Rational Inattention and Energy Efficiency

James M. Sallee   *University of Chicago*

Abstract

This paper argues that it will often be rational for consumers to pay limited attention to energy efficiency when choosing among energy-consuming durable goods like automobiles or home appliances. The reason is that the proper valuation of energy efficiency requires time and effort, but differences in efficiency across products will rarely be pivotal to choice when consumers have strong preferences regarding other product attributes. The paper first explains why proper valuation of efficiency is difficult, even in the presence of government energy labels. It next develops a model that shows how to value additional information about energy efficiency in a discrete-choice context. It then uses data on automobiles to show that consumers experience only small welfare losses when forced to choose a car without detailed information about fuel costs. Finally, the paper discusses the implications of rational inattention for both economic research and public policy.

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 on incomplete information rather than incur the cost to become perfectly informed. Rational inattention is a type of bounded rationality that has gained prominence in both microeconomic (Gabaix 2013) and macroeconomic (Reis 2006; Sims 2003) research, but its relevance for energy economics has not been explored. I argue in this paper that it will often be rational for consumers to choose among energy-consuming durables, like automobiles or home appliances, without acquiring complete information about energy efficiency.

The paper explores the following thought experiment. Suppose that a consumer must decide which durable good to purchase, given a set of options. The durable good in question consumes energy when used, so consumers should take into account the lifetime present discounted cost of buying fuel for the product when they make their choice. Consumers observe the various attributes of each product, but they have an incomplete understanding of lifetime fuel costs—that is, they have some rough idea of how much fuel will cost over the product's life, but they are uncertain about this cost. Consumers can resolve (or reduce significantly) this uncertainty by doing research and performing calculations, but this requires costly effort.

Alternatively, consumers might decide to make a purchase without exerting additional effort, thereby trading off the benefit of more precise information against the costs of effort. In this case, consumers are called rationally inattentive—to economize on effort, they choose not to become fully informed about fuel costs. Under what conditions is inattention to fuel costs rational? Are these conditions met in the real world? If rational inattention exists, what does it imply for public policy and economic research on energy?

To answer these questions, this paper takes four steps. First, a necessary condition for inattention to be rational is that there be some uncertainty about lifetime fuel costs that requires effort to resolve. If consumers were all blessed with an accurate understanding of lifetime fuel costs, or if mandatory government labels transmitted the relevant information perfectly, then there would be no scope for rational inattention. Thus, I begin in Section 2 by arguing that the accurate valuation of lifetime present discounted fuel costs is challenging, both because the calculation is cognitively difficult and because the information required is hard to obtain. Government labels aid in this task, but they do not resolve all uncertainty because the labels are incomplete and inaccurate and because heterogeneity in usage patterns implies that labels can resolve only a modest portion of the relevant uncertainty.

Second, having concluded that there are substantial effort costs involved in paying full attention to energy efficiency, I develop a heuristic model of a consumer making a discrete choice among energy-consuming durables in Section 3. Consumers are assumed to have beliefs about the energy consumption of each product that are based on readily available information and therefore require minimal effort. Consumers, however, recognize 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. Consumers must therefore make two decisions. First, they decide whether to incur the requisite effort cost to pay full attention. Second, they choose among durable goods. If they do not exert effort, then they have to choose a durable good with incomplete information. They may therefore make a mistake by choosing the "wrong" product.[1]

---

[1] This framework is similar to the one developed in Houde (2014). This paper and Houde (2014), which estimates a structural model of refrigerator choice, pursue very different empirical applications, but some of their theoretical insights overlap. This relationship is discussed further in Section 3.3.

The results of the model hinge on the fact that, in a discrete-choice framework, additional information about products is valuable to consumers only if it proves pivotal to their choice. Consumers might be incorrect about the lifetime fuel cost of the product that is their most preferred under incomplete information; but if they would choose this same product when they have full information, then the mistaken perception has no impact on their welfare because it does not lead them to make a mistake in their product choice. Thus, inattention to fuel costs is more likely to be rational whenever the resolution of uncertainty about fuel costs is less likely to be pivotal to choice. This information is less likely to be pivotal when the initial uncertainty itself is small (because mistakes in perception will be small) and when consumers have a strong preference for their first choice over the alternatives (so that mistaken perceptions must be large to affect product choice).

