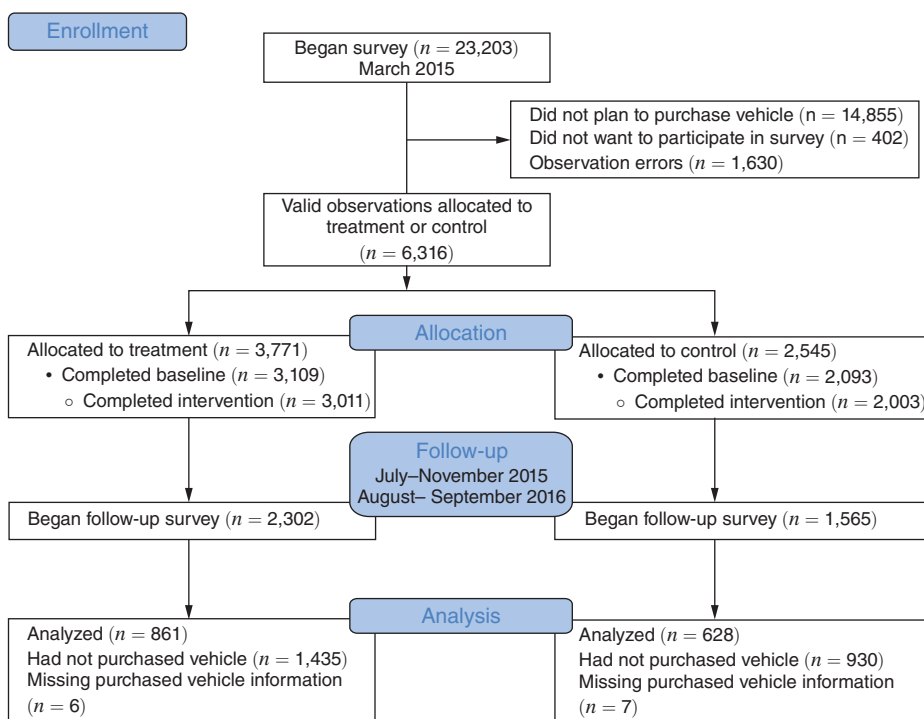


FIGURE 4. ONLINE EXPERIMENT CONSORT DIAGRAM

II. Sample Characteristics

A. Summary Statistics

Table 1 presents summary data for the samples that began the dealership and online experiments—specifically, the samples of valid observations that were randomized into treatment or control. For the dealership experiment, age, gender, and race were coded by the RAs at the end of the tablet survey, and income is the median income in the consumer’s zip code. For the online experiment, demographics are from basic demographics that the respondent provided to ResearchNow upon entering the panel. We impute missing covariates with sample means. See online Appendix A for additional details on data preparation.

Given that the dealership sample was recruited at Ford dealerships, it is not surprising that 40 percent of that sample currently drove a Ford, and 67 percent eventually purchased a Ford. By contrast, 12 percent of the online sample currently drove a Ford, and 11 percent purchased a Ford, closely consistent with the national average.

Fuel intensity (in gallons per mile (GPM)) is the inverse of fuel economy (in miles per gallon). For readability, we scale fuel intensity in gallons per 100 miles. The average vehicles use 4 to 5 gallons per 100 miles, meaning that they get 20 to 25 miles per gallon. We carry out our full analysis using fuel intensity instead of fuel economy because fuel costs are a key eventual outcome, and fuel costs scale

TABLE 1—COMPARISON OF SAMPLE DEMOGRAPHICS TO NATIONAL AVERAGES

	Dealership sample (1)	Online sample (2)	National (new car buyers) (3)
Male	0.64 (0.47)	0.60 (0.49)	0.48 (0.26)
Age	41.37 (12.87)	54.83 (13.64)	54.01 (13.14)
White	0.77 (0.41)	0.86 (0.35)	0.91 (0.29)
Income (\$000s)	73.51 (25.69)	121.93 (138.33)	82.08 (35.68)
Miles driven/year (000s)	15.29 (11.80)	11.68 (7.94)	13.38 (9.91)
Current vehicle is Ford	0.40 (0.48)	0.12 (0.32)	0.11 (0.31)
Current fuel intensity (gallons/100 miles)	4.70 (1.15)	4.57 (1.08)	4.58 (1.50)
Consideration set fuel intensity (gallons/100 miles)	4.35 (1.20)	4.15 (0.96)	—
Purchased fuel intensity (gallons/100 miles)	4.34 (1.26)	4.08 (1.00)	—
Purchased new/ late-model vehicle	0.67 (0.47)	0.68 (0.47)	—
Observations	1,989	6,316	18,053

Notes: This table shows sample means, with standard deviations in parentheses. The first two columns are the samples of valid observations that were randomized into treatment or control in the dealership and online experiments, respectively. “Purchased new/late-model vehicle” is an indicator for whether the purchased vehicle is model year 2013 (2015) or later in the dealership (online) sample. The national sample is the sample of households with model year-2008 or later vehicles in the 2009 National Household Travel Survey (NHTS), weighted by the NHTS sample weights.

linearly in GPM. “Consideration set fuel intensity” is the mean fuel intensity in the consumer’s consideration set.⁷

The final row reports that 67 to 68 percent of vehicle purchases in the two experiments were “new,” as defined by having a model year of 2013 or later (in the dealership experiment) or 2015 or later (in the online experiment). The third column in Table 1 presents the same covariates for the national sample of new car buyers from the 2009 National Household Travel Survey (NHTS), weighted by the NHTS sample weights. For the NHTS, we define “new car buyers” as people who own a model year 2008 or later vehicle in the 2009 survey. Unsurprisingly, neither of our samples is representative of the national population of new car buyers. Interestingly, however, they are selected in opposite ways for some covariates: the online sample is slightly older, significantly wealthier, and drives less than the national comparison

⁷ A small share of vehicles (0.2 to 0.3 percent of purchased and first choice vehicles) are electric. For electric vehicles, the EPA calculates MPG equivalents using the miles a vehicle can travel using the amount of electricity that has the same energy content as a gallon of gasoline. We omit electric vehicles from the descriptive analyses of gasoline cost beliefs, but we include electric vehicles in the treatment effect estimates.

group, while the dealership sample is younger, less wealthy, and drives more than the national population.

For some regressions, we re-weight the final samples to be nationally representative on observables using entropy balancing (Hainmueller 2012). We match sample and population means on the six variables in Table 1 that are available in the NHTS: gender, age, race (specifically, a White indicator variable), income, miles driven per year, whether the current vehicle is a Ford, and current vehicle fuel intensity. By construction, the mean weight is one. For the dealership and online samples, respectively, the standard deviations of weights across observations are 1.28 and 0.73, and the maximum observation weights are 12.0 and 9.2.

B. Balance and Attrition

ResearchNow allocated observations to the four treatment and four control groups using a modification of the least-fill algorithm.⁸ In the standard least-fill algorithm, a survey respondent is allocated to the group with the smallest number of completed surveys. A treatment or control group closes when it reaches the requested sample size, and the survey closes when the last group is full. In this algorithm, between the times when the groups close, group assignment is arbitrary and highly likely to be exogenous, as it depends only on an observation's exact arrival time. Over the full course of the survey, however, group assignment may be less likely to be exogenous, as some treatment or control groups close before others, and different types of people might take the survey earlier versus later. To address this possible concern, we condition regressions on a set of "treatment group closure time indicators," one for each period between each group closure time.⁹ While we include these indicators to ensure that it is most plausible to assume that treatment assignment is unconfounded, it turns out that their inclusion has very little impact on the results.

The first eight variables in Table 1 were determined before the information treatment was delivered. Online Appendix Table A2 shows that F -tests fail to reject that these eight observables are jointly uncorrelated with treatment status. In other words, treatment and control groups are statistically balanced on observables. By chance, however, several individual variables are unbalanced at conventional levels of statistical significance: current vehicle and consideration set fuel intensity in the dealership experiment, and income in the online experiment. We use the eight predetermined variables as controls to reduce residual variance and ensure conditional exogeneity in treatment effect estimates.

As we had expected, attrition rates are high. However, this does not appear to threaten internal validity. Online Appendix Table A3 shows that attrition rates are

⁸ We had instructed ResearchNow to use random assignment, but they did not do this, and we did not discover the discrepancy until we analyzed the data.

⁹ We say a "modification" of the least-fill algorithm because there were also some deviations from the above procedure. In particular, had the procedure been followed exactly, the last 20 percent of surveys would all be assigned to a treatment group, as 60 percent of observations were assigned to treatment versus 40 percent for control. However, ResearchNow modified the algorithm in several ways, and we thus have both treatment and control observations within each of the treatment group closure time indicators.

