NBER WORKING PAPER SERIES

RATIONAL INATTENTION AND ENERGY EFFICIENCY

James M. Sallee

Working Paper 19545
http://www.nber.org/papers/w19545

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2013

The author would like to thank David Austin, Hunt Allcott, Lucas Davis, Ashley Langer, and seminar participants at the University of Chicago/Resources for the Future Symposium and at the University of California at Berkeley for helpful comments. The author also would like to thank Greg Sasso, Yi Sun and Samantha Superstine for excellent research assistance, Roberts French for providing Environmental Protection Agency five-cycle test data, Hunt Allcott and Ashley Langer for providing data on second choice vehicles, and Ashley Langer for providing data used in simulation. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Rational Inattention and Energy Efficiency
James M. Sallee
NBER Working Paper No. 19545
October 2013
JEL No. D03,H23,Q48

ABSTRACT

If time and effort are required to accurately ascertain the lifetime value of energy efficiency for a durable good, consumers might rationally ignore energy efficiency. This paper argues that such inattention may be rational in the market for automobiles and home appliances. To do so, it 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 via costly effort. The model indicates under what conditions the consumer will be less likely to undertake this effort. The empirical portion of the paper argues that energy efficiency is often not pivotal to choice. This, along with a simulation of the automobile market, suggests that returns to paying attention to energy may be modest, and analysis of the information readily available to consumers suggests that the costs of being fully informed may be substantial. The paper discusses the implications of rational inattention for public policy and for empirical research on the energy paradox.

James M. Sallee
Harris School of Public Policy Studies
University of Chicago
1155 East 60th Street
Chicago, IL 60637
and NBER
sallee@uchicago.edu

1 Introduction

The purpose of this paper is to elevate consideration of rational inattention in the study of energy economics. The idea behind rational inattention is that when information is costly to acquire, decision makers may sometimes choose to act upon incomplete information, rather than incurring the cost to become perfectly informed. I argue in this paper that it will often be rational for consumers to choose among energy-consuming durables, like automobiles or refrigerators, without acquiring complete information about energy-efficiency.

To explore this possibility, the paper first presents a simple model of a consumer making a discrete choice among energy-consuming durables. The consumer is assumed to have beliefs about the energy consumption of each alternative that are based on readily available information, and therefore require minimal effort. The consumer recognizes that these beliefs may be incorrect, and it is assumed that they can resolve any such uncertainty by exerting costly effort to acquire additional information. The consumer must therefore decide whether to pay the requisite effort cost and then make their choice with the benefit of full information, or whether they would rather make their choice with the incomplete information that comes to them without cost. When consumers act on incomplete information, they make choices equivalent to the choice they would make if they ignored (were inattentive to) some component of fuel costs.

The model draws intuitive conclusions. If the variance across alternatives in the unknown cost of energy is large, then it is more likely that consumers will be attentive to efficiency. Conversely, if the variation in the value of each product is determined primarily by attributes other than energy efficiency, then energy efficiency is unlikely to be pivotal to choice and consumers will be less likely to pay attention to it. Finally (and obviously), consumers will be less likely to pay attention when the effort cost of doing so is high.

The paper then proceeds to an empirical analysis, which takes this model as a guide and makes the case that rational inattention is plausible in the market for automobiles and home appliances. For these goods, the paper documents the fuel cost variation across alternatives and shows that this variation, while substantial, is modest compared to variation in prices. To the extent that prices proxy for variation in the value of all other attributes, these descriptive statistics suggest that energy costs are unlikely to be pivotal in many cases and that consumers may therefore lose little from being inattentive to them.

For automobiles, the empirical analysis goes further by analyzing survey data on first and second choice vehicles in order to quantify the fuel cost variation among the set of vehicles that consumers actively consider, and by using the parameters from a discrete choice model to simulate the average welfare loss that consumers would experience from choosing a vehicle with incomplete information about fuel consumption. The simulation suggests average consumer loses are on the order of one to three hundred dollars per vehicle when consumers make choices with incomplete information about energy.

The paper does not directly estimate the cost of learning about energy efficiency, but it does argue that these costs are likely to be substantial and may plausibly exceed the one to three hundred dollar benchmark from the simulation. Automobiles and some home appliances come with government issued labels, which surely reduce the cost of learning about energy consumption. The paper demonstrates that significant effort costs are required even when labels are available, however, both because the labels are incomplete and sometimes biased, and because heterogeneity in usage patterns implies that labels can resolve only a modest portion of consumer uncertainty.

This framework, by necessity, concludes that inattention can only prevail when the welfare losses from it are small. Nevertheless, small individual welfare losses can still aggregate into large losses when viewed economy wide. For example, the $100 per vehicle loss from inattention estimated in the simulation is modest compared to other factors related to the purchase of a new car, but even so it amounts to $1.5 billion per year in the U.S. car market.

One reason that rational inattention is of interest to energy economics is its bearing on unresolved questions regarding the energy paradox. The energy paradox, or the energy-efficiency gap, is the hypothesis that market actors systematically fail to adopt energy efficiency improvements that are privately beneficial. Suspicion that an energy paradox exists comes in large part from the observation that many energy-saving technologies—like compact-fluorescent lightbulbs, improved building insulation and energy-efficient home appliances—have low adoption rates despite being, according to engineering estimates, cost-effective. Economists have tested for an energy paradox using revealed preference data in a variety of ways. Key early contributions were Hausman (1979) and Dubin and McFadden (1984), who analyzed cross-sectional data on home appliances to test for excessively high implied discount rates. Recent advances that use more credible panel data have been concentrated in the automobile sector, where several papers test whether new and used car prices correlate with gasoline price changes in the amount predicted by a theory in which consumers fully understand and value fuel economy (Allcott and Wozny Forthcoming; Busse, Knittel, and Zettelmeyer 2013; Sallee, West, and Fan 2009).

Despite a flurry of studies and significant improvements in methodology, the question of whether or not a paradox exists is best characterized as unsettled. Greene (2011) reviews 25 studies that test for an energy paradox in the automobile market and finds that half of them indicate that consumers undervalue fuel economy. Allcott and Greenstone (2012) review the wider literature on the energy efficiency gap and conclude that there is little evidence of a widespread paradox, but that more and better research is still needed because "this body of evidence frequently does not meet modern standards of credibility."

Rational inattention poses a key challenge to this empirical literature. Revealed preference analysis necessarily evaluates consumer choice among the products that are on offer in the marketplace. It is generally assumed that if consumers fully value energy efficiency among current products, then firms will bring to market all cost-effective innovations. Suppose, however, that consumers are attentive to some types of innovations and not others. In equilibrium, firms will bring to market only those innovations that garner attention, and consumers may choose rationally among the products on offer. To the econometrician, consumers will appear to fully value energy efficiency, but a paradox nevertheless exists because cost-effective innovations that do not garner attention are not brought to market. For example, rational inattention implies that consumers may rationally value energy efficiency across broad categories (such as vehicle class), but they may ignore small innovations (such as more efficient transmission types). Some studies, such as Busse, Knittel, and Zettelmeyer (2013), test for an energy paradox by comparing prices across broad categories. Such tests are ill equipped to detect rational inattention.

Whether or not an energy paradox exists is important for public policy. If a paradox does exist, it will influence the optimal design of public policies aimed at correcting energy-related externalities. For example, traditional economic analysis consistently finds that a gasoline tax is a superior way to reduce gasoline consumption than fuel economy standards and that fuel economy standards often fail cost-benefit analyses (see Austin and Dinan (2005) for a key example and Anderson, Parry, Sallee, and Fischer (2011) for a review of the literature). When performing their regulatory impact analyses, however, U.S. federal agencies assume a substantial energy paradox, which implies large private gains to consumers that are pivotal in allowing standards to pass the cost-benefit test (NHTSA 2010). More generally, an emerging literature on policy design has shown that the optimal externality-correcting policies are indeed altered by the existence of a paradox (Allcott, Mullainathan, and Taubinsky Forthcoming; Fischer, Harrington, and Parry 2007; Heutel 2011).

The policy implications of an energy paradox that is due to rational inattention, however, may be distinct in several ways from a paradox due to the more commonly discussed sources of present bias (myopia) or biased beliefs. If effort costs of attention represent real welfare losses, then forcing consumers to pay attention is not necessarily welfare improving. Instead, policy might focus on lowering the barriers to information processing (such as improving labels) or determining the optimal coarseness of information presented to consumers. For example, inattention might justify the use of coarse information or "notches", such as whether or not a product qualifies for an Energy Star label, alongside, or even in place of, fine-grained information. These possibilities are discussed further at the end of the paper.

This paper relates to a large literature on the energy paradox, much of which is cited above. It also relates to the broader literature on the economics of information begun by Stigler (1961), as well as to the recent theoretical literature on inattention, such as Gabaix (2013). The connection between this paper and those, as well as to the few papers that directly connect inattention to energy issues—including Howarth and Andersson (1993) and Houde (2012)—are described in greater detail after the model has been presented.

The remainder of the paper is structured as follows. Section 2 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 3 considers empirical evidence relevant from the automobile market. Section 4 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 5 discusses the implications of rational inattention for policy design, and 6 concludes.

2 Model

This paper is about a consumer who must make a discrete choice among durable goods, but must also decide whether or not it is worth paying a cost to acquire additional information before making that choice. Before proceeding to the mathematical characterization of the problem, it is useful to describe it in more general terms.

Take the example of a consumer purchasing an automobile. The automobile's color, size, body type, passenger capacity, brand and sticker price are all readily apparent to a shopper. Other features, such as the interior space, cargo volume and comfort can be determined relatively easily upon inspection. The automobile's drivability, turning radius, acceleration and performance can be, at least in good part, determined by a test drive. All of these features can be felt or experienced by the consumer, and a consumer can intuitively determine how much they are willing to pay for each feature in much the same way that they can determine their willingness to pay for produce at a grocery store. In contrast, the consumer has no way of seeing or experiencing the automobile's lifetime cost of refueling, and consequently he or she may struggle to value fuel economy.

Consumers do have some information related to fuel costs. Most consumers likely understand that bigger, heavier, more powerful cars have worse fuel economy. Mandatory labels provide the government sanctioned city and highway fuel economy ratings of every car. Thus, with relative ease, consumers can form some belief regarding the vehicle's fuel economy and its consequent impact on the cost of ownership, but this expectation will come with uncertainty. Uncertainty exists both because the true fuel economy may deviate from its expected value (this is discussed in detail below) and because, for valuation, fuel economy must be translated into lifetime operating costs. Thus, after a quick inspection of a car, the consumer may form a belief about fuel costs, but they will also have uncertainty regarding that belief. The consumer can choose a car given this easily formulated belief, or the consumer can seek additional information before making a selection.

What information might they obtain? The value of an automobile to a consumer should reflect its lifetime operating costs, which is often written in the following form:

$$\text{Present-discounted value of lifetime operating costs} = \sum_{t=0}^{T} \delta^t \frac{P_{gt} m_t}{mpg_t},$$

where t indexes time period with $t = 0$ being today, T is the lifespan of the vehicle, δ is the discount factor, P_{gt} is the price of gasoline per gallon, m_t is the number of miles driven, and mpg_t is the vehicle's fuel economy.

Ideally consumers would know this full fuel cost equation and take it under consideration when choosing a car, but they face two challenges in doing so. First, the present-discounted calculation is itself cognitively difficult. Second, none of the components of the calculation are immediately available to the consumer. A vehicle's life expectancy is uncertain. Most consumers do not know how much they drive today, and will necessarily be uncertain about their driving in the future. Consumers may know the price of gasoline today, but they will be uncertain about future prices. They may or may not have a solid grasp on the appropriate discount rate for the calculation. Driving style and conditions will affect the fuel economy that they obtain when driving a car, and fuel economy degrades over the life of a vehicle. Thus, even if label fuel economy is exactly right for the average driver, uncertainty remains. (The size of this uncertainty is estimated below.)