Guided by the model's results, in the third step I analyze variation in the fuel cost of products and estimate how often mistakes regarding them might prove pivotal. I do this for the case of automobiles in Section 4.[2] There, I document that, while there is substantial variation in fuel costs across automobiles, it is dwarfed by variation in prices. To the extent that prices proxy for the sum of all other attributes, this suggests that fuel costs will not often be pivotal and that consumers may therefore lose little from being inattentive to them. I then use 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 that only about 10 percent of new-car buyers would change their vehicle choice as a result of ignoring variation in fuel costs within vehicle categories (for example, they assume that all compact cars have the same fuel economy). As a result, the expected welfare losses from mistakes turn out to be modest—on the order of $100 or $200 per vehicle. If effort costs exceed these numbers, then inattention would be rational. Note that, by necessity, rational inattention implies that average mistakes cannot exceed the costs of effort. Even small individual mistakes, however, can be important in the aggregate; a $100-per-vehicle loss from inattention would amount to $1.5 billion per year in the U.S. car market.

The fourth step in the analysis is to speculate on how rational inattention could affect economic research and public policy (Section 5). One reason that rational inattention is of interest to energy economics is its bearing on the energy paradox. The literature on the energy paradox, or the energy efficiency gap, originates from the puzzling observation that many energy-saving technologies— like energy-efficient lightbulbs, improved building insulation, and energy-efficient home appliances—appear to be cost-effective, according to bottom-up

---

[2] Sallee (2013) reports similar exercises for home appliances.

engineering cost estimates, but suffer from very low take-up.[3] This is puzzling because it implies that Pareto improvements are left undone by the market. Many observers have thus concluded that consumers and firms are biased in their valuation of energy efficiency, perhaps because of myopia. The main alternative hypothesis is that the adoption of these technologies involves costs that are hidden to the analyst and thus have been excluded from consideration. Attention costs may be one such factor and as such may offer a partial solution to the puzzle.

The empirical literature on the energy paradox has tended to focus on revealed-preference tests to determine whether consumers appear to fully value energy efficiency.[4] If consumers are found to fully value energy efficiency—that is, if they are willing to trade $1 in purchase price for $1 in present discounted lifetime fuel costs—then it is usually presumed that producers will adopt all cost-effective technologies and bring them to market. Investment in energy efficiency will therefore be efficient.

Rational inattention threatens to break this logic. To see why, suppose that some energy-saving technologies are easily understood by consumers but others require costly effort to observe and evaluate—they are shrouded in the sense of Gabaix and Laibson (2006). In equilibrium, firms will bring to market only those innovations that garner attention, and consumers will choose rationally among the products on offer. To the econometrician, consumers will thus appear to fully value energy efficiency, but even so producers will fail to bring to market any cost-effective innovations that are hard for consumers to perceive. 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). Automakers may thus underadopt the latter, which creates inefficiency.

In the absence of inattention, present bias, or some other behavioral economic consideration, economic theory is clear in advocating Pigouvian taxes as ideal policies to correct for externalities related to energy consumption. Policy prescriptions are less clear when a behavioral bias exists, however, and an emerging literature has sought to understand how corrective policies should be designed in their presence (Allcott, Mullainathan, and Taubinsky 2014; Fischer, Harrington, and Parry 2007; Heutel 2011). Rational inattention points to several possible

---

[3] The technologies are thought to be privately cost-effective—that is, they are cost-effective even before externalities are taken into account. Throughout the discussion here, "efficiency" is meant to indicate only private efficiency. For recent literature reviews of the energy paradox, see Allcott and Greenstone (2012), Gillingham and Palmer (2013), and Greene (2010), the last of which is focused on automobiles.

[4] Allcott and Wozny (forthcoming), Busse, Knittel, and Zettelmeyer (2013), and Sallee, West, and Fan (2009) are recent examples that study this in the case of automobiles by correlating relative price movements of vehicles with variation in their lifetime fuel costs that are driven by changes in the price of gasoline. This panel methodology represents an improvement over previous literature that relied purely on cross-sectional choice data because the research design can employ fixed effects that capture unobserved attributes of vehicles that affect choice.

policy conclusions, some of which are distinct from those derived in that literature. First, consideration of inattention suggests that it is important to lower barriers to information acquisition and processing, perhaps through improvements in energy product labels. It also suggests that there might be a role for coarse information, such as Energy Star certification, which simply rates a product as efficient or inefficient, if such information is easier for consumers to digest. Surprisingly little research evaluates the impacts of alternative labels, with Newell and Siikamäki (2013) being a key exception.