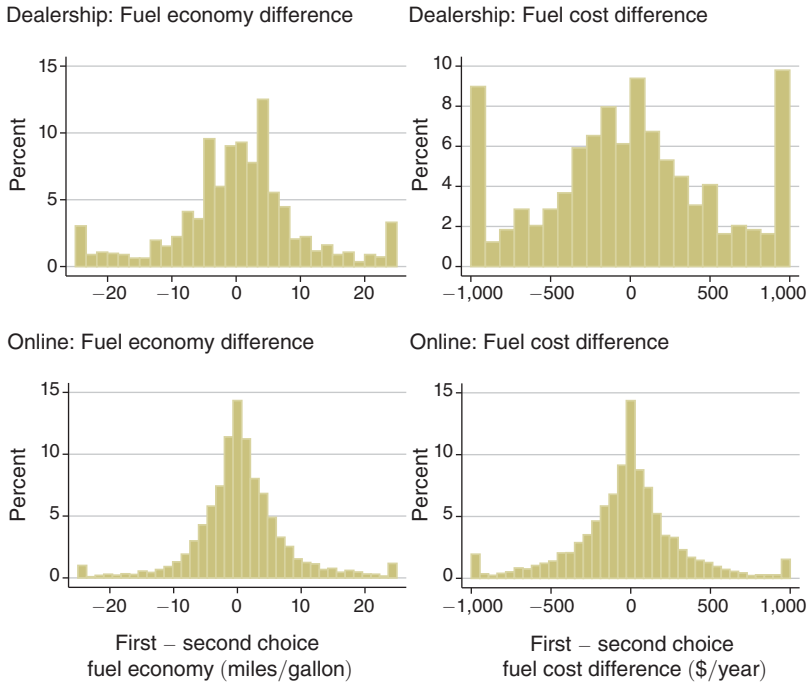


FIGURE 5. DISTRIBUTIONS OF ANNUAL FUEL COST DIFFERENCES BETWEEN FIRST- AND SECOND-CHOICE VEHICLES

Notes: The left two histograms present the distributions of fuel economy differences between consumers' first- and second-choice vehicles. The right two histograms present the distributions of fuel cost differences between consumers' first- and second-choice vehicles, given the vehicles' fuel economy ratings and consumers' self-reported miles driven, city versus highway share, and per gallon gasoline price. Outlying observations are collapsed into the outermost bars.

make as one of the two vehicles in their consideration set; this high proportion is unsurprising given that the participants were recruited from Ford dealerships. The final row of that part of the table shows a strong correlation between consideration set average fuel intensity and purchased vehicle fuel intensity.

The bottom part of Table 2 presents basic facts about the variation in fuel economy within consumers' consideration sets. The first row shows that the average consumers in the dealership and online experiments, respectively, were considering two vehicles that differed by 8.5 and 5.4 miles per gallon, or 1.1 and 0.7 gallons per 100 miles. The third row shows that the average consumers in the two experiments would have increased fuel economy by 3.9 and 2.3 MPG by switching from the first-choice vehicle to the vehicle with the highest MPG in the consideration set. This is about half of the previous number because for about half of consumers, the first-choice vehicle already is the highest-MPG vehicle in the consideration set. Finally, the average consumers in the two experiments were considering two vehicles with fuel costs that differed by \$523 and \$245 per year, at their self-reported miles driven, city versus highway share, and per gallon gasoline price.

TABLE 2—CONSIDERATION SETS

	Dealership experiment (1)	Online experiment (2)
<i>Panel A. Consideration sets versus final purchases</i>		
Share with . . .		
purchased model = first-choice model	0.42	0.30
purchased make = first-choice make	0.70	0.53
purchased model = second-choice model	0.12	0.06
purchased make = second-choice make	0.70	0.25
purchased model = first- or second-choice model	0.49	0.35
purchased make = first- or second-choice make	0.73	0.63
Correlation between consideration set average MPG and purchased MPG	0.52	0.44
<i>Panel B. Variation in consideration sets</i>		
Average of . . .		
first-choice – second-choice MPG	8.5	5.4
first-choice – second-choice gallons/100 miles	1.1	0.7
max{consideration set MPG} – First-choice MPG	3.9	2.3
max{consideration set gallons/100 miles} – First-choice gallons/100 miles	0.59	0.39
first-choice – second-choice fuel cost (\$/year)	523	245

Notes: Panel A of this table compares consideration sets (first- and second-choice vehicles) from the baseline surveys with the purchased vehicles reported in the follow-up surveys. Panel B presents variation in fuel economy and fuel costs within consumers' consideration sets.

While there is considerable variation within consideration sets, this is of course still smaller than the variation between consumers. In the dealership experiment consideration sets, the within- and between-consumer standard deviations in fuel economy are 6.5 and 9.7 MPG, respectively. For the online experiment consideration sets, the within- and between-consumer standard deviations are 5.0 and 8.7 MPG, respectively.

B. Beliefs about Consideration Set Fuel Costs

Above, we described the actual fuel costs for vehicles in consumers' consideration sets. We now examine a different question: what were consumers' beliefs about fuel costs? To do this, we follow Allcott (2013) in constructing "valuation ratios." We define \tilde{G}_{ij} as consumer i 's belief about annual gas costs of vehicle j , as elicited in the baseline survey. As above, G_{ij}^* is the "true" value given the vehicle's fuel economy rating and the consumer's self-reported miles driven, city versus highway share, and per-gallon gasoline price. For a given vehicle j , consumer i 's valuation ratio is the share of the true fuel cost that is reflected in beliefs:

$$(1) \quad \phi_{ij} = \frac{\tilde{G}_{ij}}{G_{ij}^*}.$$

For any pair of vehicles $j \in \{1, 2\}$, consumer i 's valuation ratio is the share of the true fuel cost difference that is reflected in beliefs:

$$(2) \quad \phi_i = \frac{\tilde{G}_{i1} - \tilde{G}_{i2}}{G_{i1}^* - G_{i2}^*}.$$

For both ϕ_{ij} and ϕ_i , the correct benchmark is $\phi = 1$. Note, $\phi > 1$ if the consumer perceives larger fuel costs, and $\phi < 1$ if the consumer perceives smaller fuel costs. Larger $|\phi - 1|$ reflects more “noise” in beliefs.

For example, consider two vehicles, one that gets 25 MPG (4 gallons per 100 miles) and another that gets 20 MPG (5 gallons per 100 miles). For a consumer who expects to drive 10,000 miles per year with a gas price of \$3 per gallon, the two cars would have “true” annual fuel costs $G_{i1}^* = \$1,200$ and $G_{i2}^* = \$1,500$. If on the survey, the consumer reports $\tilde{G}_{i1} = \$1,400$ and $\tilde{G}_{i2} = \$1,250$, we would calculate $\phi_i = \frac{1,400 - 1,250}{1,500 - 1,200} = 0.5$. In other words, the consumer responds as if she recognizes only half of the fuel cost differences between the two vehicles.

The fuel cost beliefs elicited in the surveys are a combination of consumers' actual beliefs plus some survey measurement error. Survey measurement error is especially important due to rounding (most responses are round numbers) and because we did not incentivize correct answers.¹⁰ Online Appendix Table A6, however, shows that elicited beliefs appear to be meaningful, i.e., not just survey measurement error: the results suggest both that ϕ_{ij} , ϕ_i , and $|\phi_i - 1|$ are correlated within individual between the baseline and follow-up surveys, and that people who perceive larger fuel cost differences (higher ϕ_i) also buy higher MPG vehicles, although the results from the dealership experiment are imprecise due to the smaller sample.

Figure 6 presents the distributions of valuation ratios in the baseline dealership and online surveys. The left panels show ϕ_{ij} from equation (1) for the first-choice vehicles, while the right panels show ϕ_i from equation (2) for the first- versus second-choice vehicles. Since there can be significant variation in ϕ_i , especially for two vehicles with similar fuel economy, we winsorize to the range $-1 \leq \phi \leq 4$.¹¹

The figure demonstrates three key results. First, people's reported beliefs are very noisy. Perfectly reported beliefs would have a point mass at $\phi = 1$. In the dealership and online experiments, respectively, 24 and 32 percent of ϕ_{ij} in the left panels are off by a factor of two or more, i.e., $\phi_{ij} \leq 0.5$ or $\phi_{ij} \geq 2$. This reflects some combination of truly noisy beliefs and survey reporting error.

Second, many people do not correctly report whether their first- or second-choice vehicle has higher fuel economy, let alone the dollar value of the difference in fuel costs. Forty-five and 59 percent of respondents in the dealership and online data, respectively, have $\phi_i = 0$, meaning that they reported the same expected fuel costs for vehicles with different fuel economy ratings. In both surveys, 8 percent have

¹⁰ Allcott (2013) shows that incentivizing correct answers does not affect estimates of belief errors in a related context.

¹¹ In the dealership experiment, this winsorization affects 5.2 and 13.2 percent of the observations of ϕ_{ij} and ϕ_i , respectively. In the online experiment, winsorization affects 5.1 and 10.2 percent of ϕ_{ij} and ϕ_i , respectively.

