

September 2012

Rational Inattention and Energy Efficiency

James M. Sallee



Rational Inattention and Energy Efficiency

James M. Sallee*

University of Chicago and the National Bureau of Economic Research

First Draft: June 3, 2011

This Draft: September 13, 2012

Abstract

If time and effort are required to ascertain the lifetime value of energy efficiency for a durable good, consumers might rationally ignore energy efficiency when making a purchase. This paper develops a heuristic model of a consumer's decision problem when purchasing an energy consuming durable good in which uncertainty about each good's energy efficiency can be resolved through costly search. The model shows that consumers will be less likely to undertake costly search when the variance of energy savings across models in a class is small, and when the variance of other attributes is large. This is important for the interpretation of the empirical literature because, if consumers are inattentive, an equilibrium could arise in which consumers make rational choices among the products on offer, but an energy paradox nevertheless exists because firms fail to bring innovations to market. The paper empirically characterizes the variance in lifetime fuel costs and prices for automobiles and several home appliances. The analysis suggests that, while variation in fuel costs is substantial, it is modest compared to the variation in prices for goods. The paper also argues that energy consumption labels are not sufficient to eliminate uncertainty. Together, these findings suggest that rational inattention may play an important role. Finally, the paper documents bunching around energy efficiency policy notches in automobiles, buildings, and home appliances. If consumers are inattentive, then providing salient signals, such as a binary energy efficiency certification, may be efficient, even if it induces supply distortions.

Keywords: energy paradox, limited rationality, rational inattention

JEL: Q48, H23, D03

*James M. Sallee, The Harris School, University of Chicago. Email: sallee@uchicago.edu; Web: <http://home.uchicago.edu/~sallee>. This paper was prepared for the 2011 University of Chicago / Resources for the Future Energy Policy Symposium. The author would like to thank the discussant, David Austin, as well as Hunt Allcott, Lucas Davis, Ashley Langer, and seminar participants at the Symposium and at the University of California at Berkeley for helpful comments. The author also would like to thank Greg Sasso and Samantha Superstine for excellent research assistance and Roberts French for providing Environmental Protection Agency five-cycle test data. ©2011 James M. Sallee. All rights reserved.

1 Introduction

An important, unresolved question in energy economics is whether consumers undervalue energy efficiency when purchasing durable goods. The hypothesis that consumers undervalue energy efficiency is sometimes called the “energy paradox.” Determining whether an energy paradox exists is important for policy because its existence could overturn the canonical preference for a Pigouvian tax on energy sources over regulatory approaches. This paper explores a novel explanation for an energy paradox, based on a model of bounded rationality. If the expected benefits from learning about energy efficiency are small relative to the costs of learning the relevant information, it may be rational for consumers to ignore energy efficiency when purchasing certain goods.

To explore this possibility, the paper develops a simple model of a consumer choosing among alternative energy-consuming durable goods. The model formalizes several intuitive points. If the variance across goods in the cost of energy is small, then it is more likely that consumers will ignore energy efficiency when making a choice because the expected benefit from learning is small. Conversely, if the variance in other attributes is large, then it is unlikely that energy cost will prove pivotal in a consumer’s choice, so ignoring energy efficiency when making a selection is rational.

The empirical portion of the paper characterizes the variance in expected energy costs and prices across models for automobiles and home appliances. These calculations monetize the expected benefit to a consumer from investigating energy efficiency benefits for different classes of goods. The paper does not directly estimate the cost of learning about energy efficiency. Rather, the aim of the empirical analysis is to describe the monetized value of search costs that would be required in order to make inattention rational.

The possibility of rational inattention has two key implications. First, if consumers rationally ignore energy efficiency, this could explain the energy paradox. In equilibrium, firms will underprovide energy efficiency if consumers ignore it. If true, this would qualitatively change the interpretation of empirical work on the energy paradox. Most empirical work

tests for the rationality of consumer choice across goods that are actually sold in the market. If rational inattention leads to an inefficient set of product offerings, consumers might choose rationally among goods in equilibrium, but a paradox still exists.

Second, if consumers are rationally inattentive to energy efficiency, this could provide direct justification for regulatory standards and “notched” policies, such as the Energy Star label system. Notched policies present coarse information to consumers, and this may create product design distortions. Notched policies are an inefficient method of raising energy efficiency because they fail to equate the marginal costs of improvement across sources. These inefficiencies could be justified if consumer inattention causes them to ignore fine grained information, whereas streamlined information, such as whether or not an appliance qualifies for the Energy Star label, may be taken into account because the cost of doing so is low. If consumers are rationally inattentive, the increased salience from the notched system might outweigh the product design distortions, creating a justification for such policies. This paper presents new estimates of distortions caused by energy-efficiency labeling systems in buildings and home appliances by examining the distribution of energy-efficiency ratings vis-à-vis certification notches.

This research relates to a large literature on the energy paradox.¹ Prior work has often mentioned imperfect information and cognitive costs as potential sources of the energy paradox (Jaffe and Stavins 1994; Howarth and Sanstad 1995), but relatively little work has sought to formalize a model of this or test its quantitative significance. Existing research that is more closely related to this paper includes Greene (2011), which estimates an energy gap arising from incomplete information and loss aversion in automobile fuel economy, but does not focus on rational inattention. Howarth and Andersson (1993) model the demand for energy efficiency when information acquisition requires a search cost. In contrast to this paper, their model does not develop a discrete choice model, nor does it allow heterogeneity

¹Much of the recent empirical work has focused on revealed preference tests of whether an energy paradox exists (Sallee, West, and Fan 2009; Allcott and Wozny 2010; Busse, Knittel, and Zettelmeyer Forthcoming), which builds on the seminal work of Hausman (1979) and Dubin and McFadden (1984).

in the preference for different models; nevertheless, the model of Howarth and Andersson (1993) arrives at similar intuition regarding the value of information acquisition as it relates to the difference in energy efficiency across alternatives.

An important distinction has been drawn by previous literature between “market barriers” and “market failures,” where only the latter justify market intervention (Metcalf 1994). Market barriers are factors that make optimal choice of energy efficiency difficult, but might be overcome by market forces. For example, Metcalf (1994) argues that informational uncertainty should give rise to market services that provide information, such as home energy audits. Such solutions are unlikely to overcome a market problem based on rational indifference, as the cost of providing such services is likely greater than the expected gains for many goods.

The remainder of the paper is structured as follows. As further motivation, section 2 shows evidence that notched policies create product design distortions. One potential justification for these distortions is rational inattention. Section 3 describes a heuristic model of a consumer choosing among alternatives in a class of durable goods, about which there is uncertainty regarding energy efficiency. Section 4 characterizes descriptive statistics about the automobile market that are motivated by the model. Section 5 does the same for household appliances. Both suggest that there is considerable scope for rational inattention, given the relative variance in prices, the barriers to ascertaining true energy efficiency, and the variance in energy costs across similar models. Section 6 concludes.

2 Notched policies

One reason to be interested in whether consumers are rationally inattentive to energy savings is that such inattention may provide a justification for the ubiquitous presence of discontinuities—sometimes called “notches”—in energy policies. Notches are points in a policy schedule at which marginal changes in behavior or attributes cause a discrete change

in policy treatment. Sallee and Slemrod (Forthcoming) argue that notches in policies aimed at correcting for externalities (such as subsidies for energy efficiency) are inefficient because they create capriciously varying incentives for the reduction in externalities.² The hallmark of a Pigouvian tax is that it gives all market participants the same marginal incentive to change behavior, which ensures that those with the lowest costs of action will act. In contrast, notches give agents close to a notch very large incentives for small changes, those far away lesser incentives, and those above the top notch no incentive at all. If notches have added salience, perhaps because of consumer inattention to energy efficiency in durable purchases, then this salience benefit could possibly outweigh inefficiencies caused by the unevenness of local incentives induced by notches.

The telltale sign of the distortion caused by local incentive variation from notches is a preponderance of products located just on the policy preferred side of notches. That is, products are “bunched” in the relevant attribute space around the critical points. To further motivate the study of rational inattention, this section provides evidence that there is product bunching around notches in several energy-efficiency policies and rating systems.

2.1 Automobiles

Policies aimed at improving energy efficiency in automobiles have historically been aimed at fuel economy, not the price of fuel, as economists would prefer. Fuel economy policies come in two main varieties: regulation and taxation. Corporate Average Fuel Economy standards are the main form of fuel economy regulation in the United States. Fuel economy taxation comes in the form of the Gas Guzzler Tax and tax credits for alternative fuel vehicles (like hybrids), which exist on both the federal and state levels.³

Historically, fuel economy taxation has always featured notches—the taxes or subsidies

²See Kleven and Slemrod (2009) for a related theoretical model of notch-driven product design distortions.

³In most cases, direct taxation of fuel is understood to be significantly more efficient at reducing externalities related to automobiles than fuel economy regulation (see for example Austin and Dinan (2005)). See Sallee (2011b) for a review of fuel economy taxation, and Anderson, Parry, Sallee, and Fischer (2011) for a review of fuel economy standards. See Sallee (2011a) and Gallagher and Muehlegger (2011) for details about hybrid vehicle tax credits.

are step functions. Thus, at certain values of fuel economy, marginal improvements in fuel economy lead to discrete changes in tax treatment. Sallee and Slemrod (Forthcoming) provide evidence—in the form of bunching in the distribution of fuel economy ratings—that automakers respond to these discrete incentives by making small modifications to fuel economy.

The Gas Guzzler Tax is a tax paid on all passenger cars rated below 22.5 miles per gallon (mpg), based on a US Environmental Protection Agency (EPA) test procedure. The amount of the tax depends on the vehicle’s fuel economy. It starts at \$1,000, and changes by several hundred dollars for every full mpg below 22.5, eventually reaching \$7,700 for vehicles that get below 12.5 mpg. Figure 1a, which modifies a figure from Sallee and Slemrod (Forthcoming), shows the distribution of all vehicles that paid the Gas Guzzler Tax from 1991 to 2009. The figure also shows, on a different vertical axis, the amount of the tax. This tax schedule is a step function, where steps occur at each 0.5 decimal place in mpg. If automakers respond to these discrete tax incentives, we would expect to find “extra” vehicles on the tax preferred (to the right) side of the notches. The data in Figure 1a, which are weighted by sales volumes, show exactly that pattern.

The response at the first notch (22.5 mpg), which may be especially important as it causes the vehicle to be labeled as a “guzzler”, is not visible when analyzing data from a list of vehicles that actually paid the tax, because vehicles just over that notch are not in the Internal Revenue Service (IRS) data. To examine the bunching around the top notch, Figure 1 shows the relevant fuel economy rating for all vehicles, using EPA fuel economy data from 1998 to 2007. Figure 1b shows the distribution of all cars, relative to the top notch at 22.5 mpg. The figure is suggestive of a significant amount of bunching just above 22.5 mpg—the distribution overall has a smooth bell shape except for the points just above and just below the Gas Guzzler Tax notch. For comparison, Figure 1c shows the same distribution for light trucks, which are exempt from the Gas Guzzler Tax and therefore have no incentive to bunch around the 22.5 value. Trucks do not show evidence of a discontinuity around the