Second, unlike in the case of present-biased consumers, it is important to recognize that overcoming rational inattention by coercing or incentivizing consumers to pay attention is not welfare improving because attention involves real costs. Ex post, myopic consumers would be grateful for a policy that had changed their choice of product, but inattentive consumers would not be. Rather than prod consumers with price incentives, it might be optimal to force producers to improve energy efficiency or to adopt certain technologies. This suggests a role for minimum efficiency standards or technology mandates in the optimal suite of policies, whereas such policies are redundant or harmful in the presence of Pigouvian taxes in a model with no limited rationality. Interestingly, product standards and inattention may create a positive feedback loop—if standards raise the minimum efficiency of products, this may compress the variance in fuel costs across products, which in turn makes inattention more likely to be rational.

In sum, this paper argues that rational inattention is relevant for economic research and public policy in the area of energy economics, thereby connecting the broader literature on the economics of information and rational inattention to energy economics. The paper does not directly test whether consumers are rationally inattentive, but rather it provides evidence that rational inattention is plausible. The hope is that subsequent research will provide additional evidence and also draw out more fully the implications of rational inattention for optimal policy design.

2. Energy Efficiency Valuation, Uncertainty, and Effort Costs

If valuing energy efficiency were sufficiently easy, there would be no room for rational inattention. Thus, I first describe the effort cost involved in valuing energy efficiency with a focus on automobiles, which is the good analyzed empirically in Section 4.

When choosing an automobile, consumers should consider its fuel costs, which can be written as

$$\text{Present discounted value of lifetime fuel costs} = \sum_{t=0}^{T} \delta^t \frac{P_{gt} m_t}{\text{mpg}_t}, \qquad (1)$$

where t indexes time period with $t = 0$ being today, T is the life span of the

Figure 1. Current fuel economy label (courtesy of the Environmental Protection Agency)

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.

With minimal or no effort, consumers likely understand coarse facts—for example, that big, heavier cars cost more to drive or that hybrids are less costly. Moreover, new vehicles come with fuel economy labels that provide some relevant information. Thus, consumers might form some belief about the fuel cost of each vehicle with minimal effort. A full calculation of equation (1), however, will require time, research, and cognitive effort. A sophisticated consumer might save on this exertion and make choices on the basis of incomplete, but easily formed, initial beliefs.

To see why fuel economy labels do not by themselves solve the valuation problem, consider the content of fuel economy labels currently in use, an example of which is provided in Figure 1. The label, which was updated in 2011, includes a city, highway, and combined fuel economy rating, which must be combined with estimated mileage, future gasoline prices, and a discount rate to calculate equation (1). It seems likely that many people do not know how much they drive, are not used to employing discount factors, do not know how long a vehicle will last, and do not have experience forecasting fuel prices. The label does include an annual fuel cost, based on an assumption about mileage and prices, but even this must be converted into a lifetime present discounted value via discounting and the choice of a time horizon. Moreover, this formula assumes full capitalization of efficiency in the used-car market; if consumers are unsure of how other people value fuel economy, there is another layer of complexity.
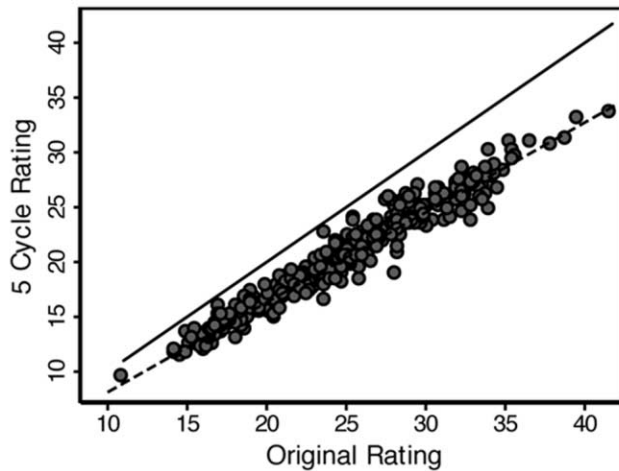
Figure 2. Fuel efficiency: new versus old Environmental Protection Agency highway ratings

Some cognitive effort is therefore required for a proper valuation. This is consistent with existing research that shows that consumers have difficulty in converting the information on labels into lifetime fuel costs.[5]

Valuation is made more difficult by the fact that official labels are not necessarily accurate. The Environmental Protection Agency (EPA) has twice made major reforms to its rating system because of concerns about accuracy, the more recent of which took effect in 2008. To show that this is important, I use data from the EPA on a sample of 380 vehicles, from model years 2003 to 2006, that were tested using both the old and new rating systems.[6] Figure 2 plots the old