How might consumers gain an estimate of fuel costs that is more accurate than their quick impression? Firstly, consumers can attempt to perform the present-discounted calculation, rather than relying on an intuitive guess. If they do not know how to do the calculation, they could seek help. Secondly, consumers could research the unknowns. They could research the reported fuel economy of a particular model to see if it deviates from the label. They could keep track of their own driving behavior more closely to determine their mileage. To calibrate the impact of their driving styles and patterns, they could try to figure out how their experienced fuel economy deviates from the label rating for their current vehicle, or attempt to separate out their city versus highway mileage and recalculate costs based on separate ratings. All of these actions would improve the information available for their decision, but such actions require time and effort, which represent real costs. And these costs are real because consumers do not already know how to do this calculation; qualitative interviews of new car buyers show that consumers cannot readily articulate any of the building blocks necessary for this calculation, save the current price of gasoline (Turrentine and Kurani 2007).

Under what conditions will a rational consumer select a vehicle without learning additional details about fuel consumption? What are the consequences for the automobile market if they do? The next section develops a heuristic mathematical model to answer these questions.

2.1 Mathematical 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 product. A risk-neutral consumer i is to choose a durable j out of a set of durables $\mathcal{J}$. The consumer will choose exactly one model.

The choice problem analyzed is standard, except that consumers are assumed to have imperfect information about each product's true fuel cost c_j^*. Consumers instead observe $\bar{c}_j$, where $c_j^* = \bar{c}_j + c_j$, so that c_j is an unobserved random error. The focus of the analysis here is on uncertainty rather than biased beliefs, so it is natural to assume that $\bar{c}_j$ is an unbiased estimate of the truth ($\mathbb{E}[\bar{c}_j] = c_j^*$). The belief $\bar{c}_j$ should be interpreted as the best guess regarding fuel costs for product j given all of the easily ascertained information about the product, including label information and all attributes that enter the utility function directly.[1] The consumer is assumed to know the distribution of c_j but not its value. This leads to the following formulation for the utility of consumer i over good j:

$$U_{ij} = \beta' \mathbf{X}_j - p_j - (\bar{c}_j + c_j) + \varepsilon_{ij} \tag{1}$$

$$\equiv V_j - c_j + \varepsilon_{ij} \equiv W_{ij} - c_j, \tag{2}$$

where U_{ij} is utility, $\mathbf{X}_j$ is a set of observed attributes, p_j is the price of the vehicle, and ε_{ij} is a random error term. V_j is defined as the representative utility of good j, without the random error component of energy costs, and W_{ij} is the idiosyncratic valuation of vehicle j absent the uncertain cost component.[2] Use of V and W will simplify notation later. All consumers are assumed to have a common valuation of attributes, prices and costs—only ε_{ij} varies across individuals. Heterogeneity in these dimensions could be allowed in the standard way without changing the analysis below, but at the cost of notational complexity.

Before choosing a product, the consumer first decides whether or not to exert effort, defined as $e \in \{0, 1\}$. If they exert effort ($e = 1$), then the vector c_j is revealed but they pay an effort cost s. If they choose not to exert effort ($e = 0$), then no cost is borne but they must choose a product j without knowing the fuel cost errors c_j. From the consumer's point of view, c_j is an unknown error term. Exerting effort reveals these error terms.

If the consumer exerts effort and gains full information, then they will simply choose the product with the highest utility (choose j if and only if $U_{ij} \geq U_{ik}\ \forall\ k \neq j \in \mathcal{J}$). Under the assumption that $\bar{c}_j$ are unbiased estimates of the truth, the consumer will choose the product with the highest W_{ij} if they choose not to acquire full information (choose j if and only if $W_{ij} \geq W_{ik}\ \forall\ k \neq j \in \mathcal{J}$). Thus, a consumer will exert effort and learn true costs if and only if the expected increase in utility from making a fully informed choice exceeds the cost s. It is equivalent to say that the consumer will exert effort if and only if the expected welfare loss from making the choice with incomplete information is larger than the costs of search. This latter formulation can be written as, consumer i will choose $e = 1$ if and only if:

$$\mathbb{E}[\max_j\ U_{ij}] - \max_j\ \mathbb{E}[W_{ij}] \geq s. \tag{3}$$

When will consumers choose rational inattention? The answer is easiest to characterize under some assumption about the distribution of the error term. The solution is particularly elegant under the assumption that c_j is a type-I extreme value random variable with zero mean and scale parameter σ (which implies a variance of $\pi^2/(6\sigma^2)$).[3] This makes the problem directly analogous to a standard logit, where now the uninformed consumer occupies the econometrician's standard position, in that they can calculate choice probabilities but do not know the error terms. Formulas familiar from the logit problem carry over, and expression 3, after rearrangement, can be rewritten

[1] In this section, I assume that all attributes in $\mathbf{X}_j$ are included in the forecast because otherwise it is illogical to assume that the error terms are independent of utility, and independence facilitates exposition. Sections 2.3 and 2.4 explore alternative structures that relax this assumption.

[2] Recall that c^*_j is the present-discounted lifetime fuel cost of the vehicle. The absence of a coefficient in front of p_j and $\bar{c}_j + c_j$ in the utility function implies full valuation—the consumer trades-off current price and lifetime fuel costs one for one. In that sense, the model assumes rationality of the type usually analyzed in the energy paradox literature, which typically tests for an equal weight of fuel costs and prices on choice. Any mistakes made in the model result from the rational decision to make the durable choice with less than full information—that is, from rational inattention.

[3] An earlier version of this paper assumed a normal distribution. All of the qualitative implications for that case are identical to those described here. Relaxation of the independence assumption is discussed below in section 2.4.

as, the consumer will choose $e = 1$ if and only if:

$$\text{Expected Welfare Loss} = \frac{1}{\sigma} \ln \left(\frac{\sum_{j \in \mathcal{J}} e^{\sigma W_{ij}}}{e^{\sigma W_{ik}}} \right) > s. \tag{4}$$

This decision rule has an intuitive interpretation, and it makes the important results of the model immediately apparent. The term inside of the parentheses is one over the probability that the consumer will choose good k under full information.[4] Choice k is the product that the consumer will choose when they are imperfectly informed. When that probability is high, the consumer is less likely to make a mistake when inattentive, and thus effort becomes less likely. As the probability rises that good k is the correct choice, the inverse of that probability will fall; its log will fall; and effort becomes less likely. The logged term approaches zero as the probability that k is the correct choice goes to one, which makes effort less beneficial.

Ceteris paribus, increasing W_{ik}, the value of the perceived top choice k, will raise the probability product k is the correct choice. If choice k is far better than the next best alternative given the information that is available for free, then effort will have little benefit because it is very unlikely that revelations about c_j will change the consumer's choice. In contrast, if choice k is similar in observable attributes to other choices (W_{ik} is close to other W_{ij}), then the consumer is likely to change his or her mind after observing c_j. More generally, as products become "more different," fewer consumers will exert effort.

As uncertainty about fuel costs rises, the expected welfare loss from inattention will rise and effort will become more likely. Recall that the variance of the unknown component c_j is inversely related to the scale parameter σ. As σ rises, the variance of uncertainty falls and the expected welfare loss falls. It is straightforward to show this by differentiating expression 4.[5] An increase in the variance of c_j increases the returns to acquiring information because the greater is the variance in unknown energy costs, the more likely the consumer is to change his or her mind about which product is optimal once information is revealed, and, conditional on changing from choice k to some other product, the larger will be the average utility gain. When the variance of c_j is high, the uninformed consumer will make more, and bigger, mistakes. This increases the return to effort.

Expression 4 also shows clearly (and obviously) that effort is less likely as the cost of effort rises. When s gets bigger, the inequality will be harder to satisfy. Note also that the effort cost bounds the expected welfare losses of inattention. When an energy paradox is due to inattention, the losses from it on average cannot exceed the cost of paying attention. This framework shows that this welfare loss can be written in terms of a standard logit model, which is a fact that may aid future empirical work.

These intuitive results constitute the main guidance for empirical work provided by the model. Consumers are more likely to be rationally inattentive if (a) effort costs are high, (b) the variance of unknown energy costs are low, and (c) products are very different, so that consumers are far from indifferent between their first choice and its alternatives. This last point implies either that the attributes of goods are quite different or that the random utility component of preferences is large. The empirical portion of the paper explores these different parameters for the cases of automobiles and home appliances. Before proceeding to the empirical analysis, however, the next several subsections describe the model's implications for producer behavior, how the model could

[4] The term in parentheses is precisely the inverse of the probability that k is chosen in a standard logit model, for any choice k.

[5] This derivative is more easily shown after multiplying both sides by σ, exponentiating and grouping terms on one side of the inequality, which must be less than zero if $e = 1$. Then, the derivative with respect to σ is $e^{\sigma W_{ik}} \frac{\sum (W_{ik} - W_{ij}) e^{\sigma W_{ij}}}{(\sum e^{\sigma W_{ij}})^2} + s e^{-\sigma s}$. Because $W_{ik} \geq W_{ij}$, both terms are positive.

be recast as a model about heterogeneity or noisy signals, the model's implications for empirical research, and its relation to existing work.

2.2 Inattention and producer behavior

What does rational inattention imply for firm behavior? Suppose that the firm that produces j (call it firm j) can implement an innovation that marginally lowers fuel costs c_j at a production cost below the informed consumer's willingness to pay. Will it adopt this improvement and increase price p_j to recoup the costs? Inattentive consumers will recognize only the price increase p_j, but not the commensurate reduction in fuel costs c_j. Thus, an innovation that improves the true value of the product will be perceived as a change that lowers value for inattentive consumers. This may cause producers to fail to bring cost-efficient innovations to market. This intuition holds, however, only in certain cases; if the *marginal* consumer is attentive, there can be rational inattention among other consumers without causing any market inefficiencies.

To see why, consider the choice of a consumer in the two good case, where there are only products j and k. Define the difference in idiosyncratic taste terms as $\Delta\varepsilon \equiv \varepsilon_{ij} - \varepsilon_{ik}$. For a given value of V_j and V_k, when $\Delta\varepsilon$ is large enough, consumer i will choose good j without exerting effort. Denote by θ^* this cutoff value of $\Delta\varepsilon$.[6] Conversely, if $\Delta\varepsilon$ is smaller than $-\theta^*$, then consumer i will choose good k without exerting effort.

The true cost difference $c_j - c_k$ determines who the marginal consumer is under full information; it is the consumer with $\Delta\varepsilon_i = c_j - c_k$. If that marginal consumer is attentive, then no inefficiency need arise. This happens when $c_j - c_k$ is between $-\theta^*$ and θ^*. It is the marginal consumer who would be influenced by the firm's decision to improve fuel economy and raise price, so if they are fully informed and firms know that, there is no impetus for underprovision of a marginal innovation.

This is illustrated in Figure 1, which plots the distribution of ε for the case where $V_j = V_k$.[7] In the figure, individuals with $\Delta\varepsilon_i < -\theta^*$ will choose model k without effort; those with $\Delta\varepsilon_i > \theta^*$ will choose model j without effort; and those in the middle with exert effort. Of those who exert effort, some will choose j and some k after learning true costs, which is shown to be point $A = c_j - c_k$. In this example, 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. If the firm knows where the marginal consumer lies, both firms j and k recognize that marginal improvements (such that c_j falls by more than p_j rises) will increase demand, and hence there is no reason for underprovision.

If, however, the true realization of c_j and c_k implies that the marginal consumer under full information is rationally inattentive in equilibrium, then underprovision may occur. In Figure 1, 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. The demand for model k is equal to the probability that $\Delta\varepsilon < \theta^*$. Raising price p_k will shift $-\theta^*$, lowering demand for k. But, lowering c_k will create no offsetting increase in demand because only consumers who exert effort recognize the change in c_k, but none of those consumers are marginal—they will choose j even if there is a marginal improvement in k. Similarly, firm j would not make the innovation because raising p_j will shift $-\theta^*$, lowering demand for j. But, the corresponding drop in c_j will again only be recognized by consumers who purchase j, with or without a marginal improvement. Thus, neither firm would have the proper incentive to adopt a cost-efficient innovation.

[6] Specifically, i will choose j and $e = 0$ iff $V_j - V_k + \Delta\varepsilon > -\ln(e^{\sigma s} - 1)/\sigma \equiv \theta^*$. Note that, when some consumers search in equilibrium, $\ln(e^{\sigma s} - 1) < 0$ so $\theta^* > 0$.

[7] If $V_j \neq V_k$, then the distribution will simply be shifted to the left or right, holding constant the values of θ.