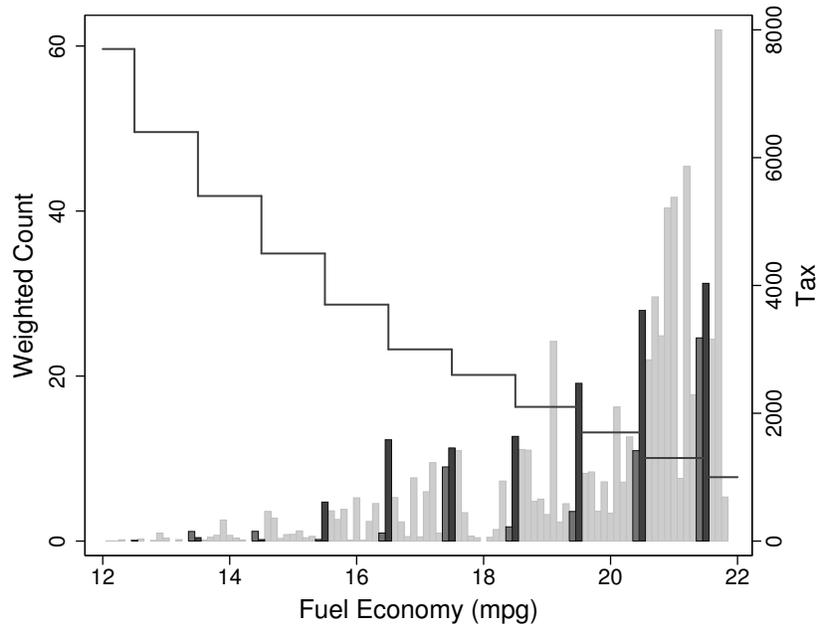
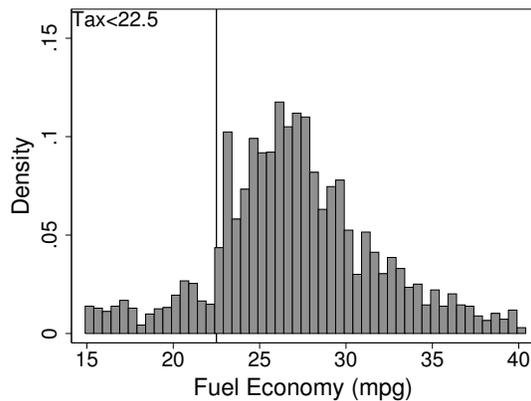
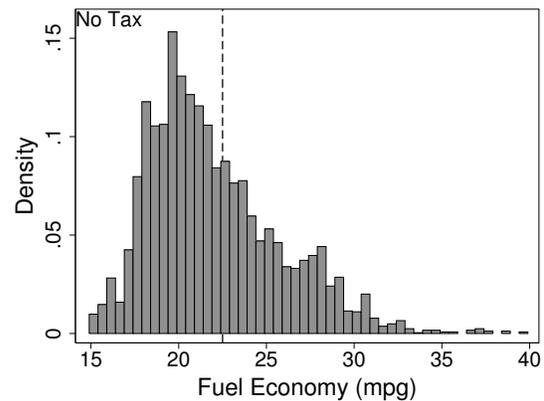
Figure 1: Gas Guzzler Tax notches**(a) Vehicles that paid the Gas Guzzler Tax****(b) Cars (taxed)****(c) Trucks (not taxed)**

Figure (a) shows a sales-weighted histogram of all vehicles that paid the Gas Guzzler Tax from 1991 to 2009, based on IRS data. The Gas Guzzler Tax schedule is shown as a line. Figure (a) is a modified version of a figure from Sallee and Slemrod (Forthcoming). Figures (b) and (c) show the unweighted distribution of Gas Guzzler Tax fuel economy ratings for all vehicles from 1998 to 2007, based on EPA data. The vertical line in (b) indicates the minimum fuel economy rating needed to escape the Gas Guzzler Tax completely. Figure (c) shows the distribution for light trucks, which are exempt from the Gas Guzzler Tax, for comparison.

Gas Guzzler Tax notch, which further bolsters the conclusion that the bunching in cars is due to the policy.

Sallee and Slemrod (Forthcoming) provide evidence of automaker responses to several other notches, including a similar program in Canada, fuel economy label ratings in the United States, and the notch in the classification of vehicles across light trucks and passenger cars. Overall, the evidence strongly suggests that automakers manipulate the fuel economy of their vehicles in response to policy notches, which implies an inefficiency in the allocation of fuel economy improvements across vehicles.

2.2 LEED certification

Another example of bunching in a rating system comes from the Leadership in Energy and Environmental Design (LEED) certification system for buildings. LEED is a non-governmental labeling program operated by the Green Building Certificate Institute (GBCI). The GBCI has separate rating systems for a number of project categories, including new construction, schools, core and shell, commercial interiors, existing buildings (operations and maintenance), and neighborhood development. Within each system, a building can achieve four levels of certification: certified, silver certified, gold certified, and platinum certified. To gain certification, projects must achieve a certain number of “points”. Each LEED system has a list of project characteristics that earn points, and a minimum number of points is required for each tier of certification. Examples range from reflective rooftops to bicycle storage rooms to reused building materials.

These certification tiers are notches. If builders are concerned about their status, they will be willing to exert considerable effort to gain an additional point if their project is close to a threshold, but much less willing if their project is far away. In terms of energy savings, however, the social external benefits are the same regardless of where the project lies vis-à-vis these notches. As such, if LEED certified projects appear to be bunched around the threshold values, it implies that strategic design plans are inefficient, because they put

different values on actions that have the same social external benefit.

The US Green Building Council (USGBC) maintains a list of all certified projects. To examine whether or not developers respond to the LEED notches, figure 2 plots the distribution of points and the certification notches for six of the most frequent LEED categories. The project list contains only projects that are certified, so it is not possible to observe the list of projects that fall short of the bottom tier. The first three are all new construction projects, with different LEED regimes.⁴ The other three include schools, commercial interiors and core and shell certification. In each diagram, the horizontal lines indicate the critical value necessary to move from certified status to silver, gold and platinum. The lines are drawn such that buildings with exactly that value qualify for the more advantageous label.

All six distributions show strong evidence of strategic bunching. There is a pronounced asymmetry around almost every threshold—buildings within two or three points of a threshold are far more likely to be on the advantageous side. The only exception is the platinum notch for schools. Note that the number of points varies across regimes. Thus, the bunching is around different cutoff values across different categories. This bunching implies inefficiency. To be justified, there must be some benefit, in terms of salience or administration, to having a system with discrete tiers as opposed to a smooth rating system that simply reports the number of LEED points earned by a project.

2.3 Energy Star certification

Most appliances in the US have minimum energy efficiency standards, which are regulated by the US Department of Energy (DOE). A subset of appliances (dishwashers, clothes washers, refrigerators and freezers, air conditioners, water heaters, furnaces and boilers, and ceiling fans) are required to carry Energy Guide labels, which provide information about estimated energy consumption from each good. The Energy Guide program is administered by the

⁴Within categories, the LEED systems have changed over time, from an original pilot phase to different regimes, named 1.0, 1.1, 2.0, etc. Figure 2 includes the three most recent regimes for new construction, which is by far the most common category.

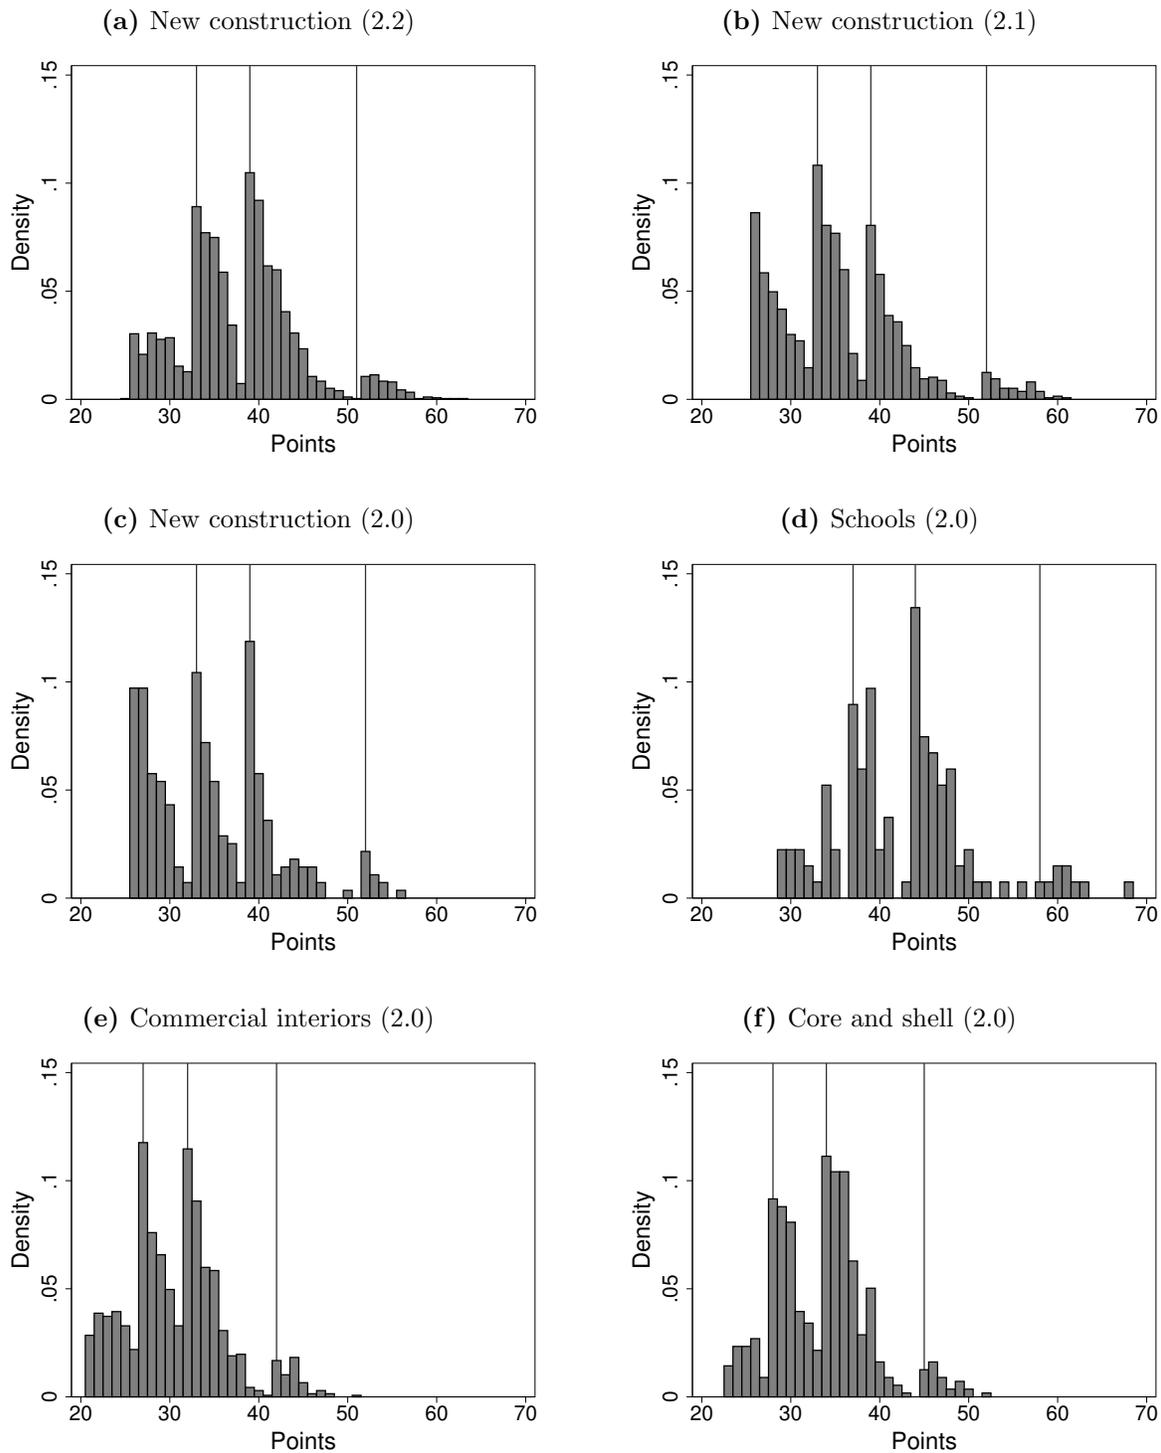
Figure 2: Distribution of LEED points

Figure plots histograms for various LEED certification systems from the USGBC registry. All buildings are completed and certified. Horizontal lines indicate the cutoffs for silver, gold, and platinum status, from left to right for each diagram. Sample sizes are (a) 2,738, (b) 1,371, (c) 282, (d) 134, (e) 172, and (f) 202.

Federal Trade Commission (FTC). Separately, the EPA administers the Energy Star program, which is a voluntary labeling program that certifies particularly efficient products and provides those that qualify with an Energy Star label. Energy Star certification covers a wide range of goods, including all of the goods labeled by the FTC, and many others, like televisions, computing and office equipment, light fixtures, air purifiers, and windows.

All Energy Star certification levels are notches (as are the mandatory efficiency standards). Appliances that meet some minimum threshold efficiency level are certified as Energy Star qualified and may carry an Energy Star label. If manufacturers believe that having a label increases demand for their product, then they will be willing to exert effort to improve the energy efficiency of models that fall short of the notch. The program does not, however, provide a similar incentive to models that already exceed the standard—there is no incentive for them to improve. Likewise, models that are “too far” from the threshold have no incentive for marginal improvements. As a result, the notched label policy provides an uneven incentive to producers—if only models that are near, but not over, the threshold are made to improve, then it is unlikely that the energy efficiency improvement is isolated to those models that have the lowest cost of improvement. Bunching in the distribution of products around the notch is an indicator of whether the uneven incentives create a distortion.