[5] Larrick and Soll (2008) document that, in a laboratory setting, consumers fail to understand the nonlinearity of costs in miles-per-gallon ratings and overestimate the pecuniary gains from increases in miles per gallon, and thus underestimate 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 Larrick and Soll (2008). The newest labels now include a combined rating in gallons-per-mile, which could aid consumer decision making because this measure is linearly related to fuel costs. 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 concluded that labels improved decision making, but consumers are sensitive to the form of the information, and responses are often inconsistent across individuals (McNeill and Wilkie 1979; Hutton and Wilkie 1980). Newell and Siikamäki (2013) give more reason for optimism, finding that, depending on which discount rate is used, the existing labels for water heaters induce rational choices, on average, though there is significant dispersion across consumers. Water heaters were chosen for the experiment because they are particularly easy to understand and thus may represent a best-case scenario for the transfer of information.

[6] The first reform was in 1986, at which time the Environmental Protection Agency (EPA) uniformly lowered city ratings by 10 percent and highway ratings by 22 percent. The second reform involved the introduction of a new testing procedure.
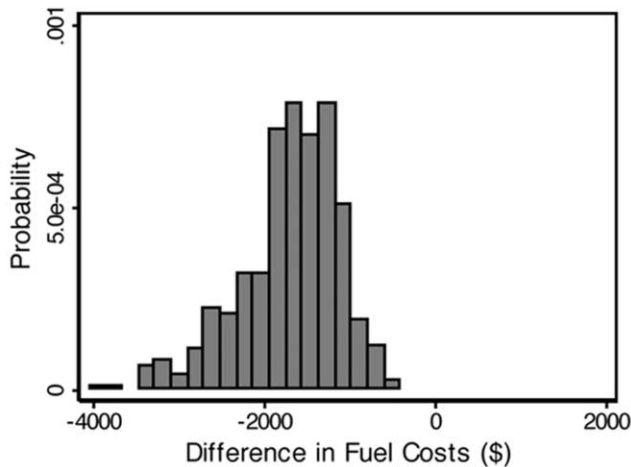
Figure 3. Change in lifetime fuel costs due to revision of Environmental Protection Agency highway ratings.

and new ratings for the highway test for this sample of cars. The dashed line shows the linear fit between the two regimes. All the data points are below the 45-degree line, which implies that the new rating is lower than the original. There is also a significant spread, which indicates that the change in fuel economy between rating systems varies substantially across vehicles.[7] This translates into significant differences in lifetime fuel costs. Figure 3 shows the distribution in the change in lifetime fuel costs implied by the ratings revision for each car in the sample. Lifetime fuel costs are calculated assuming 12,000 miles driven per year for 14 years, with a 5 percent discount rate, a $2.50-per-gallon gasoline price, and 55 percent of miles driven in the city. The old system understated fuel costs by $1,700 on average, but this ranges from $400 to $4,000. If labels are inaccurate, they cannot fully solve the consumer's information problem.

Even if labels provided perfectly accurate information for the average driver, consumers would still need to figure out how they differ from the average. Individuals have different beliefs about future gasoline prices. They drive different distances. They have different discount rates. And they have different fuel economies from driving the same car, because of driving styles. All of this implies that a single label cannot capture fuel costs accurately for all drivers.

To demonstrate the importance of heterogeneity, I adapt a simulation that is performed in Anderson, Kellogg, and Sallee (2013). The simulation draws

[7] The city ratings show a similar pattern, though both the bias and spread are less pronounced. Ordinary least squares regressions (not shown) indicate that the old rating system overestimated the cost savings of a fuel economy improvement by 7 percent for city ratings and 18 percent for highway ratings, on average.

100,000 random observations from the empirical distributions of mileage, gasoline price forecasts, and discount rates and calculates the lifetime fuel costs for each simulated observation, using equation (1). The empirical distribution of inflation-adjusted forecasted gasoline prices is from Anderson, Kellogg, and Sallee (2013).[8] The distribution of miles driven per year is from the National Highway Transportation Survey.[9] 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).[10]

For each random draw, I calculate the lifetime fuel costs at every fuel economy value in the range 16–31 mpg (which covers 90 percent of cars), assuming a starting price of gasoline of $2.50 per gallon. Figure 4 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 with the median. For example, at 20 mpg, the median fuel cost is $12,300, but the interquartile range is $8,900–$16,500.[11] 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 50 percent, on average.

