Paternalism and Energy Efficiency: An Overview

Hunt Allcott[*]

August 25, 2015

Abstract

This review paper provides an overview of the application of behavioral public economics to energy efficiency. I document policymakers' arguments for "paternalistic" energy efficiency policies, formalize with a simple model of misoptimizing consumers, review and critique empirical evidence, and suggest future research directions. While empirical results suggest that policies to address imperfect information and internalities may increase welfare in some cases, some existing policies may be mistargeted or miscalibrated.

JEL Codes: D03, D12, D83, H21, L51, L94, Q41, Q48.

Keywords: Behavioral public economics, energy efficiency.

---

[*]New York University, NBER, JPAL, ideas42, and E2e. Email: hunt.allcott@nyu.edu.

1

1 Introduction

Governments frequently intervene to protect citizens from our own choices. For example, the U.S. and many other countries tax or ban hard drugs, alcohol, and cigarettes. Motorcycle helmet and automobile seat belt laws are common, as are safety standards for food and other products. Usury laws and other regulations protect consumers from their own financial decisions. "Life-cycle myopia" is one potential justification for mandatory retirement savings programs (Feldstein and Liebman 2002).

Without models of imperfect information or misoptimization, economics would have little useful to say about these policies: with perfect information and optimizing consumers, any restriction of consumer choice mechanically reduces consumer welfare, essentially by assumption. Information economics and behavioral public economics provide the tools to formalize and test hypotheses about "mistakes" and to evaluate the welfare effects of "paternalistic" policies.[1] I use the word "mistakes" to cover any reason why our decisions don't maximize our own welfare, including that we don't know or don't properly value the consequences. In a world with mistakes, "paternalistic" is not a pejorative word - instead, it is a descriptive term that clarifies policy goals and defines empirical hypotheses.

This paper is an overview of the application of behavioral public economics to energy efficiency policy. Along with externality reduction, which of course would be best achieved through Pigouvian taxes, consumer mistakes are offered as an important justification for policies such as subsidies for energy efficient goods and minimum energy efficiency standards. I begin by reviewing the policy argument and formalizing it in a simple consumer choice model. I use the model to characterize three categories of empirical tests for mistakes: comparing demand responses to prices vs. energy costs, measuring effects of nudges, and belief elicitation. These tests have direct analogies in other domains such as consumer finance and health, and the categorization draws on similar discussions in DellaVigna (2009) and Mullainathan, Schwartzstein, and Congdon (2012). I then discuss tests between different behavioral models, outline emerging policy implications, and suggest directions for future research.

There are two important related areas that this review does not cover.[2] The first is the literature on how energy conservation marketing campaigns, social norm information, goal setting, and other "behavior levers" affect energy use.[3] Many of these interventions could act through persuasion

---

[1] In the past ten years, a growing number of papers have made empirical or theoretical contributions to behavioral public economics, including Baicker, Mullainathan, and Schwartzstein (2015), Bernheim and Rangel (2004, 2009), Carroll, Choi, Laibson, Madrian, and Metrick (2009), Gabaix and Laibson (2006), Grubb (2014), Grubb and Osborne (2015), Gruber and Koszegi (2004), Gul and Pesendorfer (2007), Gruber and Mullainathan (2005), Mullainathan, Schwartzstein, and Congdon (2012), O'Donoghue and Rabin (2006), and others.

[2] See Allcott and Greenstone (2012), Gerarden, Newell, and Stavins (2015), Gillingham, Newell, and Palmer (2009), and Gillingham and Palmer (2014) for broader reviews of energy efficiency policy. Allcott and Sunstein (2015) discuss regulation of internalities with a focus on energy efficiency.

[3] Recent work includes Allcott (2011a), Allcott and Rogers (2014), Asensio and Delmas (2015), Ayres, Raseman,

instead of correcting consumer mistakes, so the regulatory policy implications are unclear unless one can distinguish between these two channels and develop a model of socially-optimal persuasion.[4] Second, I leave aside "between-consumer" market failures that could affect energy efficiency investments, such as information asymmetries between landlords and tenants or home sellers and home buyers.[5]

2   Background

2.1   Why Energy Matters

Expenditures on energy and energy-using durables represent a significant share of the economy in industrialized countries. In 2011, US households purchased $361 billion worth of energy-using durable goods such as cars and air conditioners and another $570 billion on gasoline, electricity, natural gas, and fuel oil (BLS 2014). Even relatively small inefficiencies in such large markets could quickly add up to large welfare losses.

There is also increased policy interest due to energy use externalities, including greenhouse gas and local pollution emissions and national security concerns related to oil imports. US household energy use imposes $40 billion in carbon pollution externalities, plus additional externalities from local air pollution.[6] Any systematic consumer mistakes that increase energy use would also increase externalities.

2.2   Energy Efficiency Policies

The United States has a number of policies that encourage energy efficiency, many of which were introduced as energy prices rose in the 1970s. Table 1 presents the most significant US policies, organized into three categories. Many other industrialized countries and large developing countries have similar policies.

The first category is standards. The Energy Policy and Conservation Act of 1975 called for minimum energy efficiency standards for home appliances. The first meaningful nationwide standards were finally implemented in 1990, and since then the standards have been strengthened and additional products included. The Energy Independence and Security Act of 2007 set minimum lighting efficiency standards that banned most traditional incandescent lightbulbs between 2012

---

and Shih (2013), Costa and Kahn (2013), Delmas and Lessem (2014), Dolan and Metcalfe (2015), Harding and Hsiaw (2014), Ito, Ida, and Tanaka (2015), Yoeli, Hoffman, Rand, and Nowak (2013), and many others.

[4]See Allcott and Kessler (2015) for one initial step in this direction.

[5]Davis (2012), Gillingham, Harding, and Rapson (2012), and Myers (2014) study the landlord-tenant information asymmetry.

[6]This number is the product of the $38 social cost of carbon for 2013 estimated by the Interagency Working Group on Social Cost of Carbon (2013) and total household carbon emissions from energy use estimated by the US Energy Information Administration (2014).

and 2014 and will be tightened further in 2020. Argentina, Australia, Brazil, Canada, China, Cuba, the European Union, Israel, Malaysia, Russia, and Switzerland have also banned some or all incandescent light bulbs. Many states have building codes that mandate minimum insulation levels and other energy efficiency measures. Corporate Average Fuel Economy (CAFE) standards require that new cars and trucks meet a minimum average mile per gallon (MPG) rating.

The second category is price policies, i.e. taxes and subsidies. There were federal income tax credits of up to $3,400 for hybrid vehicles from 2006 to 2010, and credits of up to $7,500 are currently available for plug-in hybrids and electric vehicles. There are federal "gas guzzler taxes" range from $1,000 to $7,700 on the sale of low-MPG passenger cars. The Weatherization Assistance Program grants $250 million for improved insulation, air sealing, and other weatherization measures at approximately 100,000 low-income homes each year. Electricity bill surcharges fund billions of dollars of utility-managed "demand-side management" programs, which include subsidized home energy audits, energy efficiency information provision, and subsidies for energy efficient appliances, weatherization, and other investments. The 2009 American Recovery and Reinvestment Act included substantial increases in energy efficiency subsidies. In total, that legislation and related economic stimulus bills included $17 billion in energy efficiency spending, including additional weatherization subsidies and automobile and appliance cash-for-clunkers programs.

The final category is information and marketing. The US government requires all new vehicles to have fuel economy information labels, and appliances must be labeled with informational "yellow tags." Furthermore, the Environmental Protection Agency runs the Energy Star labeling and marketing program. Although cost data are not available for all programs, information and marketing are cheap compared to the fiscal cost of subsidies or the production costs of energy efficient goods.

2.3 The Paternalistic Rationale for Energy Efficiency Policy

Energy efficiency policies can address several different market failures, including energy use externalities, between-consumer asymmetric information problems, and spillovers of returns to innovation. Imperfect information and misoptimization, however, are central to the policy discussion. This section documents that with several examples.

Some of the earliest behavioral language is from Jerry Hausman's (1979) seminal analysis of the implied discount rates that rationalize consumers' tradeoffs between purchase price and energy efficiency in air conditioner purchases. On page 51, Hausman writes:

> *This finding of a high individual discount rate does not surprise most economists. At least since Pigou, many economists have commented on a "defective telescopic faculty." A simple fact emerges that in making decisions which involve discounting over time, individuals behave in a manner which implies a much higher discount rate than can be explained in terms of the opportunity cost of funds available in credit markets. Since this*

individual discount rate substantially exceeds the social discount rate used in benefit-cost calculations, the divergence might be narrowed by policies which lead to purchases of more energy-efficient equipment.

The next year, the U.S. Department of Energy relied heavily on this argument in its Regulatory Impact Analysis (RIA) of the first appliance energy efficiency standards. Using engineering analyses of appliance costs and energy savings, the RIA found that the proposed standards would generate large net private benefits for consumers. Of course, such private net benefits could only reflect welfare gains in the presence of mistakes or between-consumer market failures. The Department of Energy and other federal agencies debated whether consumers were "uninformed" and "myopic"; see Regulatory Analysis Review Group (1980).

Engineering cost calculations in subsequent RIAs have consistently found that further strengthened energy efficiency standards generate large net private benefits. For example, the RIA for the 2012-2016 CAFE standard (NHTSA 2010) found that $15 billion per year in consumer welfare gains will accrue to consumers who would be induced to buy higher-fuel economy vehicles. Without these private gains, the regulation is likely to be welfare-reducing - such stringent CAFE standards cannot be justified by externalities alone. The RIA states (page 2):

> *Although the economy-wide or "social" benefits from requiring higher fuel economy represent an important share of the total economic benefits from raising CAFE standards, NHTSA estimates that benefits to vehicle buyers themselves will significantly exceed the costs of complying with the stricter fuel economy standards this rule establishes. [...] This raises the question of why current purchasing patterns do not result in higher average fuel economy, and why stricter fuel efficiency standards should be necessary to achieve that goal. To address this issue, the analysis examines possible explanations for this apparent paradox, including discrepancies between the consumers' perceptions of the value of fuel savings and those calculated by the agency.*

Similarly, the RIA for Australia's phase-out of traditional incandescent lightbulbs (DEHWA 2008) argues that internalities and information failures justify the policy, as well as asymmetric information in housing markets. Page vii states:

> *[Incandescent lightbulbs] continue to sell remarkably well because, if their energy costs are ignored, they appear cheap ... There are significant information failures and split incentive problems in the market for energy efficient lamps. Energy bills are aggregated and periodic and therefore do not provide immediate feedback on the effectiveness of individual energy saving investments. Consumers must therefore gather information and perform a reasonably sophisticated calculation to compare the life-cycle costs of tungsten filament lamps and CFLs. But many lack the skills. For others, the amounts*

> *saved are too small to justify the effort or they do not remain at the same address long enough to benefit fully from a long lived energy saving lamp.*

Using similar arguments, lightbulb manufacturers and environmental groups lobbied together for the US lighting efficiency standards. In congressional testimony, the National Electrical Manufacturers Association argued, "New standards-setting legislation is needed in order to further educate consumers on the benefits of energy-efficient products" (US GPO 2007).[7] The Natural Resources Defense Council (2011) writes:

> *Some in Congress are considering repealing the new efficiency standards before they even take effect. That would take away $12.5 billion in consumer savings—something none of us can afford.*

The RIA for the Energy Independence and Security Act (EISA) of 2007, which included the lighting efficiency standards, argues that after accounting for incremental production costs, the lighting standards will save consumers a net present value of $27 to $64 billion over 30 years (DOE 2009). As with CAFE standards, private net benefits are considerably more important than carbon externality reduction, which the EISA RIA values at no more than $16 billion over 30 years.

Gayer (2011) succinctly summarizes the argument:

> *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.*

Additional academic articles have followed Hausman (1979). Fischer, Harrington, and Parry (2007) and Parry, Evans, and Oates (2010) use simulation models to analyze whether mistakes might justify energy efficiency policies. In a review of fuel economy policies in the Journal of Economic Literature, Parry, Walls, and Harrington (2007) conclude that:

> *Higher fuel economy standards significantly increase efficiency only if carbon and oil dependence externalities greatly exceed the mainstream estimates . . . or if consumers perceive only about a third of the actual fuel economy benefits.*
>
> *Unfortunately, there is little in the way of solid empirical (as opposed to anecdotal) evidence on this hotly contested issue ...*

---

[7]US Senator Rand Paul has strong views on this argument. During congressional testimony by Kathleen Hogan, Deputy Assistant Secretary for Energy Efficiency at the Department of Energy, Paul said (ABC 2011), "You're really anti-choice on every other consumer item that you've listed here, including light bulbs, refrigerators, toilets – you name it, you can't go around your house without being told what to buy. You restrict my choices, you don't care about my choices." Paul continues, "This is what your energy efficiency standards are. Call it what it is. You prevent people from making things that consumers want."