Figure 1: Effort and demand across consumer types

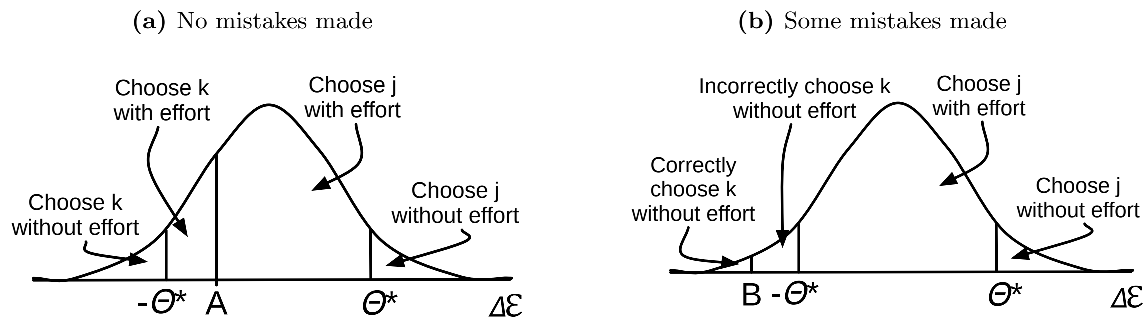

Figure shows the distribution of effort and demand across the $\Delta\varepsilon$ distribution, assuming that $V_j = V_k$ for the two good case. The actual realization of $c_j - c_k$ determines the choice cutoff for fully informed consumers. A and B in the diagrams denote this cutoff for two different realizations of $c_j - c_k$.

This model 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 an equilibrium with inattentive marginal consumers, 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. Moreover, for products with many attributes about which consumers have uncertainty, consumers will not readily know whether a higher price implies greater efficiency or an improvement in some other characteristic.

The model is also not complete in that it does not consider directly the supply of information to consumers. If there are cost-effective ways of giving consumers greater information about energy efficiency, producers (or other third-party actors) might choose to provide information that goes beyond the government's mandatory labels or the information that consumer choose to collect. It is intuitive to expect firms to provide additional information that limits the prospect for rational inattention, but as I discuss in section , this may not be feasible in some cases, especially when uncertainty stems from idiosyncrasies of consumers. In addition, there are legal restrictions on claims that producers can make regarding information in many markets—automakers, for example, are not allowed to advertise a fuel economy that exceeds that determined by the federal test. Finally, Gabaix and Laibson (2006) show that, counter to basic economic intuition, competitive forces may not force "shrouded attributes" (like unknown energy efficiency) to be revealed honestly by sellers, but rather may cause firms to intentionally obscure consumer information.

2.3 Alternative interpretations of the model

The model was cast as one in which the consumer has unbiased information about the true energy consumption of the product. This could easily be recast as a model in which the consumer has a prior belief about the product's energy consumption (formed as a function of the product's attributes perhaps) and they receive a noisy signal about energy consumption. Then, the belief $\bar{c}_j$ that appears in the model is the mean of the posterior distribution and the variance of c_j is the variance of the posterior distribution.[8]

[8] In particular, if c_j is assumed to have a normal distribution (as opposed to a type-I extreme value) then the mapping is very natural. Let the prior $p \sim \mathcal{N}(\bar{p}, \rho^2)$ and the signal $d \sim \mathcal{N}(0, \delta^2)$. Then the posterior belief $\bar{c}$ is also

The model can also be recast as a model about heterogeneity. One reason that consumers may have uncertainty about fuel consumption is that their own fuel consumption may differ from that of the average consumer. Specifically, suppose that consumers know with certainty the true fuel costs of each product for the *average* consumer $\bar{c}_j$ (such as from a government label). But, consumers differ in utilitization, so their own fuel costs are scaled functions of this mean, $\gamma_i^* \bar{c}_j$. In the case of automobiles, an example would be if consumer i drives $\gamma_i^* \bar{m}$ miles, where $\bar{m}$ is the mean miles driven, their fuel costs will be $\gamma_i^* \bar{c}_j$. If consumers are certain of $\bar{c}_j$, but uncertain about γ_i^*, then the model is very similar, and utility can be written as:

$$U_{ij} = \beta' \mathbf{X}_j - p_j - (\gamma_i + u_i)\bar{c}_j + \varepsilon_{ij} \equiv Y_{ij} - u_i \bar{c}_j. \tag{5}$$

The assumption parallel to the model above would be to assume that they have an unbiased belief about their heterogeneity γ_i, such that $\gamma_i^* = \gamma_i + u_i$ and $\mathbb{E}[\gamma_i] = \gamma_i^*$. This is exactly the same as the original setup in equation 1, except that the error term c_j from the original problem is now $u_i \bar{c}_j$. Now, errors across products cannot be independent from each other for consumer i, because they are the product of each good's average fuel consumption and the consumer's error, which is common across products. This dependence in the error structure makes the mathematics less elegant, but the main intuition goes through.

To see this, consider the case of consumer i choosing between only two goods j and k. Let k be the good with the higher utility given incomplete information ($Y_{ik} > Y_{ij}$). Then, the expected payoff to choosing k without exerting effort is

$$\int_{-\infty}^{\infty} (Y_{ik} - u_i \bar{c}_k) dF(u_i), \tag{6}$$

where $F(\cdot)$ is the distribution function of u_i. Define u_i^* as the value at which the consumer is indifferent between goods j and k, that is, $u_i^* = (Y_{ik} - Y_{ij})/(\bar{c}_k - \bar{c}_j)$.[9] Suppose that $\bar{c}_k > \bar{c}_j$ (analogous math follows for the opposite case). Then, under full information the consumer will choose k so long as $u_i < u_i^*$, but will otherwise choose j. In this case, the expected payoff under full information is:

$$\int_{-\infty}^{u_i^*} (Y_{ik} - u_i \bar{c}_k) dF(u_i) + \int_{u_i^*}^{\infty} (Y_{ij} - u_i \bar{c}_j) dF(u_i) - s. \tag{7}$$

Taking the difference between equations 6 and 7 and rearranging yields the following condition. Consumer i will exert effort if and only if:

$$\int_{u_i^*}^{\infty} u_i(\bar{c}_k - \bar{c}_j) dF(u_i) > (1 - F(u_i^*))(Y_{ik} - Y_{ij}) + s. \tag{8}$$

The left-hand side of expression 8 is the expected gain from switching to product j, conditional on wanting to do so under full information. The right-hand side is the probability that the consumer chooses k even under full information times the difference in utility between k and j under incomplete information, plus the cost of effort.

normally distributed, $\bar{c} \sim \mathcal{N}(\bar{p} + (d - \bar{p}) * \rho^2/(\rho^2 + \delta^2), (\rho^2 \delta^2)/(\rho^2 + \delta^2))$. If the posterior errors are type-I extreme value, then a different prior distribution would be required for the noisy signal interpretation to map perfectly into the model with the parametric assumptions emphasized above.

[9] Note that it makes sense to impose a restriction on the support of u_i^* so that total fuel consumption cannot be negative. With this assumption imposed, there will be values of $Y_{ik}, Y_{ij}, \bar{c}_j$ and $\bar{c}_k$ for which no realization of u_i could cause the consumer to change their mind. In such cases, the consumer will obviously not exert effort.

Straightforward analysis of equation 8 shows that the comparative statics have the same signs as those in the baseline model. As s rises, the probability of effort falls. As Y_{ik}, the desirability of product k relative to j, rises, the probability of effort falls. And, if u_i has a normal distribution, then an increase in variance will raise the probability of effort because it will raise the truncated mean of u_i above u_i^*.[10] Thus, the basic insights of the model apply when the source of uncertainty is individual heterogeneity. This is important because uncertainty from this source is especially difficult for policy makers or third parties to resolve, as detailed in section 3.4.2.

2.4 Alternative belief structures

The derivation above assumed that the consumer forms an unbiased belief, $\bar{c}_j$, which takes into account the correlation between fuel consumption and all other attributes that enter the utility function, $\mathbf{X}_j$. If beliefs are biased or they do not take all attributes into account, then the error terms c_j will not be independent across products.

For many products, an estimate of energy efficiency that takes all other attributes into account will be quite accurate. For example, as shown in section 3.3, a limited number of automobile attributes account for over 80% of the variation in fuel economy across the new car market. It makes sense to suppose that many consumers will not have such a complete belief system. Instead, they may form a belief $\bar{c}_j$ taking into account a limited set of product attributes. Less sophisticated belief formation will necessarily raise the variance of the unknown component of fuel costs, but the effect of this increased variance on rational inattention will be mitigated by the degree to which consumers are choosing between alternatives with similar attributes because those products will have correlated errors.

To see this, consider the simplest case where a product has only a single attribute X_j, which determines utility and is correlated with fuel consumption. Specifically, let the true fuel consumption of product j be $c_j^* = \alpha X_j + c_j$, where c_j is mean zero and independent of X_j. Then, the unbiased belief regarding fuel consumption that takes X_j into account would be $\bar{c}_j = \alpha X_j$. The belief of a consumer who ignored the attribute X_j of each model, but was nonetheless unbiased for the entire set of vehicles would be $\hat{c}_j = \alpha \bar{X}$, where $\bar{X}$ is the mean of the attribute among the set of alternatives.

In this case, the true utility of good j is:

$$U_{ij} = \beta X_j - p_j - (\alpha X_j + c_j) + \varepsilon_{ij},$$

whereas the perceived utility is:

$$\beta X_j - p_j - \hat{c}_j + \varepsilon_{ij} = \beta X_j - p_j - \alpha \bar{X} + \varepsilon_{ij}.$$

The difference between true utility and perceived utility is then $\alpha(X_j - \bar{X}) + c_j$, whereas this difference is simply c_j when the belief accurately takes X_j into account. The misperception about the value of good k will likewise be $\alpha(X_k - \bar{X}) + c_k$, and the difference in errors across j and k is therefore $\alpha(X_j - X_k) + (c_j - c_k)$. Thus, if a consumer's first and second choice products have very similar (or identical) values of X, then ignoring the role of X in determining fuel costs will have a small (or zero) effect on the consumer's rank ordering of the goods, which is all that matters for choice. That is, as β gets larger compared to α, failing to account for X_j in fuel cost beliefs will have less impact on choice.

[10]More generally, any transformation that increases the truncated mean of u_i above u_i^* will have this effect, which will be true of a broad class of variance increasing transformations.

For example, suppose that a consumer has a strong preference for large vehicles, but they ignore the role of weight in determining their beliefs about fuel consumption. They will underestimate the fuel costs of large vehicles, but their mistakes will be similar across vehicles of similar sizes, so their rank ordering among large vehicles may not change very much. Such misinformation may raise the effort cost of becoming informed without necessarily raising the benefits of information.

2.5 Implications for empirical research

Rational inattention has three important implications for empirical research on the energy paradox. First, it is possible for consumers to make on average rational choices given the products available in the market, and yet an energy paradox exists because firms decline to make cost-effective innovations when the marginal consumer is inattentive. If firms only bring innovations to market when they are likely to garner attention, then the natural outcome will be that consumers will do a reasonably good job of choosing among products that actually exist, but a paradox (and hence social inefficiency) still exists.

By necessity, all revealed-preference empirical work on this topic examines how consumers choose among the products actually in the marketplace. This is true for research on appliances, as in Hausman (1979) or Dubin and McFadden (1984), and automobiles, as in Allcott and Wozny (Forthcoming), Busse, Knittel, and Zettelmeyer (2013), Goldberg (1998), Kilian and Sims (2006), Sallee, West, and Fan (2009), and Sawhill (2008). The idea in this literature is to test how consumers trade-off fuel costs and up-front purchase prices, among available products. This approach will detect an energy paradox that is due to myopia or systematic mis-valuation, but it might not detect a paradox due to inattention.

Much of the consternation regarding the energy paradox owes to a dissonance between "bottom up" engineering estimates and these types of revealed preference tests. Engineering studies tend to conclude that there are a host of innovations that appear to be cost effective that have not been brought to market. Often these innovations are small in nature. For example, in the automobile market, NRC (2011) concludes that the cost-effective technology improvements that appear to pay for themselves but are left undone are small features that are not easily visible to the consumer, like the use of low-friction lubricants, engine friction reduction, and variable valve timing. In contrast, visible changes like the deployment of advanced diesel engines, hybrid technologies or even turbocharging are far less cost-effective. This is typical of engineering cost studies, which generally find that a host of small cost-effective features are left undone by the market (see also NRC (2002) and NHTSA (2010)).