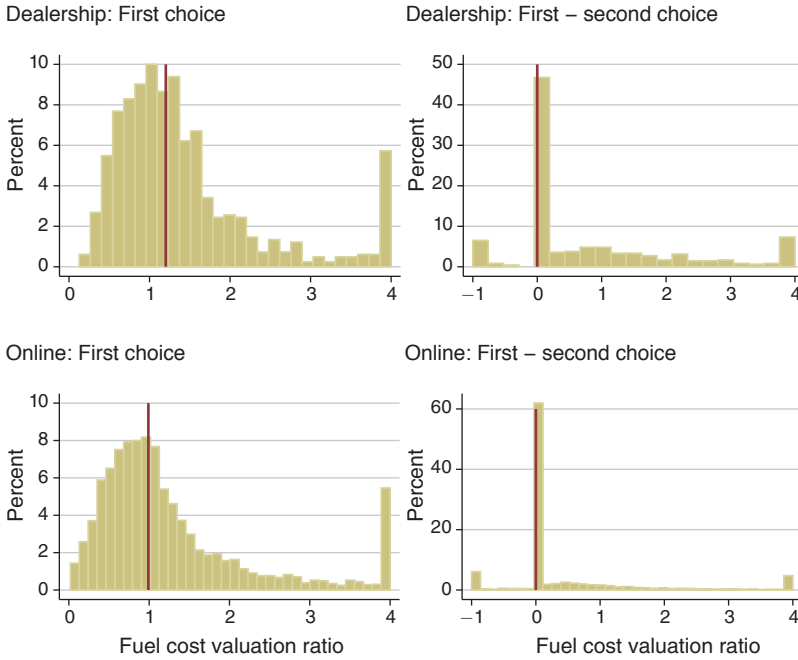


FIGURE 6. DISTRIBUTIONS OF FUEL COST BELIEFS: VALUATION RATIOS

Notes: These figures present the distribution of valuation ratios in the baseline surveys for the dealership and online experiments. The left panels present the valuation ratio from equation (1) for the first-choice vehicles. The right panels present the valuation ratios from equation (2) for the first- versus second-choice vehicles. In the right panels, a valuation ratio of zero means that the consumer reported the same expected fuel costs for both vehicles.

$\phi_i < 0$, meaning that they have the MPG rankings reversed. Thus, in the dealership and online surveys, respectively, only 47 and 33 percent of people correctly report which of their first- versus second-choice vehicle has higher fuel economy. This result also reflects some combination of incorrect beliefs and survey reporting error.

Third, it is difficult to argue conclusively whether people systematically overstate or understate fuel costs. The thin vertical lines in Figure 6 mark the median of each distribution. The top left figure shows that the median person in the dealership survey overestimated fuel costs by 20 percent ($\phi_{ij} = 1.2$), which amounts to approximately \$200 per year. The median person in the online survey, by contrast, has $\phi_{ij} = 0.99$. In the histograms on the right, the median ϕ_i is zero in both surveys, reflecting the results of the previous paragraph. All four histograms show significant dispersion, making the means harder to interpret.

IV. Empirical Results

We estimate the effects of information by regressing purchased vehicle fuel intensity on a treatment indicator, controlling for observables. Define Y_i as the fuel intensity of the vehicle purchased by consumer i , measured in gallons per 100 miles. Define T_i as a treatment indicator, and define \mathbf{X}_i as a vector of controls for the eight predetermined variables in Table 1: gender, age, race, natural log of income, miles

driven per year, an indicator for whether the current vehicle is a Ford, current vehicle fuel intensity, and consideration set average fuel intensity. The latter two variables soak up a considerable amount of residual variance in Y_i . For the online experiment, \mathbf{X}_i also includes the treatment group closure time indicators. The primary estimating equation is

$$(3) \quad Y_i = \tau T_i + \beta \mathbf{X}_i + \varepsilon_i.$$

We first study effects on stated preference questions in the online experiment, both immediately after the intervention and in the follow-up survey. The immediate stated preference questions are useful because they show whether the intervention had any initial impact. By comparing effects on the exact same questions asked months later during the follow-up, we can measure whether the intervention is forgotten. We then estimate effects on the fuel economy of purchased vehicles, for the full sample and then for subgroups that might be more heavily affected.

A. Effects on Stated Preference in the Online Experiment

We first show immediate effects on stated preference questions asked just after the online intervention. To increase power, we use the full sample available from the baseline survey, which includes many participants who do not appear in the follow-up survey. Table 3 reports results for three sets of questions. Panel A reports estimates of equation (3) where the dependent variable is the response to the question, “How important to you are each of the following features? (Please rate from 1–10, with 10 being “most important.)” Panel B reports estimates where the dependent variable is the answer to the question, “Imagine we could take your most likely choice, the [first-choice vehicle], and change it in particular ways, keeping everything else about the vehicle the same. How much additional money would you be willing to pay for the following?” In both panels, the feature is listed in the column header. Panel C presents the expected fuel intensity, i.e., weighted average of the first- and second-choice vehicles, weighted by the post-intervention reported purchase probability. In panel C, the R^2 is very high, and the estimates are very precise. This is because \mathbf{X}_i includes the consideration set average fuel intensity, which is the same as the dependent variable except that it is not weighted by post-intervention reported purchase probability.

Results in panels A and B show that the information treatment actually *reduced* the stated importance of fuel economy. The treatment group rated fuel economy 0.56 points less important on a scale of 1–10 and was willing to pay \$92.18 and \$237.96 less for five and 15 MPG fuel economy improvements, respectively. The treatment also reduced the stated importance of price, although the effect size is less than half of the effect on fuel economy. Preferences for power, leather interior, and sunroof are useful placebo tests, as the intervention did not discuss these issues. As expected, there are no effects on preferences for these attributes.¹²

¹²We thank a referee for pointing out that a WTP of only \$242 for a one-second improvement in 0–60 time would suggest that automakers’ large investments in engine power may be misguided. This could reinforce the

TABLE 3—IMMEDIATE EFFECT OF INFORMATION ON STATED PREFERENCE IN ONLINE EXPERIMENT

	Power (1)	Fuel economy (2)	Price (3)	Leather interior (4)	Sunroof (5)
<i>Panel A. Importance of features, from 1 (least important) to 10 (most important)</i>					
Treatment	−0.04 (0.06)	−0.56 (0.06)	−0.24 (0.05)	−0.06 (0.09)	0.10 (0.08)
Observations	5,036	5,036	5,036	5,036	5,036
R^2	0.04	0.13	0.06	0.07	0.04
Dependent variable mean	6.62	7.68	8.31	4.65	3.80
	Leather interior (1)	5 MPG improvement (2)	15 MPG improvement (3)	Power: 0–60 MPH 1 second faster (4)	
<i>Panel B. Willingness-to-pay for additional features</i>					
Treatment	4.49 (16.77)	−92.18 (15.81)	−237.96 (35.14)	16.89 (19.35)	
Observations	4,609	4,512	4,512	4,609	
R^2	0.06	0.06	0.07	0.05	
Dependent variable mean	380	409	1,043	242	
	Expected fuel intensity (gallons/100 miles) (1)				
<i>Panel C. Expected fuel intensity</i>					
Treatment	−0.032 (0.004)				
Observations	5,018				
R^2	0.97				
Dependent variable mean	4.12				

Notes: This table presents estimates of equation (3). The dependent variables in panel A are responses to the question, “How important to you are each of the following features? (Please rate from 1–10, with 10 being “most important).” Dependent variables in panel B are responses to the question, “Imagine we could take your most likely choice, the [first choice vehicle], and change it in particular ways, keeping everything else about the vehicle the same. How much additional money would you be willing to pay for the following?” In both panels, the feature is listed in the column header. In panel C, the dependent variable is the weighted average fuel intensity (in gallons per 100 miles) of the two vehicles in the consideration set, weighted by post-intervention stated purchase probability. Data are from the online experiment, immediately after the treatment and control interventions. All columns control for gender, age, race, natural log of income, miles driven per year, an indicator for whether the current vehicle is a Ford, current vehicle fuel intensity, consideration set average fuel intensity, and treatment group closure time indicators. Robust standard errors are in parentheses.

Why might the intervention have reduced the importance of fuel economy? One potential explanation is that people initially overestimated fuel costs and fuel cost differences, and the quantitative information in the treatment helps to correct these biased beliefs. As we saw in Figure 6, however, there is no clear evidence that this is the case for the online experiment sample. Furthermore, we can calculate the

usual concerns about taking unincentivized stated preference questions too seriously; our main focus is the effects on actual purchases in Tables 5 and 6.

actual annual savings from 5 and 15 MPG fuel economy improvements given each consumer's expected gasoline costs and driving patterns and the MPG rating of the first-choice vehicle. The control group has average willingness-to-pay of \$464 and \$1,186 for 5 and 15 MPG improvements, respectively. The actual *annual* savings are \$266 and \$583. This implies that the control group requires a remarkably fast pay-back period—approximately two years or less—for fuel economy improvements. It therefore seems unlikely that the control group overestimated the value of fuel economy improvements. Notwithstanding, the results in panels A and B are very robust: for example, they are not driven by outliers, and they don't depend on whether or not we include the control variables \mathbf{X}_i .

Panel C of Table 3 shows that the treatment shifted purchase probabilities toward the higher MPG vehicle in consumers' consideration set. This effect is small: a 25-MPG car has a fuel intensity of 4 gallons per 100 miles, so a decrease of 0.032 represents only a 0.8 percent decrease. In units of fuel economy, this implies moving from 25 to 25.2 miles per gallon.