The rest of this paper formalizes these policy assertions in a model and provides an overview of recent empirical evidence. Such empirical evidence is crucial: an alternative explanation for the RIAs' findings of large private net benefits is that their engineering cost calculations are overly optimistic. If this is true, then energy efficiency policies based on these RIAs might actually be welfare reducing.

3   Model

3.1   Setup

This section formalizes policy assertions about consumer mistakes in a very simple model. It roughly follows Allcott and Taubinsky (2015, "AT") and Allcott, Mullainathan, and Taubinsky (2014, "AMT"), but it also draws on work by Heutel (2011), Sallee (2014), and the frameworks in DellaVigna (2009) and Mullainathan, Schwartzstein, and Congdon (2012).

Consumers have unit demand and choose between two energy-using durable goods, indexed $j \in \{E, I\}$. Good E is more energy efficient than good I, with energy requirement $e_E < e_I$ per unit of usage. The goods are produced at marginal cost c_j, and markets are perfectly competitive, with $p_j = c_j$ in the absence of a subsidy. Time is indexed by t, and the durables have fixed lifetime T. Consumers have exogenous utilization of m units per unit time, the energy cost is g, and consumers use discount factor δ. Expected total lifetime energy costs for good j are thus:

$$G_j \equiv \sum_{t=0}^{T} \delta^t m e_j g. \tag{1}$$

Consumers receive usage utility v_j from owning good j. I define differences $p \equiv p_E - p_I$, $G = G_E - G_I$, and $v = v_E - v_I$.

3.2   Specific Behavioral Models of Mistakes

Fully informed and optimizing consumers purchase good E if and only if the net benefits outweigh the relative purchase price, i.e. if $v - G > p$. Since $G < 0$, the left hand side adds the energy cost savings to usage utility difference. By contrast, a biased consumer might misperceive energy cost savings by amount b, purchasing good E if and only if $v - G - b > p$. Bias b could be positive. negative, or zero. As defined, $b > 0$ implies that a consumer undervalues energy cost savings and is thus less likely to buy good E, while $b < 0$ implies the opposite. In the language of Herrnstein, Loewenstein, Prelec, and Vaughan (1993), b is an "internality": an externality that decisionmakers impose on themselves by not correctly valuing all consequences of a decision.

Mistakes might arise through the following mechanisms:

Present bias. If consumers are present biased with β, δ preferences as in Laibson (1997), and

if they cannot save or borrow between periods, they downweight the future energy cost savings G by factor $\beta < 1$, purchasing good E if $v - \beta G > p$.

Bias toward concentration. In Koszegi and Szeidl (2013), consumers underweight future cash flows that, like energy bill savings, accrue in small amounts over many future dates. While many predictions differ from those of the present bias model, the implication is the same in this simple framework: consumers downweight energy cost savings by β.

Biased beliefs. Consumers' beliefs about energy costs could be biased by factor ϕ. For example, they could perceive that energy prices are $\widetilde{g} = \phi g$, or that the total energy cost difference is $\widetilde{G} = \phi G$. In this case, they purchase good E if $v - \phi G > p$.

Costly information acquisition. Consumers might learn G_E and G_I only by incurring some cost. Otherwise, they remain imperfectly informed. In some rational expectations models, they might assume that G_E and G_I both equal the average G, purchasing good E if $v > p$.

Exogenous inattention. Similar to Chetty, Looney, and Kroft (2009), share λ of the population could be attentive to energy costs G, while share $1 - \lambda$ are exogenously fully inattentive, purchasing E if $v > p$.

Endogenous inattention. In rational inattention models such as Gabaix (2014) and Sallee (2014), consumers are more likely to pay attention to an attribute if it is more likely to matter in their purchase decision. A simple way to capture this is to allow $\lambda = \lambda(G)$, with $\frac{\partial \lambda}{\partial G} > 0$.

With present bias and bias toward concentration, $b = (1 - \beta)G$. In the biased beliefs model, $b = (1 - \phi)G$. Imperfectly informed or inattentive consumers have $b = G$.

3.3   Policy Implications of Mistakes

If factors such as the six above cause mistakes, the first-best policy would be some "nudge" that eliminates the distortion. For example, if consumers are imperfectly informed or have biased beliefs, one natural policy is to provide information. In practice, however, the ideal nudge may not be feasible, and the policymaker must resort to a price or quantity instrument. The logic of O'Donoghue and Rabin (2006) suggests that a tax or subsidy could increase welfare by offsetting the bias, even if the first best is not achieved.

To see this, consider a social planner that sets subsidy s for good E using public funds raised from lump-sum taxation, so $p = c - s$. Define the triple $\theta = \{v, G, b\}$, and assume that the joint distribution $f_\theta(\theta)$ is continuous with support such that there is always positive demand for both goods. For expositional simplicity, assume that there are no distortions other than consumer mistakes.[8]

---

[8]AT, AMT, Heutel (2015), and Tsvetanov and Segerson (2013, 2014) consider the case with both internalities and externalities. In these models, the optimal policy is a combination of product subsidies and energy taxes, although the energy tax will not necessarily equal the externality. One result of these models is that externality taxes can actually increase consumer welfare (even without including externality benefits) if they also reduce internalities; AMT call this the "internality dividend."

If all consumers have homogeneous bias $b^\dagger$, then subsidy $s = b^\dagger$ causes consumers to purchase good E if and only if $v - G - b^\dagger > c - s$. The s and $b^\dagger$ cancel, and the first best is achieved: consumers make the same purchase that they would make if they were rational and perfectly informed and had to pay marginal cost. Thus, the welfare-maximizing subsidy is $s^* = b^\dagger$: a subsidy that exactly offsets the bias.

When consumers have heterogeneous bias, AT show that the average marginal bias function $B(p)$ is a sufficient statistic for the optimal subsidy. Formally, $B(p)$ is the average b of all consumers that are marginal at price p: $B(p) = \frac{\int_\theta b \cdot M(\theta, p) \cdot f_\theta(\theta) d\theta}{\int_\theta M(\theta, p) \cdot f_\theta(\theta) d\theta}$, where $M(\theta, p) = 1(v - G - b = p)$ as an indicator function for whether a consumer of type θ is on the margin at price p. Relative to no subsidy, the welfare effect of a subsidy of amount s is:

$$\Delta W(s) = \int_{p=c-s}^{c} \int_\theta \underbrace{(v - G - c)}_{Allocative\ gain} \cdot \underbrace{M(\theta, p) \cdot f_\theta(\theta) d\theta}_{Number\ of\ consumers\ marginal\ at\ p} \quad dp \tag{2}$$

The term $(v - G - c)$ is the allocative gain when a consumer purchases good E instead of good I. It measures the difference between a consumer's true utility gain $v - G$ and the relative production cost c. If and only if the allocative gain is positive, then the social planner would like the consumer to purchase good E. Equation (2) can also be re-written as:

$$\Delta W(s) = \int_{p=c-s}^{c} \left(\underbrace{B(p)}_{Internality\ reduction} - \underbrace{s}_{Harberger\ distortion} \right) \cdot \underbrace{\int_\theta M(\theta, p) \cdot f_\theta(\theta) d\theta}_{Number\ of\ consumers\ marginal\ at\ p} \quad dp \tag{3}$$

AT use an analogous equation to highlight the two effects of a corrective subsidy. First, the subsidy increases allocative efficiency through internality reduction, measured by $B(p)$. Second, however, it distorts consumers' perceived optimal decisions by amount s, which would integrate to a standard Harberger deadweight loss triangle. If the average marginal bias $B(p)$ is larger than the marginal "Harberger distortion" s, a marginal increase in the subsidy increases welfare because on average, marginal consumers' true relative utility is larger than relative cost.

Taking the derivative of Equation (3) gives that the optimal subsidy s^* must equal the average marginal bias:

$$s^* = B(c - s^*) \tag{4}$$

AT, AMT, Baicker, Mullainathan, and Schwartzstein (2013), and Mullainathan, Schwartzstein, and Congdon (2012) derive this basic result that the optimal internality tax equals the average marginal internality. This directly parallels Diamond (1973), who shows that the optimal externality tax equals the average marginal externality.

Although this framework is highly stylized, it generates insights for a broad class of policies: at least in simple models, the most important policies listed in Table 1 are isomorphic to a subsidy s. Subsidies for Energy Star appliances and weatherization can be analyzed directly, although typically it would be necessary to extend the model to incorporate more than two products. With unit demand and lump-sum taxation or revenue recycling, a gas guzzler tax or analogous tax on good I is isomorphic to a subsidy, and a minimum energy efficiency standard is an infinite tax on good I. CAFE standards impose a shadow cost on vehicle sales that can be approximated as a subsidy for high-MPG vehicles.

Equations (3) and (4) emphasize the importance of empirical work that can identify $B(p)$. Some issues matter for welfare analysis, while others do not. First, $B(p)$ could naturally be heterogeneous at different p, and different types of consumers could be marginal to different policies. Thus, to set and evaluate subsidies and standards, it is insufficient to know just the population average bias. Second, $B(p)$ is a "reduced form" description of the behavioral models described above in Section 3.2. Once we know $B(p)$, the optimal subsidy and welfare evaluation does not depend on whether consumers are present biased, inattentive, or biased in some other way.

4   Empirical Tests

This section discusses three empirical tests: comparing demand responses to prices vs. energy costs, measuring effects of nudges, and belief elicitation. These tests are directly analogous to the three empirical studies in Chetty, Looney, and Kroft (2009) and are parallel to categories of tests described in Mullainathan, Schwartzstein, and Congdon (2012). The first two tests also parallel two of DellaVigna's (2009) proposed methods of testing for inattention, although as we shall see, they are joint tests of inattention and other behavioral models.

4.1   Comparing Demand Responses

A general test for systematically misperceived costs is to compare demand response to a correctly-perceived cost vs. a potentially misperceived cost. This is the approach followed by Abaluck and Gruber (2011), Barber, Odean, and Zheng (2005), Chetty, Looney, and Kroft (2009), Finkelstein (2009), Hossain and Morgan (2006), and others when studying sales taxes, shipping and handling charges, highway tolls, and other potentially shrouded costs.

The energy efficiency version of this test is to compare demand response to price vs. energy cost. Consumers should care only about the total lifetime cost of using a car, not the individual components of that cost. Thus, consumers should be indifferent between a dollar of upfront price and a present discounted dollar of energy costs. Many papers have implemented this test, including Dubin (1992), Dubin and McFadden (1984), Goldberg (1998), Grigolon, Reynaert, and Verboven (2014), Hausman (1979), and Houde and Spurlock (2015). A simple representation of the approach

is to imagine cross-sectional data on market shares and product attributes in a logit model. Define s_j, X_j, and ξ_j as the market share, observed characteristics, and unobserved characteristic of product j. In the representative consumer logit model, the market share equation is

$$ln(s_j) = -\eta(p_j + \gamma G_j) + \alpha X_j + \xi_j. \tag{5}$$

If the estimated $\hat{\gamma}$ is smaller (larger) than one, this implies that consumers undervalue (overvalue) one dollar of energy cost G relative to purchase price p.

A crucial part of the approach is estimating G. This requires expectations of future energy costs, utilization patterns, and discount rates.[9] Of course, estimating the "implied discount rate" that gives $\gamma = 1$ does not eliminate the need to take a stand on discount rates, as once an implied discount rate is estimated, it must be compared to some benchmark rate to determine if consumers are behaving according to some benchmark model. The need to make assumptions to calculate G makes it more complicated to implement this test with energy efficiency relative to sales taxes, shipping and handling fees, or other costs that are easily quantified and are paid at the time of purchase. On the other hand, many other product attributes - a car's horsepower, the sweetness of strawberry jam, etc. - cannot be monetized, and this test could not be implemented at all.

One immediate problem with the cross-sectional test is the endogeneity of both p and G. As Berry, Levinsohn, and Pakes (1995) and others have pointed out, prices are likely to be correlated with unobserved characteristics ξ. Furthermore, energy efficiency, which generates the cross-sectional variation in G_j, is also likely to be correlated with ξ. For example, an "economy" car both has high fuel economy and fewer amenities, and not all amenities are observed by the econometrician. Allcott and Wozny (2014) show that fuel economy is highly negatively correlated with price in a cross section of vehicles, suggesting the low-fuel economy vehicles have more observed and unobserved amenities.

To address these issues, Allcott and Wozny (2014, "AW"), Busse, Knittel, and Zettelmeyer (2013, "BKZ"), and Sallee, West and Fan (2015, "SWF") exploit variation in g and p from panel data on used auto markets. Allow jat to index a model j of age a at time t, for example a three-year-old Honda Civic DX in 2006. In used markets, s_{ja} is fixed over time by definition because vehicles of a given model year have already been produced.[10] Defining ξ_{ja} as a model fixed effect and including time indicators τ_t, Equation (5) can be re-arranged as:

---

[9] See Allcott (2011), Allcott and Wozny (2014), and Anderson, Kellogg, and Sallee (2013) for evidence on consumers' energy price predictions. Anderson, Kellogg, and Sallee (2013) use Michigan Survey of Consumers (MSC) data from 1993-2010 to show that on average, consumers believed that the future real price of gasoline would equal the current price. However, both oil futures markets and MSC consumers believed that the price shocks of 2008 were temporary. Allcott and Wozny (2014) show that used vehicle prices move as if consumers' forecasts are some combination of current and futures prices.