The tension in the literature has been to square analyses of this type with credible revealed preference studies, like Allcott and Wozny (Forthcoming) and Busse, Knittel, and Zettelmeyer (2013), which find that consumers fully (or nearly fully) value fuel economy. To the extent that engineering estimates point to cost-plus innovations that are likely to be invisible to the inattentive consumer, rational inattention can potentially reconcile these two sets of facts.

A second, closely related, lesson for empirical research is that the level of the comparison matters when there is rational inattention. That is, under rational inattention, consumers may correctly recognize and properly value fuel cost differences across broad product categories, but they may not value small differences correctly. In the automobile market, this might mean that consumers correctly value the fuel cost differences across SUVs and compact cars, but they do not understand the differences in fuel costs across two small compact cars. If true, then empirical work that focuses on broad category differences will be less likely to find evidence of an energy paradox.

This is relevant for interpreting recent work in the automobile sector. Busse, Knittel, and Zettelmeyer (2013), for example, relates variation in gasoline prices to variation in average used car prices across fuel economy *quartiles*. That is, they measure how a change in gasoline prices affects the relative price of the 25% of vehicles with the worst fuel economy relative to the 25% of vehicles with the best fuel economy. They conclude that price variation is consistent with full valuation, which is interpreted as evidence against an energy paradox. Their results do not rule out a paradox due to inattention, however, so long as consumers can sort vehicles into quartiles when forming their low-cost beliefs $\bar{c}_j$.

Third, in the presence of rational inattention, empirical studies that use actual fuel consumption in choice models are biased by measurement error. Rationally inattentive consumers use $\bar{c}_j$ to make choices, but the econometrician typically uses c_j^*. Econometric estimates will thus be biased by mismeasurement. When the classical measurement error model holds, the bias will attenuate the coefficient on fuel costs, which drives the estimate towards zero. This coefficient is typically the coefficient analyzed in order to test for the presence of an energy paradox.

Bias need not follow the classical measurement error model in the context of discrete choice models, however. Such mismeasurement will still likely bias coefficients, as evidenced by the analogy between this problem and the one analyzed by Bento, Li, and Roth (2012), who show that heterogeneity in the valuation of fuel economy leads to bias in discrete choice models. Their simulations suggest that such bias will be downwards and could easily be of sufficient magnitude to explain some of the findings in the literature of undervaluation. Rational inattention could have a similar effect on coefficient estimates.

2.6 Relationship to existing literature

The model developed here relates to a growing literature on inattention, as well as to an older tradition on search decisions. The model has roots in the optimal search literature begun by Stigler (1961), in that information is costly and the consumer must decide how much information is worth obtaining before making a purchase. It is also closely related to recent theories of inattention, including Bordalo, Gennaioli, and Shleifer (2013), Gabaix (2013), and Gabaix and Laibson (2005). The model of Gabaix (2013) is particularly relevant. That model is a much more general statement of the consumer choice problem, where a consumer decides how much to buy of many different goods. Gabaix (2013) develops a sparsity-based decision rule in which consumers have a prior over the price of each good and they know the price variation. Then, they decide how much costly attention to allocate towards learning the price of each good with precision. The fundamental insight is that, in this setting, consumers will pay more attention to goods that have more price variation (and which they buy more of). This result is similar to the result obtained here regarding the importance of variance in garnering attention.

The inattention model in this paper differs from those discussed above by focusing on a discrete choice problem. In that regard the model is similar to that of Matejka and McKay (2013), which explores how uncertainty in a discrete choice model can rationalize the logit framework. Finally, a strain of research in macroeconomics uses models of costly information acquisition and rational inattention to solve dynamic consumption problems (Reis (2006) and Sims (2003) are key examples).

In the literature on energy efficiency, 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) discuss the demand for energy efficiency when information acquisition requires an effort cost. Their model does not develop a discrete choice model, nor does it allow heterogeneity in the preference for different models, so the overlap in insights with this paper are modest.

The work most similar in spirit to this paper is a working paper, Houde (2012). Houde (2012) presents a model of costly search for information about energy efficiency in a consumer's discrete choice among refrigerators. The model in that paper has some technical differences from the one presented here, but it captures many of the same insights. Houde (2012) then estimates the model using data on refrigerators and concludes that a substantial fraction of consumers choose not to search in equilibrium and subsequently choose a refrigerator as if they are indifferent to, or ignorant of, energy efficiency. That paper and this one are complementary: whereas this paper explores *why* inattention may be rational for several types of goods, but does not estimate a demand model, Houde (2012) estimates a structural demand model for a single good, but does not delve empirically into why inattention might be rational.

3 Automobile energy costs

The next two sections use the model as a guide for assessing the plausibility of rational inattention in the marketplace. I begin by considering automobiles here, and then discuss home appliances in section 4. The theoretical framework tells us that rational inattention is more likely when the variance of the unknown component of fuel costs c_j is small. But in the absence of data on how consumers form beliefs, what I do below (in section 3.1) is document the variation in fuel costs across all cars and within vehicle class, which I argue gives a conservative estimate of how much variation inattentive consumer beliefs would capture. Of course, what really matters is how fuel costs vary within the set of vehicles that a consumer might choose between. To better measure that object, I analyze data from surveys that ask consumers about their first and second choice vehicles in section 3.2.

The model indicates that it is not only variation in fuel costs that determines rational inattention, but also the degree to which a consumer's first choice vehicle is superior in all attributes to other vehicles. To proxy for variation in all other attributes, I document variation in prices and compare this to variation in fuel costs in section 3.1. Both aspects of the problem are combined in section 3.3, where I use the parameters from a discrete choice model to simulate the choice that consumers would make under full information and under rational inattention and quantify the welfare loss from inattention. It is shown that welfare loss is modest.

The model indicates also that rational inattention is more likely when effort costs are high. Lacking any quantitative estimates of effort costs, instead I argue in qualitative terms that effort costs are likely to be substantial, both because readily available information on fuel economy labels is biased and incomplete and because heterogeneity across consumers makes resolving uncertainty difficult.

Finally, in section 3.4 I also document that price variation *within* model is large compared to fuel cost uncertainty. This means that if a consumer faced an effort budget, they might do better to exert effort towards getting the best possible price on their most preferred model, rather than exerting effort in learning about fuel costs *across* models.

3.1 Fuel cost and transaction price variation

To quantify fuel cost variation, I estimate the distribution of lifetime fuel costs implied by transaction data from a large random sample of new vehicle purchases for model year 2006 vehicles. I compare this variation to price variation, which both serves to contextualize the fuel cost magnitudes and acts as a proxy for variation in all other attributes. 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

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
Midsize 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.

subsidies calculated relative to the Federal Reserve's survey of 48-month car loan interest rates from commercial banks.[11]

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.

Table 1 shows the average actual transaction price, net of incentives, and the average fuel cost across vehicle class, as well as their standard deviations. The relevant object for determining inattention is the variation in fuel costs that is unknown to the consumer. This is determined by the information that consumers use to form their beliefs ($\bar{c}_j$). A conservative approach is to assume that they recognize only vehicle class—a coarse industry measure that delineates cars into compact, midsize, luxury, and sports cars, and light-duty trucks into SUVs, pickups, and vans. Table 1 shows price and cost variation by class, as well as for the whole fleet. The table also shows 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 equal to roughly half of the purchase price for most categories, reaching nearly two-thirds for compact cars. Variation in average fuel costs *across* class is also large, with mean differences of several thousand dollars across car and light-truck classes. It therefore makes little sense for consumers to be rationally inattentive to fuel costs when comparing models across classes, unless effort costs are extraordinarily high.

Consumers are likely to be drawn to vehicles that are similar in attributes, which will generally

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

imply that consumers are comparing vehicles in the same class. Table 1 shows that, even *within* class, fuel cost variation is large, ranging between $1,100 for vans and $2,500 for SUVs. This variation is substantial, and on its face does not suggest rational inattention is likely.

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 prices 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. Thus, fuel costs are a modest part of the variation in vehicle features overall. Hence fuel costs may not be pivotal to choice in many cases, which suggests that inattention could be rational. To address this further, I next consider data on second choices.

3.2 Second-choice data

Consumers report actively considering only a few models before buying a car (Ratchford and Srinivasan 1993; Moorthy, Ratchford, and Talukdar 1997; Ratchford, Lee, and Talukdar 2003). It is the fuel cost variation among these choices that determines whether inattention is rational, rather than variation across entire classes of vehicles. Most sources of automobile data, including the transaction data used here, do not include information on the vehicle that a consumer would have chosen if their preferred model was unavailable. Some such data do exist, including the Vehicle Ownership and Alternatives Survey (VOAS), an Internet survey of about three thousand consumers that is analyzed in Allcott (2011) and Allcott (2013), as well as survey data from about thirteen thousand consumers used in Langer (2012).

I use these data to estimate the distribution of fuel cost savings that a consumer faces between their first and second choice vehicles.[12] To make results comparable to the fuel cost numbers from the other tables, I calculate lifetime *differences* in fuel costs between first and second choice vehicles, assuming the same usage pattern and discounting used above: cars are driven 12,000 miles a year for 14 years, with a $2.50 per gallon price of gasoline and a 5% discount rate. This difference can be positive (if the second choice vehicle was more fuel economic) or negative (if the alternative was less fuel economic). In the VOAS, the mean of the difference is just $170, and the median is zero. The standard deviation is $2,890; the 25th percentile is -$1,350; and the 75th percentile is $1,620. In Langer's data, the mean of the difference is -$225, and the median is zero. The standard deviation is $3,112; the 25th percentile is -$1,472; and the 75th percentile is $984.

These standard deviations are somewhat larger than the standard deviations of within vehicle class fuel costs reported in table 1, but they are skewed heavily by a modest fraction of consumers who report having chosen between very different vehicles. To show the full distribution more clearly, I plot the cumulative density function of the lifetime fuel cost difference in figure 2. In the VOAS, around 18% of all consumers are choosing between two vehicles with exactly the same fuel economy rating, which means that the difference in lifetime fuel costs is zero. In Langer's data, only 6% report their first and second choice vehicles having identical fuel economy, but the distributions are otherwise fairly similar. In particular, around 40% of consumers in both surveys are choosing between two vehicles with lifetime fuel costs within $1,000 of each other, and around three-quarters of all consumers would save less than $1,000 by switching to their second choice vehicle.

These differences are of the same order of magnitude as the within class standard deviations described in table 1. They indicate that there are many consumers who report choosing among vehicles with very similar fuel economies. For these consumers, inattention might be rational. There

[12]The VOAS includes choices over used cars as well as new ones, whereas my interest is only in new car purchases. The survey does not ask if the vehicle was new, so instead I keep cases where the model year is within 1 year on either side of the year in which the vehicle was purchased. This makes results comparable to the tabulations from the transaction data in table 1, which include only new vehicles. This restriction produces a sample of 798 observations.