It need not be surprising that the intervention shifted stated preference toward higher MPG vehicles in the consideration set while also reducing the stated general importance of fuel economy. As we saw in Figure 6, about two-thirds of online survey respondents do not correctly report which vehicle in their consideration set has higher MPG. Thus, even if the treatment makes fuel economy less important in general, it is still a positive attribute, and the treatment can shift preferences toward higher-MPG vehicles by clarifying which vehicles are in fact higher MPG. Furthermore, even consumers who do correctly report which vehicle in their consideration has lower fuel costs may be uncertain, and the treatment helps make them more certain.

We also asked the same stated preference questions from panels A and B on the follow-up survey, which respondents took 4 to 18 months later. Table 4 parallels panels A and B of Table 3, but using these follow-up responses. Only one of the nine variables (importance of price from 1–10) demonstrates an effect that is statistically significant with 90 percent confidence. For the fuel economy variables, there are zero remaining statistical effects, and we can reject effects of the sizes reported in Table 3. This suggests that the effects of information wear off over time, perhaps as people forget.

B. *Effects on Vehicle Purchases*

Did the interventions affect only stated preference, or did they also affect actual purchases? Table 5 presents treatment effects on the fuel intensity of purchased vehicles. Columns 1–3 present dealership experiment results, while columns 4–6 present online experiment results. Columns 1 and 4 omit the \mathbf{X}_i variables, while columns 2 and 5 add \mathbf{X}_i ; the point estimates change little. Columns 3 and 6 are weighted to match US population means, as described in Section II. In all cases, information provision does not statistically significantly affect the average fuel intensity of the vehicles consumers buy.

The bottom row of Table 5 presents the lower bounds of the 90 percent confidence intervals of the treatment effects. Put simply, these are the largest

TABLE 4—EFFECT OF INFORMATION ON STATED PREFERENCE IN ONLINE EXPERIMENT FOLLOW-UP SURVEY

	Power (1)	Fuel economy (2)	Price (3)	Leather interior (4)	Sunroof (5)
<i>Panel A. Importance of features, from 1 (least important) to 10 (most important)</i>					
Treatment	0.12 (0.12)	-0.10 (0.11)	-0.17 (0.10)	0.15 (0.17)	0.07 (0.16)
Observations	1,542	1,544	1,543	1,542	1,541
R^2	0.03	0.07	0.03	0.05	0.03
Dependent variable mean	6.90	7.76	8.49	4.95	4.02
	Leather interior (1)	5 MPG improvement (2)	15 MPG improvement (3)	Power: 0–60 MPH 1 second faster (4)	
<i>Panel B. Willingness-to-pay for additional features</i>					
Treatment	-37.41 (29.38)	2.66 (23.97)	20.31 (56.25)	13.48 (27.76)	
Observations	1,359	1,329	1,329	1,359	
R^2	0.06	0.04	0.04	0.03	
Dependent variable mean	316	346	940	168	

Notes: This table presents estimates of equation (3). The dependent variables in panel A are responses to the question, “How important to you are each of the following features? (Please rate from 1–10, with 10 being “most important).” Dependent variables in panel B are responses to the question, “Imagine we could take your most likely choice, the [first choice vehicle], and change it in particular ways, keeping everything else about the vehicle the same. How much additional money would you be willing to pay for the following?” In both panels, the feature is listed in the column header. Data are from the follow-up survey for the online experiment. All columns control for gender, age, race, natural log of income, miles driven per year, an indicator for whether the current vehicle is a Ford, current vehicle fuel intensity, consideration set average fuel intensity, and treatment group closure time indicators. Robust standard errors are in parentheses.

TABLE 5—EFFECTS OF INFORMATION ON FUEL INTENSITY OF PURCHASED VEHICLES

	Dealership			Online		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.07 (0.13)	0.11 (0.11)	-0.21 (0.17)	0.05 (0.05)	0.03 (0.04)	0.01 (0.06)
90% confidence interval lower bound	-0.15	-0.06	-0.49	-0.03	-0.04	-0.08
Observations	375	375	375	1,489	1,489	1,489
R^2	0.00	0.39	0.29	0.00	0.39	0.38
Dependent variable mean	4.33	4.33	4.33	4.09	4.09	4.09
Controls	No	Yes	Yes	No	Yes	Yes
Weighted	No	No	Yes	No	No	Yes

Notes: This table presents estimates of equation (3). The dependent variable is the fuel intensity (in gallons per 100 miles) of the vehicle purchased. All columns control for gender, age, race, natural log of income, miles driven per year, an indicator for whether the current vehicle is a Ford, current vehicle fuel intensity, and consideration set average fuel intensity. Columns 4–6 also control for treatment group closure time indicators. Samples in columns 3 and 6 are weighted to match the national population of new car buyers.

statistically plausible effects of information on fuel economy. With equally weighted observations in columns 2 and 5, the confidence intervals rule out fuel intensity decreases of 0.06 and 0.04 gallons per 100 miles in the dealership and

online experiments, respectively. When re-weighted to match the national population, the confidence intervals rule out decreases of 0.49 and 0.08 gallons per hundred miles, respectively. For comparison, for a 25-MPG car, a decrease of 0.1 gallons per 100 miles represents a decrease from 4 to 3.9 gallons per 100 miles, i.e., an increase from 25 to 25.64 miles per gallon.

Power and Economic Significance.—Should we think of these estimates as precise zeros, with enough statistical power to rule out any economically significant effects? Or are these imprecise zeros, meaning that there could be economically significant effects that we cannot statistically distinguish from zero? We consider five benchmarks of economic significance, focusing on the primary unweighted estimates in columns 2 and 5 of Table 5.

First, we can compare our effect sizes to the variation in the dependent variable, purchased vehicle fuel intensity. This variation reflects the variation in consumers' full choice sets. For the dealership and online samples, respectively, Table 1 reported that 1 standard deviation in purchased vehicle fuel intensity is 1.26 and 1.00 gallons per 100 miles. Thus, using the lower bounds of the 90 percent confidence intervals in columns 2 and 5 of Table 5 for the dealership and online experiments, respectively, we can rule out that the treatment decreased fuel intensity by more than $0.06/1.26 \approx 0.05$ and $0.04/1.00 \approx 0.04$ standard deviations.

Second, we can compare our effect sizes to the variation in consumers' consideration sets that was documented in Section III. As reported in Table 2, the average absolute difference in fuel intensity between consumers' first- and second-choice vehicles is 1.1 and 0.7 gallons per 100 miles in the dealership and online experiments, respectively. Again comparing these to the 90 percent confidence bounds, we can rule out that the intervention decreased fuel intensity by more than $0.06/1.1 \approx 0.05$ and $0.04/0.7 \approx 0.06$, i.e., about 6 percent, of the average difference between consumers' two most preferred vehicles.

Third, we can benchmark against the effect sizes that would be expected if the intervention moved all consumers from their initially preferred vehicle (i.e., their first-choice vehicle) to the highest fuel economy vehicle in the consideration set. This benchmark is naturally smaller than the average absolute differences discussed above, because for about half of consumers, their first-choice vehicle is already the highest-MPG vehicle between their first- and second-choice vehicles. As reported in Table 2, the average differences between the highest-MPG vehicle in the consideration set and the first-choice vehicle are 3.9 and 2.3 MPG in the two experiments, or 0.59 and 0.39 gallons per 100 miles. With 90 percent confidence, we thus can rule out effects larger than $0.06/0.59 \approx 0.10$ and $0.04/0.39 \approx 0.10$, i.e., 10 percent, of that benchmark.

A fourth way to benchmark the effect sizes is to compare them to how consumers respond to changes in gasoline prices. This benchmark could make any MPG effect seem small, as fleet fuel economy is relatively inelastic to gas price changes: Klier and Linn (2010) find that a \$1 gasoline price increase would increase the average fuel economy of new vehicles sold by only 0.8 to 1 MPG. Using this result, we can reject effect sizes equivalent to more than a gas price increase of about \$1.08–\$1.35 (in the dealership experiment) and \$0.29–\$0.37 (in the online experiment).

Of course, the fact that we have two experiments instead of just one adds confidence to these results—both because this provides additional evidence of generalizability, and because statistically combining the results would make the estimates even more precise. By these first four benchmarks, we have enough power to conclude that the information treatments did not have *economically* significant effects on average fuel economy.

Our fifth benchmark is whether our estimates are precise enough to be policy-relevant: can we reject the effect sizes that would be needed to justify the Corporate Average Fuel Economy standards currently in place in the United States? Section V considers that question in more depth.

Alternative Estimates.—As discussed in Section I, the online intervention actually had four separate sub-treatments. Online Appendix Table A10 presents estimates of equation (3) for stated preference fuel intensity immediately after the intervention, paralleling panel C of Table 3, and for fuel intensity of purchased vehicles, paralleling column 5 of Table 5. For both outcomes, Wald tests fail to reject that the coefficients on the four sub-treatments are jointly equal. Interestingly, the “Base + Climate” treatment, which included information about both fuel costs and climate change damages, has a statistically positive treatment effect on purchased vehicle fuel intensity, meaning that it caused people to buy statistically lower MPG vehicles. It would be useful to test whether this replicates in other samples.