[10] AW exclude vehicles older than 15 years, as they reject that scrappage is exogenous to gasoline prices for these older vehicles.

$$p_{jat} = \gamma G_{jat} + \xi_{ja} + \tau_t + \epsilon_{jat}. \tag{6}$$

Figure 1 gives graphical intuition for the identification, using data from AW. It is constructed by grouping above-median and below-median fuel economy vehicles into composite groups, analogous to goods E and I. Each point on the graph is a month of the sample, from 1999 through the end of 2008. The y-axis plots the mean price difference $p_I - p_E = -p$, while the x-axis plots the mean gas cost difference $G_I - G_E = -G$. Notice that low-MPG cars are both more expensive and have higher G, so $-p$ and $-G$ are both positive. This graph plots raw data, before removing fixed effects or including other controls. The graph shows that the line slopes downward: as gas prices increase, the relative gas costs of good I increase, and relative prices drop in response. Thus, consumers are clearly highly responsive to gas prices. The more responsive are relative prices to gasoline costs, the more that we infer that consumers value fuel economy.

AW estimate a 15 percent implied discount rate, meaning that $\gamma = 1$ when they assume $\delta = \frac{1}{1.15}$. They calculate consumers' weighted average intertemporal opportunity cost of capital based on used car loan interest rates from the Survey of Consumer Finances (SCF) and returns to the S&P 500, and weight these two by the share of consumers that pay with loans vs. cash. This weighted average is 6 percent, which gives $\delta = \frac{1}{1+6\%}$. At this δ, the $\hat{\gamma}$ in AW's primary specification is 0.76. They also estimate $\hat{\gamma}$ for a range of alternative assumptions and specifications.

Using an analogous approach that compares vehicles in different quartiles of the MPG distribution, BKZ report a range of implied discount rates from -6.8 percent to 20.9 percent, depending on the assumptions used to calculate G and the quartiles being compared. When using assumptions that correspond most closely to AW's primary specification (vehicle-miles traveled and survival probabilities from the National Highway Transportation Safety Administration), the average implied discount rate for used vehicles is 13 percent. At $\delta = \frac{1}{1+6\%}$, this gives $\hat{\gamma} = 0.78$. However, BKZ's implied discount rates drop under alternative estimates of vehicle miles-traveled, and the auto loan interest rates in BKZ's data appear to be higher than in the SCF, so their benchmark δ is different. Thus, BKZ conclude that they find little evidence that auto consumers undervalue future gasoline costs.

SWF have a slightly different empirical specification: they include separate model-age fixed effects for every month in their sample, as well as flexible controls for each transacted vehicle's odometer reading. Autos with less past mileage have a longer future life, and so their transaction prices should be more responsive to changes in gas prices. SWF are thus testing a different hypothesis than AW and BKZ: whether consumers recognize that gas price changes should differentially affect *older vs. newer vehicles*, instead of *high-MPG vs. low-MPG vehicles*. Although the AW and BKZ question is one step closer to the policy question of how consumers value fuel economy, SWF can include a larger set of controls, and their test is plausibly a nested test: if consumers fully value how gas prices differentially affect vehicles of the same model and model year with different

odometer readings, it's likely that they fully value how gas prices differentially affect models with different MPG ratings. SWF's baseline model gives robust estimates consistent with full valuation, but they emphasize that different assumptions around the construction of G can shift the coefficient estimate in either direction.

As I discuss more in Section 6.2 below, while these three papers differ, their results are consistent in showing that consumers are highly responsive to gas prices, and they can rule out the magnitude of mistakes that would be required to justify the stringency of the current CAFE standard.

Houde (2014a) estimates a model of the refrigerator market that can also be categorized as "comparing demand responses." He first estimates an analogue to Equation (5) with model fixed effects, identifying the γ parameter using the fact that consumers in different counties and states pay different electricity prices. This variation consistently identifies γ as long as geographic variation in models stocked and consumer preferences is not correlated with electricity prices. The base specification gives an implied discount rate of 62 percent. Houde then posits a structural model in which some consumers are fully inattentive to electricity costs, some are fully attentive, and a third group forms expectations of electricity costs using only the Energy Star label. The model is identified by changes in the choice set over his 2007-2011 sample: models enter and exit, prices change, and some models lost their Energy Star certification due to the discovery of a testing error. The share of consumers that substitute equally to models regardless of electricity use represents the share of inattentive consumers, and the share that substitute equally to all Energy Star or non-Energy Star models regardless of electricity use is the share that attend to Energy Star. Consistent with his high implied discount rate, Houde finds that the modal consumer is inattentive, and many other consumers are attentive only to the Energy Star label and not the exact level of electricity use.

4.1.1   What Parameter Do the Used Vehicle Market Studies Identify?

In the choice models from Section 3, what does the $\hat{\gamma} = -\frac{dp}{dG}$ from Equation (6) represent? For expositional simplicity, assume for this section that v is the only source of heterogeneity across consumers. For the exogenous and endogenous inattention models, also assume that $b \perp v$, and take a first-order approximation to demand, so that the marginal distribution of v, denoted $f_v(v)$ with CDF $F_v(v)$, is locally uniform.

Consider first the case with informed and rational consumers. Because only v varies across consumers, there is some marginal consumer with $v = v^*$, and all consumers with $v < v^*$ purchase good I. Given that market share is constant in used vehicle markets, v^* must remain the same as p and G change. Mathematically, $v^* = p + G$. Taking the total derivative gives $dv^* = dp + dG$. Given that $dv^* = 0$, rearranging gives $-\frac{dp}{dG} = -1$. This says that relative prices and relative gasoline costs must offset each other one-for-one in the standard model in order to keep the same marginal consumer indifferent.

Similarly taking total derivatives, we can see that $-\frac{dp}{dG}$ equals β and ϕ in the present bias and biased beliefs models, respectively. In the exogenous inattention model, the market share of good I is $s_I = \lambda F_v(p + G) + (1 - \lambda)F_v(p)$. Taking the total derivative, $ds_I = \lambda\left(f_v(dp) + f_v(dG)\right) + (1 - \lambda)f_v(dp)$, which gives $-\frac{dp}{dG} = \lambda$. Thus, the approach of comparing demand responses identifies the average bias among consumers with a combination of present bias, biased beliefs, and exogenous inattention.

In the endogenous inattention model, however, taking the total derivative gives $-\frac{dp}{dG} = \lambda(G) + \lambda'\frac{G}{dG}$: relative price moves *more* than would be predicted in the exogenous inattention model. Intuitively, as gas prices increase, the relative prices of high-MPG vehicles increase both because they are more valuable to attentive consumers and because more consumers become attentive and thus are willing to pay for fuel economy. Thus, when attention (or information acquisition) is endogenous, the comparing demand responses identification strategy will overstate the share of attentive consumers λ at initial prices. Proposition 3 of AMT derives this result in a much more general setting, showing that under plausible assumptions, the comparing demand responses approach gives a lower bound for the average marginal bias.[11] This is an issue not just for AW, BKZ, SWF, and other energy papers, but for papers in other domains such as Chetty, Looney, and Kroft (2009) that use an analogous approach to identify the magnitude of inattention.

4.2 Estimating Effects of Nudges

Building on Thaler and Sunstein (2008), AT define a "pure nudge" as an intervention that eliminates bias but has no other effects. For example, providing information on energy costs can cause consumers who were previously uninformed or inattentive to make privately-optimal decisions. In theory, pure nudges can identify the average marginal bias function. To see this, consider the set of consumers who have baseline willingness-to-pay (WTP) of $v - G - b = p$ and are thus marginal at price p. Since the pure nudge causes consumers to make the same decisions they would make if they were informed and rational, their average WTP *after* being given a pure nudge is just their average $(v - G)$. Thus, $\tau(p)$, the average WTP change from a pure nudge for consumers marginal at price p, equals the average marginal bias:

$$\tau(p) = \frac{\int_\theta (v - G) \cdot M(\theta, p) \cdot f_\theta(\theta)d\theta}{\int_\theta M(\theta, p) \cdot f_\theta(\theta)d\theta} - (v - G - b) = B(p) \tag{7}$$

Carrera and Villas-Boas (2014), Chetty, Looney, and Kroft (2009), Choi, Laibson, and Madrian (2010), Kling *et al.* (2012), and others measure the effects of information in choices between generic and branded drugs, sales taxes, retirement savings, health insurance, and other domains.

---

[11]AT also show that even when each consumer's attention or bias is exogenous, comparing demand responses only approximates the average marginal bias $B(p)$ when b is heterogeneous across consumers and not independent of v. They show that the approximation can be highly inexact, which provides a further challenge to this approach.

DellaVigna (2009) points to Chetty, Looney, and Kroft (2009) as an example of how treatment effects from experimentally-provided information can be a measure of bias from inattention. Of course, nudges only identify the magnitudes of the specific biases they target: information provision can eliminate imperfect information and inattention, but it does not affect or identify present bias. Offering commitment contracts could identify present bias for sophisticates, but this does not affect or identify imperfect information.

In the context of energy efficiency, the nudge identification strategy might take the form of measuring effects of durable good energy cost information. Allcott and Taubinsky (2015) run two informational field experiments with lightbulb buyers. The AT study is motivated by the fact that only 28 percent of light sockets in American homes that could accommodate compact fluorescent lightbulbs (CFLs) actually had them as of 2010, despite the fact that a standard 60-watt-equivalent CFL saves \$40 in electricity and bulb replacement costs over its rated eight-year life relative to traditional incandescent bulbs. Is this low market share due to preferences, i.e. the distribution of v in the model above? Or is this due to imperfect information or inattention?

Only very specific experimental designs can plausibly use Equation (7) to identify $B(p)$. First, experimental or quasi-experimental designs are obviously needed to consistently identify a treatment effect τ. Second, one needs a design that allows the measurement of $\tau(p)$ at different p. To do this, one must observe consumers' WTP *before* being given a pure nudge to know which consumers are marginal at a given p and then measure the nudge's effect on WTP for these marginal consumers. Third, the treatment must be a pure nudge, without affecting consumers through persuasion, demand effects, or other mechanisms. This highlights perhaps the crucial drawback of this identification strategy: it is difficult to design interventions that plausibly approximate a pure nudge, and it is not clear how such an assumption could be decisively tested.

The AT experiment using Time-Sharing Experiments for the Social Sciences (TESS) attempts to satisfy these three criteria. A crucial feature is the within-subject design. All consumers in the experiment were given a \$10 shopping budget and asked to make choices between CFLs (good E) and incandescents (good I). Each consumer made baseline choices using a 15-part multiple price list. The treatment group was then given hard information about cost savings from CFLs, while the control group received information that was parallel in form but vacuous in content. Each consumer then made endline lightbulb choices on an identical 15-part multiple price list. Consumers' decisions were incentive-compatible: one of the 30 choices was randomly selected to be the consumer's "official purchase," and consumers were shipped the lightbulbs that they had chosen at that relative price and given the remainder of their \$10 shopping budget. This within-subject design identifies $\tau(p)$ by measuring the average treatment effect on WTP for consumers at each level of baseline WTP. In this context, WTP refers to incremental willingness-to-pay for the CFL relative to the incandescent.

The information treatment was designed to plausibly approximate a pure nudge. It provided

15

only hard information on lightbulb costs, without persuasive cues such as social comparisons or environmental framing. Care was taken to avoid demand effects, where experimental subjects might be more or less likely to purchase energy efficient goods not due to new information but because they wish to comply with (or perhaps defy) the perceived wishes of the experimenter. In addition, AT take steps to ensure and document that treated consumers understood, believed, and internalized the information. Despite this, they view their information treatments as only approximations to a pure nudge, and they carry out analyses under alternative assumptions.

Figure 2 gives the intuition for AT's main results. The graph presents a discrete version of Equation (3) above. The dashed line is the demand curve from the baseline multiple price list. Because the CFL and incandescent lightbulb packages were chosen to have the same unsubsidized retail price, the unsubsidized market equilibrium is at a relative price of $0. At this price, the CFL market share is about 70 percent. The shaded rectangle above the x-axis is the average treatment effect for consumers whose baseline relative WTP for the CFL is between $0 and $-1. By the assumptions for Equation (7), this treatment effect equals the average marginal bias $B(p)$. The triangle below the x-axis is the "Harberger distortion" from a $1 subsidy, in which the subsidy moves consumers away from their perceived private optima. A subsidy of $1 moves CFL market share to about 0.82, generating internality reduction equal to the rectangle but Harberger distortion equal to the triangle. Subtracting the triangle from the rectangle, it is clear that this $1 subsidy increases welfare. Incremental increases in the subsidy increase welfare until the incremental internality reduction above the x-axis is larger than the incremental Harberger trapezoid below the x-axis. Using this logic, AT find that the optimal subsidy to address imperfect information and inattention is $3.