Figure 2: Cumulative distribution function of the absolute value of the lifetime difference in fuel costs across first and second choice models

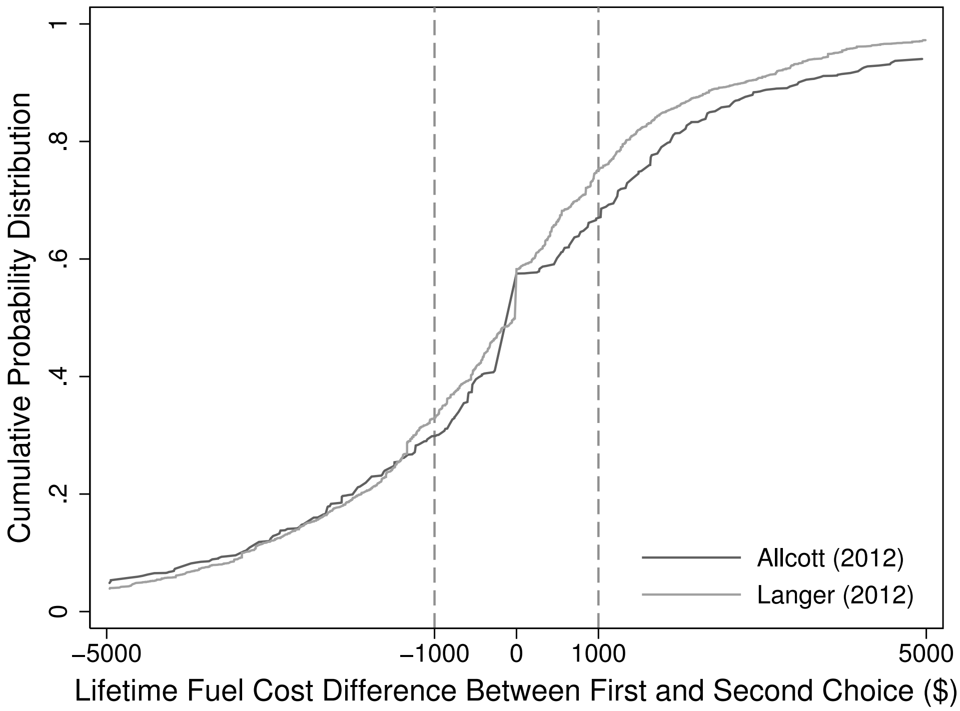

Data are from Allcott (2013) and Langer (2012). See text for details.

are many others, however, choosing among vehicles with far greater cost variation. The problem with relying on second choice data is that these close second choice pairings may be the result of attention to fuel economy, and they do not indicate how consumer welfare differs across first and second choices. When the second choice is more fuel efficient, changing to that choice would save fuel costs but may offer lower utility from other attributes associated with fuel economy. To better address these issues, I turn next to a simulation.

3.3 Simulated choices with rational inattention

This section uses a discrete choice model to simulate how vehicle choices would be influenced if consumers were inattentive to fuel costs. The idea is to calculate how much consumer utility would fall if consumers ignored fuel costs when making a vehicle purchase by calculating their optimal choice, and then calculating the choice that they would make with limited information (due to inattention). The difference between utility values of these two choices (calculated with full information) represents the utility lost from making a choice with incomplete information. The distribution of this welfare loss statistic could be directly compared to effort costs in order to assess the rationality of inattention.

This exercise requires estimates of the utility value of many different cars for a sample of consumers. The estimates from a discrete choice model of the vehicle market, of which there are many, provide coefficients that can be combined with data to calculate utility values, but neither the data nor the coefficients provide the unobserved error terms. The error terms are not known, but their distribution is estimated as part of a discrete choice model estimation, so I simulate choices using actual data, estimated coefficients, and randomly drawn error terms.

There are a variety of discrete choice models of the vehicle market. Many model the car market using coarse categories—i.e., "Ford midsize car". These are not adequate for present purposes because they would hide the many cases where inattention caused consumers to switch between very similar models that fall within the same aggregated category. Models also vary in how flexibly they allow preferences over attributes, including fuel economy, to vary. In this exercise, it is important to allow as much flexibility in preferences as possible, which is best done in a mixed logit that estimates the shape of a heterogeneous distribution of taste for attributes, including fuel economy. For the mixed logit, research has shown that having stated second choice data as well as the actual car purchased by a consumer is critical for precision (Train and Winston 2007). Some models of vehicle demand also allow consumers to have incomplete valuation of fuel economy (e.g., Allcott and Wozny (Forthcoming)), whereas for the current exercise it is best to impose full valuation when the consumer has full information because the goal is to determine how a fully rational consumer will behave with different information sets.

One paper that is ideal for the exercise is Langer (2012), which estimates a mixed logit demand system for a fine-grained version of the car market that includes 213 different models in the choice set. The estimation uses actual car purchases and stated second choices. Langer estimates the utility function as follows:

$$U_{ij} = \delta_j + \alpha p_j z_i + \sum_k \beta_k x_{jk} z_i + \tilde{\alpha}\nu_{ip} p_j + \sum_k \tilde{\beta}_k \nu_{ik} x_{jk} z_i + \epsilon_{ij}$$
$$\text{where } \delta_j = \bar{\alpha} p_j + \sum_k x_{jk}\bar{\beta}_k + \xi_j \text{ for each } j = 1, 2, ..., J,$$

where i denotes a consumer, j denotes a vehicle, p_j is price, z_i are consumer characteristics, and x are the attributes (denoted $k = 1, 2, ..., K$). The ν terms are standard normal random variables, unique to each person i and characteristic k (but not car j), and ϵ_{ij} is a type-I extreme value error. The parameters to be estimated are α, β, δ and ξ. Each vehicle is modeled as having a common utility value δ_j, which is a function of observable attributes p and x_k (not interacted with individual characteristics or the random ν). ξ is the value of unobserved attributes of the vehicle, which is estimated through a standard contraction mapping to fit market demand, as in Berry (1994). Heterogeneity in demand for each vehicle is allowed both through the interaction of consumer characteristics with prices and attributes and by allowing individual's to have random variation in tastes around the mean taste. The random variation is assumed to be distributed normally with mean zero, and the variance is estimated in parameters $\tilde{\alpha}$ and $\tilde{\beta}$.[13]

In the heterogeneous portion of the utility function, Langer includes price divided by income, and random coefficients for price, vehicle classes, horsepower, fuel consumption and curb weight. Langer includes price, class dummies, curb weight, number of passengers, turning radius and a dummy for domestic production in her estimates of δ_j. Note that this second list does not include fuel consumption, which means that fuel consumption is included in the ξ_j terms. This is ideal because it allows me to impose that the average consumer fully values fuel consumption, as described below.

The simulation works as follows. I take the demographic characteristics of each of the 13,454 observations in Langer's data set. I then draw a random ν_k for each attribute that has a random

[13]Langer estimates such equations for each of four demographic groups in order to study price discrimination. See Langer (2012) for details.

component (there are i x k, that is 13,454 x 9 such errors). These random coefficients ν_k create correlation in the demand across similar car types for each consumer—e.g., consumers who get a high draw in their taste for horsepower will like all powerful cars more than the average consumer does. I then draw a type-I extreme value error term for each of the 213 vehicle choices in the model (there are i x j, that is 13,454 x 213 such errors). Given these random draws and the coefficient estimates reported in Langer (2012), I then calculate the utility (consumer surplus) of each of the 213 vehicles. I identify the vehicle with the highest utility as the choice that the consumer would make if they had full information. (The model does include an outside good, the value of which is normalized to zero, so some consumers will choose not to buy a car.)

To determine choices and utility under inattention, I recalculate a perceived utility for each vehicle using the same random error draws and coefficients, but after substituting a perceived fuel consumption variable for the true one. That is, rather than assuming that consumers observe true fuel costs c_j^*, I assume they make decisions based on three alternative forms of perceived fuel costs $\bar{c}_j$. This substitution has two effects on the perceived utility of each vehicle. First, the perceived fuel cost enters into the random coefficient term on fuel consumption and changes the apparent utility of each vehicle. Vehicles with below average fuel economy will see a rise in their perceived utility, and those with above average fuel economy will see their perceived utility fall. Second, I raise or lower the vehicle's average utility δ_j by the difference between true and perceived lifetime fuel costs, multiplied by the average coefficient on price $\bar{\alpha}$.[14] This imposes that the average consumer fully values lifetime fuel costs, but the random coefficient component allows some consumers to value it more or less. Once the perceived utility of each vehicle is calculated, I identify the vehicle that the consumer would have chosen given these perceptions. Having identified the vehicle that consumer i would have chosen under full information and under partial information, the difference in actual utility can be calculated by taking the difference in consumer surplus across the two choices (based on true fuel consumption for both vehicles).

I consider three scenarios that represent different levels of information and generate different perceived fuel consumption values. To represent total ignorance, the first scenario replaces the actual fuel consumption term of all vehicles with the global mean value of fuel consumption, weighted by sales, across the sample. This preserves the average utility measure across all consumers, but it wipes away all perceived differences in fuel consumption. This necessarily will improve the perceived merit of all cars with below average fuel economy, and weaken the appeal of all cars with above average fuel economy. Each consumer is assumed to have the same misperception, but because there is a random coefficient on fuel consumption, this misperception will lead to varying changes in perceived utility. Given these perceived utilities, I identify the vehicle that consumers would choose under total ignorance. I then calculate the welfare loss for each consumer of choosing under imperfect information by comparing the true utility of this choice to the true utility of the choice made with full information.

In the second scenario, rather than using the global mean of fuel consumption, I use the class mean. Classes in Langer's data are luxury cars, sports cars, all other cars, pickup trucks, sport-utility vehicles and vans, which differs only slightly from the class definitions used above. This represents the case in which consumers have an unbiased estimate of the lifetime fuel costs of each type of car, but they know nothing about fuel cost variation within type. They assume, for example, that all pickup trucks have the same fuel economy, which is equal to the class average. Class accounts for around 50% of the variation in fuel costs of a vehicle, so this scenario represents a considerable improvement over pure ignorance.

[14] I use the same parameters as in the lifetime fuel cost calculations above to transform fuel economy ratings into lifetime fuel costs.

Table 2: Estimated welfare impacts of choice under incomplete information

	Perceived Fuel Economy Used for Choice		
	Global Mean	Class Mean	Attribute Predicted
Average welfare lost (per vehicle purchased)	\$552	\$291	\$89
	(33)	(21)	(11)
Percent who change vehicle	19%	14%	7%
	(0.8%)	(0.7%)	(0.7%)
Average welfare lost conditional on changing vehicle	\$2957	\$2092	\$1285
	(141)	(119)	(123)
Standard deviation of welfare lost	\$1725	\$1112	\$539
	(96)	(79)	(63)

Statistics are averages over 500 simulations, based on the data and coefficients from Langer (2012). The values in parentheses are the standard deviations of the test statistics over the 500 trials. In the global mean scenario, choices are determined as if all vehicles have the fleet average fuel economy. In the class mean scenario, choices are determined as if all vehicles have the average fuel economy for their class (luxury car, sports car, all other cars, pickup trucks, sport-utility vehicles and vans). Class predicts 41% of the variation in fuel consumption. In the attribute predicted scenario, vehicles are assigned their predicted fuel consumption from a regression of fuel consumption on class dummies, horsepower, curb weight, number of passengers and a dummy for being domestically produced. These variables predict 82% of the variation in fuel consumption.

In the third scenario, I use a more sophisticated prediction of a vehicle's fuel economy that takes into account horsepower, weight, number of passengers and location of production (foreign or domestic), in addition to class. Specifically, I regress fuel consumption on class dummies, horsepower, a dummy for whether the vehicle is made by a domestic automaker, curb weight and number of passengers, and I assume that consumers' perception is the predicted value from this regression. The R^2 on this regression is around .82, with most of the explanatory power coming from vehicle classes and weight. The attributes used here to predict fuel consumption are all in some sense "visible". A vehicle's class, brand and the number of passengers are easy to perceive. A vehicle's horsepower and weight are commonly cited statistics, but they may not be as easy to perceive. Thus, it likely makes sense to think that the average consumer has an easily determined fuel cost perception $\bar{c}_j$, the accuracy of which lies somewhere between that used in the second and third scenarios.

The procedure draws all of the random terms once, and then calculates four choices that would emerge, given the same random draws, under the four information scenarios. I then repeat this procedure, using a new set of random draws, 500 times and report average welfare changes over all trials. For each trial and in each information scenario, each consumer's welfare loss from misperception is calculated based on the choice they would make given incomplete information but the utility they would derive given the vehicle's true characteristics. Consumers who choose the same vehicle under complete and incomplete information will have a utility loss of zero. All other consumers will lose utility; by definition, it is not possible to make a better choice than that which would be made with full information.

Table 2 summarizes the results. The first row reports the average loss in consumer surplus per vehicle across all vehicle purchases over all 500 trials.[15] This statistic is an average of zero utility

[15] The model includes an outside good, but the statistics reported exclude cases in which the consumer chose the outside good (did not buy a car) in both cases of a pairwise comparison. Welfare losses therefore represent average utility losses, conditional on a vehicle being purchased. Including cases where the consumer chooses the outside good in both scenarios would drive down all of the estimated losses but also change the interpretation.

loss for many consumers whose choice is not changed by misperception and some utility losses from those who do change their choice. The second row of table 2 reports the percentage of consumers who change their vehicle choice because of misperception. The third row shows the average welfare loss conditional on the consumer changing vehicles as a result of misperception. The fourth row shows the standard deviation of the average welfare losses reported in the first row. The table also reports the standard deviation across trials in each of the welfare statistics. (These standard deviations are fairly small because each trial has a large sample.)