The characteristics required to earn an Energy Star label vary across the goods. Typically, the Energy Star label requires a certain percentage efficiency beyond the DOE regulatory minimum. The percentage required, however, changes over time. Often, the rating used is different from the rating that appears on FTC Energy Guide labels. For example, the Energy Star criteria for a clothes washer is that the item have a Modified Energy Factor (MEF) above 2.0 and a Water Factor (WF) below 6.0. The MEF is equal to the capacity of the clothes container divided by the energy used per load from running the machine, from heating the hot water, and from removing moisture after the cycle is complete. The WF is equal to the water used per cycle, divided by the machine’s capacity. In contrast, the FTC reports only a kilowatt-hour per year (kWh/year) measure in its labels, which does not

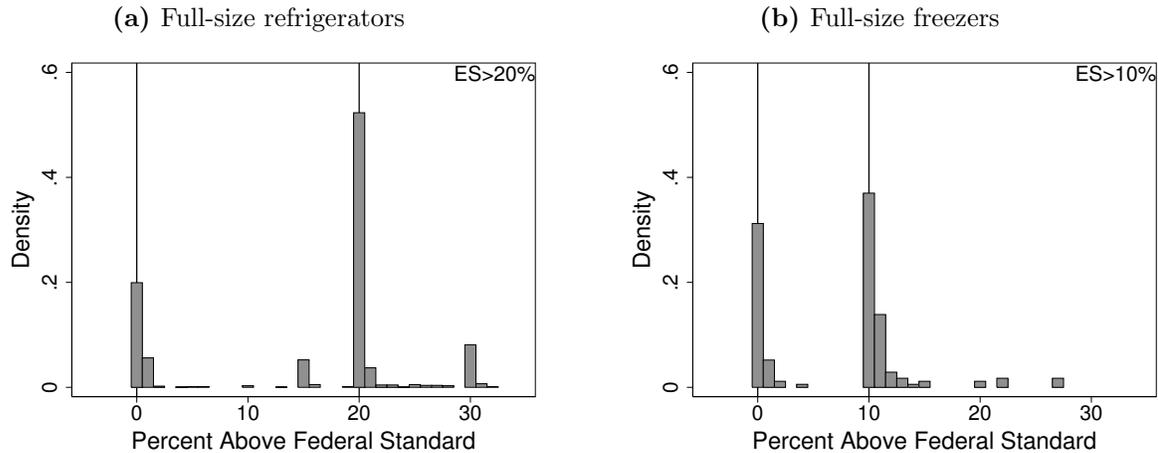
Figure 3: Distribution of refrigerator and freezer energy efficiency

Figure plots histograms for refrigerator and freezer energy efficiency from Association of Household Appliance Manufacturers data. Appliances rated at 0 exactly meet the federal standard. Appliances rated at or above 20% for refrigerators and 10% for freezers are eligible for Energy Star labels. Sample sizes are (a) 1,358 and (b) 187.

include all of the electricity factors or account for water use.

As a result, for many goods, the only data readily available that have the relevant metrics to compute Energy Star status are the Energy Star approved lists themselves. But these lists do not include the appliances that just failed to meet the Energy Star standard—the EPA only keeps track of the list of models that attain certification. Thus, the data to test for bunching is not immediately available from federal sources for many goods.

Third-party data are available in some cases, including for refrigerators and freezers. The Association of Household Appliance Manufacturers (AHAM) provides a variety of appliance market data, including a directory of refrigerators and freezers. To gain an Energy Star label, refrigerators must be 20% more efficient than the federal standard, which is a function of each model’s attributes, including its adjusted volume (a weighted function of fresh food and freezer volume), the location of the freezer (top, bottom, side by side), manual versus automatic defrost, and whether the model has through-the-door ice. For freezers, a model must exceed the minimum efficiency by 10% in order to earn a label.

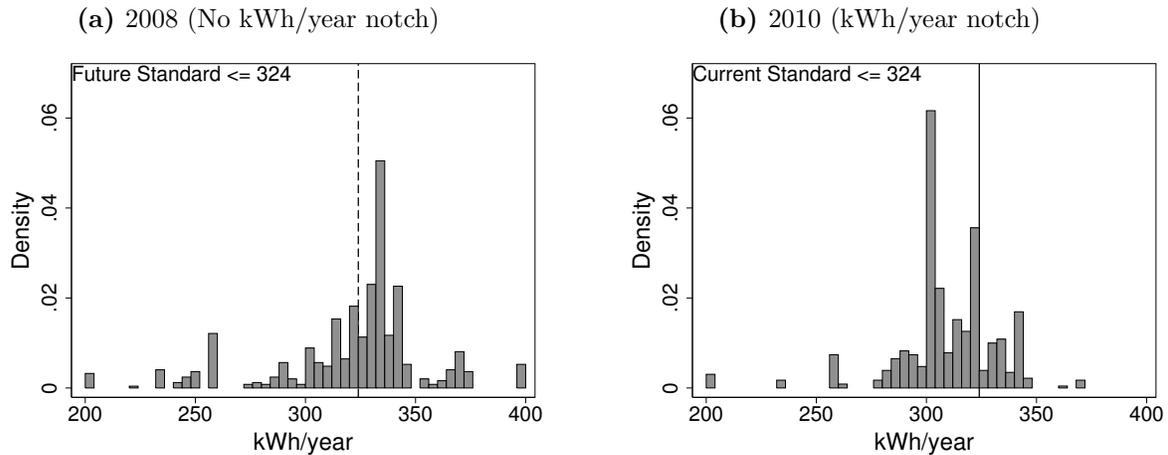
Figure 4: Distribution of dishwasher energy efficiency

Figure plots histograms for dishwasher energy efficiency from FTC data. Appliances rated *below* the Energy Star standard are certified as energy efficient. The EPA Energy Star standard changed in 2009. Before 2009, certification was a function of both kWh/year and water use. Starting in 2010 a 324 kWh/year standard went into effect. The 2008 data are included for comparison—there is no single cutoff in kWh/year in that year.

Figure 3 plots the distribution of the percentage by which refrigerators and freezers exceed the regulatory minimum, using the AHAM data. Models that are 0% above the minimum just qualify for sale. Vertical lines indicate the federal minimum (at 0 in both cases) and the Energy Star minimum (at 20% for refrigerators and 10% for freezers). Both appliances show dramatic bunching at both thresholds. Nearly half of all refrigerators are located exactly at the Energy Star notch, and over one-third of freezers are located there. Another 20% of refrigerators and another one-third of freezers are located at the regulatory minimum, leaving roughly one-third of each good located more than 1% away from the two notches.

Another important appliance with available data is dishwashers. For dishwashers, the only metric needed to calculate Energy Star status in the current regime (effective January 1, 2009) is kWh/year, which is also available on the FTC Energy Guide label. To qualify for a label, dishwashers must use less than 324 kWh/year. Before 2009, the Energy Star label was based on an adjusted Energy Factor, for which comprehensive data are not available. To

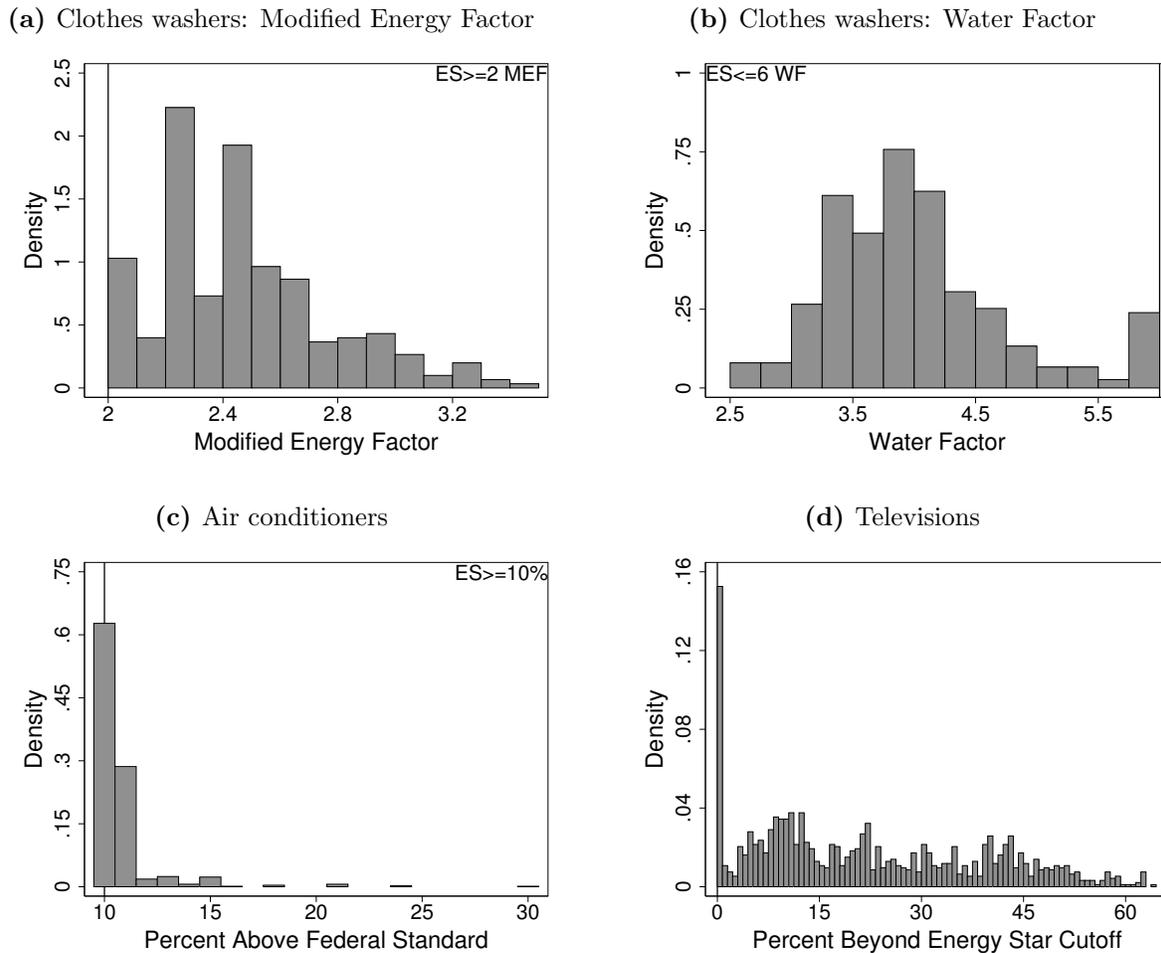
Figure 5: Distribution of Energy Star certified appliances

Figure plots histograms for appliances that qualify for Energy Star labels in several product categories. All data are from the EPA. Vertical lines indicate the critical point for qualification.

show the impact of the label notch, Figure 4 plots the distribution in 2008—when the notch was not in effect—and the distribution in 2010, from the FTC’s list of all dishwashers. The data for 2010 show significant bunching just under 324; many models just qualify for the label. The distribution looks quite different in 2008—the bulk of the dishwashers are above the notch (on the not qualifying side). The distribution shows a marked shift between 2008 and 2010, and the shift is particularly pronounced right next to the notch.

The full list of current qualifying Energy Star appliances is available for a broad set of

categories. These data show only the set of models that qualify, so they provide only a partial picture of bunching. Nevertheless, it is possible to see whether many appliances “just qualify” for the label. Figure 5 includes clothes washers, which have two standards, an energy consumption per cycle standard and a water use standard. Both show evidence that many models “just qualify.” The same is true for air conditioners and televisions. The distribution for televisions is especially pronounced. The strength of this effect is driven largely by a single manufacturer, Samsung, a majority of whose products exactly qualify for the label. Even after removing Samsung (which accounts for around one-fifth of all observations) from the sample, there are many observations very close to the threshold.

Despite the data limitations, the appliance market shows consistent evidence of bunching near the Energy Star notches across many different goods. Combined with the evidence from LEED certification and automobiles, these data suggest that there is indeed product bunching in response to energy efficiency notches. This implies a welfare distortion, because energy efficiency improvements are concentrated in models that happen to be near, but not above, notch points. Unless these models happen, by coincidence, to be the models with the lowest costs for improvement, concentrating improvements in those models is not efficient. If notched policies are optimal, then there must be some benefit that offsets these costs. One possibility is increased salience, which may be important if consumers are rationally inattentive to energy efficiency when making purchases. The remainder of the paper explores the case for such rational inattention, beginning in section 3 with a heuristic model of consumer choice, followed by empirical analyses related to automobiles in section 4 and appliances in section 5.

3 Model

To gain insight into the conditions that determine when inattention might be rational, this section describes a simple model of a consumer making a discrete choice who faces a cost to

learn the exact energy consumption of each model. Consumer i is to choose an appliance j out of a set of appliances \mathcal{J} . The consumer will choose exactly one model. The utility from owning model j for consumer i is written as:

$$\begin{aligned} U_{ij} &= \beta' \mathbf{X}_j + k_j - p_j - c_j + \varepsilon_{ij} \\ &\equiv V_j - c_j + \varepsilon_{ij}, \end{aligned}$$

where U_{ij} is utility, X_j is a set of observed attributes, k_j is a model-specific fixed effect, p_j is the price, c_j is the present discounted value of lifetime energy costs, and ε_{ij} is a random error term. V_j is the representative utility of good j , without the energy cost component.