In the model with unit demand and lump sum revenue recycling, a ban on good I is equivalent to subsidy $s = \infty$. Any incremental subsidy increases beyond $3 decrease welfare, and the ban on incandescent lightbulbs is so stringent as to fully reverse the possible welfare gains from the $3 subsidy. This result is driven by the fact that there are a large group of consumers who strongly prefer incandescent lightbulbs, even after being informed about the CFL's cost advantages.

This nudge-based identification strategy has several important advantages. First, the nudge approach provides a consistent estimate of the baseline level of bias even if the bias is endogenous or heterogeneous. As discussed above, this is not necessarily true for the comparing demand responses approach. Second, the nudge approach requires none of the assumptions on discount rates, energy price forecasts, and utilization that are required to calculate G in order to compare demand responses. Third, when paired with the multiple price lists and within-subject design of the lightbulb experiment, the nudge approach allows the analyst to infer the joint distribution of baseline demand and $B(p)$, which is required both for setting the globally-optimal subsidy and for evaluating the welfare effects of a ban. An important future research area would be to design a natural field experiment that has the conceptual advantages of the AT experiment without the

artefactual setting.

Because their nudge identifies only the effects of imperfect information and inattention, AT cannot rule out that other mistakes might justify a ban on incandescent lightbulbs. Indeed, there are two interesting results that suggest other potential mistakes. First, AT point out that the demand curve in Figure 2 is highly elastic around zero relative price. What's interesting about this is that G is typically very large and negative, because CFLs can save significant amounts of money on electricity costs. For $v - G$ to have significant mass around zero, the distribution of v must be closely symmetric to the distribution of G, i.e. also typically very large and negative. While many people dislike CFLs, the required symmetry of this disutility might not be plausible. One alternative explanation is that there is some alternative bias, such as bias toward concentration with β close to zero, such that the distribution of WTP largely reflects a distribution of v centered close to zero.

The second result comes from the fact that the AT survey participants are a nationally-representative sample with at least one person in every state other than Alaska. There is substantial variation in electricity prices across states: using data from the U.S. Energy Information Administration, the mean across states is 12.7 cents per kilowatt-hour, with a standard deviation of 4.3. Using Equation (1), I calculate $G_s = G_{Es} - G_{Is}$ for CFLs versus incandescents in each state s using a six percent discount rate over the eight-year rated life of a CFL; this averages \$39. If $v \perp G_s$, then the coefficient on G_s in a regression of WTP on G_s should be one. In other words, as long as people's non-monetary preferences for lightbulbs are not correlated with state electricity prices, people should be willing to pay \$1 more for a CFL if their state electricity prices imply an additional \$1 savings.

Table 2 presents results. Regardless of whether or not I control for additional covariates that might be correlated with v and electricity prices, the coefficient is statistically indistinguishable from zero and easily distinguishable from one. This is again consistent with the idea that there is some additional bias that would not be addressed by an information provision nudge, such as bias toward concentration with β close to zero.

4.3   Belief Elicitation

Starbucks customers tend to overestimate the calories in drinks and underestimate the calories in food (Bollinger, Leslie, and Sorensen 2011). People signing up for gyms are overconfident about their future attendance and about their likelihood of canceling automatically renewed memberships (DellaVigna and Malmendier 2006). More than 70 percent of seniors choosing between Medicare plans underestimate potential cost savings from switching (Kling *et al.* 2012). Spinnewijn (2015) shows that the unemployed overestimate how quickly they will find work. Could consumers have systematically biased beliefs about the financial benefits of energy efficiency?

There is some evidence of misperceived energy costs from the psychology literature. Attari *et*

17

al. (2010) show that consumers underestimate the energy use from heating, cooling, and other large energy uses relative to the energy use of a lightbulb. Larrick and Soll (2008) document MPG Illusion: people intuitively think as if automobile fuel costs scale linearly in miles per gallon (MPG), whereas they in fact scale linearly in gallons per mile. Turrentine and Kurani (2007) show that even well-educated auto owners with quantitative jobs have trouble with the net present value calculations required to estimate G_j.

While these results are interesting and important, this evidence that consumers misperceive energy costs does not directly test whether consumers systematically overestimate or underestimate the value of energy efficiency, i.e. whether $B(p)$ is non-zero. The Vehicle Ownership and Alternatives Survey (VOAS) (Allcott 2011b, 2013) is a nationally-representative survey on TESS platform, designed to test whether consumers systematically underestimate the fuel cost savings from higher-MPG vehicles. The VOAS asks consumers what vehicle they currently drive, how much they spend on gasoline, and how much they would spend on gasoline if they had bought their second-choice vehicle or another "replacement vehicle" with randomly-selected MPG difference. These questions are used to construct "valuation ratios": the share of true fuel cost differences that consumers perceive. This was denoted ϕ in the model in Section 3. Indexing the current vehicle as $j = o$ and an alternative vehicle as $j = a$, the valuation ratio for consumer i is

$$\phi_{ia} = \frac{\widetilde{G}_{ia} - \widetilde{G}_{io}}{G_{ia} - G_{io}}. \tag{8}$$

Consumers were told to assume that they drove the alternative vehicles the same amount as their current vehicle, which allows m to be assumed constant. To measure and limit potential confusion, various response frames were randomized across consumers. For example, some were told to report annual gasoline cost estimates, while others were told to report estimates over the full vehicle life. Some consumers were asked for absolute levels of gasoline costs, while others were told to report cost savings or additional costs relative to their current vehicle.

Figure 3 plots the distribution of valuation ratios between current vehicles and second choice vehicles. If all consumers were perfectly informed and reported correctly on the survey, the distribution would have point mass at $\phi = 1$. In fact, the figure shows that valuation ratios are quite dispersed, likely reflecting a combination of truly dispersed beliefs and reporting error. The large mass at zero represents consumers who incorrectly report that their current vehicle and second choice vehicle have "exactly the same" fuel economy. Allcott (2013) shows that on average, consumers correctly underestimate or perhaps slightly underestimate the financial benefits of higher-MPG vehicles.

There are two concerns with the belief elicitation approach to measuring bias. The first is that because the goal is to measure beliefs as they existed at the time of purchase, the survey must induce accurate recall without additional calculation. Making the belief elicitation fully incentive

compatible could induce survey respondents to look up fuel economy ratings and calculate answers on a calculator or spreadsheet. On the other hand, providing no incentives could lead to thoughtless answers. The VOAS offered moderate incentives (up to $10) with a vague criteria for payout: if an answer "makes sense" given answers to other questions. Allcott (2013) shows that results are indistinguishable between consumers randomly assigned to be offered incentives vs. not offered incentives.

The second concern is that stated beliefs have wide variance, and it is not clear whether this reflects true variation in beliefs or reporting error. The average marginal ϕ represents an estimate of $B(p)$, but the standard error on the estimated mean can be large. Allcott (2013) also estimates median regressions as an alternative measure of central tendency.

One benefit of belief elicitation is that it provides a direct measure of biased beliefs, whereas the comparing demand responses approach could estimate a combination of biases, and the informational nudge could estimate effects of biased beliefs, costly information acquisition, and inattention.

5   Testing Between Behavioral Models

Section 3 proposed several behavioral models, and the three empirical approaches discussed in Section 4 identify different subsets of these models. This section provides additional evidence for and against two of the proposed models.

5.1   Evidence of Endogenous Attention

Sallee (2014) describes reasons to believe that auto buyers might be rationally inattentive to gas costs. First, he argues that some cognitive cost is required to calculate the lifetime present discounted value of gas costs, both because fuel economy information labels can be incomplete and inaccurate, and because the labels provide average cost information that consumers must then adjust to their own energy costs and usage patterns. Second, he shows that in many cases, fuel economy may not be worth paying attention to. The variation in lifetime gas costs across vehicles is dwarfed by the variation in prices, and only about 10 percent of new vehicle buyers would change their decisions if they ignored vs. fully attended to within-class fuel economy differences, generating a welfare loss of only about $100 to $200 per purchase.

Below, I provide three pieces of empirical evidence in support of endogenous attention models. I use the phrase "endogenous attention" as shorthand - these tests measure both endogenous information acquisition and endogenous attention to that information. I use the word "endogenous" instead of "rational" because these tests do not necessarily imply that attention responds optimally.

In rational attention models such as Gabaix (2014) and Sallee (2014), consumers are more likely to pay a cost to attend to an attribute if it is more likely to affect their decisions. In the context of the model in Section 3, this implies that λ is increasing in energy price g and utilization m.

Furthermore, attention might vary across consumers due to variation in cognitive costs, not just gains from attention. These models guide the evidence below.

The first piece of evidence comes from the VOAS survey, originally reported in Allcott (2011b). The median completion time for this survey was ten minutes. The final question was: *In this survey, we asked you to calculate fuel costs fairly mathematically and precisely. Think back to the time when you were deciding whether to purchase your vehicle. At that time, how precisely did you calculate the potential fuel costs for your vehicle and other vehicles you could have bought?* There were five possible responses, ordered from "I did not think about fuel costs at all when making my decision" to "I calculated more precisely than I did just now during this survey." Forty percent of Americans reported that they did not think about fuel costs at all during their most recent purchase. I construct a measure of calculation effort by coding the five responses as integers from 1 to 5 and normalizing to mean zero, standard deviation one.

Table 3 reports regressions of calculation effort on potential explanatory factors available in the VOAS. "Gas Price at Purchase" is the US average retail gasoline price in the month that the vehicle was purchased from the US Energy Information Administration, inflated to 2014 dollars. Implied miles/year is backed out from self-reported gasoline expenditures using the current vehicle MPG rating and gasoline prices. Environmentalist is the consumer's response to the question, "Would you describe yourself as an environmentalist?" "Yes, definitely," "Yes, somewhat," and "No" are coded as 1, 1/2, and 0, respectively.

Column 1 shows that a $1 increase in gasoline prices is associated with a 0.148 standard deviation increase in calculation effort. Consumers that drive more miles per year also calculate more. These two results are consistent with a model in which consumers pay more attention to fuel economy when it has larger financial implications, either because they drive more miles or because gas prices are higher. Column 2 repeats the regression with the addition of the interaction between gas prices and miles per year. As the same model would predict, gas prices have an especially large association with calculation effort for consumers who drive more. Calculation effort is also positively associated with education, which is consistent with a model in which education is correlated with lower cognitive costs. Self-described environmentalists also are more attentive to fuel costs, a result that I return to in Section 6.1.

A second category of evidence for endogenous attention comes from the used vehicle markets studied by AW, BKZ, and SWF. As gas prices rose between 2003 and 2008, endogenous attention predicts that the estimated γ parameter should increase. Figure 1 from above shows graphical evidence of this using data from AW. From 1999-2003, gasoline prices were relatively low, and the slope of $\frac{dp}{dG}$ is relatively flat. From 2004-March 2008, gasoline prices were at a higher average, and the slope becomes steeper. From April-December 2008, as gas prices reached almost unprecedented highs, the slope of $\frac{dp}{dG}$ is close to -1. More formally, AW show that the estimated γ is larger when excluding 1999-2003 and smaller when excluding April-December 2008.

A third piece of evidence comes from direct observation of information acquisition: internet search volumes for terms related to fuel economy and page visits to the fueleconomy.gov informational website. This analysis parallels Hoopes, Reck, and Slemrod (2015), who show that Google searches for tax-related terms respond to stock market activity, policy changes, and news events.

For search volume data, I use Google Trends, which provides weekly data beginning in January 2004. I consider searches for five fuel economy-related phrases (gas mileage, fuel economy, fuel efficiency, miles per gallon, and mileage calculator) in the Vehicle Shopping category.[12] Search counts are normalized by total Vehicle Shopping category search volume in order to control for time variation both in the number of consumers currently in the auto market and in the use of web search. To corroborate the Google Trends data, I also acquired monthly counts of visits to fueleconomy.gov, the official U.S. government fuel economy information website, beginning in October 2002. A "visit" is as defined by Google Analytics: a series of one or more page views by a unique visitor which ends after 30 minutes of inactivity or when the browser is closed. There are an average of 2.3 million visits per month.

Panel A of Figure 4 plots weekly Google searches for fuel economy on the left axis and gas prices on the right axis. Search volumes track gas prices extremely closely. Panel B plots the monthly data on visits to fueleconomy.gov, which also closely track gas prices. Unlike the Google Trends data, fueleconomy.gov visits are not normalized by any measure of web traffic or vehicle consumers in the market, and there is a general upward trend over time. Both graphs show a sharp uptick in search volume in June-August 2009, as the Cash for Clunkers program made government subsidies conditional on fuel economy ratings.