In the scenario of total ignorance, where consumers assign the global mean fuel consumption to all vehicles and choose accordingly, the average welfare loss is \$552. Even in this scenario, 81% of consumers will choose the same vehicle under full and partial information. The modest number of switchers is driven by two factors. One is that a top choice is sufficiently preferred to other choices that modifying one characteristic does not change the discrete choice made in many situations. Another is that, if a consumer's preferences are such that their first and second (and higher order) choices are all similar cars, then the misperception will have similar impacts on all of those cars, which preserves the ordering and does not change choice. When misperception does cause a consumer to switch vehicles, the welfare costs are large, at an average of \$2,946.

These welfare losses are reduced significantly when the consumer's perceived fuel economy is set to the class average, as shown in column 2. The average welfare loss from misperception when consumers take class into account is around half of the losses that arise from total ignorance. In the class scenario, the average welfare lost is \$291, and this reduced misperception causes only 14% of consumers to change their choice of vehicle.

If consumers make a more sophisticated prediction and accurately account for the impact of curb weight, horsepower, domestic production and the number of passengers, then the average welfare lost falls to a modest \$89 per vehicle, as shown in column 3. In this scenario, only 7% of consumers change their vehicle choice. As mentioned above, the sophisticated prediction model likely overstates the amount of information that consumers have at no or low cost, while it is intuitive to believe that consumers do know the class of vehicles they are considering and can perceive cost differences across classes on average. A reasonable conjecture is therefore that the true welfare loss from rational inattention in the car market would lie somewhere between the estimates in columns 2 and 3.

The next section will discuss why effort costs may be substantial, but there are no direct quantitative estimates of these costs. But, one can imagine that effort would exceed \$89 or even \$291. Another way to frame the issue is to say that, when buying a car, if paying attention to some feature provides only a few hundred dollars of value on average, many people will focus their efforts elsewhere.

Rational inattention is more likely as the welfare costs of it get smaller. But, even small welfare costs add up across the market. In the past decade, the new automobile market in the United States has averaged 14.9 million new cars sold, fluctuating between a low of 10.6 million in 2009 during the financial crisis and a high of 17.4 million in 2005. If consumers were rationally inattentive in scenarios two or three, that would correspond to total welfare losses of \$4.3 billion or \$1.3 billion per year, on average. Thus, while the individual welfare costs of inattention are modest, because the market is large, the total welfare impact on society is non-trivial.

3.4 Effort costs and fuel economy labels

The simulation suggests that the gains to effortful attention may be low. But, even if these gains 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 fuel economy.

Figure 3: Current fuel economy label

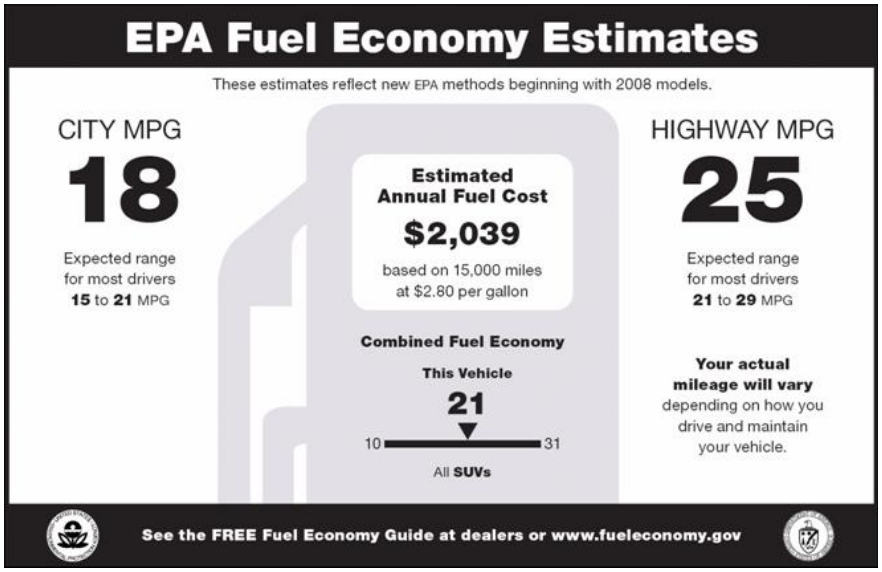

Given that, is it reasonable to suppose that consumers face nontrivial costs in calculating lifetime fuel costs? This section argues that the costs of becoming fully informed are likely high, both because label information is incomplete and possibly inaccurate, and because heterogeneity is an important source of uncertainty. One way to think about effort costs is simply to multiply the wage rate by the time required. The material below suggests that several hours would be required to overcome information limitations, which means that effort costs could easily be in the hundreds of dollars, even when fuel economy labels exist.

3.4.1 Labels are incomplete and inaccurate

Figure 3 provides an example of the current fuel economy label for new vehicles. It includes three fuel economy estimates (city, highway, and combined), as well as an estimated annual fuel cost and a comparison of the combined fuel economy to other vehicles in the same class. 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 miles per gallon (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.

Research shows that consumers find difficulty in converting the information on labels into lifetime fuel costs. 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 (2013) 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. 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).

The effort required of consumers to become fully informed is increased by the fact that the EPA's fuel economy ratings, which are based on a laboratory test, are imprecise and sometimes biased. 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%. Assuming that these adjustments were accurate, the original ratings were biased by a noticeable amount.

More recently, the EPA determined that the test ratings were inaccurate because tests no longer reflected typical driving behavior. (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.) The EPA determined that it was necessary to replace the old procedure, which conducted one test for the highway rating and one for the city, with a procedure that involves five tests that are combined to create the same two ratings with greater accuracy. New ratings were generated as of 2008.

To understand how the transition would affect labels, the EPA analyzed a sample of 615 vehicles for which they determined the old and new ratings. They have generously provided these data to me. Of the original sample, 380 are gasoline-only vehicles with complete information on all five tests. (The new test procedure has a much more pronounced impact on gas-electric hybrids vehicles, which I omit from my analysis.) Using these data, I reconstruct the original fuel economy label ratings and the new five-cycle ratings and compare them. Under the assumption that the new ratings are correct, the difference between the two rating systems indicates the degree of error in the prior rating system.

Figure 4 shows scatter plots of the old and new tests, for both city and highway values, along with the 45-degree line for reference. The vast majority of the data points are below the 45-degree line, implying that the new rating is lower than the original. 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 substantially across vehicles. OLS regressions (not shown) confirm that the slopes differ from one by a statistically significant amount, and the point estimates indicate that 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.

The apparent mismeasurement of the official fuel economy ratings translates into significant differences in lifetime fuel costs. Figure 5 shows the distribution of the difference in lifetime fuel costs between the two label ratings—using the same assumptions about mileage, gasoline prices, discounting, vehicle lifetime, and ratio of highway to city miles as employed elsewhere in this paper—in two ways. Figure 5a calculates lifetime fuel costs using the old test ratings and subtracts cost using the new ratings. Figure 5b instead calculates this difference using the predicted new fuel economy ratings from an OLS regression of the new ratings on the old; this removes the systematic bias so that differences are centered around zero and leaves variation that reflects only the idiosyncratic effects of the test change on particular vehicles.

The mean mismeasurement is -\$1,700—that is, on average, the old EPA ratings understated lifetime fuel costs by \$1,700—and mismeasurement ranges from -\$400 to -\$4,000. This is large relative to the welfare costs of inattention estimated in section 3.3 and the savings from switching vehicles estimated in section 3.1. Even once the systematic bias is removed, the variation in the

Figure 4: New and old EPA fuel economy ratings, in mpg

(a) City rating **(b)** Highway rating

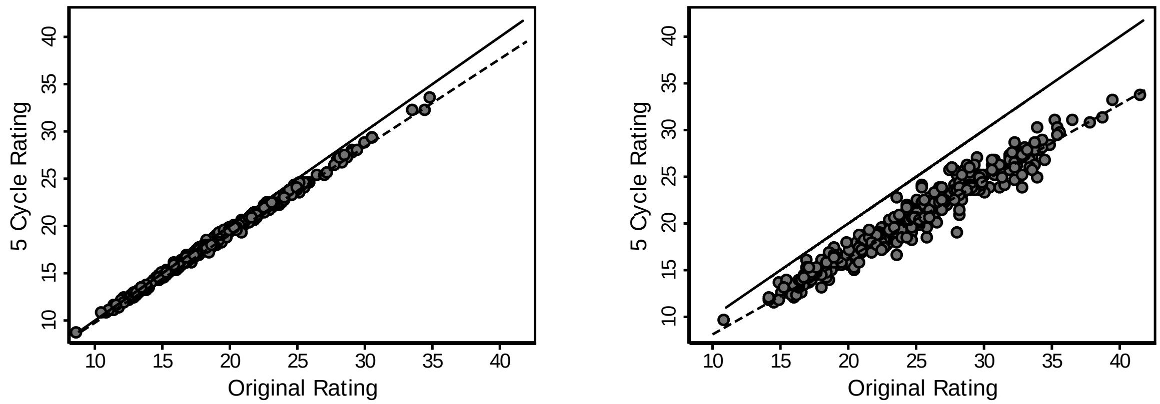

Figures based on 380 test vehicles for which data for both rating systems are available. Test vehicles come from model years 2003 to 2006. Dashed lines show a linear fit.

Figure 5: Difference in lifetime fuel costs across new and old EPA ratings

(a) Old - actual new **(b)** Predicted new - actual new

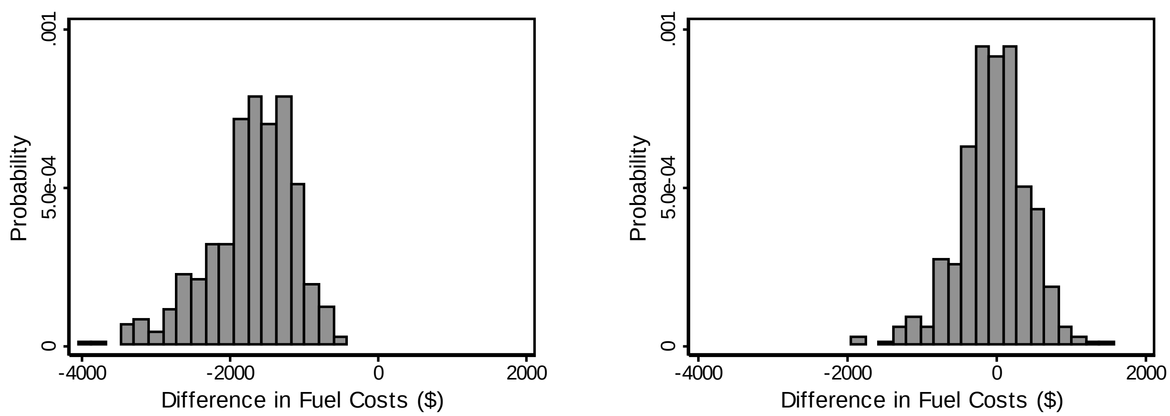

Figures based on 380 test vehicles for which both rating systems are available. Lifetime fuel costs are constructed assuming 12,000 miles driven per year for 14 years, with a 5% discount rate, a $2.50 per gallon gasoline price, and 55% of miles driven in the city. Predicted new ratings are calculated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

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 effort cost of ascertaining true lifetime fuel costs must be substantial.

3.4.2 Heterogeneity and effort costs

Even if labels provided perfectly accurate information of the lifetime fuel costs for the *average* driver, significant effort would still be required for consumers to figure out the value of fuel economy to them, if they differ from the average in their consumption behavior. Recall the formula used above for the lifetime fuel costs of a vehicle, $\sum_{t=0}^{T} \delta^t \frac{P_{gt} m_t}{mpg_t}$. Each parameter in this formula engenders some heterogeneity, including forecasted future gasoline prices P_{gt}, miles driven per period m_t, the discount factor δ^t and fuel economy for a given vehicle mpg_t (based on driving behavior and style). As a result, the lifetime fuel cost of the same car will be different for different drivers. Importantly, this type of uncertainty is difficult to resolve with labels or other information provision.