The absence of a coefficient in front of p_j and c_j implicitly assumes full valuation. For simplicity, all consumers are assumed to have a common valuation of attributes, prices and costs. Heterogeneity in these dimensions could be allowed, but it is not central to the point of this exercise. If the consumer knows the fuel cost for all models, then the consumer will choose model j if and only if $U_{ij} \geq U_{ik} \forall k \neq j \in \mathcal{J}$.

To introduce the possibility of rational inattention, suppose that c_j is not known for certain, unless a search cost $s > 0$ is incurred. Costs are known to come from a distribution $f(c_j; X_j, p_j)$, which could depend on the observable attributes. Without loss of generality, assume that $f(\cdot)$ has a mean of zero, which indicates that c_j is the deviation from the mean in energy costs, where the mean costs are loaded onto the fixed effect k_j . Consumers are aware of the attributes X_j and k_j and prices p_j of each item in all cases. The distribution of c is assumed to be independent of ε —all individuals face the same distribution of appliance characteristics. There is only one search decision: if a consumer pays s , he or she will learn about all costs simultaneously. If a consumer does not pay search costs s to learn true energy costs, he or she will choose the model with the highest expected value.

The consumer's problem has two steps. First, the consumer decides whether to search. Then the consumer selects the model with the highest utility, with or without information

about energy costs. In the first step, an expected utility maximizing consumer i will decide to search if and only if the expected increase in utility from making an enlightened choice exceeds the cost s . That is, the consumer will search if and only if:

$$\mathbb{E}[\max_j U_{ij}] - s \geq \max_j \mathbb{E}[V_j + \varepsilon_{ij}].$$

To gain additional intuition, it is useful to reduce the problem by assuming there are only two alternatives, j and k , and that the variance of c_{ij} and c_{ik} are the same. Then, with search, consumer i will choose model j if and only if $V_j - c_j + \varepsilon_{ij} \geq V_k - c_k + \varepsilon_{ik}$. Without search, the consumer will choose j if and only if $V_j + \varepsilon_{ij} \geq V_k + \varepsilon_{ik}$. How does the consumer decide whether to search?

From the point of view of individual i , ε_{ij} and ε_{ik} are fixed; the uncertain random variables are c_j and c_k . Define the difference between the two as $\tilde{c} \equiv c_j - c_k$. Define θ_i as the preference of person i for model j over model k , given no search, $\theta_i \equiv V_j - V_k + \varepsilon_{ij} - \varepsilon_{ik}$. For an individual i with $\theta_i > 0$, someone who prefers j without search, the expected gain from searching is the product of the probability that the consumer would have chosen k instead, had they searched, and the expected utility gain from choosing k , conditional on wishing to have chosen k :

$$\begin{aligned} \mathbb{E}[\text{Gain}|\theta_i > 0] &= \Pr(V_j - c_j + \varepsilon_{ij} < V_k - c_k + \varepsilon_{ik}|\theta_i > 0) \\ &\quad \cdot \mathbb{E}[V_k - c_k + \varepsilon_{ik} - V_j + c_j - \varepsilon_{ij}|U_{ij} < U_{ik} \text{ and } \theta_i > 0] \\ &= \Pr(\theta_i < \tilde{c}|\theta_i > 0) \cdot \mathbb{E}[\tilde{c} - \theta_i|\theta_i < \tilde{c} \text{ and } \theta_i > 0]. \end{aligned}$$

The second term is a truncated mean. If we assume a functional form for \tilde{c} , then we can write out the truncated mean as a function of the underlying parameters. Let c_j and c_k be identically, independently normally distributed, with mean zero (which comes from the normalization of the fixed effects above) and standard deviation σ_c . Then $\tilde{c} \sim \mathcal{N}(0, 2\sigma_c^2)$.

The expected gain from searching is then:

$$\mathbb{E}[\text{Gain}|\theta_i > 0] = (1 - \Phi(\theta_i/\sigma_c)) \cdot \left[\sqrt{2}\sigma_c \frac{\phi(\theta_i/\sigma_c)}{1 - \Phi(\theta_i/\sigma_c)} - \theta_i \right],$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the normal cumulative density function and probability density function, respectively. A person i who would have chosen model j without search ($\theta_i > 0$) will pay search costs s if and only if:

$$(1 - \Phi(\theta_i/\sqrt{2}\sigma_c)) \cdot \left[\sqrt{2}\sigma_c \frac{\phi(\theta_i/\sqrt{2}\sigma_c)}{1 - \Phi(\theta_i/\sqrt{2}\sigma_c)} - \theta_i \right] > s.$$

The gain from search is falling in θ and rising in σ_c , both of which are intuitive. A consumer with a higher θ_i —whose preference for j over k is greater—is less likely to wish to change upon learning the true energy efficiency. This means that as models become “more different,” fewer consumers will wish to search, because their preference for their choice will be stronger.

The net returns to search are falling in s . As the cost of search rises, fewer people will find it worth searching. An increase in σ_c , the standard deviation of the unexpected portion of energy costs, will increase the return to search because it raises the probability that, upon searching, the consumer will switch to the other model. This parameter represents the “surprise” element of energy costs. Intuitively, this means that, as uncertainty about fuel costs increases, search will be more valuable and thus more frequent in the population.

For any given σ_c and s , there will be a cutoff value of θ_i , call it θ^* , such that a consumer with $\theta_i > 0$ will search if and only if $\theta_i < \theta^*$. There will be an analogous threshold value for those with $\theta_i < 0$, who will choose k without search. Call this threshold $\underline{\theta}$. If $V_j = V_k$, then the search decision will be symmetric, so $\underline{\theta} = -\theta^*$. Looking across consumers, the probability that a consumer searches is a function of the distribution of θ . In particular, the fraction of consumers who do not search is equal to the area of the tails of the θ distribution

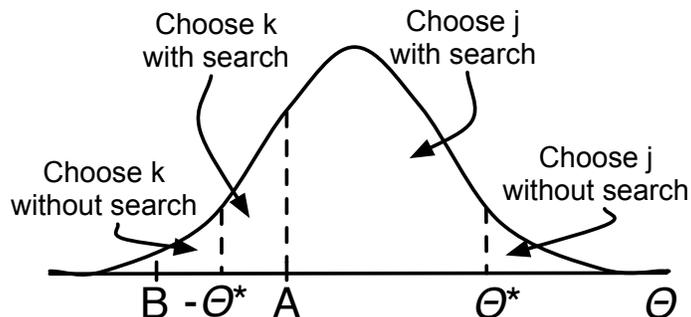
Figure 6: Search and demand across types

Figure shows the distribution of search and demand across types.

beyond the cutoff:

$$\text{Fraction who do not search} = 1 - G(\theta^* - V_j + V_k) + G(\underline{\theta} - V_j + V_k).$$

When $V_j = V_k$, this reduces further to $1 - G(\theta^*) + G(-\theta^*)$. This case is illustrated in Figure 6. Consumers who have very strong preferences—those in the upper and lower tails of the θ distribution—will not search. Those in the middle, who are relatively close to indifference, will search. If we raise s , θ^* will fall, and the slice of the distribution where search happens will shrink. The opposite is true of σ_c ; as it increases θ^* will rise, expanding the slice of the distribution that searches.

These intuitive results constitute the main guidance for empirical work provided by the model. Consumers are more likely to be rationally inattentive if (a) search costs are high, (b) the variance of unmeasured energy costs are low, and (c) products are very different, so that consumers are far from indifferent between their first choice and the alternatives. This last point implies either that the attributes of goods are quite different or that the random utility component of preferences is large.

3.1 Inattention and underprovision

What does rational inattention imply about firm behavior? Suppose that firm j has a choice to increase energy efficiency (lower c_j) at some cost below the informed consumer's willingness to pay. Will it adopt this improvement, and increase price to recoup the costs? It is intuitive to expect that firms might underprovide energy efficiency if some consumers are unaware of such improvements because they are inattentive, as they would perceive only the price increase and not the commensurate benefit. It is worth noting, however, that this intuition holds only in certain cases. The reason is that, in a situation in which some consumers search and some do not, the ones that do not search are the ones that are very far from being indifferent, so that marginal consumers are all searching. An explication of the market demand for each product makes this clear.

Assume that the distribution of consumers is measure 1. The demand for model j is the sum of two components. The first is the fraction of consumers who prefer model j without searching, and choose not to search—those for whom $\theta_i > \theta^*$. The second component is the fraction of consumers who search, and prefer model j after searching. Together, the demand for j , labeled D_j , is:

$$D_j(\mathbf{V}, \mathbf{c}, \mathbf{s}) = \Pr(\theta_i > \theta^*) + \Pr(-\theta^* < \theta_i < \theta^*) \cdot \Pr(\theta_i - \tilde{c} > 0 \mid -\theta^* < \theta_i < \theta^*), \quad (1)$$

where the symmetry assumption that $\underline{\theta} = -\theta^*$ is maintained from above. The formula for the conditional probability in the second term depends on the size of the difference between c_j and c_k . This is illustrated in Figure 6, which plots the distribution of θ . In the figure, individuals with $\theta_i < -\theta^*$ will choose model k without searching; those with $\theta_i > \theta^*$ will choose model j without searching; and those in between will search. Of those who search, some will choose j and some k , after learning true costs. The example is drawn assuming that $c_j - c_k = A$, so all searchers who have $\theta_i > c_j - c_k = A$ will choose j . As drawn, A lies between $-\theta^*$ and θ^* , which means that marginal consumers search. In such a situation,

there will be no “mistakes”. Consumers who prefer j , accounting for energy costs, will end up purchasing j .

As a result, there will be no underprovision of energy efficiency. If firm j has an opportunity to decrease c_j (to improve energy efficiency) at a cost below the reduced operating cost, it will do so because all consumers who will buy its model in equilibrium will be aware of the true energy cost. In this case equation 1, the demand for j , becomes:

$$\begin{aligned} D_j(\mathbf{V}, \mathbf{c}, \mathbf{s}) &= \Pr(\theta_i > \tilde{c}) \\ &= 1 - G(-V_j + V_k + \tilde{c}) \\ &= 1 - G(-\beta' \mathbf{X}_j + p_j + c_j + \beta' \mathbf{X}_k - p_k - c_k). \end{aligned}$$

Therefore, D_j is unchanged from an increase in p_j that is offset by a commensurate change in c_j . Demand will increase if the firm can lower c_j by less than the price increase passed on to consumers. Thus, there is no distortion that creates underprovision of energy efficiency, even if some consumers are inattentive.

The same is not true, however, when the difference between c_j and c_k is large enough so that the marginal consumer, under full information, is in the set of consumers who choose not to search. In Figure 6, this would be true if $c_j - c_k$ were at point B . Given these cost realizations, some of the consumers who choose to purchase k without searching would in fact prefer to have model j . In this situation, the demand for model j (equation 1) is:

$$\begin{aligned} D_j(\mathbf{V}, \mathbf{p}, \mathbf{c}, \mathbf{s}) &= \Pr(\theta_i - c_j + c_k) \\ &= 1 - G(\underline{\theta}). \end{aligned}$$

If firm j improves energy efficiency (lowers c_j) and raises the price to cover the additional

marginal cost, demand will fall:

$$\left. \frac{\partial D_j}{\partial p_j} \right|_{\Delta p_j = -\Delta c_j} = -g(\underline{\theta}) \frac{\partial \underline{\theta}}{\partial p_j} < 0.$$

The reason is that an increase in p_j will raise $\underline{\theta}$, which will increase the number of people who are not searching who would in fact prefer j . This will dampen the willingness of firm j to invest in energy efficiency.

More generally, consumers who do not search will perceive only an increase in price from a product change that decreases c and raises p . Thus, an increase in p will reduce the number of searchers. If some of those consumers would have switched products after search, then raising p involves a reduction in demand from marginal consumers, and firms will not always increase energy efficiency, even when their cost of doing so is below the value of increased efficiency to consumers (i.e., there will be an energy paradox). This will always be true for the low cost firm in an equilibrium in which search costs are high enough so that no consumer searches. It will also be true more often when the variance in energy costs is large, so that the average gap between realizations is greater.

This explanation for the energy paradox has distinct implications for empirical work, compared to more commonly discussed notions of consumer myopia or other biases. Even if consumers are, on average, making rational choices given the products that are offered to them in equilibrium—a paradox may nevertheless exist. Most empirical tests of the energy paradox are tests of whether consumers rationally substitute between goods that are on offer. Such a test cannot detect a market failure in the set of products offered.