To formally test the relationship, I regress search volume in week t (or website visits in month t) S_t on gas price as G_t, controlling for a cash-for-clunkers indicator for June-August 2009 C_t, linear time trend L_t, and week-of-year (or month-of-year) indicators μ_t:

$$\ln S_t = \eta \ln G_t + \psi C_t + \zeta L_t + \mu_t + \varepsilon_t. \tag{9}$$

Table 4 presents results, with Newey-West standard errors allowing for one year of lags (52 weeks or 12 months). Columns 1 and 3 show that information acquisition is highly statistically significantly correlated with gas prices in both the Google Trends and fueleconomy.gov data, with very similar elasticities of approximately 1.9. Column 2 shows that Google searches for four alternative vehicle attributes (horsepower, torque, warranty, and safety ratings) are not correlated with gas prices, suggesting that attention to gas costs may not draw attention away from other attributes, but instead may increase total decisionmaking effort.

---

[12]These search terms were chosen to measure information acquisition, not product demand. While one could also show that search terms like "hybrid vehicles" and "Toyota Prius" covary positively with gasoline prices, this effect could be driven by informed consumers searching for information about vehicles already known to have higher MPG, instead searches motivated by determining which vehicles have higher MPG.

5.2 Is Present Bias a Plausible Model?

Across many non-energy domains, there is strong evidence consistent with present bias and temptation models. Heutel (2015) and Tsvetanov and Segerson (2013, 2014) theoretically model the welfare effects of energy efficiency policy in present bias and temptation models, respectively. I suggested the possibility of a present bias model above, as do AT, AMT, Gillingham and Palmer (2014), and others. Is there empirical evidence consistent with such models?

There is one theoretical reason to question the relevance of present bias models for choices of energy-using durables. Taken literally, agents in these models are present biased over consumption, not cash flows. Because consumers in developed countries often can save and borrow to buy energy-using durables, paying additional money for an energy efficient good does not affect consumption in the present. Instead, it affects consumption over many future periods - perhaps the same future periods in which energy cost savings accrue. Similarly, the actual consumption value from energy using durables, by which I mean the enjoyment of driving a large SUV or a sporty car with lots of torque, also accrues over many future periods. Home buying is a clear example: because of mortgages, the purchase price, home energy costs, and consumption value of the home are all flows that occur in the future. Thus, while present bias might cause consumers to delay the effort of buying a home or car, it should not cause consumers to underweight energy costs relative to purchase prices. This suggests that any mistakes in this domain would be more likely driven by imperfect information, inattention, or bias toward concentration.

One might propose an alternative present bias model where agents are present biased over cash flows. Bradford, Courtemanche, Heutel, McAlvanah, and Ruhm (2014) estimate time-consistent discounting and present bias parameters using multiple price lists that trade of cash flows at different times. They find that the present bias parameter β is statistically associated in the predicted direction with driving a fuel-efficient car, having a well-insulated home, and the temperature setting on one's thermostat, along with other non-energy behaviors such as smoking, binge drinking, and savings. Interestingly, however, both Bradford *et al.* (2014) and AT find no correlation between the β from multiple price lists and demand for energy efficient lightbulbs.

Present bias and temptation models both predict that sophisticated agents will want to restrict their choice sets, perhaps by taking up commitment devices. Casual observation suggests that this is uncommon: many people seem to want to get their future selves to lose weight or quit smoking, but one rarely hears someone saying that he hopes to be able to get his future self to buy an Energy Star air conditioner. There are two ways to corroborate this observation with data. First, one highly-publicized commitment device is stickK.com, which allows people to individually tailor commitment contracts. stickK.com shared with me their data through mid-2014. Only 0.3 percent of their commitment contracts are categorized as "Green Initiatives," while 0.8 percent are "Home Improvement and DIY" and 3 percent are "Money & Finance." Energy efficiency contracts are likely to be a small subset of these three categories. By contrast, other categories are much

more popular, such as weight loss (35 percent), exercise (21 percent), other health and wellness (17 percent), career (8 percent), and education and knowledge (7 percent).

A second piece of evidence can be found in Harding and Hsiaw (2014), who study a program that invited electricity consumers to set energy conservation goals and also provided conservation information and usage feedback. In the Harding and Hsiaw (2014) model, failing to reach a goal imposes a utility cost due to reference-dependent preferences, and present-biased consumers thus view the program as a commitment device. While it is difficult to measure takeup rates conditional on being aware of the program, the authors present two statistics that suggest that these rates are low. First, while the program was marketed throughout the greater Chicago metropolitan area, only 2487 households enrolled. Second, when the authors took a random sample of 10,000 households in the zip codes where there was any enrollment, only 36 had actually enrolled.

In summary, commitment device takeup rates from stickK.com and Harding and Hsiaw (2014) suggest that few consumers are sophisticated in a present bias or temptation model. Unless the ratio of sophisticates relative to naifs is somehow very large for energy conservation, this suggests that these models are not very relevant for energy-using durables. While this does not say that low-cost commitment devices should not be offered, it does imply that present bias and temptation are unlikely to justify energy efficiency subsidies and standards.

6   Policy Implications

6.1   Targeting

Mistakes, when they occur, are almost certainly heterogeneous. For example, Figure 3 suggests that some consumers overestimate the private value of energy efficiency, even as other consumers underestimate. The optimal subsidy to address heterogeneous bias would be consumer-specific subsidies tailored to each consumer's bias b. Given that consumer-specific subsidies are not practically feasible, Equation (4) shows that the optimal uniform subsidy s is a compromise between (i.e., the average of) the optimal consumer-specific subsidies for the marginal consumers. However, this uniform subsidy distorts decisions of less biased consumers even as it improves decisions by more biased types.

A "well-targeted" policy approximates a pure nudge by preferentially affecting decisions by biased types without affecting the already-optimal decisions of rational types. A poorly-targeted policy, by contrast, preferentially affects decisions by less-biased types. At best, poorly targeted subsidies are disadvantageous because they might address only a small share of the total allocative inefficiency caused by consumer mistakes. At worst, a poorly-targeted subsidy could reduce welfare if $s > 0$ but the marginal consumers have $B(p) < 0$. Bernheim and Rangel (2004) give an intuitive example: if addicts are more biased but are highly inelastic to sin taxes, then a sin tax addresses little of the welfare losses from addiction. Furthermore, if rational types are relatively elastic, then

a sin tax of any non-trivial magnitude might reduce welfare by inducing many rational types to underconsume the sin good.

Allcott, Knittel, and Taubinsky (2015, "AKT") provide suggestive evidence on the targeting of energy efficiency subsidies. They propose two mechanisms that might cause some energy efficiency efficiency subsidies to be poorly targeted. First, many consumers are unaware of subsidies such as the weatherization incentives offered by local utilities, and the consumers who are most aware of and attentive to energy costs should also be more likely to be aware of energy efficiency subsidies. Because consumers won't respond to a subsidy if they are unaware of it, this mechanically makes the marginal consumers better informed and more attentive. Second, some energy efficiency investments are niche goods, such as hybrid cars and weatherization, with small market shares and only environmentalists near the margin. Thus, environmentalists are more likely to be marginal to anything other than a large subsidy, and environmentalists might have less bias against energy efficiency (i.e. lower b).

AKT first confirm that environmentalism is correlated with lower b. Environmentalist consumers believe that energy efficient lightbulbs and water heaters save more money, and as shown in Table 3, they pay more attention to fuel costs when buying autos.[13] AKT also show that environmentalists are more likely to be aware of energy efficiency subsidies offered in their area.

AKT then test what observable characteristics are correlated with takeup of major energy efficiency subsidies: the federal Residential Energy Credit, hybrid vehicle subsidies, and energy efficiency subsidies from a major utility. They find that environmentalism (as measured by self reports, owning a home solar energy system, and voluntarily paying for green energy) and the calculation effort variable from Table 3 are correlated with takeup, which strongly suggests that the average adopter is less biased against energy efficiency (i.e. has lower b) than the average consumer. They also show that wealthy people and homeowners are more likely to take up subsidies, which suggests that these policies are also not well-targeted to address credit constraints and landlord-tenant information asymmetries. Using much more extensive tax microdata, Borenstein and Davis (2015) fully develop this result that high-income people are more likely to take up the Residential Energy Credit.

These results suggest that some energy efficiency subsidies may not target the market failures they were designed to target. Instead, they primarily pay wealthy environmentalists to become even more green. Notwithstanding, this empirical evidence is only suggestive. In particular, AKT's regressions show that the *average* (instead of marginal) subsidy adopters have characteristics *likely to be correlated* with lower b. It would be an interesting and important contribution to (1) convincingly identify biased consumers and (2) convincingly test whether consumers marginal to a policy

---

[13]The differences in beliefs are interesting *per se*, because energy cost savings are a factual issue, not an ideological viewpoint, and the lightbulb and water heater survey instruments were apolitical. This finding relates to earlier work by Kahn (2007), Costa and Kahn (2013), and Gromet, Kunreuther, and Larrick (2013) showing political differences in energy efficiency behaviors.

are more or less biased.[14]

6.2 Calibrating Magnitudes

Sections 3 and 4 walked through the process of defining a socially-optimal policy and calibrating it using data. By contrast, the processes that most utilities and government agencies undertake when setting subsidies for weatherization or energy efficient goods bear little resemblance to this theoretically-grounded process. Instead, the subsidies are often set to exhaust the available budget or to maximize effects per subsidy dollar spent. As a result, the subsidies in effect are almost certain to be inefficiently large or small.[15]

As a stark example of the importance of calibration, imagine that $f(v)$ is uniform (so demand is linear) and $B(p) = b^{\dagger}$ at all p. In the model from Section 3, the optimal subsidy is $s = b^{\dagger}$, but any subsidy $s' > 2b^{\dagger}$ distorts prices so much that s' is worse than no subsidy at all. In other words, even if the policymaker correctly infers the sign of consumer bias, a subsidy can easily be worse than no policy if the subsidy is not optimally calibrated. This highlights the importance of not only measuring market failures, but also ensuring that policymakers correctly use the estimates.

A large set of papers (AMT, Fischer, Harrington, and Parry (2007), Heutel (2011), and Parry, Evans, and Oates (2010)) use different models to show that CAFE standards are much more stringent than can be justified by the largest estimates of consumer bias from multiple papers (AW, BKZ, SWF, and Allcott (2013)). This poses a significant challenge to CAFE standards: the RIAs have relied heavily (although not exclusively) on consumer mistakes to justify the policy, so they will likely have to come up with some alternative explanation for why the engineering models predict such large consumer welfare gains from regulation. One potential explanation is that the engineering models understate the development and manufacturing costs and/or vehicle performance tradeoffs of fuel economy improvements.[16]

---

[14]This point about optimal targeting is different than the point made by Boomhower and Davis (2014), Joskow and Marron (1992), and Ito (2015) about additionality. These papers argue that a policy's cost effectiveness (subsidy disbursed per kilowatt-hour saved) is worse if the subsidy has many inframarginal consumers. Thus, from a cost-effectiveness perspective, it is better if marketing or program design can target a subsidy at consumers who are more likely to be marginal. By contrast, in the simple framework from Section 3, inframarginal consumers do not matter for welfare because subsidy funds are raised through non-distortionary lump-sum taxes.

[15]By contrast, the Department of Energy sets energy efficiency standards so as to equate the marginal cost of more energy efficient (lower e_j) goods with the marginal reduction in G_j. While this is non-trivial to implement in practice, at least the process is theoretically optimal in a simple model.

[16]One scenario that would be useful to see in future RIAs is to adjust the fuel economy cost curves upward until the models correctly predict the current market equilibrium. This explicitly loads the discrepancy between model predictions and market equilibrium onto modeling error instead of an unknown market failure. Then the adjusted engineering model would be used to predict the socially-optimal CAFE standard to address uninternalized externalities. The optimal CAFE standard under these assumptions would be much less stringent.

6.3 Information Provision

As discussed in Section 3, the US has several major information provision programs: "yellow tags" for appliances, MPG labels for autos, and the Energy Star marketing campaign. There are now a number of papers on energy cost information provision for durable goods, largely in artefactual settings with stated preference, but some in field settings. What are we learning from a policy perspective?

First, while it is hard to object to the basic idea information provision, existing policies can be improved. Regulatory implementation costs and firms' compliance costs are low relative to subsidies and standards, so as long as some consumers benefit, it is plausible that benefits outweigh costs. However, Davis and Metcalf (2015), point out one way that energy information labels could be improved: because electricity costs vary substantially across states, presenting national average costs could mislead many consumers. In an stated choice experiment, they find that better information leads to better choices: consumers in high (low) energy cost states are more (less) likely to choose an energy efficient air conditioner with state-specific energy cost information relative to the national average. Larrick and Soll's (2008) finding of MPG Illusion is another example of how information can be presented more clearly and understandably, and fuel economy labels have since been redesigned to include information in gallons per mile.