Many consumers will have some sense of whether their fuel costs are higher or lower than the average. But to quantify this they must gather detailed idiosyncratic information on their own driving behavior. It will be especially difficult for the government, producers, or another third party to resolve this type of uncertainty because they too will lack the necessary facts.

In sum, consumers can use fuel economy labels and basic intuition to form a belief about lifetime fuel costs of different vehicles with minimal effort. Resolving the remaining uncertainty will require additional time, research, and

---

[8] The data used in Anderson, Kellogg, and Sallee (2013) are from the Michigan Survey of Consumers and report a 5-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 5 years, after which time gasoline prices are assumed to be constant. This assumption of constant prices after 5 years likely understates forecast heterogeneity.

[9] 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 on the basis of the prior period's accumulated miles. This parameter of the simulation is identical to that used in Anderson, Kellogg, and Sallee (2013), which includes more detail.

[10] 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) 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 percent lower than highway ratings (equivalent to the difference between a Volkswagen Jetta and a Ford Crown Victoria). Drivers who drive mostly on the highway or mostly in the 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.

[11] This corresponds to a median value of improving fuel economy by 1 mpg of $585, with an interquartile range of $426–$786.
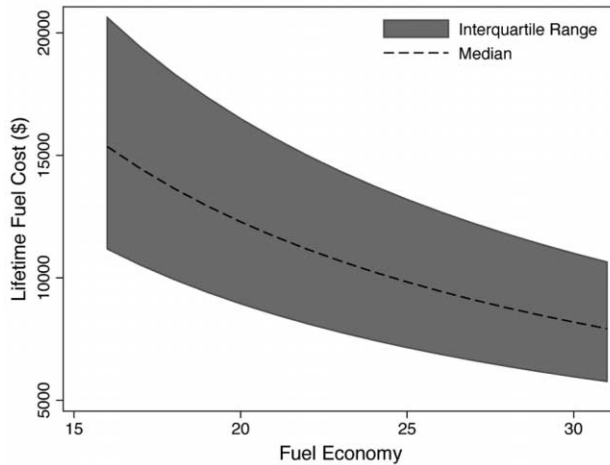
Figure 4. Heterogeneity in the lifetime fuel costs of vehicles

cognitive effort. Consumers can resolve uncertainty by learning the algebraic formula and doing the proper calculation. They can also resolve uncertainty by researching the reported fuel economy of a particular model to see if it deviates from the label, tracking their own driving behavior more closely to determine their mileage and the impact of their driving style, researching each model's reliability and expected life span, and so on. Any of these actions requires time and effort, and consumers may alternatively decide to forgo these costs and make a vehicle choice on the basis of incomplete information. The model that is laid out in Section 3 provides a way of understanding what the benefits of the additional information might be and thereby demonstrates when these costs are worth paying and conversely when attention would be rational.

3. A Model of Inattention

Having established that there are costs to determining the value of energy efficiency in Section 2, this section turns to a formal description of the benefits of acquiring additional information through the use of a discrete-choice model. In a discrete-choice context, additional information is beneficial only if it changes the consumer's choice. As a result, new information will be less valuable as initial uncertainty is smaller and as the consumer's first-choice product under limited information is more appealing relative to the alternatives—both of which make it less likely that new information about energy efficiency will cause a consumer to switch products.

To see this formally, consider a model with the following assumptions. 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. At no cost, consumers observe the relevant attributes for each product j, but for lifetime fuel costs they have imperfect information. They observe $\bar{c}_j$ for each product j, whereas the true fuel cost is $c_j^* = \bar{c}_j + c_j$, so c_j may be thought of as a random error term.

Consumers make two choices. They first decide whether to exert effort, defined as $e \in \{0, 1\}$, to learn about fuel costs. After that, they choose a product to purchase. If they exert effort ($e = 1$), then c_j is revealed, but they pay an effort cost s, which is the monetized value of time and energy spent learning about fuel costs. If they choose not to exert effort ($e = 0$), then no cost is borne, but in the second step they must choose a product j without knowing the fuel cost error c_j, which could lead them to make a mistake.