To show how wide is the dispersion in lifetime fuel costs for the same vehicle when drivers are heterogeneous, I simulate the value of increasing fuel economy by one mile-per-gallon using estimates of the distribution for these parameters taken from the literature. This is similar to the simulation in Anderson, Kellogg, and Sallee (Forthcoming). That paper provides an empirical distribution of inflation-adjusted forecasted gasoline prices from a sample of nationally-representative households.[16] The distribution of miles driven per year comes from the National Highway Transportation Survey.[17] The distribution of discount factors is calculated using the loan rate on new automobile purchases from a nationally-representative sample of transactions analyzed in Anderson and Sallee (2011).[18]

Specifically, I draw a sample of 100,000 observations from these empirical distributions, assuming independence across the parameters for each individual. I then calculate the lifetime fuel costs for each observation at every fuel economy value in the range of 16 to 31 mpg, assuming a starting price of gasoline of \$2.50 per gallon. The range of 16 to 31 mpg is the 10th and 90th percentiles in the transaction data analyzed above.

Figure 6 plots the median value of fuel economy and the interquartile range from this simulation. The downward slope of the values reflects the nonlinear relationship between fuel costs and fuel economy. The key result is that the interquartile range is very wide, in absolute dollar terms and as compared to the median. For example, at 20 miles-per-gallon, the median fuel cost is \$12,300,

[16]The data used in Anderson, Kellogg, and Sallee (Forthcoming) come from the Michigan Survey of Consumers and report a five-year ahead forecast of gasoline prices. I convert those forecasts into an annual growth rate and model heterogeneity by assuming growth at that rate for five years, after which time gasoline prices are assumed to be constant. This assumption of constant prices after five years likely understates forecast heterogeneity.

[17]I do not use the survey data directly, but instead use the mean and variance of annual miles driven calculated from the survey that are reported in Li, Timmins, and von Haefen (2009), adjusted for vehicle survival rates as reported in Lu (2006), to calculate the survival-probability weighted mileage of vehicles in each time period, based on the prior period's accumulated miles. This parameter of the simulation is identical to that used in Anderson, Kellogg, and Sallee (Forthcoming), which includes more detail.

[18]I do not model variation in fuel economy across drivers of the same car because the empirical distribution required is unavailable. Langer and McRae (2013) do show that there is wide dispersion in experienced fuel economy of a set of drivers using identical vehicles. 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). Drivers who drive mostly highway versus mostly city will therefore have very different fuel consumption per mile, but the location of miles driven is unavailable. Thus, idiosyncratic fuel economy for a given vehicle is likely to be significant, and omitting it from the simulation therefore understates fuel consumption heterogeneity.

Figure 6: Heterogeneity in the lifetime fuel cost of vehicles

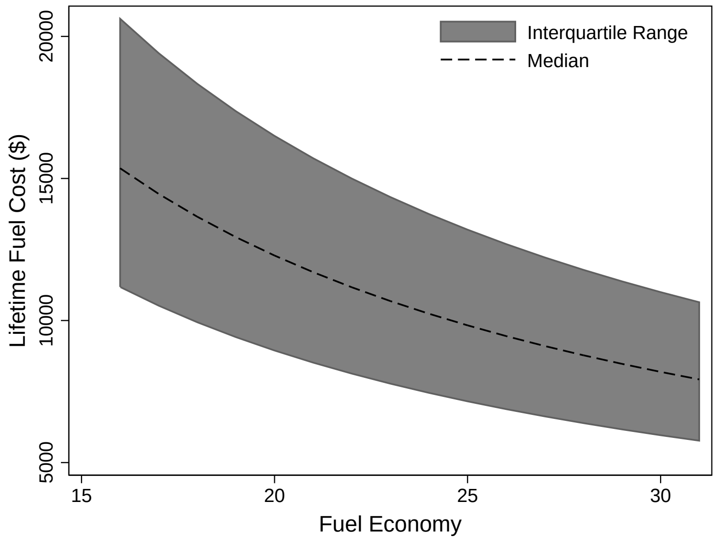

Estimates are based on a simulation using random draws from the empirical distribution of gasoline price forecasts, interest rates, and lifetime mileage. See text for details.

but the interquartile range is \$8,900 to \$16,500.[19] The standard deviation is \$6,200. This means that even if a fuel economy label explained the lifetime fuel costs accurately for the median driver, that estimate will be too high or too low by \$6,200, or 38%, on average.

Of course, many consumers will have some sense of whether their fuel costs are higher or lower than the average. But, it is not obvious how they would know by how much they differ from the average, short of gathering detailed idiosyncratic information on their own driving behavior and performing the necessary calculation. To the degree that this uncertainty stems from consumer's not knowing individual specific parameters—how much they drive, or whether they are especially aggressive drivers, what they think future gasoline prices will be or what discount rate they wish to apply to the choice problem—an outside party will have trouble overcoming the information problem for them.

This type of uncertainty can greatly raise the effort cost of resolving uncertainty without causing a commensurate increase in the returns to those efforts. As discussed in section 2.3, if a consumer's driving patterns, gasoline price forecasts and discount rates are independent of the vehicle they chose, then the errors they make will be closely correlated across models and will therefore have a mitigated impact on choice.

3.5 Price variation within model

Until now, the paper has emphasized a comparison of the costs and benefits of exerting effort to resolve uncertainty about the value of fuel economy. Another way to think about the consumer's decision to pay attention to fuel economy is to ask what the opportunity cost of that attention is.

[19] This corresponds to a median value of *improving* fuel economy by one mpg of \$585, with an interquartile range of \$426 to \$786.

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
Midsize 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 type. Column (3) shows the variation in fuel costs. Column (4) shows the ratio of within-VIN price variation to fuel costs.

Suppose that the consumer has a fixed amount of time to devote to researching a car purchase. Should they devote more time to studying fuel economy, or spend their time on something else?

One answer comes from comparing the welfare costs of fuel consumption estimated above to variation in retail prices that different consumers pay for identical cars. In the automobile market, haggling over the final price of the good is important. One can attempt to get price quotes for identical cars from multiple dealerships, or spend effort in negotiation. How does this price variation compare to the consequences of ignoring fuel consumption?

To quantify the value of negotiation, I use the transaction data detailed above to calculate price variation for vehicles with the same vehicle identification number (VIN). The VIN indicates a vehicle's 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 type. 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 type 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 VIN types 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 VIN types.

For reference, column 1 in Table 3 reproduces from Table 1 the total variation in price both between and within VIN for each vehicle class. Column 2 shows the standard deviation in price within each VIN, which is an estimate of 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 than across, as would be expected, but the remaining variation is still quite large. 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 within a class from Table 1, and column 4 shows the ratio of within-VIN price variation to this across-VIN fuel cost variation. In all categories, the within-VIN price variation exceeds the across-VIN fuel cost variation. In other words, consumers would save more, on average, by improving the transaction price for the vehicle he or she has chosen to buy by one standard deviation than he or she would

Figure 7: Vehicle prices over model year cycle

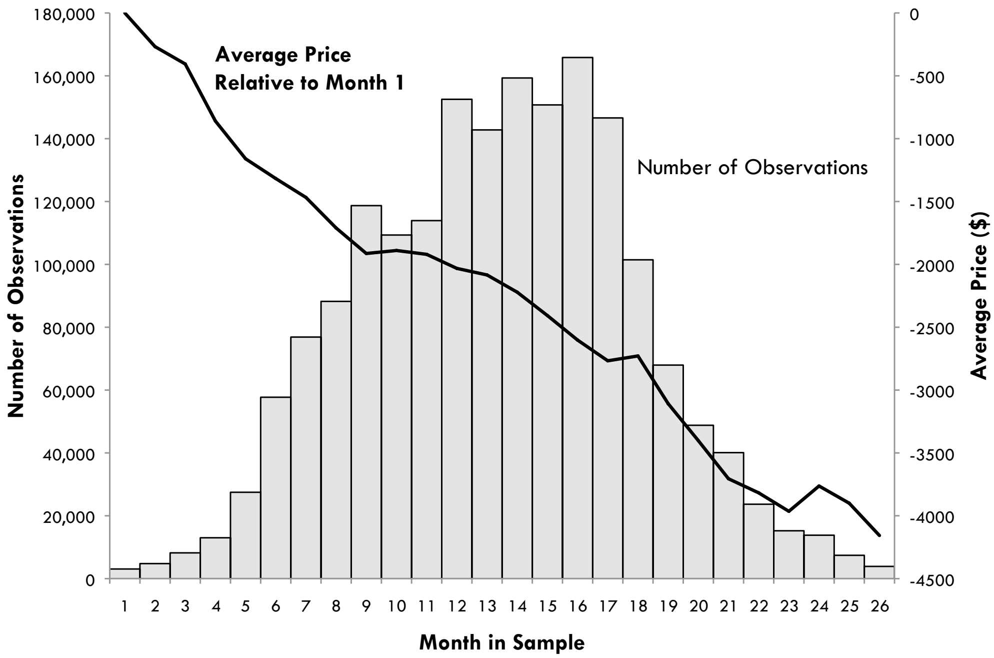

Solid line plots the month coefficients from a regression of vehicle price (in levels) on a set of month dummies for every month in the sample with VIN fixed effects. Bar graph plots total sample size over the model year cycle. Month 1 is April 2005.

gain from making a one-standard deviation improvement in fuel economy, within their vehicle class, by changing vehicle.[20]

It may be difficult to negotiate a better price, but some transaction price variation is predictable and can be taken advantage of by the consumer. 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:

$$p_{ijt} = \alpha + \sum_{s=2}^{26} 1(t = s)\delta_t + \sum_{v=2}^{1,705} 1(v = j)\gamma_j + u_{ijt}, \tag{9}$$

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

[20]Adjusting the price data for outliers and dropping transactions that have interest rate subsidies (the calculation of which requires use of a market benchmark and therefore may exaggerate variance for individuals) reduces price variation by only modest amounts. The standard deviation for a sample that drops the highest and lowest 2% of transactions within VIN is $2,100 across all classes, and the sample that excludes interest rate subsidized vehicles has a standard deviation of $2,344, as compared to the $2,435 for the full sample.

Table 4: 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.

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. Thus, a consumer can expect to gain more, on average, by waiting one month to buy an identical vehicle than they stand to lose by making the wrong choice due to inattention, as estimated in the simulation in section 3.3. 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 (9) 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).

4 Appliance energy costs

The preceding analysis suggests that even though fuel costs and their variation are substantial in the automobile market, inattention at some level may be rational. This section briefly considers related data for home appliances. The necessary data on household appliances are not uniformly available from the federal government, but a patchwork of sources enables an analysis of the most important appliances, including dishwashers, clothes washers, ovens, ranges, refrigerators, and freezers.

Table 4 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 Federal Trade Commission (FTC). FTC data for other appliances are incomplete for varied reasons, so other sources are used. Association of Home Appliance Manufacturers (AHAM) data are used for refrigerators and freezers, and data from

Natural Resources Canada (NRC) are used for ovens and ranges.

Table 4 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 4 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. Unlike the data for automobiles, these price variations reflect only variation in suggested retail prices across models. Price variation from sales and promotions and differences across retailers will make the full price variation much larger.

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.

4.1 Effort 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 effort cost required of consumers to become fully informed. Prior to 2008, labels included a 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 8a shows an example for a refrigerator label. The Energy Policy Act of 2005 mandated a review of the labels, and the labels were changed, effective in 2008. An example of a refrigerator label from the new regime is provided in figure 8b.