A second lesson is that while some have argued that information provision programs are sufficient to address consumer mistakes, this may not be true. Gayer and Viscusi (2013) write,

> *How can it be that consumers are leaving billions of potential economic gains on the table by not buying the most energy-efficient cars, clothes dryers, air conditioners, and light bulbs?... If the savings are this great, why is it that a very basic informational approach cannot remedy this seemingly stunning example of completely irrational behavior? It should be quite simple to rectify decisions that are this flawed.*

Similarly, Mannix and Dudley (2015) argue that energy efficiency policies should be designed to address only what they call "classical" market failures. They then write,

> *Information asymmetry is an example of such a market failure, and we have long had fuel economy and energy-efficiency labels on cars and appliances to remedy it. This is a reasonable policy "nudge," and ought to be sufficient to the task.*

In other words, Gayer and Viscusi (2013) and Mannix and Dudley (2015) argue first that imperfect information is the primary (or only) consumer mistake that merits regulatory intervention, and second that existing information provision regulations should completely address this. These arguments are hard to square with the results of Houde (2014a), who estimates that large shares of refrigerator buyers do not value electricity costs despite having both yellow tags and Energy Star labels. Houde's evidence implies that either some non-informational mistake causes these consumers

to undervalue electricity costs or the existing information provision regulations do not completely address imperfect information. This second possibility is consistent with casual observation - in some appliance showrooms, the yellow tags are hidden inside washing machines or not posted, and when buying online, energy cost information often requires extra clicks that many consumers might not make.

A third potential lesson is that regulators should think carefully about the distinction between "information provision" and "persuasion." One view is that regulators should only be involved in the former. But the psychology literature provides many examples of how people respond differently to subtly-different information that a neoclassical model would classify as "hard information." For example, Camilleri and Larrick (2014) point out that presenting fuel economy information in gallons per 100,000 miles increases preferences for high-MPG vehicles relative to providing information in gallons per 100 miles. They suggest that policymakers should use the metric that induces more demand for high-MPG vehicles, but it is not clear that this increases allocative efficiency.

Indeed, two recent papers find that some energy information labels cause consumers to overvalue energy efficiency. Houde (2014a) estimates that the implied value that refrigerator buyers place on the Energy Star label is 15 to 83 percent larger than the true average electricity cost savings from Energy Star models. Newell and Siikamaki (2014) find that showing the Energy Star label causes 50 percent overvaluation of energy savings for water heater purchases in a stated choice experiment, when discounting future savings at individual-specific discount rates. They also find that a letter grade similar to the European Union's energy cost labels causes similar or greater overvaluation. These findings reinforce that policymakers should carefully consider whether current and future information labels are and should be used for marketing and persuasion instead of to help consumers minimize monetary costs. Glaeser (2006,2014) expresses concern about using persuasion as a policy tool, pointing out that many nudges are in fact "emotional taxes" that decrease consumer utility but do not raise revenues.

7   Open Research Questions

Many research questions remain. In this section, I sketch five of them.

First, what discount factor δ should be used in the "comparing demand responses" empirical tests? Many empirical papers, such as AW, BKZ, and SWF, use market average discount rates inferred from different estimates of the opportunity costs of capital. These approaches ignore consumer heterogeneity and may not reflect the correct opportunity cost of the marginal dollar, given that many consumers have savings, car loans, and credit card debt with widely-varying interest rates. Newell and Siikamaki (2015) make progress on this issue by eliciting preferences over hypothetical cash payments of $1000 in one month vs. larger amounts in in 12 months. They show that the elicited discount rates vary significantly across consumers and strongly predict

stated preference for energy efficiency. Ideally, this approach could be extended to be incentive compatible, perhaps by paying respondents with some small probability, and to cover a longer time horizon relevant for most energy-using durables.[17]

Second, if consumers misoptimize, how does this affect firms' incentives to develop, offer, and price energy efficient goods? Gabaix and Laibson (2006) study firms' incentives to debias consumers about add-on costs like energy use, but a key feature of that model is that the firm selling the base good also sells the add-on at some profit. By contrast, retailers of durable goods do not sell energy, so depending on markups, they may have stronger incentives to debias compared to the Gabaix and Laibson (2006) "curse of debiasing." There is a growing literature on firms' incentives in marketing energy efficiency. Allcott and Sweeney (2015) use a field experiment to study one retailer's ability to inform consumers about energy efficiency. Houde (2014b) shows how firms bunch product characteristics around the minimum Energy Star eligibility cutoffs, and Houde and Spurlock show how energy efficiency standards can increase welfare by increasing competition (and thus decreasing markups) among products that comply with the standard. Fischer (2010) studies how internalities interact with market power in firms' decisions to provide higher-fuel economy vehicles and how this affects the optimal determination of minimum standards and gasoline taxes.

Third, could bias toward concentration be relevant for some energy efficiency decisions? Because the benefits of energy efficiency come in small flows every time an electricity bill is paid or the gas tank is filled, this is a natural application of the Koszegi and Szeidl (2013) model. Even a lab experiment framed in an energy use context would be interesting if it measured and distinguished between bias toward concentration from other potential mistakes.

Fourth, many utilities have "on-bill financing" programs that essentially offer a loan to consumers for energy efficiency upgrades, which they pay back monthly on their energy bills. In theory, such a program could address present bias, bias toward concentration, and credit constraints. How effective are these programs, and why?

Fifth, in an endogenous inattention model, how do consumers form initial beliefs about the importance of energy efficiency when deciding whether to attend to this attribute? In Sallee (2014), for example, consumers have beliefs about the variance of energy costs across products, and in Gabaix (2014), consumers must have analogous initial beliefs about how much any attribute matters. Conlisk (1996), Gabaix, Laibson, Moloche, and Weinberg (2006), Lipman (1991), and others provide insight from other contexts.

---

[17]Laibson, Repetto, and Tobacman (2007) estimate that consumers use much higher discount rates over shorter horizons, which also poses a challenge to the idea of using one time-consistent δ.

8   Conclusion

As documented in Section 2.3, paternalism is an important factor used to justify energy efficiency policies. In recent years, models like the one in Section 3 have formalized this rationale, and three categories of empirical tests from Section 4 have been used to estimate consumer bias. Although the results are far from ironclad or comprehensive, there is some evidence that some consumers make systematic mistakes when purchasing some energy-using durables. In cases such as the vehicle market, however, the weight of the evidence suggests that if there are any systematic mistakes, they are small. In the absence of other market failures, the estimated consumer bias cannot justify the stringency of the current CAFE standards, and other energy efficiency subsidies and standards may also be miscalibrated. Furthermore, some existing policies do not appear to effectively target the market failures that they were designed to address. There is much work still to be done, and as rigorous evidence grows, it will be crucial to bring these results to policymakers.

References

[1] ABC (2011). Rand Paul's Toilet Tirade. http://abcnews.go.com/blogs/politics/2011/03/rand-pauls-toilet-tirade/

[2] Allcott, Hunt. 2011a. Social Norms and Energy Conservation. *Journal of Public Economics*. 95(9-10):1082-1095.

[3] Allcott, Hunt. 2011b. Consumers' Perceptions and Misperceptions of Energy Costs. *American Economic Review*. 101(3):98-104.

[4] 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.

[5] Allcott, Hunt and Michael Greenstone. 2012. Is There an Energy Efficiency Gap? *Journal of Economic Perspectives*. 26(1):3-28.

[6] Allcott, Hunt and Michael Greenstone. 2015. Measuring the Welfare Effects of Energy Efficiency Programs. Working Paper, NYU.

[7] Allcott, Hunt, and Judd Kessler. 2015. The Welfare Effects of Nudges. Working Paper, New York University (August).

[8] Allcott, Hunt, Christopher Knittel, and Dmitry Taubinsky. 2015. Tagging and Targeting of Energy Efficiency Subsidies. *American Economic Review*. 105(5):187-191.

[9] Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. 2014. Energy Policy with Externalities and Internalities. *Journal of Public Economics*. 112:72-88.

[10] Allcott, Hunt, and Todd Rogers. 2014. The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation. *American Economic Review*. 104(10):3003-3037.

[11] Allcott, Hunt, and Cass Sunstein. 2015. Regulating Internalities. *Journal of Policy Analysis and Management*. 34(2): 698-705.

[12] Allcott, Hunt, and Richard Sweeney. 2015. Information Disclosure through Agents: Evidence from a Field Experiment. *Management Science*, forthcoming.

[13] Allcott, Hunt, and Dmitry Taubinsky. 2015. Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market. *American Economic Review*, 105(8):2501-2538.

[14] Allcott, Hunt, and Nathan Wozny. 2014. Gasoline Prices, Fuel Economy, and the Energy Paradox. *Review of Economics and Statistics*, 96(10):779-795.

[15] Anderson, Soren, Ryan Kellogg, and James Sallee. 2013. What Do Consumers Believe about Future Gasoline Prices? *Journal of Environmental Economics and Management*, 66(3):383-403.

[16] Asensio, Omar, and Magali A. Delmas. 2015. Nonprice Incentives and Energy Conservation. *Proceedings of the National Academy of Sciences*. 112(6):E510-E515.

[17] Attari, Shahzeen, Michael DeKay, Cliff Davidson, and Wandi Bruine de Bruin. 2010. Public Perceptions of Energy Consumption and Savings. *Proceedings of the National Academy of Sciences*. 107(37):16054-16059.

[18] Ayres, Ian, Sophie Raseman, and Alice Shih. 2013. Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage. *Journal of Law, Economics, and Organization*. 29(5):992-1022.

[19] Baicker, Kate, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. Behavioral Hazard in Health Insurance. *Quarterly Journal of Economics*, forthcoming.

[20] Bernheim, B. Douglas, and Antonio Rangel. 2004. Addiction and Cue-Triggered Decision Processes. *American Economic Review*. 90(5):1558-1590.

[21] Bernheim, B. Douglas, and Antonio Rangel. 2009. Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics. *Quarterly Journal of Economics*. 124(1):51-104.

[22] Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. Automobile Prices in Market Equilibrium. *Econometrica*. 63(4):841-890.

[23] Boomhower, Judson, and Lucas W. Davis. 2014. A credible approach for measuring inframarginal participation in energy efficiency programs. *Journal of Public Economics*. 113(C):67–79.

[24] Borenstein, Severin, and Lucas Davis. 2015. The Distributional Effects of U.S. Clean Energy Tax Credits. NBER Working Paper No. 21437.

[25] Bradford, David, Charles Courtemanche, Garth Heutel, Patrick McAlvanah, and Christopher Ruhm. 2014. Time Preferences and Consumer Behavior. NBER Working Paper No. 20320.

[26] Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer. 2013. Are Consumers Myopic? Evidence from New and Used Car Purchases. *American Economic Review*. 103(1):pages 220-256.

[27] Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin. 2003. Regulation for Conservatives: Behavioral Economics and the Case for Asymmetric Paternalism. *University of Pennsylvania Law Review*. 151:1211-1254.

[28] Camilleri, Adrian R., and Richard P. Larrick. 2014. Metric and Scale Design as Choice Architecture Tools. *Journal of Public Policy and Marketing*. 33(1):108-125.

[29] Chetty, Raj, Adam Looney, and Kory Kroft. 2009. Salience and Taxation: Theory and Evidence. *American Economic Review*. 99(4):1145-1177.

[30] Conlisk, John. 1996. Why Bounded Rationality? *Journal of Economic Literature*. 34(2):669-700.

[31] Costa, Dora, and Matthew Kahn. 2013. Energy Conservation "Nudges" and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment. *Journal of the European Economic Association*. 11(3):680-702.

[32] Davis, Lucas W. 2012. Evaluating the Slow Adoption of Energy Efficient Investments: Are Renters Less Likely to Have Energy-Efficient Appliances? In The Design and Implementation of U.S. Climate Policy, Don Fullerton and Catherine Wolfram, Eds. Chicago, IL: University of Chicago Press.

[33] Davis, Lucas W., and Gilbert E. Metcalf. 2014. Does Better Information Lead to Better Choices? Evidence from Energy-Efficiency Labels. NBER Working Paper No. 20720.

[34] DEHWA (Australian Department of the Environment, Water, Heritage, and the Arts). 2008. Regulatory Impact Statement: Proposal to Phase-Out Inefficient Incandescent Light Bulbs. http://www.energyrating.gov.au/wp-content/uploads/Energy_Rating_Documents/Library/Lighting/Incandescent_Lamps/200808-ris-phaseout.pdf

[35] DellaVigna, Stefano. 2009. Psychology and Economics: Evidence from the Field. *Journal of Economic Literature*. 47(2):315-372.

[36] Delmas, Magali A., and Neil Lessem. 2014. Saving Power to Conserve Your Reputation? The Effectiveness of Private versus Public Information. *Journal of Environmental Economics and Management*. 67:353:370.