The utility of consumer i choosing product j is written as

$$U_{ij} = \beta' X_j - p_j - (\bar{c}_j + c_j) + \varepsilon_{ij} \tag{2}$$

$$\equiv W_{ij} - c_j, \tag{3}$$

where U_{ij} is utility, X_j is a set of observed attributes, p_j is the price of the vehicle, and ε_{ij} is a random error term. The term W_{ij} is the idiosyncratic valuation of vehicle j absent the uncertain cost component—for the inattentive consumer W_{ij} is the observable portion of utility—and c_j is a random error term.[12] All consumers are assumed to have a common valuation of attributes (captured by β), 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.

The initial belief $\bar{c}_j$ is interpreted as the consumer's 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. 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; that is, $\mathbb{E}[\bar{c}_j] = c_j^*$.[13] The consumer is assumed to know the distribution of c_j but not its value.

In this case, where $\bar{c}_j$ is an unbiased estimate of the truth, the uninformed consumer ($e = 0$) will choose the product with the highest W_{ij} in the second step. The informed consumer ($e = 1$) will choose the product with the highest

[12] Because c_j^* is the present discounted lifetime fuel cost of the vehicle, the absence of a coefficient for 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.

[13] In this section, I assume that all attributes in X_j are included in forming $\bar{c}_j$ because otherwise it is illogical to assume that the error terms are independent of utility, and independence facilitates exposition. Section 3.2 discusses alternative structures that relax this assumption.

U_{ij}. In the first step, the consumer will exert effort if the expected return from information—which is the difference between the expected value of the choice made under complete versus incomplete information—exceeds the search costs. Formally, consumer i will choose $e = 1$ if and only if

$$\mathbb{E}[\max_j U_{ij}] - \max_j \mathbb{E}[W_{ij}] \geq s. \tag{4}$$

To proceed, it is useful to make a distributional assumption, so that the expectations of U_{ij} and W_{ij} can be expressed in terms of underlying parameters. Results are particularly elegant if c_j is a type I extreme-value random variable.[14] A type I extreme-value random variable has a mean, assumed here to be 0 so that $\bar{c}_j$ is unbiased, and a scale parameter σ, which is inversely related to variance and is assumed to be known to the consumer.[15]

This parametric assumption makes the problem directly analogous to a standard logit model, in which now the uninformed ($e = 0$) consumer occupies the econometrician's standard position of knowing the expected utility of each choice, which can be used to calculate choice probabilities, but not observing the error term. The uninformed consumer knows W_{ij} and knows that the probability that j has the highest true utility is $e^{\sigma W_{ik}}/\sum_{j \in \mathcal{J}} e^{\sigma W_{ij}}$, which is the standard expression for a choice probability in a logit problem. Thus, when c_j is a type I extreme-value random variable, the decision rule in expression (4) can be rewritten as consumer i will choose $e = 1$ if and only if

$$\frac{1}{\sigma} \ln\left(\frac{\sum_{j \in \mathcal{J}} e^{\sigma W_{ij}}}{e^{\sigma W_{ik}}}\right) \geq s, \tag{5}$$

where k is the product with the highest W_{ij}. The term in parentheses is the inverse of the probability that k has the highest true utility U_{ij}, and the entire left-hand-side expression is equal to the expected difference in choice utility between the informed and uninformed consumer. That is, it equals the expected welfare gain in product value from being fully informed. The consumer will choose $e = 1$ if this gain exceeds s but will choose $e = 0$ otherwise.

Examination of expression (5) makes clear the main comparative statics of interest. Ceteris paribus, increasing W_{ik}, which is the value of the perceived top choice k, will increase the probability that product k is the correct choice, thereby lowering the parenthetical term. 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 value to other choices (that is, W_{ik} is close to W_{ij}), then the consumer is likely to change his or her mind after observing c_j. Loosely speaking, as products become more different, fewer consumers will exert effort.

---

[14] An earlier version of this paper assumed a normal distribution. All of the qualitative implications for that case are identical to those described here.

[15] The variance of a type I extreme-value distribution is $\pi^2/6\sigma^2$.

As σ rises, the variance of c_j falls, and the expected welfare gain from information falls.[16] An increase in the variance of c_j increases the returns to acquiring information because the greater 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 (5) 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. If consumers make mistakes because of inattention, the losses from these mistakes on average cannot exceed the cost of paying attention. 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 effort costs are high, the variance of unknown energy costs are low, and 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.