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, know how their local elecriticy price varies from the national average, 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.[21]

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

Figure 8: Energy guide labels

(a) Before 2008 **(b)** After 2008

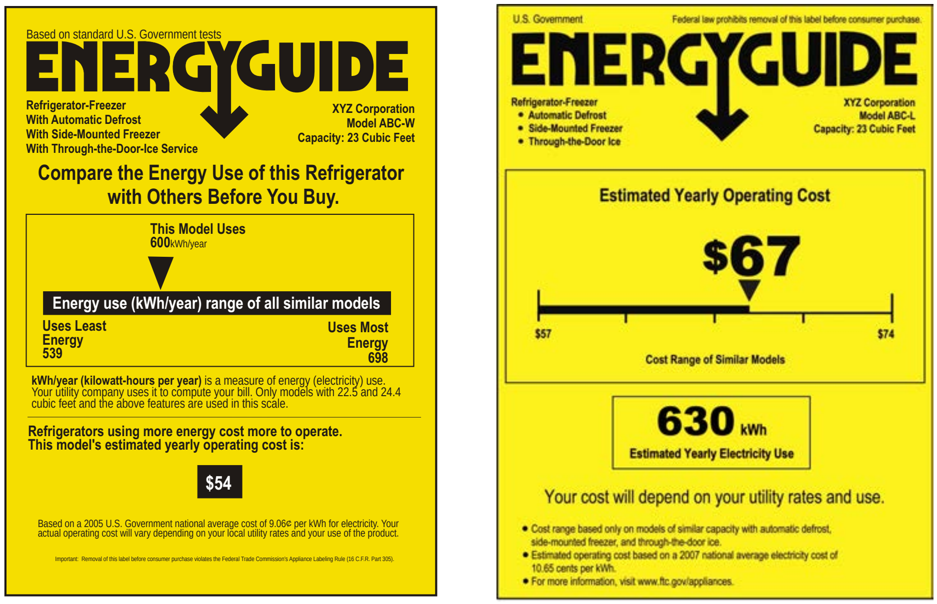

Table 5: NRC lifetime cost inaccuracy

	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.

The Canadian government does not include a lifetime cost on its labels, but NRC does 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 5 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. Knowing the annual cost is not in itself a solution to the need for costly effort to estimate the correct costs when purchasing a durable good. The Canadian government made a sizable error when doing the calculation itself.

4.2 Summary 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 Department of Energy 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.

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.

5 Implications for policy

In a wide range of applications, the first-best solution to a market failure caused by an externality is a Pigouvian tax. In energy economics, Pigouvian taxes typically take the form of a direct tax on fuel or emissions, and in the presence of a Pigouvian tax there is no role for taxes or regulations on the product market. Theoretical research has found, however, that product taxes or subsidies can be used in conjunction with fuel or emissions taxes to improve welfare when both externalities and an energy paradox are present (Allcott, Mullainathan, and Taubinsky Forthcoming; Fischer, Harrington, and Parry 2007; Heutel 2011). This logic has been used by policy makers to promote alternatives to Pigouvian taxes. For example, the government credits recent Corporate Average Fuel Economy regulations with large improvements in the *private* welfare of consumers. These benefits, derived under the assumption that consumers exhibit significant myopia, are pivotal in making the standards pass the government's formal cost-benefit analysis (NHTSA 2010). More generally, the energy paradox is frequently cited as a reason for reliance on policies other than Pigouvian taxes.

Rational inattention is not necessarily a market failure, however. If making better choices requires costly effort, then a policy that, through some incentive, induced consumers to exert effort has both benefits (from improved choice) and costs (from exerted effort). If inattention is rational, then consumers have demonstrated by revealed preference that costs exceed benefits, and a policy that gets them to pay more attention will be welfare decreasing, at least initially. If consumer attention changes the products on offer in the marketplace by changing producer innovations, then the welfare impacts of inducing consumers to pay attention are less clear. This is an important question for future research.

What does seem clear is that if policy can lower the cost of attention (rather than incentiving it) then that would improve welfare. The most obvious way of lowering attention costs is to direct more effort towards the appropriate design of labels, where there is almost certainly room for improvement. For example, research has shown clearly that consumers are confused by fuel economy labels written in miles-per-gallon (Larrick and Soll 2008), but after extensive review the EPA has decided to keep the mpg rating, though they have added the gallons-per-mile rating also. Allcott (2013) discusses this reform process and concludes that label redesign is highly cost effective. Unfortunately, there is little to no research on how effective these labels are in practice, and thus more work is needed. (Allcott and Mullainathan (2010) also argue that more research on labels is warranted.)

A closely related question is what factors might make information easier for consumers to digest. Take as an example "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 (Sallee and Slemrod 2012). Empirical evidence that these notches do indeed create product distortions is presented in Houde (2013), Sallee and Slemrod (2012), and the working paper version of this paper. These inefficiencies could be justified, however, if inattention causes consumers 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 possibility is explored in more detail in Houde (2013).

The energy paradox literature has long made an important distinction between "market failures" and "market barriers". Whereas market failures are immediately suggestive of a role for policy, market barriers are factors that make optimal choice of energy efficiency difficult, but they might be overcome by market forces even without government intervention. Incomplete information is typically categorized as a market barrier, and Metcalf (1994), for example, argues that informational uncertainty should therefore give rise to market services that provide information, such as home energy audits. This suggests perhaps a muted role for policy.

In the case of rational inattention, however, market forces may not be able to provide a way around the barrier. For some goods, like microwaves or laptops, even if information provision costs can be lowered from economies of scale, it may still be the case that consumers will be rationally inattentive if the other attributes of those goods are so much more important that energy concerns are rarely pivotal to discrete choice. Even for goods with larger costs, like automobiles, if the uncertainty stems from heterogeneous parameters—like how much one drives or idiosyncratic driving styles that affect fuel economy—that consumers themselves do not know, it will be challenging for a third party (or the government) to solve the information problem for the consumer.

Finally, if consumers demand too little energy efficiency because of inattention, then any action that draws attention to efficiency may help reduce externalities. An interesting implication of this is that programs that create "buzz" around a product or technology can be important and change the traditional evaluation of programs. For example, Sallee (2011a) demonstrates that the federal tax credit for the hybrid Prius was effective during a time period when there was excess demand for the Prius, and therefore the subsidy had no effect on demand. Normally, this would be damning to any cost-benefit analysis of the subsidy, but if the credit created public awareness of the new technology (brought attention to hybrids), then it may have been helpful in overcoming welfare losses due to inattention. A catch to this implication, however, is that if attention is costly then the welfare implications of driving people's attention towards something are ambiguous. That is, lowering information acquisition costs is obviously beneficial, but causing someone to pay attention to energy efficiency may not be welfare improving, even if it helps them avoid making a product choice mistake, because they must incur the cost of effort.

6 Conclusion

The goal of this paper is to elevate consideration of rational inattention in the study of energy economics. The heuristic model aids this task by providing a framework for thinking about conditions under which rational inattention is likely to emerge. Specifically, rational inattention is more likely when uncertainty around variation in energy costs is small, when variation in the differentiation of products in other dimensions is large, and when effort costs of paying attentions are large.

The empirical portions of the paper argued that inattention is quite plausible in the markets for automobiles and home appliances. In particular, a simulation of the automobile market indicates that the welfare cost to consumers of being inattentive is modest. The paper does not, however, develop a direct test of inattention or prove its existence. This is a key step for future research. The analysis in Houde (2012), which integrates costly attention into a discrete choice model for refrigerators, is an important step in this direction. 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.

Greater consideration of rational inattention also has implications for both empirical and theoretical research in energy economics. In terms of empirics, an energy paradox that is driven by rational inattention may not be detected by standard empirical methodologies because these methods take as given the set of products consumers face. Researchers may need to develop additional tests of the paradox that use data on revealed preferences but nevertheless allow for the possibility that certain products may not be brought to market. This will likely require tighter integration of models of consumer demand and product design decisions.

In terms of theory, rational inattention may indicate how notched program policies, simple labels, or other policies that present coarse information to consumers may be welfare improving. These types of public policies are ubiquitous, yet standard economic models suggest that they are redundant or even welfare decreasing (Sallee and Slemrod 2012). Further study of rational inattention could provide an explanation as to why these policies exist in the form that they do. More broadly, the presence of rational inattention complicates the welfare analysis of externality correcting policies, and additional research is essential to help guide optimal policy design.

References

Allcott, Hunt. 2011. "Consumers' Perceptions and Misperceptions of Energy Costs." *American Economic Review Papers & Proceedings* 101 (3):98–104.

———. 2013. "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market." *American Economic Journal: Economic Policy* 5 (3):30–66.

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

Allcott, Hunt and Sendhil Mullainathan. 2010. "Behavior and Energy Policy: Online Supporting Material." *Science* 327 (5970):1204–1205.

Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. Forthcoming. "Energy Policy with Externalities and Internalities." *Journal of Public Economics* .

Allcott, Hunt and Nathan Wozny. Forthcoming. "Gasoline Prices, Fuel Economy, and the Energy Paradox." *Review of Economics and Statistics* .

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. Forthcoming. "What Do Consumers Believe About the Future Price of Gasoline?" *Journal of Environmental Economics and Management* .

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.

Anderson, Soren T. and James M. Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." *American Economic Review* 101 (4):1375–1409.

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.

Bento, Antonio M., Shanjun Li, and Kevin Roth. 2012. "Is There an Energy Paradox in Fuel Economy? A Note on the Role of Consumer Heterogeneity and Sorting Bias." *Economics Letters* 115 (1):44–48.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics* 25 (2):242–262.

Bordalo, Peter, Nicola Gennaioli, and Andrei Shleifer. 2013. "Competition for Attention." NBER Working Paper 19076.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." *American Economic Review* 103 (1):220–256.

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.

Fischer, Carolyn, Winston Harrington, and Ian W.H. Parry. 2007. “Should Automobile Fuel Economy Standards Be Tightened?” *The Energy Journal* 28 (4):1–29.

Gabaix, Xavier. 2013. “A Sparsity-Based Model of Bounded Rationality, Applied to Basic Consumer and Equilibrium Theory.” Manuscript: New York University.

Gabaix, Xavier and David Laibson. 2005. “Bounded Rationality and Directed Cognition.” Manuscript: New York University.

———. 2006. “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics* 121 (2):505–540.

Goldberg, Pinelopi Koujianou. 1998. “The Effects of the Corporate Average Fuel Efficiency Standards in the U.S.” *The Journal of Industrial Economics* 46 (1):1–33.

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.

Heutel, Garth. 2011. “Optimal Policy Instruments for Externality-Producing Durable Goods Under Time Inconsistency.” NBER Working Paper 17083.

Houde, Sébastien. 2012. “How Consumers Respond to Product Certification: A Welfare Analysis of the Energy Star Program.” Manuscript: Stanford University.

———. 2013. “Bunching with the Stars: How Firms Respond to Environmental Certification.” Manuscript: University of Maryland.

Howarth, Richard B. and Bo Andersson. 1993. “Market Barriers to Energy Efficiency.” *Energy Economics* 15 (4):262–272.

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.

Kilian, Lutz and Eric R. Sims. 2006. “The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data.” Manuscript, University of Michigan.

Langer, Ashley. 2012. “Demographic Preferences and Price Discrimination in New Vehicle Sales.” Manuscript: University of Michigan.

Langer, Ashley and Shaun McRae. 2013. “Step on It: Evidence on the Variation in On-Road Fuel Economy.” Manuscript: University of Arizona.

Larrick, Richard P. and Jack B. Soll. 2008. “The MPG Illusion.” *Science* 320 (5883):1593–1594.

Li, Shanjun, Christopher Timmins, and Roger H. von Haefen. 2009. “How Do Gasoline Prices Affect Fleet Fuel Economy?” *American Economic Journal: Economic Policy* 1 (2):113–137.

Lu, S. 2006. “Vehicle Survivability and Travel Mileage Schedules.” Technical Report DOT HS 809 952, National Highway Traffic Safety Administration.

Matejka, Filip and Alisdair McKay. 2013. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." Manuscript: Boston University.

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.

NHTSA. 2010. "Final Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2012-MY 2016 Passenger Cars and Light Trucks." Washington, DC: National Highway Traffice Safety Administration, U.S. Department of Transportation.

NRC. 2002. "Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards." National Research Council. Washington, DC: National Academies Press.

———. 2011. "Assessment of Fuel Economy Technologies for Light-Duty Vehicles." National Research Council. Washington DC: National Academies Press.

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.

Reis, Ricardo. 2006. "Inattentive Consumers." *Journal of Monetary Economics* 53 (8):1761–1800.

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. 2012. "Car Notches: Strategic Automaker Responses to Fuel Economy Policy." *Journal of Public Economics* 96 (11-12):981–999.

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* .

Sawhill, James. 2008. "Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the US Automobile Market." Working Paper, University of California at Berkeley.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3):665–690.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3):213–225.

Train, Kenneth E. and Clifford Winston. 2007. "Vehicle Choice Behavior and the Declining Market Share of U.S. Automakers." *International Economic Review* 48 (4):1469–1496.

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.