Online Appendix Table A11 presents alternative estimates of equation (3), except with G_{ij}^* , the purchased vehicle annual fuel costs (using consumers’ self-reported miles driven, city versus highway share, and per gallon gasoline price) as the dependent variable. The treatment effect is now in units of annual fuel costs saved, which is in some senses more directly relevant than fuel intensity. Furthermore, consumers who expect to drive more receive more weight in the estimation, which is useful in that these consumers should theoretically be more affected by information. As in Table 5, the estimated effects are still statistically zero, and we can reject that information caused consumers to save more than \$28 and \$18 per year in fuel costs in the dealership and online experiments, respectively.

Online Appendix B explores whether the treatment makes fuel cost beliefs meaningfully more precise, or whether baseline beliefs meaningfully moderate the treatment effect. Because people’s fuel cost beliefs are so dispersed, the estimates deliver imprecise zeros.

C. Effects in Subgroups

Several hypotheses predict specific subgroups where the treatment effects might be larger or smaller. First, information might have smaller effects on people who are considering vehicles only in a narrow fuel economy range: fuel economy information will likely have smaller effects for a consumer deciding between 22- and 23-MPG vehicles compared to a consumer deciding between a Hummer and a Prius. Second, as suggested by comparing the stated preference results between baseline and follow-up in Tables 3 and 4, the treatment’s possible impact may have worn off as people forgot the information. Consumers who bought their new cars

TABLE 6—TREATMENT EFFECTS FOR SUBGROUPS HYPOTHESIZED TO BE MORE RESPONSIVE

	Full sample (1)	\geq Median consideration set MPG variance (2)	\leq Median time until purchase (3)	Less sure (4)	\leq Median research time (5)
<i>Panel A. Dealership experiment</i>					
Treatment	0.11 (0.11)	0.03 (0.17)	0.09 (0.14)	0.24 (0.15)	0.24 (0.16)
Observations	375	188	168	185	168
R^2	0.39	0.29	0.47	0.39	0.41
Dependent variable mean	4.33	4.15	4.27	4.28	4.22
<i>Panel B. Online experiment</i>					
Treatment	0.03 (0.04)	-0.03 (0.06)	0.01 (0.06)	0.03 (0.05)	0.00 (0.06)
Observations	1,489	745	745	1,095	743
R^2	0.39	0.36	0.44	0.35	0.42
Dependent variable mean	4.09	3.93	4.06	4.10	4.07

Notes: This table presents estimates of equation (3), with samples limited to the subgroups indicated in the column headers. The dependent variable is the fuel intensity (in gallons per 100 miles) of the vehicle purchased. All columns control for gender, age, race, natural log of income, miles driven per year, an indicator for whether the current vehicle is a Ford, current vehicle fuel intensity, and consideration set average fuel intensity. Panel B also includes treatment group closure time indicators. Robust standard errors are in parentheses.

sooner after the intervention are less likely to have forgotten. Third, information might be more powerful for people who have done less research and are less sure about what car they want to buy.

Table 6 presents estimates in specific subgroups that, per these hypotheses, might be more responsive. Column 1 reproduces the treatment effect estimate for the full sample. Column 2 considers only consumers with above-median variance of fuel intensity in their consideration set. Column 3 considers only the consumers with below-median time between the intervention and the date of vehicle purchase reported in the follow-up survey. Column 4 drops the approximately half of consumers who report being “almost certain” what vehicle they will purchase, using only consumers who are “fairly sure,” “not so sure,” or “not at all sure.” Column 5 considers only consumers who report having spent less than median time researching what vehicle to buy. In all of these eight subgroup analyses, the effects are statistically zero at conventional levels.

V. Theoretical Model: Implications for Optimal Policy

Corporate Average Fuel Economy standards are a cornerstone of energy and environmental regulation in the United States, Japan, Europe, China, and other countries.¹³ The US government’s Regulatory Impact Analysis (RIA) for CAFE

¹³There is a large literature on various aspects of fuel economy standards in the United States—see Austin and Dinan (2005), Goldberg (1998), Jacobsen (2013), and Jacobsen and van Benthem (2015)—and other countries, including Japan (Ito and Sallee 2014), Europe (Reynaert and Sallee 2016), and China (Howell 2018).

standards finds that they generate a massive win-win: not only do they reduce externalities, but they also save consumers money. Over 2011–2025, the standards are projected to cost \$125 billion, reduce externalities (mostly from climate change, local air pollution, and national energy security) by \$61 billion, and reduce private costs (mostly from buying gasoline) by \$540 billion (NHTSA 2012). Thus, even ignoring externalities, the regulation generates \$415 billion in net private benefits, with a private benefit/cost ratio of better than three-to-one. Net private benefits are almost seven times more important than externalities in justifying the regulation. The large net private benefit implies that there must be some large non-externality market failure that is keeping the private market from generating these results in the absence of CAFE.¹⁴

While some possible market failures are on the supply side—for example, cross-firm spillovers from research and development of fuel economy-improving technologies—significant attention has been focused on demand-side market failures. The US government’s RIA argues that information, inattention, “myopia,” and other behavioral biases might keep consumers from buying higher fuel economy vehicles that would save them money in the long run at reasonable discount rates. For example, the RIA argues that

*“consumers might lack the information necessary to estimate the value of future fuel savings,” and “when buying vehicles, consumers may focus on visible attributes that convey status, such as size, and pay less attention to attributes such as fuel economy” (EPA 2012, 8–7).*¹⁵

Thus, our information treatments could conceivably address some—although not all—of the alleged classes of imperfect information and inattention that are used to justify this regulation.

In this section, we use a theoretical model to formalize the argument that imperfect information and inattention cause systematic misoptimization, and that CAFE standards can help address these distortions. We then show how our empirical estimates can be relevant for evaluating this argument.

¹⁴“By non-externality,” we more precisely mean market failures other than the specific environmental and energy security externalities comprising the \$61 billion.

¹⁵There are many other examples of this argument. For example, the CAFE standard final rule (EPA 2010, 25510) argues, “In short, the problem is that consumers appear not to purchase products that are in their economic self-interest. There are strong theoretical reasons why this might be so,” including that “consumers might lack information” and “the benefits of energy-efficient vehicles may not be sufficiently salient to them at the time of purchase, and the lack of salience might lead consumers to neglect an attribute that it would be in their economic interest to consider.”

As another example, Fischer, Harrington, and Parry (2007, 3) concludes, “The bottom line is that the efficiency rationale for raising fuel economy standards appears to be weak unless carbon and oil dependency externalities are far greater than mainstream economic estimates, or consumers perceive only about a third of the fuel saving benefits from improved fuel economy.” Gayer (2011) summarizes the arguments, “Energy-efficiency regulations and fuel economy regulations are therefore justified by [cost-benefit analyses] only by presuming that consumers are unable to make market decisions that yield personal savings, that the regulator is able to identify these consumer mistakes, and that the regulator should correct economic harm that people do to themselves.”

A. Model Setup

In our theoretical model, a social planner wants to set the socially optimal fuel economy standard. Consistent with the current policy of tradable CAFE credits, we model the standard as creating a tradable credit market with credit price t dollars per vehicle-GPM. This means that when an auto manufacturer sells a vehicle with fuel intensity e_j gallons per mile, it must also submit credits valued at te_j for each unit sold.

Auto manufacturing firms produce a choice set of J vehicles, indexed $j \in \{1, \dots, J\}$. Marginal production cost is c_j , price is p_j , and fuel intensity in GPM is e_j . In the model, supply is perfectly competitive, so price equals total marginal cost: $p_j = p_j(t) = c_j + te_j$.¹⁶ Like some prior literature, we assume that the choice set is fixed, so automakers comply with fuel economy standards by increasing the relative price of low MPG vehicles, instead of by introducing more hybrid vehicles or MPG-improving technologies.

Consumers choose exactly one option from the J vehicles or an outside option indexed $j = 0$. There are L consumer types, each with different preferences; l indexes types and i indexes consumers within a type. We normalize each consumer type to have measure one consumer. Here, G_{lj} is the present discounted value of fuel cost for vehicle j given fuel intensity e_j and consumer type l 's utilization patterns. Consumer i of type l who buys vehicle j enjoys true utility $U_{ij} = \eta_l(Z_l - p_j - G_{lj}) + \xi_{lj} + \epsilon_{ij}$, where Z_l is income, ξ_{lj} is utility from vehicle use (i.e., utility from vehicle attributes other than price and fuel cost), and ϵ_{ij} is a logit taste shock. Notice that although we assume this particular distribution of ϵ_{ij} to simplify the derivations, preferences are very general because η_l , G_{lj} , and ξ_{lj} can vary arbitrarily across types.

Consumers are potentially biased: when choosing a vehicle, imperfect information or inattention cause them to perceive fuel costs $(1 + b_{lj})G_{lj}$ instead of G_{lj} . Their vehicle choices thus maximize decision utility $\tilde{U}_{ij} = \eta_l(Z_l - p_j - (1 + b_{lj})G_{lj}) + \xi_{lj} + \epsilon_{ij}$. Note, $b_{lj} = 0$ implies no bias. Positive b_{lj} means that the consumer overestimates fuel costs and thus would get more utility than expected because there is additional money left to buy more units of the numeraire good. Conversely, negative b_{lj} means that the consumer underestimates fuel costs and thus would get less utility than expected. Define \mathbf{b}_l as type l 's vector of biases for each of the J vehicles.