3.1. Inattention and Producer Behavior

If consumers are rationally inattentive, how might this affect firm behavior? One interesting possibility is that consumer inattention might lead to underadoption of technologies that boost a product's energy efficiency but do so in a way not salient to consumers. That is, suppose that a firm can add a technology to a product that would lower the lifetime present discounted fuel costs by more than the incremental production cost. If consumers are attentive, the firm can increase profit by adding the technology and raising the price to recoup the incremental cost. If consumers are inattentive, however, they might recognize the price increase but fail to recognize the fuel cost decrease, so that the revised product looks worse (when in fact it is better). In that case, the adoption of a cost-effective energy-efficient technology could lead to a decrease in demand for the product, which in turn implies that firms may choose not to adopt it.

The existence of rational inattention among some consumers, however, is not sufficient to imply that such a social inefficiency will prevail. What matters is the attention of the marginal consumer, which can be easily illustrated in the

---

[16] This can be seen by differentiating expression (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 0 if $e = 1$. Then the derivative with respect to σ is $e^{\sigma W_{ik}}\{[\sum(W_{ik} - W_{ij})e^{\sigma W_{ij}}]/(\sum e^{\sigma W_{ij}})^2\} + se^{-\sigma s}$. Because $W_{ik} \geq W_{ij}$, both terms are positive.

two-good case, in which the consumer chooses only between goods j and k. For expositional ease, suppose that the two goods have the same average perceived value, which is to say that $\beta'X_j - p_j - \bar{c}_j = \beta'X_k - p_k - \bar{c}_k$, and that the error terms c_j and c_k are two independent draws from a common distribution. In this case, the value of the two goods differs only because of the idiosyncratic taste parameters ε_{ij} and ε_{ik}, which vary across individuals, and the fuel cost error terms c_j and c_k. I assume that there are many heterogeneous consumers with different realizations of $\Delta\varepsilon_i$. For the inattentive consumer, the difference in expected values of the two products is simply the difference in their idiosyncratic error terms, which I denote $\Delta\varepsilon_i \equiv \varepsilon_{ij} - \varepsilon_{ik}$.

In this setup, $\Delta\varepsilon_i$ summarizes how much better consumer i's first choice is compared with the alternative, when he or she is inattentive. As this statistic gets larger (in absolute value), rational inattention becomes more likely because the realized fuel cost differences are less likely to be pivotal. It is straightforward to show that there will be a cutoff value, denoted $\theta^\star$, such that if $\Delta\varepsilon_i$ exceeds $\theta^\star$, the consumer will choose product j while being inattentive ($e = 0$).[17] Conversely, if $\Delta\varepsilon_i$ is below $-\theta^\star$, then consumer i will choose product k while being inattentive. If the consumer has a realization of $\Delta\varepsilon_i$ in between these cutoffs, he or she will exert effort ($e = 1$) and then choose a product after becoming fully informed.

To return to the firm's decision, suppose that the technological change in question is small, so that it can be represented as a marginal change to cost and price. In this case, technological adoption will influence the product choice of only the marginal consumer—the one who is indifferent between product j and product k at original prices and costs. The realizations of the fuel cost error term difference, $c_j - c_k$, determines which consumer is marginal—it is the consumer for whom $\Delta\varepsilon_i = c_j - c_k$. If this marginal consumer is attentive, then either firm would wish to adopt a small cost-effective technological improvement because the only consumer whose choice will be affected, the marginal one, will perceive the change as a net improvement in the product. Thus, even if there are many inattentive consumers buying each product, the firms would each make adoption choices that are cost-effective from a societal perspective.

This is illustrated in Figure 5A, which plots a hypothetical distribution for $\Delta\varepsilon_i$ and shows the attention and choice cutoff areas. Those in the middle of the $\Delta\varepsilon_i$ distribution will exert effort and learn the true fuel cost difference $c_j - c_k$, which is the choice cutoff point for attentive consumers. In Figure 5A, the cost realization is labeled point A, which lies in between $-\theta^\star$ and $\theta^\star$, so that the marginal consumer is attentive. This means that there are no mistakes. Every individual i ends up purchasing the product that he or she would purchase under full information, even though some consumers are inattentive.