[37] Dolan, Paul, and Robert Metcalfe. 2015. "Neighbors, Knowledge, and Nuggets: Two Natural Field Experiments on the Role of Incentives on Energy Conservation." Working Paper, University of Chicago (April). Available from http://media.wix.com/ugd/fe9abe_bd09ded18b584eaf9230678e24d7cd4e.pdf.

[38] Dubin, Jeffrey, and Daniel McFadden. 1984. An Econometric Analysis of Residential Electric Appliance Holdings and Consumption. *Econometrica*. 52(2):345-362.

[39] Feldstein, Martin, and Jeffrey Liebman. 2002. Social Security. In Auerbach, Alan, and Martin Feldstein, Eds., Handbook of Public Economics, Volume 4. Amsterdam: Elsevier.

[40] Finkelstein, Amy. 2009. E-ZTAX: Tax Salience and Tax Rates. *Quarterly Journal of Economics*. 124(3):969-1010.

[41] Fischer, Carolyn. 2005. On the Importance of the Supply Side in Demand-Side Management. *Energy Economics*. 27(1):165-180.

[42] Fischer, Carolyn. 2010. Imperfect Competition, Consumer Behavior, and the Provision of Fuel Efficiency in Light-Duty Vehicles. Resources for the Future Discussion Paper 10-60.

[43] Fischer, Carolyn, Winston Harrington, and Ian Parry. 2007. Do Market Failures Justify Tightening Corporate Average Fuel Economy (CAFE) Standards? *The Energy Journal*. 28(4):1-30.

[44] Gabaix, Xavier, and David Laibson. 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics*. 121(2):505-540.

[45] Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model. *American Economic Review*. 96(4):1043-1068.

[46] Gayer, Ted. 2011. A Better Approach to Environmental Regulation: Getting the Costs and Benefits Right. Hamilton Project Discussion Paper 2011-06.

[47] Gayer, Ted, and Kip Viscusi. 2013. Overriding Consumer Preferences with Energy Regulations. *Journal of Regulatory Economics*. 43:248-264.

[48] Gerarden, Todd D., Richard G. Newell, and Robert N. Stavins. 2015. Assessing the Energy Efficiency Gap. Resources for the Future Discussion Paper No. 15-07.

[49] Gillingham, Kenneth, Matthew Harding, and David Rapson. 2012. Split Incentives and Household Energy Consumption. *The Energy Journal*. 33(2):37-62.

[50] Gillingham, Kenneth, Richard G. Newell, and Karen Palmer. 2009. Energy Efficiency Economics and Policy. *Annual Review of Resource Economics*. 1(1):597–620.

[51] Gillingham, Kenneth, and Karen Palmer. 2014. Bridging the Energy Efficiency Gap: Policy Insights from Economic Theory and Empirical Evidence. *Review of Environmental Economics and Policy*. 8(1):18-38.

[52] Glaeser, Edward. 2006. Paternalism and Psychology. *University of Chicago Law Review*. 73:133-156.

[53] Glaeser, Edward. 2014. The Supply of Environmentalism: Psychological Interventions and Economics. *Review of Environmental Economics and Policy*. 8(2):208-229.

[54] Goldberg, Pinelopi. 1998. The Effects of the Corporate Average Fuel Economy Standards in the US. *Journal of Industrial Economics*. 46:1-33.

[55] Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2014. Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market. Working Paper, University of Leuven.

[56] Gromet, Dena M., Howard Kunreuther, and Richard P. Larrick. 2013. Political Ideology Affects Energy-Efficiency Attitudes and Choices. *Proceedings of the National Academy of Sciences*. 110(23):9314-9319.

[57] Grubb, Michael. 2014. Consumer Inattention and Bill-Shock Regulation. *Review of Economic Studies*, forthcoming.

[58] Grubb, Michael, and Matthew Osborne. 2015. Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock. *American Economic Review*. 105(1):234-271.

[59] Gruber, Jonathan, and Botond Koszegi. 2004. Tax Incidence when Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes. *Journal of Public Economics*. 88:1959-1987.

[60] Gruber, Jonathan, and Sendhil Mullainathan. 2005. Do Cigarette Taxes Make Smokers Happier? *B.E. Journal of Economic Analysis and Policy*. 5(1).

[61] Gul, Faruk, and Wolfgang Pesendorfer. 2007. Harmful Addiction. *Review of Economic Studies*. 74(1):147-172.

[62] Harding, Matthew, and Alice Hsiaw. 2014. Goal Setting and Energy Conservation. *Journal of Economic Behavior and Organization*. 107:209-227.

[63] Hausman, Jerry. 1979. Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables. *Bell Journal of Economics*. 10(1):33-54.

[64] Heutel, Garth. 2015. Optimal Policy Instruments for Externality-Producing Durable Goods under Present Bias. *Journal of Environmental Economics and Management*, forthcoming.

[65] Herrnstein, R. J., George Loewenstein, Drazel Prelec, and William Vaughan, Jr. 1993. Utility Maximization and Melioration: Internalities in Individual Choice. *Journal of Behavioral Decision Making*. 6:149-185.

[66] Hoopes, Jeffrey L., Daniel H. Reck, and Joel Slemrod. 2015. Taxpayer Search for Information: Implications for Rational Attention. *American Economic Journal: Economic Policy*. 7(3):177-208.

[67] 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.

[68] Houde, Sebastien. 2014a. How Consumers Respond to Environmental Certification and the Value of Energy Information. Working Paper, University of Maryland.

[69] Houde, Sebastien. 2014b. Bunching with the Stars: How Firms Respond to Environmental Certification. Working Paper, University of Maryland.

[70] Houde, Sebastien, and C. Anna Spurlock. 2015. Do Energy Efficiency Standards Improve Quality? Evidence from a Revealed Preference Approach. Working Paper, University of Maryland.

[71] Ito, Koichiro. 2015. Asymmetric Incentives in Subsidies: Evidence from a Large-Scale Electricity Rebate Program. *American Economic Journal: Economic Policy*. 7(3):209-37.

[72] Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2015. The Persistence of Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand. Working Paper, University of Chicago.

[73] Jessoe, Katrina and David Rapson. 2014. Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use. *American Economic Review*. 104(4):1417-1438.

[74] Joskow, Paul L., and Donald B. Marron. 1992. What Does a Negawatt Really Cost? Evidence from Utility Conservation Programs. *Energy Journal*. 13(4):41–74.

[75] Kahn, Matthew. 2007. Do Greens Drive Hummers or Hybrids? Environmental Ideology as a Determinant of Consumer Choice. *Journal of Environmental Economics and Management*. 54(2):129-145.

[76] Kling, Jeffrey, Sendhil Mullainathan, Eldar Shafir, Lee Vermeulen, and Marian Wrobel. 2012. Comparison Friction: Experimental Evidence from Medicare Drug Plans. *Quarterly Journal of Economics*. 127(1):199-235.

[77] Koszegi, Botond, and Adam Szeidl. 2013. A Model of Focusing in Economic Choice. *Quarterly Journal of Economics*. 128(1):53-107.

[78] Laibson, David. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*. 112(2):443-477.

[79] Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. Estimating Discount Functions with Consumption Choices over the Lifecycle. NBER Working Paper No. 13314.

[80] Lipman, Barton. 1991. How to Decide How to Decide How To ... : Modeling Limited Rationality. *Econometrica*. 59(4):1105-1125.

[81] Mannix, Brian F., and Susan E. Dudley. 2015. The Limits of Irrationality as a Rationale for Regulation. *Journal of Policy Analysis and Management*. 34(3):705-712.

[82] Mullainathan, Sendhil, Joshua Schwartzstein, and William J. Congdon. 2012. A Reduced-Form Approach to Behavioral Public Finance. *Annual Review of Economics*. 4:17.1-17.30.

[83] Myers, Erica. 2014. "Asymmetric Information in Residential Rental Markets: Implications for the Energy Efficiency Gap. Working Paper, UC Berkeley (January).

[84] Newell, Richard, and Juha Siikamaki. 2014. Nudging Energy Efficiency Behavior: The Role of Information Labels. *Journal of the Association of Environmental and Resource Economists*. 1(4):555-598.

[85] Newell, Richard, and Juha Siikamaki. 2015. Individual Time Preferences and Energy Efficiency. *American Economic Review*. 105(5):196-200.

[86] NHTSA (National Highway Traffic Safety Administration). 2010. Final Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2012-MY 2016 Passenger Cars and Light Trucks. Office of Regulatory Analysis and Evaluation, National Center for Statistics and Analysis.

[87] O'Donoghue, Ted, and Matthew Rabin. 2006. Optimal Sin Taxes. *Journal of Public Economics*. 90:1825-1849.

[88] Parry, Ian, David Evans, and Wallace Oates. 2010. Are Energy Efficiency Standards Justified? Resources for the Future Discussion Paper 10-59.

[89] Regulatory Analysis Review Group. 1980. Department of Energy's Proposed Energy Efficiency Standards for Consumer Appliances. http://cwps.mercatus.org/wp-content/uploads/161501.pdf

[90] Sallee, James. 2014. Rational Inattention and Energy Efficiency. *Journal of Law and Economics*. 57(3):781-820.

[91] Sallee, James, Sarah West, and Wei Fan. 2015. Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations. NBER Working Paper No. 21441.

[92] Spinnewijn, Johannes. 2015. Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs. *Journal of the European Economic Association*. 13(1):130-167.

[93] Sudarshan, Anant. 2014. Nudges in the Marketplace: Using Peer Comparisons and Incentives to Reduce Household Electricity Consumption. Working Paper. http://www.anantsudarshan.com/uploads/1/0/2/6/10267789/nudges_sudarshan_2014.pdf.

[94] Sunstein, Cass and Richard Thaler. 2003. Libertarian Paternalism is not an Oxymoron. *University of Chicago Law Review*. 70(4):1159-1202.

[95] Thaler, Richard, and Cass Sunstein. 2008. Nudge: Improving Decisions about Health, Wealth, and Happiness. New York: Penguin Books.

[96] Tsvetanov, Tsvetan, and Kathleen Segerson. 2013. Re-evaluating the Role of Energy Efficiency Standards: A Behavioral Economics Approach. *Journal of Environmental Economics and Management*. 66:347-363.

[97] Tsvetanov, Tsvetan, and Kathleen Segerson. 2013. The Welfare Effects of Energy Efficiency Standards When Choice Sets Matter. *Journal of the Association of Environmental and Resource Economists*. 1(1/2): 233-271.

[98] US GPO (Government Printing Office). 2007. Testimony on the Status of Energy Efficient Lighting Technologies. http://origin.www.gpo.gov/fdsys/pkg/CHRG-110shrg39385/html/CHRG-110shrg39385.htm

[99] Yoeli, Erez, Moshe Hoffman, David G. Rand, and Martin A. Nowak. 2015. Powering Up with Indirect Reciprocity in a Large-Scale Field Experiment. *Proceedings of the National Academy of Sciences*. 110(2):10424-10429.

Tables

Table 1: **Major US Energy Efficiency Policies**

Policy	Years	Magnitude
Standards		
Appliance efficiency standards	1990-	$2.9 billion annual cost
Building codes	1978-	
CAFE standards	1978-	$10 billion annual cost
Prices		
Federal Hybrid Vehicle Tax Credit	2006-2010	$426 million annual credit
Gas guzzler tax	1980-	$200 million annual revenues
Weatherization Assistance Program	1976-	$250 million annual cost
Demand-Side Management	1978-	$7.6 billion annual cost
2009 Economic Stimulus	2009-	$17 billion total
Information and Marketing		
Fuel economy labels	mid-1970s	
Appliance "yellow tags"	1980	
Energy Star program	1992	$50 million annual cost

Notes: Data in this table are primarily from Allcott and Greenstone (2012).

Table 2: **Correlation of WTP for a CFL with Electricity Cost Savings**

	(1)	(2)
CFL electricity cost savings	-0.031	-0.025
	(0.031)	(0.031)
Income ($ millions)		1.903
		(6.167)
Education (years)		-0.032
		(0.147)
Age		-0.012
		(0.019)
Male		0.707
		(0.537)
Liberal		-0.071
		(0.385)
Party		0.572
		(0.346)*
Environmentalist		1.168
		(0.847)
1(Homeowner)		0.262
		(0.695)
N	1,229	1,217
R^2	0.00	0.01

Notes: The dependent variable is incremental willingness-to-pay for a compact fluorescent lightbulb relative to an incandescent lightbulb, using data from Allcott and Taubinsky (2015). Liberal and Party are self-reported political variables, normalized to mean zero, standard deviation one, with more positive meaning more liberal and more Democratic. Environmentalist is the consumer's response to the question, "Would you describe yourself as an environmentalist?" "Yes, definitely," "Yes, somewhat," and "No" are coded as 1, 1/2, and 0, respectively. Robust standard errors in parentheses. *,**, ***: Statistically different from zero with 90, 95, and 99 percent confidence, respectively.