3.2 Summary and limitations

This description of a durables market is not complete, as it has not specified fully the form of competition among firms or the consistency of consumer beliefs with that strategy. Nevertheless, in a no-search equilibrium, firms will not be able to signal higher energy efficiency

simply by raising price, as that would give all firms an incentive to inflate prices, so it cannot be a signal of quality in equilibrium, holding constant all other attributes. Even without fully closing the model, the description here is sufficient to motivate the empirical exercise. It illustrates both the conditions that describe rational inattention, and it explicates how such inattention could lead to underprovision of energy efficiency.

One advantage of the model setup here is that it leads directly to a standard discrete choice model that could be estimated with data. Unfortunately, the necessary data, which include sale shares, prices, and energy efficiency of various appliance categories, are not immediately available. In future work, I am pursuing the necessary data to implement this model directly. In the remainder of this paper, I examine facts about the market for automobiles and appliances that quantify the variance of energy costs within similar products. As a loose way of comparing the variance in consumer utility from different products, I focus on variation in prices and compare the size of that variation to the variation in energy costs. I do not pursue methods for directly measuring search costs, but I do make note of some barriers to accurate discernment of energy costs. Thus, the goal of the empirical portion is to ask whether the variation in energy costs is small compared to the variance in preferences, as proxied by price. If so, then a modest search cost may be sufficient to make inattention rational.

4 Automobile energy costs

The fuel costs of new vehicles are substantial—lifetime fuel costs vary by thousands of dollars across new vehicles. But, this variation is relevant to consumer choice only to the extent that such variation would prove pivotal in choices, meaning that consumers are close enough to indifferent between alternatives to be swayed by fuel costs. To put fuel costs into this context, this section compares the variation in fuel costs to the variation in transaction prices for similar vehicles. It then considers how much uncertainty about fuel costs remains, given

information from fuel economy labels.

4.1 Fuel cost and transaction price variation

To quantify fuel costs and prices, I analyze transaction data from a large random sample of new vehicle purchases for model year 2006 vehicles. Transaction data come from an industry source that directly samples a large, representative sample of dealers across the country. The data contain information about the transaction price, trade-in allowance, cash rebates and financing. The final prices are adjusted for incentives, including cash rebates and interest rate subsidies calculated relative to the Federal Reserve's survey of 48-month car loan interest rates from commercial banks.⁵

The sample used here ranges from May 2005, when the very first model year 2006 vehicles appear in the data, to May 2007, when the very last are sold. I focus here on a single model year in order to provide a snapshot of the market that a consumer wishing to purchase a vehicle at a particular time would face. The model year 2006 was chosen as the most recent year for which a complete cycle is available in the data.

To calculate fuel costs, I assume that all vehicles are driven 12,000 miles a year for 14 years. Both estimates are close to their national averages (US Department of Transportation 2008). I use a 5% annual discount rate and a gasoline price of \$2.50 per gallon, which is approximately the average price over the months from which the sample is drawn. For fuel economy, I merge EPA fuel economy ratings onto the transaction data. For fuel costs, I use the combined EPA fuel economy rating, which is a weighted harmonic average of the city and highway ratings.

To begin, Table 1 shows the average actual transaction price, net of incentives, and the average fuel cost across vehicle categories, as well as their standard deviations. The use of categories is a coarse industry measure that delineates cars into compact, midsize, luxury, and sports cars, and light-duty trucks into SUVs, pickups, and vans. The table also shows

⁵The same methodology is employed in Sallee (2011a), which provides additional detail.

Table 1: Median fuel cost and transaction prices across vehicle categories

	Median Price	SD Price	Median Fuel Cost	SD Fuel Cost	N	No. VINs
Compact Car	16,829	(3,890)	9,899	(1,680)	372,802	192
Midsized Car	22,053	(4,676)	11,878	(1,347)	344,974	320
Luxury Car	33,642	(12,909)	13,498	(1,397)	151,555	238
Sports Car	24,882	(13,150)	14,141	(1,647)	86,738	61
SUV	26,612	(10,162)	15,629	(2,458)	442,409	432
Pickup	24,450	(5,879)	17,468	(1,623)	309,424	372
Van	24,539	(5,298)	14,141	(1,131)	153,535	90
Total	23,405	(9,506)	13,498	(3,101)	1,861,437	1,705

Table shows median transaction price, accounting for customer rebates, trade-in allowances, and interest rate subsidies. Fuel costs are calculated assuming 12,000 miles driven per year for 14 years, with a 5% discount rate and a \$2.50 per gallon price of gasoline. The fuel economy used is the EPA estimated combined fuel economy. The fifth column shows the total sample size, and the sixth shows the number of distinct VIN codes included in the sample.

information about the overall sample size, including total observations and the number of different Vehicle Identification Numbers (VINs), which distinguish across models and engine sizes and trim levels.

Table 1 shows that fuel costs are significant, ranging between \$10,000 and \$18,000. The total fuel cost is also a substantial fraction of the total vehicle price, with fuel costs roughly half of the purchase price for most categories, reaching nearly two-thirds for compact cars. Fuel cost also varies substantially within categories, with the standard deviation ranging from \$1,100 to \$2,500 from vans to SUVs. Importantly, it is the variation in fuel cost across substitutes that provides consumer with an incentive to be attentive to fuel economy during the purchase decision. Even if fuel costs are high, a consumer stands to gain little by studying fuel economy unless there is variation among the set of vehicles that the consumer might choose. In absolute terms, the variation in lifetime fuel costs across vehicles within a category is large, but it is considerably smaller than the variation in transaction costs across the same set of vehicles, particularly for light trucks, where the standard deviation in fuel costs is around 20% of the variation in transaction prices.

The variation in fuel costs *across* categories is substantial. To make this easier to see, Table 2 shows the change in lifetime fuel cost that would occur if a consumer changed from

Table 2: Fuel cost savings from switching among median vehicle across categories

	Compact Car	Midsize Car	Luxury Car	Sports Car	SUV	Pickup	Van
Compact Car	.	1,980	3,600	4,242	5,731	7,570	4,242
Midsize Car	-1,980	.	1,620	2,263	3,751	5,590	2,263
Luxury Car	-3,600	-1,620	.	643	2,131	3,970	643
Sports Car	-4,242	-2,263	-643	.	1,489	3,327	0
SUV	-5,731	-3,751	-2,131	-1,489	.	1,839	-1,489
Pickup	-7,570	-5,590	-3,970	-3,327	-1,839	.	-3,327
Van	-4,242	-2,263	-643	0	1,489	3,327	.

Table shows the change in lifetime fuel costs that would occur if a consumer changed from purchasing the sales-weighted median fuel economic vehicle from the category in the row to the median vehicle from the category in the column. Fuel costs are calculated assuming 12,000 miles driven per year for 14 years, with a 5% discount rate and a \$2.50 per gallon price of gasoline. The fuel economy used is the EPA estimated combined fuel economy. A zero value occurs where the median fuel economy of the vehicle in the two categories is the same. The table is symmetric (with opposite sign) about the main diagonal.

the median fuel economic vehicle in the vehicle category in each row into the median vehicle in the column. Table 2 shows that a consumer stands to lose (or gain) thousands of dollars when changing from fuel economic compact or midsize cars to light trucks or sports cars (or vice versa).

The variation in fuel costs across categories is not, however, likely to be the factor that motivates a consumer to scrutinize fuel economy when purchasing a vehicle. Marketing studies suggest that, before explicitly shopping for a vehicle, consumers have in mind a rather short “shopping list” of vehicles that they might purchase. In the end, they only end up examining and test driving a small number of vehicles, with a median of just 3 models, before making a purchase (Ratchford and Srinivasan 1993; Moorthy, Ratchford, and Talukdar 1997; Ratchford, Lee, and Talukdar 2003). Even if consumers respond to fuel price levels by switching their category of search, unless they search across multiple categories in this second stage of intensive search, then the relevant variance for attention is within category, not across.

Thus, the standard deviations in fuel costs within categories reported in table 1 provide a more relevant metric for gauging the magnitude of the gains to consumer attention. How do these measures, which range from \$1,000 to \$2,500, compare to other sources of variation

Table 3: Standard deviation in vehicle transaction prices and fuel costs

	(1) Price (All)	(2) Price (Within VIN)	(3) Fuel Costs	(4) Ratio: (2)/(3)
Compact Car	3,890	1,707	1,680	1.0
Midsized Car	4,676	2,107	1,347	1.6
Luxury Car	12,909	3,099	1,397	2.2
Sports Car	13,150	2,562	1,647	1.6
SUV	10,162	2,667	2,458	1.1
Pickup	5,879	2,745	1,623	1.7
Van	5,298	2,443	1,131	2.2
Total	9,506	2,435	3,101	0.8

Table shows the standard deviations. Column (1) shows variation between and within models in transaction price. Column (2) shows the variation in transaction price within VIN. Column (3) shows the variation in fuel costs. Column (4) shows the ratio of within-VIN price variation to fuel costs.

in vehicle costs? Is this variation large?

To further contextualize these costs, Table 3 shows the amount of the variation in the price of vehicles, within a particular type of vehicle, where a vehicle type is a unique VIN. The first 10 characters of a VIN indicate the vehicle's characteristics, including manufacturer, model name, model year, engine cylinders, engine displacement, drive type, body style, trim level, fuel type, transmission, and aspiration (e.g., turbo charged). For example, a 2006 flexible-fuel Ford F150 extended-cab pickup with a 5.4-liter V8 engine and manual transmission is a unique VIN. The VIN does not indicate differences in options packages, such as carpeted floor mats, roof racks or satellite radio. The variation in prices that remains after removing the mean price for each VIN in the sample therefore reflects differences in bargaining outcomes and financial incentives from manufacturers, as well as differences in final options. There are 1,705 different VINs in the sample for model year 2006. I demean all prices by VIN, and then calculate the variation that remains. This is equivalent to measuring the standard deviation of the residuals from a fixed effects regression, where the fixed effects are the VINs.

Column 1 in Table 3 reproduces the total variation in price, including between and within variation, from Table 1 by category. The variation measured in this column includes comparisons across very different vehicles. For example, luxury SUVs and compact SUVs

Table 4: Standard deviation in vehicle transaction prices within VIN

	(1) All	(2) Windsorized	(3) Cash Only	(4) Base Price	(5) Rebates	(6) Financing
Compact Car	1,707	1,468	1,483	1,427	487	753
Midsize Car	2,107	1,815	1,899	1,832	703	864
Luxury Car	3,099	2,701	3,145	2,859	996	878
Sports Car	2,562	2,183	2,515	2,298	491	938
SUV	2,667	2,312	2,646	2,476	1,185	1,096
Pickup	2,745	2,349	2,474	2,497	1,503	1,162
Van	2,443	2,094	2,330	2,281	1,154	951
Total	2,435	2,100	2,344	2,205	1,024	970

Table shows the standard deviation in vehicle prices within VIN. Column (1) shows total variation within VIN for the full sample. Column (2) eliminates the top and bottom 2% of prices for each VIN to Windsorize the data. Column (3) drops vehicles financed through the dealerships from the sample. The final three columns show the variation coming from the three components of price—base price (4), manufacturer rebates (5) and low-interest financing incentives (6).

are grouped together in the SUV category. This likely overstates the variation in prices relevant to consumer choice. Column 2 shows the standard deviation in price within each VIN. This number shows the amount of money that a consumer could save by reducing the transaction price by one standard deviation, conditional on having already selected a vehicle to purchase.

There is substantially less variation within VIN, as would be expected, but the remaining variation is still quite large. The within-VIN variation is only a quarter of the total variation for the most expensive vehicle categories—luxury cars, sports cars, and SUVs. But, in all categories, a one standard deviation reduction in transaction price, within VIN, still represents thousands of dollars in savings.

For comparison, column 3 reproduces the variation in fuel costs across vehicles from Table 1, and column 4 shows the ratio of within-VIN price variation and fuel costs. In all categories, the within-VIN variation exceeds the fuel cost variation. In other words, this consumer will save more, on average, by improving the transaction price for the vehicle he or she has chosen by one standard deviation than he or she would gain from making a one-standard deviation improvement in fuel economy, within their vehicle class.

Recall that fuel cost variation is necessarily across vehicles within a category. As a result, outliers are of minimal concern for fuel costs, but unusual transaction prices, and the method used to impute financing subsidies, could cause an overstatement in the variation within VIN. To examine this, Table 4 shows the variation in transaction prices across several subsamples and components. Column 1 reproduces the within-VIN standard deviation for the full sample. Column 2 eliminates, for each VIN, the top and bottom 2% of transactions, to ensure that outliers are not influencing the estimate. Mechanically, this procedure must reduce the variation. If outliers play a large role, then this reduction will be significant. The reduction in variation is notable, but it is on the order of 10 to 20%, and the remaining variation is still roughly as large as, or larger than, variation in fuel costs.