Given decision utility \tilde{U}_{ij} , the representative decision utility and choice probabilities are standard for the logit model. For any credit price t and any bias b_j , representative decision utility is $V_{lj}(t, b_j) = \eta_l(Z_l - p_j(t) - (1 + b_{lj})G_{lj}) + \xi_{lj}$, and the logit choice probability for any vector of biases \mathbf{b} is $P_{lj}(t, \mathbf{b}) = \frac{\exp(V_{lj}(t, b_j))}{\sum_k \exp(V_{lk}(t, b_k))}$, where j and k both index vehicles.

The aggregate value of fuel economy credit revenues is $T(t) = \sum_l \sum_j te_j P_{lj}(t, \mathbf{b}_l)$. If credits must be bought from the government, we assume that these revenues are recycled to consumers in lump-sum payments, which would enter utility in the same

¹⁶ In reality, the vehicle market is of course not perfectly competitive. The propositions below also hold with markups that are nonzero but identical across vehicles. When markups vary across vehicles, the optimal fuel

way as income Z_l . If credits are grandfathered to auto manufacturers, as is essentially the case under the current policy, then these revenues enter as producer surplus.

We define the “stringency” of the fuel economy standard as $S(t) \equiv \sum_l \sum_j e_j [P_{lj}(t, \mathbf{b}_l) - P_{lj}(0, \mathbf{b}_l)]$. In words, S is the required change in sales-weighted average fuel intensity relative to the baseline with no standard. A value of $S < 0$ reflects a decrease in fuel intensity, i.e., an increase in fuel economy. Because higher t increases the relative price of higher fuel intensity vehicles, there is a unique and monotonically decreasing relationship between S and t : the more stringent the required fuel-intensity reduction, the higher the credit price. The policymaker sets t (or equivalently, S) to maximize social welfare, which is the sum of true utility across consumer types:

$$(4) \quad W(t) = \underbrace{T(t)}_{\text{Credit revenue}} + \sum_l \left[\underbrace{\frac{1}{\eta_l} \ln \left(\sum_j \exp(V_{lj}(t, b_{lj})) \right)}_{\text{Perceived consumer surplus}} + \underbrace{\sum_j b_{lj} G_{lj} P_{lj}(t, \mathbf{b}_l)}_{\text{Bias}} \right].$$

The first terms $T(t)$ reflects credit revenues. The second term is perceived consumer surplus, from the standard Small and Rosen (1981) formula. The final term is the bias: the expected difference $b_{lj} G_{lj}$ between perceived and true consumer surplus, summing over vehicles and weighting by choice probability P_{lj} .

Ideally, the policymaker could achieve the first best through some perfect informational intervention that fully removes all bias, causing all consumers to have $\mathbf{b}_l = \mathbf{0}$. Alternatively, the first best would obtain under a hypothetical system of type-by-vehicle-specific taxes that exactly offset each type’s bias in evaluating each vehicle: $\tau_{lj}^* = -b_{lj} G_{lj}$. Of course, such individually tailored taxes are not practical. Furthermore, a perfect information provision intervention seems both unrealistic and costly; our information provision intervention took a meaningful amount of consumers’ time to deliver, and it only provided information about a few vehicles. For this reason, the social planner is constrained to considering the second-best social optimum under a fuel economy standard.

B. Results

We use this framework to derive a proposition that demonstrates the potential policy implications of our treatment effect estimates. In online Appendix C, we first derive a result that parallels results in Diamond (1973) and Allcott and Taubinsky (2015): the socially optimal fuel economy standard imposes a credit price t^* that equals the average marginal bias—that is, the average misperception of fuel costs across types l , weighted by each type’s responsiveness to the tax.

standard also depends on the covariance between markup and fuel economy, and the optimal policy formula has an additional term reflecting this. If, as is likely to be the case, markups are higher for low-fuel economy vehicles, this would suppress sales of these vehicles. To offset this, the optimal fuel economy standard would thus be less stringent than under perfect competition. On the other hand, if high-fuel economy vehicles had relatively higher markups, it would be optimal to offset this through a more stringent fuel economy standard.

For our key proposition, we define Q as the effect of a pure nudge on sales-weighted average fuel intensity: $Q \equiv \sum_l \sum_j e_j [P_{lj}(0, \mathbf{0}) - P_{lj}(0, \mathbf{b}_l)]$. We further assume that b and χ are either homogeneous or heterogeneous in a way such that the “mismatching” of the fuel economy standard—that is, the difference between a vehicle’s CAFE credit cost and consumers’ bias in evaluating the vehicle—is orthogonal to fuel intensity and true preferences across vehicles. Under this assumption, online Appendix C derives the following proposition.

PROPOSITION 1: *The socially optimal fuel economy standard reduces fuel intensity by the same amount as a pure nudge:*

$$(5) \quad S(t^*) = Q.$$

C. Using Treatment Effects in the Context of the Model

In Section VD, we will discuss the real-world interpretation of our results, including important caveats. In this section, we first present a mechanical implementation of our treatment effect estimates in the context of Proposition 1. This section answers the following question: *how stringent of a fuel economy standard would be justified by the classes of imperfect information and attention addressed by our interventions?*

Table 7 illustrates these mechanical implications. The top panel presents estimates of stringency $S(t)$ for the current and proposed CAFE standards. The objective of the “counterfactual” is to establish the average fuel intensity that would arise in the absence of CAFE standards, or $\sum_l \sum_j e_j P_{lj}(0, \mathbf{b}_l)$ in our model. The appropriate counterfactual depends on assumptions about technological change, consumer demand, and gas prices. As a simple benchmark, we use the sales-weighted average fuel economy for model year 2005 vehicles. We choose 2005 both because gas prices were very similar to their current (2016 average) levels and because it just precedes the modern increase in the stringency of the CAFE regulation. Using later years as a counterfactual would incorrectly include increasing effects of the regulation in the no-regulation counterfactual, whereas using earlier years would involve increasingly outdated vehicle technologies and consumer preferences. The 2005 stringency may be too high, as CAFE standards were already binding for some automakers in 2005, or too low, as technological change and consumer preferences could have evolved since then in the absence of the regulation. An alternative possible counterfactual is the baseline fleet assumed in the 2012 Regulatory Impact Analysis (NHTSA 2012), which delivers a similar number.¹⁷

We calculate stringency of the CAFE regulation as of 2016 and 2025 by subtracting the regulatory requirement in each year from the 2005 counterfactual.

¹⁷This counterfactual would be 20.5 MPG, which is comparable to our 2005 benchmark of 19.9 MPG. 20.5 MPG is the 25.9 MPG (unadjusted) fuel economy from Table 15 of NHTSA (2012), multiplied by 0.790 to transform to adjusted MPG using the 2010 model year adjustment factor in Table 10.1 of EPA (2016).

TABLE 7—TREATMENT EFFECTS VERSUS ACTUAL CAFE STANDARDS

	Gallons per 100 miles (1)	Miles per gallon (2)
<i>Panel A. Current CAFE standards</i>		
“Counterfactual” (2005 sales)	5.03	19.9
2016 sales	3.91	25.6
2025 CAFE standard	2.77	36.1
“2016 stringency”: 2016 sales – counterfactual	–1.12	5.7
“2025 Stringency”: 2025 CAFE standard – counterfactual	–2.26	16.2
<i>Panel B. Treatment effects of information</i>		
Stated preference (point estimate; Table 3, panel C)	–0.032	0.20
Revealed preference (90% confidence bound; Table 5)		
Dealership experiment, equally weighted (column 2)	–0.06	1.08
Dealership experiment, re-weighted (column 3)	–0.49	3.14
Online experiment, equally weighted (column 5)	–0.04	0.29
Online experiment, re-weighted (column 6)	–0.08	0.62

Notes: Panel A details the CAFE standards currently in effect for light-duty vehicles. Sales-weighted adjusted fuel economy for model years 2005 and 2016 are from Table 2.1 of EPA (2016). The 2025 CAFE standard is the “achieved” unadjusted sales-weighted MPG of 46.2 from NHTSA (2012), multiplied by 0.782 to transform to adjusted MPG; the 0.782 adjustment factor reflects data for the most recent year in Table 10.1 of EPA (2016). Panel B presents the treatment effects of information, as estimated in Tables 3 and 5. In the bottom panel, the miles per gallon estimates in column 2 are calculated by re-estimating equation (3) with fuel economy in miles per gallon as the dependent variable.

For the 2016 regulatory requirement, we directly use sales-weighted fuel economy of model year 2016 vehicles from EPA (2016). For 2025, we use the fuel economy that the NHTSA (2012) projects would be achieved under the presumptive standard, after accounting for various alternative compliance strategies. Subtracting the counterfactuals from the regulatory requirements gives fuel intensity decreases of 1.12 and 2.26 gallons per 100 miles in 2016 and 2025, respectively, or increases of 5.7 and 16.2 MPG.

The bottom panel recaps our key treatment effect estimates from Section IV. Column 1 is restated directly from previous tables, while the results in units of MPG in column 2 are from re-estimating the same regressions with fuel economy in MPG as the dependent variable. The stated preference results from panel C of Table 3 would justify a required decrease of 0.032 gallons per 100 miles, or equivalently an increase of 0.20 MPG.