Table 3: **Correlates of Fuel Cost Calculation Effort**

	(1)	(2)
Gas price at purchase ($/gallon)	0.148	-0.801
	(0.046)***	(0.431)*
ln(Implied miles/year)	0.069	-0.205
	(0.034)**	(0.124)*
Environmentalist	0.254	0.251
	(0.115)**	(0.115)**
Income ($ millions)	0.201	0.216
	(0.859)	(0.860)
Education (years)	0.051	0.052
	(0.015)***	(0.015)***
Age	-0.006	-0.006
	(0.002)**	(0.002)***
1(Male)	0.211	0.204
	(0.070)***	(0.070)***
1(Rural)	-0.132	-0.137
	(0.082)	(0.083)*
Gas price x ln(Implied miles/year)		0.106
		(0.049)**
N	1,444	1,444

Notes: The dependent variable is a measure of fuel cost calculation effort from the Vehicle Ownership and Alternatives Survey, normalized to a standard deviation of one. Environmentalist is the consumer's response to the question, "Would you describe yourself as an environmentalist?" "Yes, definitely," "Yes, somewhat," and "No" are coded as 1, 1/2, and 0, respectively. Robust standard errors in parentheses. *, **, ***: Statistically different from zero with 90, 95, and 99 percent confidence, respectively.

Table 4: **Fuel Economy Information Acquisition as a Function of Gas Prices**

	(1)	(2)	(3)
	Google Trends	Alternative Attributes	Fueleconomy.gov Website
ln(Gas price)	1.925	-0.114	1.905
	(0.320)***	(0.116)	(0.409)***
1(Cash for Clunkers)	0.398	-0.099	1.615
	(0.077)***	(0.048)**	(0.114)***
Time (years)	-0.049	-0.038	0.147
	(0.014)***	(0.009)***	(0.023)***
Number of Obs.	565	565	144
Data Frequency:	Weekly	Weekly	Monthly

Notes: The dependent variable in column 1 is the natural log of relative popularity of five fuel economy-related search terms (gas mileage, fuel economy, miles per gallon, fuel efficiency, and mileage calculator) from Google Trends data. The dependent variable in column 2 is the natural log of relative popularity of four search terms for alternative vehicle attributes (horsepower, torque, warranty, and safety ratings). The dependent variable in column 3 is the natural log of the count of visits to the Fueleconomy.gov website. Columns 1 and 2 include 52 week-of-year indicators, while column 3 includes 12 month-of-year indicators. Newey-West standard errors, with 52 lags (columns 1 and 2) and 12 lags (column 3), in parentheses. *,**, ***: Statistically different from zero with 90, 95, and 99 percent confidence, respectively.
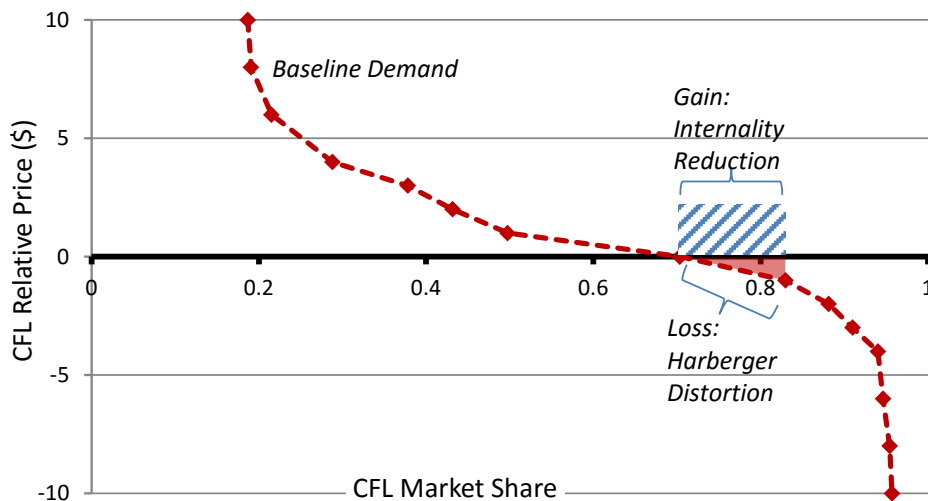
Figures

Figure 1: **Comparing Demand Responses: Graphical Results**
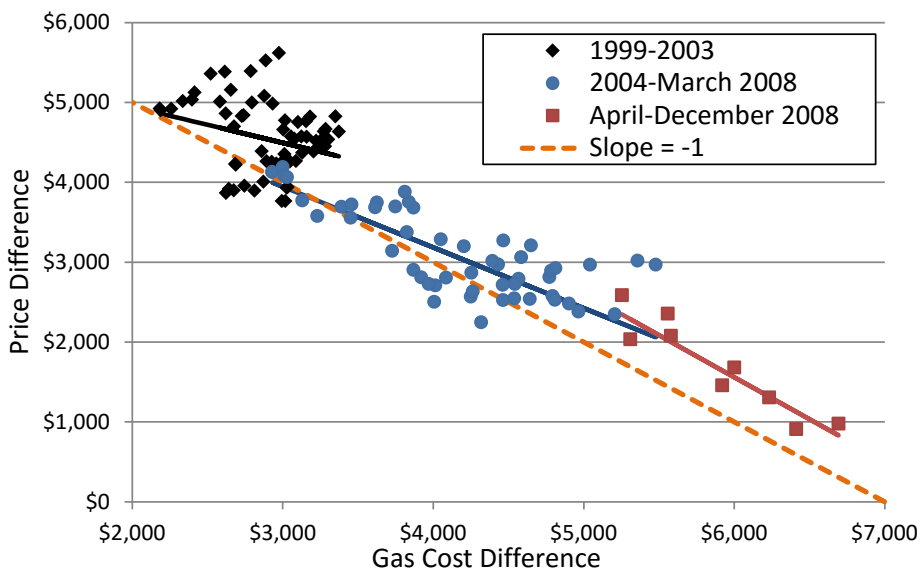
Notes: The vertical axis presents the difference in average prices between below-median MPG and above-median MPG used vehicles. The horizontal axis presents the difference in average present discounted value of gasoline costs. Each dot represents a month of the sample, from January 1999 through December 2008. Data are from Allcott and Wozny (2014).

Figure 2: **Welfare Effects of Compact Fluorescent Lightbulb Subsidies**

Notes: This figure illustrates welfare effects of a $1 CFL subsidy. Observations are weighted for national representativeness. Data are from Allcott and Taubinsky (2015).
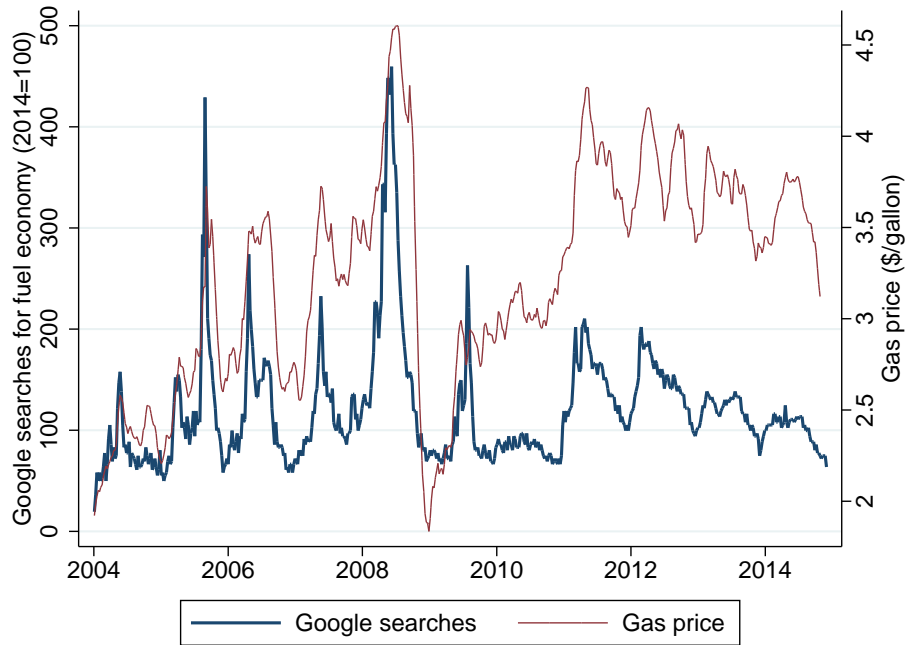
Figure 3: **Second Choice Vehicle Valuation Ratios**
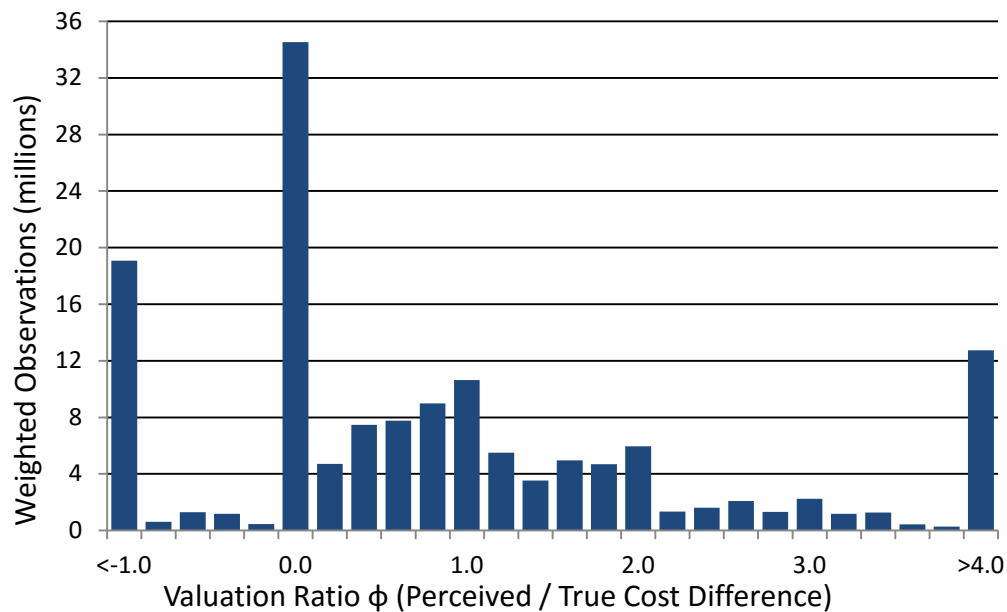
Notes: This figure shows the distribution of valuation ratios for second choice vehicles from the Vehicle Ownership and Alternatives Survey. Observations are weighted for national representativeness. Data are from Allcott (2013).

Figure 4: **Gasoline Prices and Search Volume**

(a) **Google Searches (Weekly Data)**

(b) **Fueleconomy.gov Visits (Monthly Data)**
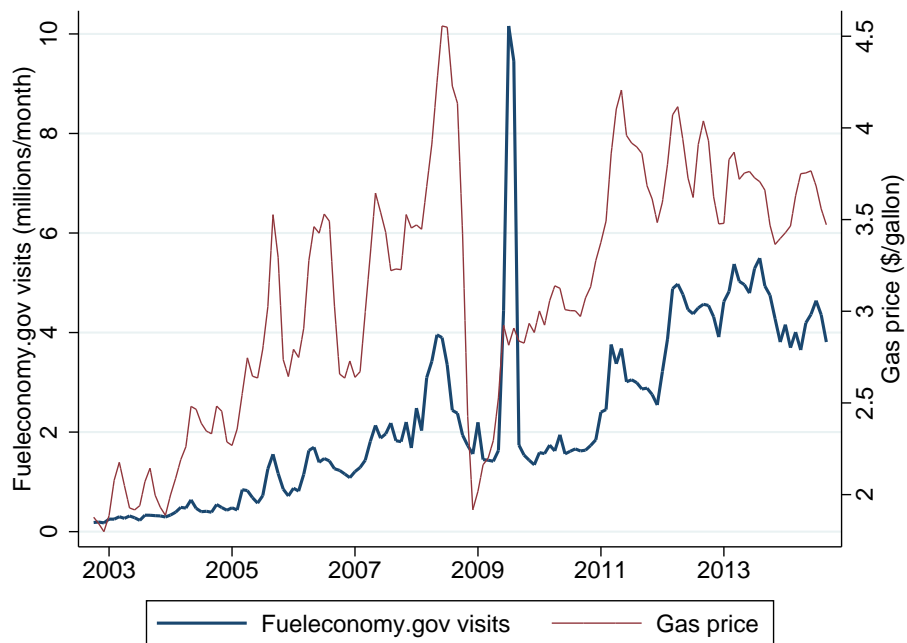
Notes: Google searches is the relative popularity of five fuel economy-related search terms (gas mileage, fuel economy, miles per gallon, fuel efficiency, and mileage calculator) from Google Trends data. Fueleconomy.gov visits is unique sessions as defined by Google Analytics. Gas price is the "U.S. All Grades All Formulations Retail Gasoline Price" from the U.S. Energy Information Administration.