The transaction prices used here account for customer cash rebates, as well as an imputed estimate of the value of interest rate subsidies for transactions that use financing at the dealership, based on a comparison of actual interest rates and a benchmark from the Federal Reserve. This imputation procedure could introduce additional variation into prices. To see if that is driving the results, column 3 of Table 4 limits the sample to cash transactions—that is, transactions that are not financed by the dealership. (Around half of transactions are cash transactions; customers need not actually have cash on hand, they may have borrowed from a commercial bank or other source.) The price variation in this subsample is still quite large, with modest reductions in most categories.

Finally, the table also breaks out the three components of the final price—the base price (including trade-in over or under allowance), direct cash rebates, and financing subsidies. The biggest portion of the variation comes from differences in the price of the vehicle before subsidies, but roughly a quarter comes from rebates and another quarter from financing. Overall, what this suggests is that there is substantial variation in several dimensions, and it appears that the variation in transaction prices is not an artifact of subsidy calculations or outliers. Instead, it appears that transaction prices vary substantially, even *within* a particular type of vehicle, enough so that this variation is as large as the variation in fuel

costs *across* models within a vehicle class.

4.2 Model year cycle variation

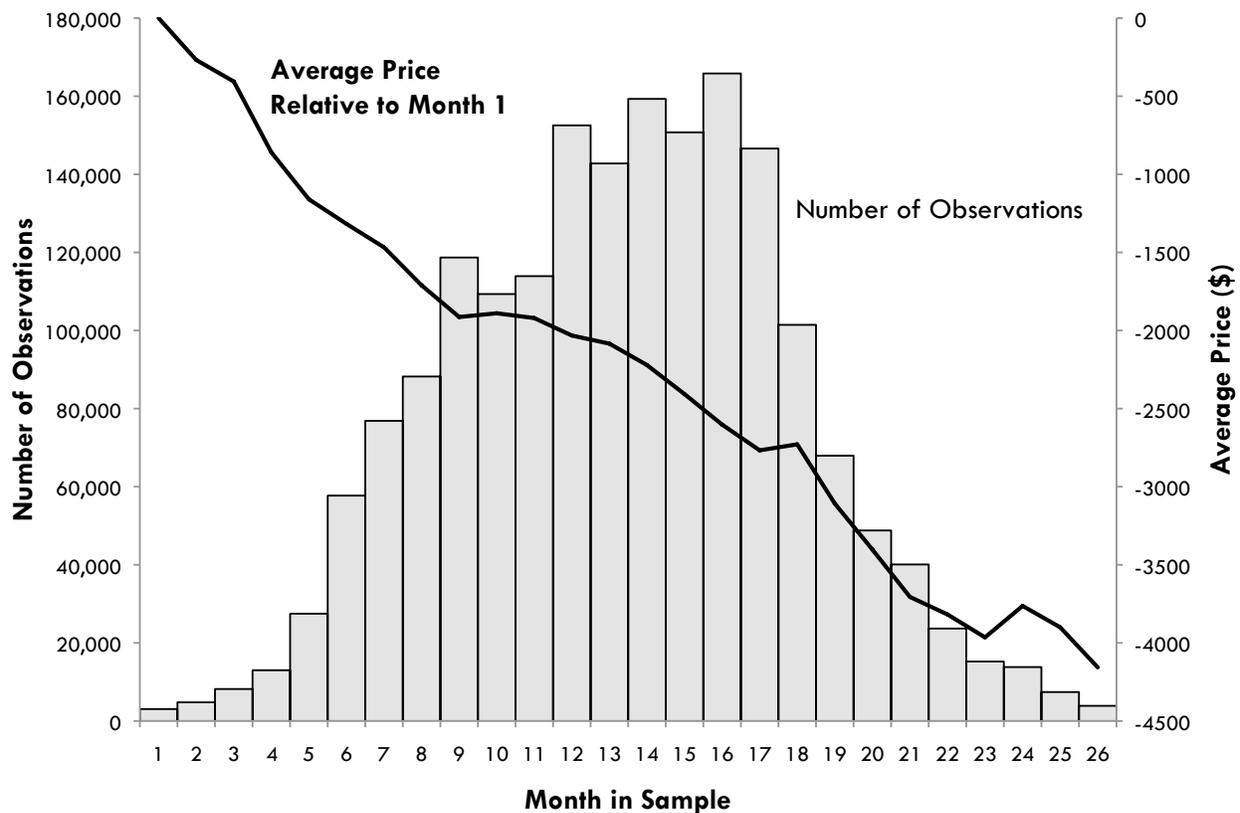
Some transaction price variation is predictable. The most pronounced example is variation over the model year cycle. When a new model year for a vehicle is first introduced, the price of the vehicle is at its highest point. In following months, the price declines steadily, but slowly. To quantify this variation in prices, I run a regression of the following form:

$$p_{ijt} = \alpha + \sum_{s=2}^{26} 1(t = s)\delta_t + \sum_{v=2}^{1,705} 1(v = j)\gamma_j + u_{ijt}, \quad (2)$$

where p_{ijt} is the transaction price of observation i of VIN type j in month t , δ_t are month dummies for the 26 months in the sample (with month 1 the omitted category), γ_j are dummies for each of the 1,705 VIN types in the sample, $1(\cdot)$ is the indicator function, and u_{ijt} is an error term. The month dummy variables represent the average price in each sample month, controlling for VIN fixed effects.

Figure 7 plots these month coefficients to show the average price decline over the model year cycle. The figure also shows the sample size in each month, to show the distribution of sales over the time period. The price decline is relatively smooth, with a slope around $-\$156$ per month. The total price decline over the cycle is around $\$4,000$, which is 17% of the median vehicle price over the entire sample. Estimating equation (2) can also be written with logged price on the left-hand side. Then, the month coefficients represent average percentage price declines relative to the first month in the sample. The coefficients from this regression have a very similar shape to the one shown in Figure 7. They indicate that prices decline by 13% over the 25 months in the sample, which is broadly similar to the average annual model year price decline of 9.0% estimated by Copeland, Dunn, and Hall (2011).

The standard deviation of the timing of purchases in the data is 4.5 months. Delaying a transaction by 4.5 months creates, on average, a $\$750$ savings for a consumer. This is around

Figure 7: Vehicle prices over model year cycle

Solid line plots the month coefficients from a regression of vehicle price (in levels) on a set of month dummies for every month in the sample with VIN fixed effects. Bar graph plots total sample size over the model year cycle. Month 1 is April 2005.

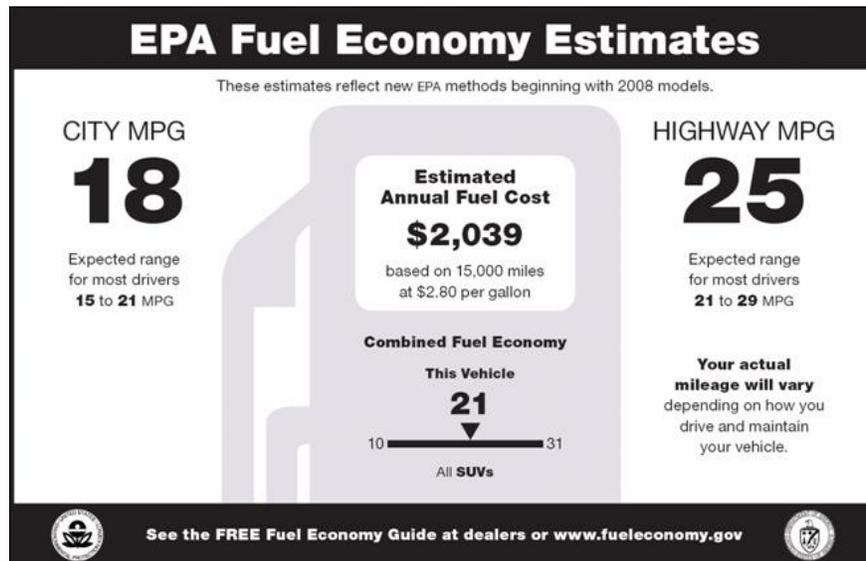
half of the savings from making a one-standard deviation improvement in fuel costs within vehicle class, for most classes.

4.3 Search costs: automobile label inaccuracies

Even if the gains to search are relatively low, consumers will become informed unless the cost of doing so is significant. All new automobiles come with fuel economy labels that show the EPA's official estimates of mpg. Given that, is it reasonable to suppose that consumers face nontrivial costs in calculating lifetime fuel costs?

Figure 8 provides an example of the current fuel economy label for new vehicles. The label includes three fuel economy estimates (city, highway, and combined), as well as an

Figure 8: Current fuel economy label



estimated annual fuel cost and a comparison of the combined fuel economy to other vehicles in the same class. Without this information, consumers would be hard pressed to make reasonable conjectures on lifetime fuel cost; nevertheless, the label does not provide all of the information necessary to fully inform the consumer. The label contains an annual fuel cost, but this must be transformed via a present discounted value calculation, which requires information about the life of the vehicle, a discount rate, the schedule of mileage over time, and future fuel costs. The comparison to other vehicles is in mpg, not dollars. Also, the comparison shows the extreme values in each category, without providing any sense of the distribution between the points, which would be necessary for the consumer to judge how likely they would be to find a similar vehicle with substantially different fuel costs. Finally, heterogeneity looms large. To the degree that consumers drive different numbers of miles in a year, drive more or less aggressively, or drive more or less of their time in highway settings, their fuel economy may vary a great deal.

Prior research gives several reasons to suspect that consumers find converting the information on labels into lifetime fuel costs difficult. Larrick and Soll (2008) document that, in a laboratory setting, consumers fail to understand the nonlinearity of costs in mpg and

overestimate the pecuniary gains from increases in high mpg vehicles, underestimating improvements in inefficient models. Allcott (2010) runs a stated preferences experiment and shows that consumer beliefs about the value of fuel economy are inaccurate and biased in ways consistent with the findings of Larrick and Soll (2008). Qualitative interviews documented in Turrentine and Kurani (2007) show that consumers lack information on all of the building blocks necessary for a lifetime fuel cost calculation, save the current price of gasoline. Sallee (2011b) shows that city ratings are, on average, 19% lower than highway ratings (equivalent to the difference between a Volkswagen Jetta and a Ford Crown Victoria), so that different proportions of highway and city driving can have a significant impact on fuel costs. In sum, recent research suggests that consumers do face significant costs for ascertaining the lifetime cost of fuel, even when fuel economy labels are present. Early marketing research on energy efficiency labels came to a similar conclusion—labels improved decision making, but consumers are sensitive to the form of the information and responses are not entirely consistent (McNeill and Wilkie 1979; Hutton and Wilkie 1980).

Consumers face an additional obstacle in making fuel economy valuation calculations because the official fuel economy estimates have exhibited significant bias. EPA fuel economy rates are estimates of typical driving behavior based on results from vehicle tests on specific courses. Actual consumer driving will differ from test patterns, so the tests will not be accurate measures for any individual consumer. Aside from the inability to capture heterogeneity, the tests may be poor measures of the actual fuel economy of particular vehicles if they do not reflect *average* driving behavior correctly.

Since its inception in 1978, the fuel economy label program has undergone two major changes. First, in 1986, in response to consumer complaints that EPA ratings significantly overestimated on-road fuel economy, the EPA adjusted the test ratings. Rather than devising a new procedure to improve accuracy, the EPA simply scaled down the original values. City values were reduced by 10% and highway values were reduced by 22%.

In recent years, the EPA became concerned that the tests were again inaccurate, on

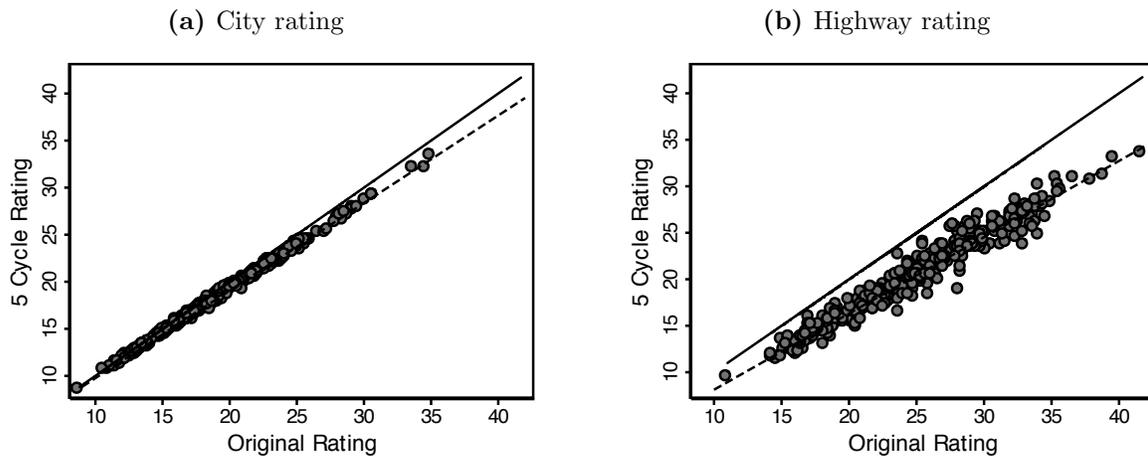
average, due to changes in typical driving. For example, the top speed on the highway test was 60 miles per hour, and the tests were conducted without running the vehicle's air conditioner. Stopping and starting were also more gentle than was believed to be typical of current drivers. The EPA decided that a new testing procedure was necessary, which would take estimates from not two but five different tests for each vehicle. The new tests are being phased in currently, with revised estimates to appear for models beginning in 2008.