The revealed preference estimates from Table 5 show statistically zero effect. The 90 percent confidence intervals for the dealership and online experiments, respectively, reject fuel intensity decreases of more than -0.06 and -0.04 gallons per 100 miles in sample, and -0.49 and -0.08 when re-weighted for national representativeness on observables. When reestimated with the dependent variable in MPG, the confidence bounds for the two experiments are 1.08 and 0.29 MPG, respectively, or 3.14 and 0.62 MPG when re-weighted. Thus, the current and proposed CAFE standards are significantly more stringent than would be optimal to address the classes of imperfect information and inattention addressed by our interventions.

D. Interpretation and Caveats

Having stated the mechanical results, we now discuss the interpretation and real-world implications.

First, it is important to consider what classes of imperfect information and inattention these treatments would address (if they exist), versus what classes of imperfect information and inattention could still be present even with our zero treatment effects. There are at least four models in which our interventions would *not* address informational and attentional distortions even if they do exist. First, information provision can be ineffective in models such as Sims (2010), in which agents face cognitive costs in *using* information to make a decision, even if they have previously *seen* all relevant information. Second, our treatments may be ineffective in models that include imperfect memory, in which consumers must not only receive information, but must receive it at the right time. Indeed, we have reported some evidence that consumers forgot the information we provided, as the immediate effects on stated preference in the online experiment are no longer evident in identical stated preference questions in the follow-up survey. Third, our treatments may be ineffective in models where consumers need to receive information for more than just the vehicles that they are considering most closely. Fourth, one can always propose more nuanced models where the *presentation* or *trust* of information matters, not just the fact that it was provided, and such models can always be constructed ad hoc to argue that any particular treatment should have been ineffective.

On the other hand, these treatments would mechanically address at least two standard types of imperfect information and inattention, if they exist. First, it is mechanically true that our treatments drew attention to fuel economy for at least a short period, and so would address the distortion in any model where consumers simply fail to think about fuel economy at all. This type of model is often discussed in the literature: for example, Gabaix and Laibson (2006, 506) introduce their analysis of “shrouded attributes” with discussions of consumers who “do not think about add-ons.” Prior survey evidence suggests that this type of model could be highly relevant in this context: a remarkable 40 percent of American car buyers report that “I did not think about fuel costs at all when making my decision” (Allcott 2011a). Second, by providing individually tailored fuel cost information, our treatments address the distortion in models such as Sallee (2014) in which consumers observe product attributes, can foresee their driving patterns, and can form some imprecise understanding of how this translates into total fuel costs, but they face cognitive costs to precisely do that calculation.¹⁸ Prior literature suggests that this model could also have been relevant: Davis and Metcalf (2016) show that individually tailored energy cost information has significant effects on stated choices between energy-using durables when hypothetical choices are made immediately after the information is provided.

¹⁸ In the model of Sallee (2014, 782), “Consumers observe the various attributes of each product, but they have an incomplete understanding of lifetime fuel costs that is, they have some rough idea of how much fuel will cost over the products life, but they are uncertain about this cost. Consumers can resolve (or reduce significantly) this uncertainty by doing research and performing calculations, but this requires costly effort.” Our treatments provide this information, personally tailored to the consumer’s consideration set, gas price beliefs, and driving patterns.

Second, on a practical level, one might question how participants in our experiments engaged with the information provided. For example, some people in our treatment groups might have wanted to ignore or speed through the intervention. To help mitigate this, we had the dealership RAs record whether people had completed the intervention, and we required online experiment participants to answer quizzes before completing the intervention. As another example, the interventions could have induced experimenter demand effects, in which participants changed their vehicle purchases to conform to what they perceived the researchers wanted. To address this, we clearly communicated to the dealership RAs that “our explicit goal is not to influence consumers to pursue fuel-efficient vehicles. Rather, we are exploring the ways in which the presentation of information affects ultimate purchasing behavior.” It seems unlikely that experimenter demand effects would meaningfully influence such large purchases, especially given that experiment participants typically did not make purchases the same day as the intervention and were probably uncertain as to whether they would ever hear from us again. Any experimenter demand effects would likely increase the treatment effects, which biases against our result of zero effect.

Third, as we have documented above, imperfect information and inattention are only *part* of the potential rationale for fuel economy standards: externalities and other market failures, plus political constraints against raising gasoline taxes, are also important motivations. Thus, our analysis can be viewed as evaluating these biases in isolation as a justification for CAFE. This is still relevant, because as described earlier, the Regulatory Impact Analyses rely largely on consumers’ private net benefits—not externalities—to justify the stringency of the policy.¹⁹ Our results suggest that the classes of imperfect information and inattention addressed by our interventions should not be used as principal justifications for stringent CAFE standards.

Fourth, our samples are not representative of the US population, both because of selection into the original randomized sample and attrition from that sample to the final sample for which we have vehicle purchase data. To help mitigate this issue, we ran two experiments in very different populations and reweighted on observables. Of course, both of our samples likely still differ in unobservable ways from the policy-relevant target population.

While each of these concerns is important, the results imply that the true effects of an ideal informational intervention would have to be dramatically different than our estimates for imperfect information and inattention to be valid as a significant justification for the current CAFE standards.

VI. Conclusion

It has long been argued that consumers are poorly informed, inattentive, or otherwise cognitively constrained when evaluating fuel economy, and that this

¹⁹ Allcott, Mullainathan, and Taubinsky (2014) presents a model that includes externalities, as well as other extensions, such as a vehicle utilization margin (the decision of how much to drive) and gas taxes, as a potential policy instrument.

causes them to buy systematically lower-fuel economy vehicles than would be optimal. We tested this hypothesis with two information provision field experiments. In both experiments, we find that our treatments did not have a statistically or economically significant effect on the average fuel economy of purchased vehicles.

Qualitatively, there are perhaps two main interpretations of these results. The first is that while our interventions did draw attention to fuel economy for a few minutes, the information we provided was not very useful, and/or people soon forgot it. Put simply, the interventions did not come close to fully informing people about fuel economy. This still points to a deeply interesting implication. New cars already have fuel economy information labels prominently posted in the windows, and the Environmental Protection Agency has a useful fuel economy information website, www.fueleconomy.gov. Then, in addition, our dealership intervention provided in-person, individually tailored fuel economy information via a well-designed tablet computer app. If, after all of these experimental and official government efforts, we still need stringent fuel economy standards to address lack of information about fuel economy, this is a striking testament to the challenges to providing information to consumers.

The second interpretation is to take the empirical estimates more seriously in the context of our optimal policy model, arguing that imperfect information and inattention do not have a significant systematic effect on vehicle markets. This would imply either that some other market failure or behavioral failure must justify the CAFE standard, or that the large net private benefits projected in the CAFE Regulatory Impact Analyses do not actually exist. The latter possibility would arise if the RIAs' engineering models did not account for the full fixed costs, production costs, or performance reductions from fuel economy-improving technologies. In this case, there would still be an economic justification for fuel economy standards as a second-best externality policy—albeit a highly inefficient one, as shown by Jacobsen (2013). But if fuel economy is more expensive than the RIA models assume, the socially optimal CAFE standard would likely be significantly less stringent than the current or proposed levels.

REFERENCES

- Abaluck, Jason, and Jonathan Gruber.** 2011. "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review* 101 (4): 1180–1210.
- Allcott, Hunt.** 2011a. "Consumers' Perceptions and Misperceptions of Energy Costs." *American Economic Review* 101 (3): 98–104.
- Allcott, Hunt.** 2011b. "Social norms and energy conservation." *Journal of Public Economics* 95 (9–10): 1082–95.
- Allcott, Hunt.** 2013. "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market." *American Economic Journal: Economic Policy* 5 (3): 30–66.
- Allcott, Hunt.** 2015. "Site Selection Bias in Program Evaluation." *Quarterly Journal of Economics* 130 (3): 1117–65.
- Allcott, Hunt, and Christopher Knittel.** 2019. "Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments: Dataset." *American Economic Journal: Economic Policy*. <https://doi.org/10.1257/pol.20170019>.
- Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky.** 2018. "Regressive Sin Taxes, with an Application to the Optimal Soda Tax." Unpublished. <https://sites.google.com/site/allcott/research>.

- Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. 2014. "Energy Policy with Externalities and Internalities." *Journal of Public Economics* 112: 72–88.
- Allcott, Hunt, and Richard L. Sweeney. 2017. "The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment." *Management Science* 63 (1): 21–39.
- Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market." *American Economic Review* 105 (8): 2501–38.
- Allcott, Hunt, and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *Review of Economics and Statistics* 96 (10): 779–95.
- Austin, David, and Terry Dinan. 2005. "Clearing the Air: The Costs and Consequences of Higher CAFE Standards and Increased Gasoline Taxes." *Journal of Environmental Economics and Management* 50 (3): 562–82.
- Barber, Brad M., Terrance Odean, and Lu Zheng. 2005. "Out of Sight, Out of Mind: The Effects of Expenses on Mutual Fund Flows." *Journal of Business* 78 (6): 2095–2119.
- Bhargava, Saurabh, and Dayanand Manoli. 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11): 3489–3529.
- Bollinger, Bryan, Phillip Leslie, and Alan Sorensen. 2011. "Calorie Posting in Chain Restaurants." *American Economic Journal: Economic Policy* 3 (1): 91–128.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* 103 (1): 220–56.
- Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–77.
- Choi, James J., David Laibson, and Brigitte C. Madrian. 2010. "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds." *Review of Financial Studies* 23 (4): 1405–32.
- Davis, Lucas W., and Gilbert E. Metcalf. 2016. "Does Better Information Lead to Better Choices? Evidence from Energy-Efficiency Labels." *Journal of the Association of Environmental and Resource Economists* 3 (3): 589–625.
- Diamond, Peter A. 1973. "Consumption Externalities and Imperfect Corrective Pricing." *Bell Journal of Economics and Management Science* 4 (2): 526–38.
- Dolan, Paul, and Robert Metcalfe. 2013. "Neighbors, Knowledge, and Nuggets: Two Natural Field Experiments on the Role of Incentives on Energy Conservation." London School of Economics Centre for Economic Performance (CEP) Discussion Paper 1222.
- Dranove, David, and Ginger Zhe Jin. 2010. "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature* 48 (4): 935–63.
- Dubin, Jeffrey A., and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52 (2): 345–62.
- Dufo, Esther, and Emmanuel Saez. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 118 (3): 815–42.
- Dupas, Pascaline. 2011. "Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics* 3 (1): 1–34.
- Environmental Protection Agency (EPA). 2010. "Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards; Final Rule." *Federal Register* 75 (88): 25324–25728.
- Environmental Protection Agency (EPA). 2012. *Regulatory Impact Analysis: Final Rulemaking for 2017–2025 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards*. Environmental Protection Agency, Washington, DC, August.
- Environmental Protection Agency (EPA). 2016. *Light-Duty Automotive Technology, Carbon Dioxide Emission, and Fuel Economy Trends: 1975 Through 2016*. Environmental Protection Agency, Washington, DC, November.
- Farhi, Emmanuel, and Xavier Gabaix. 2015. "Optimal Taxation with Behavioral Agents." <https://scholar.harvard.edu/files/farhi/files/br-taxes.pdf>.
- Ferraro, Paul J., and Michael K. Price. 2013. "Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment." *Review of Economics and Statistics* 95 (1): 64–73.
- Fischer, Carolyn, Winston Harrington, and Ian W. H. Parry. 2007. "Should Automobile Fuel Economy Standards be Tightened?" *Energy Journal* 28 (4): 1–29.
- Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121 (2): 505–40.
- Gayer, Ted. 2011. "A Better Approach to Environmental Regulation: Getting the Costs and Benefits Right." Hamilton Project Discussion Paper 2011-06.

- Goldberg, Pinelopi Koujianou.** 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the U.S." *Journal of Industrial Economics* 46 (1): 1–33.
- Greene, David L., Philip D. Patterson, Margaret Singh, and Jia Li.** 2005. "Feebates, rebates and gas-guzzler taxes: A study of incentives for increased fuel economy." *Energy Policy* 33 (6): 757–75.
- Grigolon, Laura, Mathias Reynaert, and Frank Verboven.** 2017. "Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the European car market." University of Leuven Working Paper 17-836.
- Grubb, Michael D.** 2009. "Selling to Overconfident Consumers." *American Economic Review* 99 (5): 1770–1807.
- Gruber, Jonathan, and Botond Köszegi.** 2004. "Tax incidence when individuals are time-inconsistent: The case of cigarette excise taxes." *Journal of Public Economics* 88 (9–10): 1959–87.
- Hainmueller, Jens.** 2012. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis* 20 (1): 25–46.
- Handel, Benjamin R., and Jonathan T. Kolstad.** 2015. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–2500.
- Hastings, Justine S., and Jeffrey M. Weinstein.** 2008. "Information, School Choice, and Academic Achievement: Evidence from Two Experiments." *Quarterly Journal of Economics* 123 (4): 1373–1414.
- Hausman, Jerry A.** 1979. "Individual discount rates and the purchase and utilization of energy-using durables." *Bell Journal of Economics* 10 (1): 33–54.
- Heutel, Garth.** 2015. "Optimal policy instruments for externality-producing durable goods under present bias." *Journal of Environmental Economics and Management* 72: 54–70.
- Hossain, Tanjim, and John Morgan.** 2006. "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay." *Advances in Economic Analysis and Policy* 6 (2): 1–30.
- Howell, Sabrina T.** 2018. "Joint ventures and technology adoption: A Chinese industrial policy that backfired." *Research Policy* 47 (8): 1448–62.
- Ito, Koichiro, and James M. Sallee.** 2014. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards." National Bureau of Economic Research (NBER) Working Paper 20500.
- Jacobsen, Mark R.** 2013. "Evaluating US Fuel Economy Standards in a Model with Producer and Household Heterogeneity." *American Economic Journal: Economic Policy* 5 (2): 148–87.
- Jacobsen, Mark R., and Arthur A. van Benthem.** 2015. "Vehicle Scrappage and Gasoline Policy." *American Economic Review* 105 (3): 1312–38.
- Jensen, Robert.** 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* 125 (2): 515–48.
- Jessoe, Katrina, and David Rapson.** 2015. "Commercial and Industrial Demand Response Under Mandatory Time-of-Use Electricity Pricing." *Journal of Industrial Economics* 63 (3): 397–421.
- Jin, Ginger Zhe, and Alan T. Sorensen.** 2006. "Information and consumer choice: The value of publicized health plan ratings." *Journal of Health Economics* 25 (2): 248–75.
- Kempton, Willett, and Laura Montgomery.** 1982. "Folk quantification of energy." *Energy* 7 (10): 817–27.
- Klier, Thomas, and Joshua Linn.** 2010. "The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data." *American Economic Journal: Economic Policy* 2 (3): 134–53.
- Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel.** 2012. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *Quarterly Journal of Economics* 127 (1): 199–235.
- Larrick, Richard P., and Jack B. Soll.** 2008. "The MPG Illusion." *Science* 320 (5883): 1593–94.
- Mullainathan, Sendhil, Joshua Schwartzstein, and William J. Congdon.** 2012. "A Reduced-Form Approach to Behavioral Public Finance." *Annual Review of Economics* 4 (1): 511–40.
- National Highway Traffic Safety Administration (NHTSA).** 2012. *Final Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2017-MY 2025 Passenger Cars and Light Trucks*. Office of Regulatory Analysis and Evaluation, National Center for Statistics and Analysis. Washington, DC, March. https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/fria_2017-2025.pdf.
- Newell, Richard G., and Juha Siikamäki.** 2014. "Nudging Energy Efficiency Behavior: The Role of Information Labels." *Journal of the Association of Environmental and Resource Economists* 1 (4): 555–98.
- Oak Ridge National Laboratory.** 2010. *How Consumers Value Fuel Economy: A Literature Review*. U.S. Environmental Protection Agency. Washington, DC, March.

- O'Donoghue, Ted, and Matthew Rabin.** 2006. "Optimal sin taxes." *Journal of Public Economics* 90 (10–11): 1825–49.
- Reynaert, Mathias, and James M. Sallee.** 2016. "Corrective Policy and Goodhart's Law: The Case of Carbon Emissions from Automobiles." National Bureau of Economic Research (NBER) Working Paper 22911.
- Sallee, James M.** 2014. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics* 57 (3): 781–820.
- Sallee, James M., Sarah E. West, and Wei Fan.** 2016. "Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations." *Journal of Public Economics* 135: 61–73.
- Sanstad, Alan H., and Richard B. Howarth.** 1994. "'Normal' markets, market imperfections, and energy efficiency." *Energy Policy* 22 (10): 811–18.
- Scanlon, Dennis P., Michael Chernew, Catherine McLaughlin, and Gary Solon.** 2002. "The impact of health plan report cards on managed care enrollment." *Journal of Health Economics* 21 (1): 19–41.
- Sims, Christopher A.** 2010. "Rational Inattention and Monetary Economics." In *Handbook of Monetary Economics*, Vol. 3, edited by Benjamin Friedman and Michael Woodford, 155–81. Amsterdam: North-Holland.
- Small, Kenneth A., and Harvey S. Rosen.** 1981. "Applied Welfare Economics with Discrete Choice Models." *Econometrica* 49 (1): 105–30.
- Stern, Paul C., and Elliot Aronson, eds.** 1984. *Energy Use: The Human Dimension*. New York: W. H. Freeman and Co.
- Tsvetanov, Tsvetan, and Kathleen Segerson.** 2013. "Re-evaluating the role of energy efficiency standards: A behavioral economics approach." *Journal of Environmental Economics and Management* 66 (2): 347–63.
- Turrentine, Thomas S., and Kenneth S. Kurani.** 2007. "Car buyers and fuel economy?" *Energy Policy* 35: 1213–23.

This article has been cited by:

1. Davide Cerruti, Claudio Daminato, Massimo Filippini. 2019. The Impact of Policy Awareness: Evidence from Vehicle Choices Response to Fiscal Incentives. *SSRN Electronic Journal*. [[Crossref](#)]
2. Cass R. Sunstein. 2019. Back to Mill? Behavioral Welfare Economics. *SSRN Electronic Journal*. [[Crossref](#)]