To determine the implications of a new test procedure, the EPA used a sample of 615 vehicles. For these vehicles, they calculated both the old label values and the new label values. They have generously provided these data to me for analysis here. Of the original sample, 380 are gasoline-only vehicles with complete information on all five tests. Using the original test data, I reconstruct the original fuel economy label ratings and the new five-cycle ratings for this sample.⁶

Analysis of these test vehicles allows me to answer the question: what is the expected miscalculation of lifetime fuel costs for a consumer who uses the old, less accurate measure of fuel economy? This exercise presumes that the new five-cycle test is the true average fuel economy. Under that assumption, how inaccurate were old fuel economy measures, in terms of lifetime fuel cost?

There are two possible way to answer this question. The first is to simply compare the estimated lifetime fuel costs using the old rating system to the lifetime fuel costs using the new rating system. If the old fuel economy ratings were biased on average, then this bias will be reflected in the fuel cost difference. A second method is to adjust the old test number for any systematic bias, and then measure only the residual—the unexpected portion of the difference. This method measures the uncertainty facing a consumer who is aware that the old numbers did not reflect average driving behavior and made a scaled adjustment, but was

⁶My calculation is not exactly the same as the EPA's. One test, the us06 test, is now run in two distinct versions, one for use in the highway calculation and the other for use in the city calculation. Only one unified version is available in the data. The EPA imputed two different values for the us06 test in its analysis, but I do not have the information necessary to replicate EPA's imputation and therefore I use the single test value for both ratings. The difference resulting from this discrepancy should be quite small.

Figure 9: New and old EPA fuel economy ratings, in mpg

Figures based on 380 test vehicles for which data for both rating systems are available. Test vehicles come from model years 2003 to 2006. Dashed lines show a linear fit.

not able to determine the variation from the average mistake for each particular vehicle.

When looking only at gasoline vehicles, the relationship between the old test rating and the new test rating is linear for both city and highway. (The new test procedure has a much more pronounced impact on gas-electric hybrid vehicles, which I omit from the sample.) Figure 9 shows scatter plots of the old and new tests, for both city and highway values, along with the 45 degree line for reference.

In each figure, the deviation from the 45 degree line is indicative of the difference between the two measures, and the spread in the data shows the degree of idiosyncratic differences. That is, if the data fall on a straight line, then a scalar multiple of the old formula would produce the new estimate. In both diagrams, the vast majority of the data points are below the 45 degree line, implying that the new rating is lower than the original. Moreover, it is clear that the slope is less than 1, indicating that the error in fuel economy measurement rises with fuel economy. The effect is more pronounced for the highway rating. The spread is also much more pronounced for the highway rating, indicating that, even after accounting for systematic (average) bias, the change in fuel economy between rating systems varies

Table 5: Relationship between new and old EPA ratings

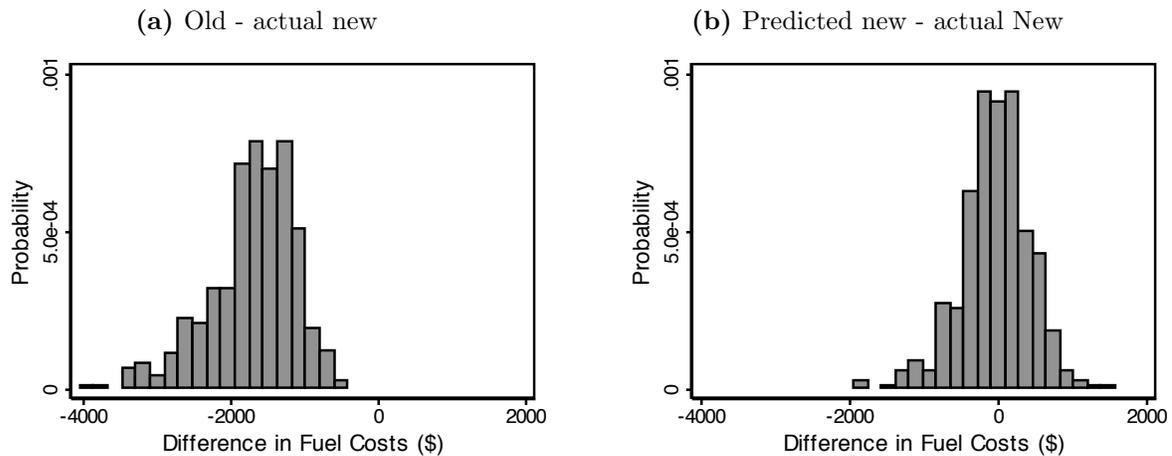
	City	Highway
Old Rating	0.930 (0.003)	0.820 (0.010)
Constant	0.455 (0.056)	-0.087 (0.251)
N	380	380
R^2	.996	.948

Standard errors in parentheses. Dependent variable is the new five-cycle EPA test rating.

substantially across vehicles.

To quantify this relationship in more detail, I regress the new rating on the old. Table 5 reports ordinary least squares (OLS) estimates from a bivariate regression of the old city rating on the new city rating and the old highway rating on the new highway rating. Consistent with the graphical evidence, both city and highway regressions show a slope estimate that is statistically distinguishable from 1. The difference between the coefficient estimates and 1 is a measure of the degree to which the old system overstated the true improvement in fuel economy that arose from a 1-unit increase in the old measure. By that measure, the old rating system overestimated the cost savings of a fuel economy improvement by 7% for city ratings and 18% for highway ratings. As a result, a consumer basing his or her fuel economy calculation on the official ratings would have misstated his or her potential fuel savings by a substantial amount, on average.

Another way to illustrate the implications for consumers is to calculate and compare the estimated lifetime fuel costs—using all of the same assumptions about mileage, gasoline prices, discounting, vehicle lifetime, and ratio of highway to city miles—for both the old and the new measures for each vehicle in the sample. Figure 10 shows the distribution of this difference in two ways. Figure 10a calculates lifetime fuel costs using the old test ratings and subtracts cost using the new ratings. Figure 10b instead calculates this difference using the predicted new fuel economy ratings (the predicted values from the regressions in Table 5). This procedure removes the systematic bias so that differences are centered around zero, and

Figure 10: Difference in lifetime fuel costs across new and old EPA ratings

Figures based on 380 test vehicles for which both rating systems are available. Lifetime fuel costs are constructed assuming 12,000 miles driven per year for 14 years, with a 5% discount rate, a \$2.50 per gallon gasoline price, and 55% of miles driven in the city. Predicted new ratings are calculated using the predicted values from an OLS regression of the new ratings on the old ratings.

it also removes any variation in fuel cost differences that is a linear function of the ratings.

The average bias from the old measures is very large relative to the potential fuel economy savings that consumers might gain from improving the fuel economy of their vehicles. In this sample, the mismeasurement in lifetime fuel costs ranges from -\$400 to -\$4,000. The mean mismeasurement is -\$1,700—that is, on average, the old EPA ratings understated lifetime fuel costs by \$1,700. Even once the systematic bias is removed, the variation in the idiosyncratic mismeasurement is substantial. The standard deviation of the difference between the old and new ratings is \$597, and the standard deviation using the predicted values is not much smaller, at \$467.

Imprecise measurement of fuel costs should not prevent consumers from valuing fuel economy; rational consumers facing uncertainty should have made expected value calculations and based their decision on those. There is little reason, however, to believe that consumers had a reliable way of estimating the bias in the old fuel economy labels, and it is less likely that they had the ability to determine the variation in mismeasurement across models. As

such, the cost of ascertaining true lifetime fuel costs are substantial.

4.4 Summary of automobile costs

The comparisons made here are not strictly derived from the model of section 3, but they do have a clear intuition within a framework of limited rationality. Even once a consumer has chosen a particular car to purchase—defined not just by model, but also by engine type, transmission, and trim level—that consumer still faces choices in financing, time of purchase within the model year cycle, options packages and haggling that have cost implications as large as making fuel economy improvements of one standard deviation within class. If consumers have an attention budget, it stands to reason that they might, rationally, devote their attention to these other margins. Moreover, even if consumers wish to make well informed fuel economy choices, they face uncertainty about true values as a result of imprecision in fuel economy ratings. Consumers who used the old fuel economy rating faced bias roughly as large as the savings from a one-standard deviation improvement in fuel economy within class, and idiosyncratic variation roughly one-third that size. Overall, this analysis suggests the plausibility of rational inattention, even in the automobile market, where the stakes are the largest and when consumers are well acquainted with fuel economy measures that are presented on mandatory labels.

5 Appliance energy costs

The preceding analysis suggests that, while the level and variation in lifetime fuel costs are significant in the automobile market, there is equally large variation in other attributes and search costs may be substantial. This section reviews similar statistics for home appliances. The necessary data on household appliances are not uniformly available from the federal government, but a patchwork of sources allows for an analysis of the most important appliances, including dishwashers, clothes washers, ovens, ranges, refrigerators, and freezers.

Table 6: Appliance lifetime energy costs

Appliance	Mean Lifetime Cost	SD Lifetime Cost	Mean Retail Price	SD Retail Price
Dishwasher	267	21	807	495
Clothes washer (top loading)	282	96	590	244
Clothes washer (front loading)	144	32	919	377
Oven	280	66	2,056	1,300
Range	591	88	962	568
Refrigerator (top, auto)	523	74	709	265
Refrigerator (side, auto)	734	67	2,368	1,548
Refrigerator (bottom, auto)	577	55	1,481	553
Refrigerator (side, auto, TTD)	729	75	2,368	1,548
Refrigerator (bottom, auto, TTD)	680	40	1,481	553
Freezer (upright, manual)	558	81	618	42
Freezer (upright, auto)	776	119	651	126
Freezer (chest)	451	110	391	114

Dishwasher data from FTC, with 10 year life assumed. Clothes washer data from FTC, with 11 year life assumed. Oven and range data from NRC, with 15 year life assumed. Refrigerator data from AHAM, with 17 year life assumed. For refrigerator/freezer data, auto and manual refer to defrost modes. Freezer location is above fresh food refrigerator (top), below (below), or side-by-side (side). TTD indicates through-the-door ice. All retail price data are from *Consumer Reports*. Price data do not distinguish between TTD and non-TTD refrigerators.

Table 6 quantifies the mean lifetime cost and the standard deviation in lifetime cost for these appliances, assuming appliance lifetimes as cited by the DOE, a 5% discount rate, and a 10.66 cent per kWh price of electricity (the price used on current Energy Guide labels). Data for dishwashers and clothes washers come directly from the FTC. FTC data for other appliances are incomplete for varied reasons, so other sources are used. AHAM data are used for refrigerators and freezers, and data from Natural Resources Canada (NRC) are used for ovens and ranges.

Table 6 shows that the lifetime fuel costs for most of these appliances is substantial, ranging from a low of \$144 for efficient front-loading clothes washers to \$776 for freezers. The standard deviations in these costs across models is much smaller. Only freezers show a standard deviation above \$100. This means that, for most major appliances, consumers stand to gain modest amounts from a one-standard deviation improvement in energy efficiency for a given appliance category.

To compare these gains to the prices of the goods, Table 6 also includes the mean retail price and the standard deviation in the retail price across models, as calculated from *Consumer Reports* data. Looking within categories, the lifetime energy costs of many of the cheaper appliance categories is a substantial fraction of the purchase price. The *variation* in energy costs, however, is dwarfed by the variation in retail prices for all appliances except for freezers.

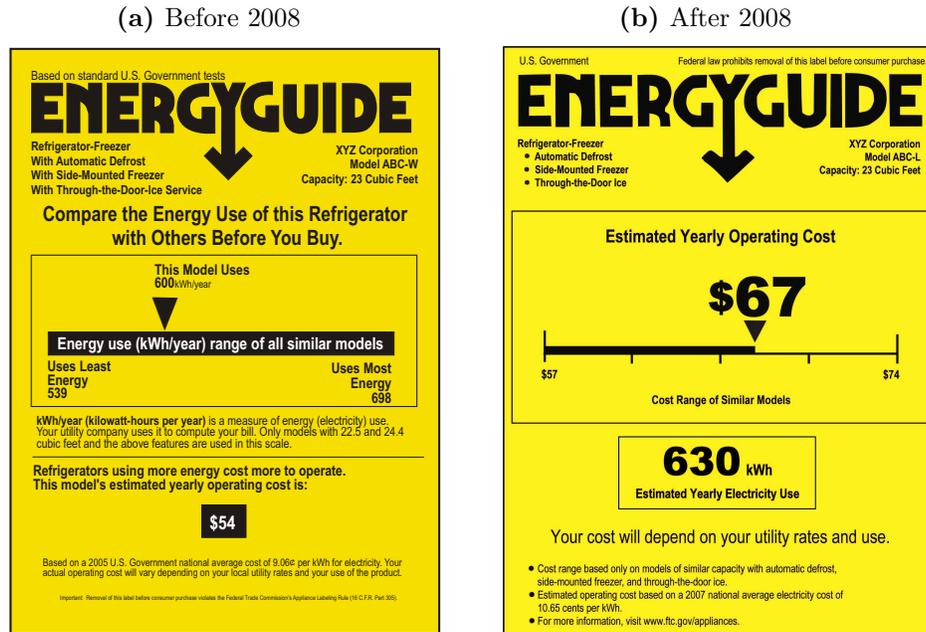
This leads to a conclusion similar to that for automobiles. On an absolute level, energy costs are important, and variation in energy costs is non-trivial. At the same time, there is far greater variation in prices, which implies both that (a) differences in attributes may be sufficient to imply that consumers are far from indifferent between different choices and (b) that a consumer with limited attention may choose to focus on attributes other than energy efficiency.

5.1 Search costs: appliance labels

Some appliances carry Energy Guide labels in the United States, which are mandated by federal law. These include refrigerators and freezers, air conditioners, clothes washers, dishwashers, fluorescent lighting, furnaces, boilers, and heaters. This leaves out a variety of goods, including ovens, ranges, televisions, VCR and DVD players, incandescent lighting, computers, monitors, audio equipment, printers, fax machines, scanners, air purifiers, and dehumidifiers. When present, Energy Guide labels include categorical information about the appliance, along with an estimate of the kWh used per year, the annual operating cost, and the range of costs for similar models.

Labels greatly reduce, but do not eliminate, the search cost of consumers. Prior to 2008, labels included the comparison with other models in terms of kWh per year, instead of dollars, meaning that consumers had to do additional calculations in order to monetize the difference between their model and its alternatives. Figure 11a shows an example for a refrigerator label. The Energy Policy Act of 2005 mandated a review of the labels, and

Figure 11: Energy guide labels



the labels were changed, effective in 2008. An example of a refrigerator label from the new regime is provided in figure 11b.

Even under the new design, labels do not provide all the necessary information to calculate the lifetime cost of operating the appliance, which is the relevant dollar amount for consumers. To do so, consumers still need to estimate the lifetime of the appliance, posit future electricity prices, and calculate the present discounted value using the appropriate discount rate. Thus, labels lower search costs but do not eliminate them. In fact, labels may cause consumers to *undervalue* energy efficiency differences by presenting estimates from only a single year, and only for comparable models that share the same features, which makes a lifetime cost difference of \$100 show up as an annual difference of a few dollars. An alternative is to present an estimate of the lifetime costs of energy on the label, which is partly done in Canada.⁷

The Canadian government does not include a lifetime cost on its labels, but NRC does

⁷Early marketing research suggested that lifetime fuel costs created a bigger response in consumers than labels with annual consumption (Hutton and Wilkie 1980).

Table 7: NRC lifetime cost inaccuracy

	Average kWh/year	Lifetime	Average Second Price Tag	Overstatement: 5% Discount Rate	Overstatement: 10% Discount Rate
Dishwashers	295	13	383	32%	66%
Clothes Washers	215	14	302	42%	80%
Refrigerators	486	17	827	73%	118%
Clothes Dryers	933	18	1,681	83%	131%
Ranges	509	18	916	82%	130%
Freezers	576	21	1,210	113%	169%

Data from Natural Resources Canada. NRC calculates a “second price tag” which is the undiscounted cost of energy given a 10 cent per kWh price of electricity and the appliance lifetime in the table. Overstatement of costs is calculated as the difference between the undiscounted and discounted lifetime costs divided by the discounted costs.

calculate a lifetime cost on its website, which it calls the “second price tag”. The second price tag, which is described as the lifetime cost of operating the appliance, is included along with other energy efficiency statistics in the NRC’s guide to appliances. Unfortunately, NRC calculates this as the *undiscounted* cost of energy over the life of an appliance, given a set of assumptions about the lifetime of each type of appliance, the price of electricity, and the estimated annual energy consumption of the appliance. The lack of discounting causes a significant overstatement of lifetime energy costs.

Table 7 shows the estimated lifetime operating costs, assuming a 10 cent per kWh cost of electricity and the lifetime of appliances indicated on the NRC website. The first column shows the average annual kWh for each appliance type, according to the NRC’s microdata on appliance efficiency. Column three calculates the average second price, which NRC includes in its data for consumers. The fourth and fifth columns show how much the NRC’s undiscounted measure overstates the lifetime costs of the appliances, assuming a 5% and 10% discount rate, respectively, calculated as a percentage exaggeration over the true discounted cost. The overstatement is significant. Given the long estimated lifetime of appliances, a failure to discount leads to an exaggeration of more than 100% of true costs for many appliance categories when the true discount rate is 10%. In other words, NRC’s estimation more than doubles the true value of an increase in energy efficiency in these cases.

The lesson here is that, even when label information is available, it may not be perfect. And, knowing the annual cost is not in itself a solution to the need for costly search to estimate the correct costs when purchasing a durable good. The Canadian government made a sizable error when doing the calculation itself.

5.2 Summary and caveats for appliances

As was the case for automobiles, an analysis of appliance data indicates that, while fuel cost variation is important, it is substantially smaller than the variation in prices. In addition, energy labels seem inadequate for resolving consumer uncertainty, which leaves open the possibility that consumers will decide that figuring out lifetime energy costs is not worth the trouble.

Missing from the preceding analysis are a variety of appliances, including microwave ovens, televisions, computers, monitors, fax machines, dehumidifiers, and air purifiers. They are missing from the prior calculations because data on their energy consumption is much more difficult to come by. The Energy Star data files contain lists of approved appliances, but neither the DOE nor the NRC provide a list of non-qualified items. Moreover, for most of these goods, the information in the Energy Star label does not include an annual cost, but instead relies on metrics that depend on the intensity of use, which means that consumers have an additional burden of calculating their use in order to estimate the value of increased efficiency. It is these goods that are most likely to be the victims of rational inattention.

In other words, the empirical analysis here has focused on the goods for which it is least likely that consumers will exhibit rational inattention—automobiles and the largest home appliances. Even for these goods, there seems to be room for rational inattention because variation in energy costs is quite modest compared to variation in price, and the barriers to assessing lifetime cost, even when labels are present, is significant. The case is much stronger for the variety of goods that lack any official labels.

6 Conclusion

The analysis provided here leads to several conclusions. First, there is significant evidence of bunching around energy efficiency notches in automobiles, buildings, and home appliances. This bunching is prima facie evidence of inefficient product design distortions from uneven policy incentives. For such policies to be justified, there must be some benefit to discrete labeling or signaling, such as the ability of a single energy efficiency certification label to overcome rational inattention.

Second, the heuristic model developed above shows that such rational inattention will be more likely as the variation in energy costs for models of a given type is small, as the variation in other attributes is large, and as the difficulty in acquiring complete information about energy is large. For both automobiles and home appliances, energy costs are substantial, both in absolute terms and as a fraction of the purchase price. Variation in these energy costs is, however, much smaller than the variation in other attributes, as captured by prices. This suggests a possible role for rational inattention in the market for many goods.

Finally, while many goods have energy efficiency labels that provide information on the operating cost of goods, these labels are not sufficient to resolve all of the cognitive costs around estimating lifetime costs. Labels do not provide lifetime cost estimates, so consumers still must perform a present discounted value calculation involving parameters not given on the labels. In addition, for some goods labels are inaccurate on average, and they do not easily account for heterogeneity.

The analysis here is largely suggestive. It indicates that rational inattention might play an important role, but it does not provide a direct test. Future work would do well, not only to test directly for the role of rational inattention, but also to devote more directed attention to the exact design of the information transmitted to consumers, both by the government and third-party certifiers.

References

- Allcott, Hunt. 2010. "Beliefs and Consumer Choice." Manuscript: New York University.
- Allcott, Hunt and Nathan Wozny. 2010. "Gasoline Prices, Fuel Economy, and the Energy Paradox." Manuscript: Princeton University.
- Anderson, Soren T., Ian W.H. Parry, James M. Sallee, and Carolyn Fischer. 2011. "Automobile Fuel Economy Standards: Impacts, Efficiency and Alternatives." *Review of Environmental Economics and Policy* 5 (1):89–108.
- Austin, David and Terry Dinan. 2005. "Clearing the Air: The Costs and Consequences of Higher CAFE Standards and Increased Gasoline Taxes." *Journal of Environmental Economics and Management* 50 (3):562–582.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. Forthcoming. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* .
- Copeland, Adam, Wendy Dunn, and George Hall. 2011. "Inventories and the automobile market." *RAND Journal of Economics* 42 (1):121–149.
- Dubin, Jeffrey A. and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52 (2):345–362.
- Gallagher, Kelly Sims and Erich J. Muehlegger. 2011. "Giving Green to Get Green? The Effect of Incentives and Ideology on Hybrid Vehicle Adoption." *Journal of Environmental Economics and Management* 61 (1):1–15.
- Greene, David L. 2011. "Uncertainty, Loss Aversion, and Markets for Energy Efficiency." *Energy Economics* 33 (4):608–616.
- Hausman, Jerry A. 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *The Bell Journal of Economics* 10 (1):33–54.
- Howarth, Richard B. and Bo Andersson. 1993. "Market Barriers to Energy Efficiency." *Energy Economics* 15 (4):262–272.
- Howarth, Richard B. and Alan H. Sanstad. 1995. "Discount Rates and Energy Efficiency." *Contemporary Economic Policy* 13 (3):101–109.
- Hutton, R. Bruce and William L. Wilkie. 1980. "Life Cycle Cost: A New Form of Consumer Information." *The Journal of Consumer Research* 6 (4):349–360.
- Jaffe, Adam B. and Robert N. Stavins. 1994. "The Energy-Efficiency Gap." *Energy Policy* 22 (10):804–810.
- Kleven, Henrik Jacobsen and Joel Slemrod. 2009. "A Characteristics Approach to Optimal Taxation and Tax-Driven Product Innovation." Working Paper: London School of Economics.

- Larrick, Richard P. and Jack B. Soll. 2008. "The MPG Illusion." *Science* 320 (5883):1593–1594.
- McNeill, Dennis L. and William L. Wilkie. 1979. "Public Policy and Consumer Information: Impact of the New Energy Labels." *The Journal of Consumer Research* 6 (1):1–11.
- Metcalf, Gilbert E. 1994. "Economics and Rational Conservation Policy." *Energy Policy* 22 (10):819–825.
- Moorthy, Sridhar, Brian T. Ratchford, and Debabrata Talukdar. 1997. "Consumer Information Search Revisited: Theory and Empirical Analysis." *The Journal of Consumer Research* 23 (4):263–277.
- Ratchford, Brian T., Myung-Soo Lee, and Debabrata Talukdar. 2003. "The Impact of the Internet on Information Search for Automobiles." *Journal of Marketing Research* 40 (2):193–209.
- Ratchford, Brian T. and Narasimhan Srinivasan. 1993. "An Empirical Investigation of Returns to Search." *Marketing Science* 12 (1):73–87.
- Sallee, James M. 2011a. "The Surprising Incidence of Tax Credits for the Toyota Prius." *American Economic Journal: Economic Policy* 3 (2):189–219.
- . 2011b. "The Taxation of Fuel Economy." *Tax Policy and The Economy* 25:1–37.
- Sallee, James M. and Joel Slemrod. Forthcoming. "Car Notches: Strategic Automaker Responses to Fuel Economy Policy." *Journal of Public Economics* .
- Sallee, James M., Sarah E. West, and Wei Fan. 2009. "Consumer Valuation of Fuel Economy: A Microdata Approach." *Proceedings of the National Tax Association Annual Conference on Taxation* .
- Turrentine, Thomas S. and Kenneth S. Kurani. 2007. "Car Buyers and Fuel Economy?" *Energy Policy* 35 (2):1213–1223.
- US Department of Transportation. 2008. "Corporate Average Fuel Economy for Model Years 2011-2015 Passenger Cars and Light Trucks." Preliminary Regulatory Impact Analysis.