If, however, the true realization of c_j and c_k implies that the marginal consumer

---

[17] Specifically, i will choose j and $e = 0$ if and only if $\Delta\varepsilon > -\ln(e^{gs} - 1)/\sigma \equiv \theta^\star$. Note that when some consumers search in equilibrium, $\ln(e^{gs} - 1) < 0$, so $\theta^\star > 0$.
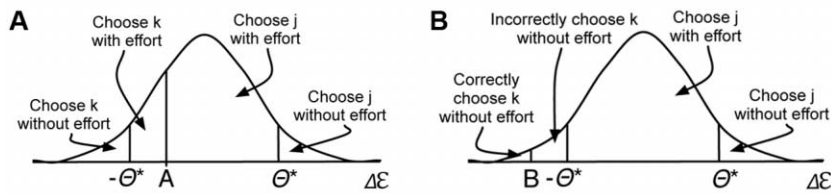
Figure 5. Effort and demand across consumer types

under full information is rationally inattentive in equilibrium, then underprovision may occur. This is illustrated in Figure 5B, which labels a realization of $c_j - c_k$ at point B, which is below $-\theta^*$. 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 $-\theta^*$, which lowers 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 $-\theta^*$, which lowers demand for j. But the corresponding drop in c_j will again be recognized only 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. In such a situation, social welfare could be improved by forcing both producers to adopt the technology, which suggests that rational inattention might provide a justification for energy efficiency standards or technology mandates.

This model is not complete, as it has not specified fully the form of competition among firms, the formation of consumer beliefs and their consistency with firm behavior, or the endogenous provision of information. A richer exploration of these issues is a promising area for future research, but I briefly discuss here a few possible concerns. A first concern is whether firms can overcome investment inefficiencies by signaling increases in nonsalient energy efficiency through changes in salient features, like price. In an equilibrium with inattentive marginal consumers, firms will not be able to signal greater energy efficiency simply by raising prices, 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.

A second concern is that consumer beliefs may be biased, perhaps because of strategic firm behavior. If inattentive consumers assume that firms deploy a set of cost-effective technologies that are nonsalient, then firms might reduce their provision of those technologies to reduce costs. Inattentive consumers will be

unaware of this, so their initial beliefs would become biased. Note that this can still be an equilibrium, albeit one in which beliefs are incorrect. A race-to-the-bottom equilibrium, in which firms adopt no technologies that are nonsalient, is consistent with unbiased beliefs and inattention.

A third concern is that all inefficiencies might be overcome through endogenous information provision by producers (perhaps via advertising), the government, or some third-party information provider. This may not be feasible in some cases, especially when uncertainty stems from consumer heterogeneity in usage, expectations, or discounting. 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. In addition, 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.[18]

3.2. Model Variations

The baseline model makes a number of simplifying assumptions and presents one interpretation of fuel cost uncertainty. Many variations of the model would deliver the same qualitative conclusions about inattention, though they change the mathematical details. Three such variations are briefly described here.

First, the model can be recast as a Bayesian updating process. In this version, the consumer has a prior belief about the product's energy consumption (formed as a function of the product's attributes perhaps) and receives 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.[19]

Second, the model can be recast so that uncertainty about fuel costs derives entirely from consumer heterogeneity. In this case, uncertain cost components cannot be independent across products (as the consumer's heterogeneity in usage and beliefs presumably carries over, at least partially, between products). As a

---

[18] Note also that even if advertising can substitute for consumer attention, this would not necessarily alleviate the problem of underadoption of cost-effective technologies because advertising involves real costs as well. Thus, there would be an analogous problem for firms deciding whether to advertise a feature. The gain to producer surplus from the information would have to exceed the cost of advertising, and the advertisement would have to be pivotal to consumer choice in order for it to affect producer surplus.

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

result, the mathematics are less elegant, but the same basic insights emerge. Details are provided in the Appendix.

Third, the baseline model assumes that the consumer's initial belief about fuel costs, $\bar{c}_j$, is unbiased and takes into account all attributes X_j that enter the utility function. If beliefs are biased or do not take all attributes into account, then the error term c_j will not be independent across products, which makes the mathematical solution less elegant. Of course, consumer beliefs may be biased, and errors may be correlated with vehicle attributes, but again the basic insights of the model are robust. The Appendix provides further details of this alternative.

3.3. 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). Gabaix (2013), for example, develops a model in which consumers have a prior belief about the price of each good and know the variance of prices, but costly attention is required to learn prices exactly. Consumers will pay more attention to goods that have greater price variation (and that they buy more of), which 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 Matějka 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 Houde (2014), which 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 (2014) 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 but does not estimate a demand model, Houde (2014) estimates a structural demand model for a single good but does not delve empirically into why inattention might occur.

4. Inattention in the Automobile Market

Section 2 argued that the effort costs of obtaining an accurate assessment of lifetime fuel costs are likely to be significant. Section 3 established a framework within which the benefits of obtaining accurate information about fuel costs can be understood. This section provides several types of empirical evidence on the benefits of information in the market for automobiles. If these benefits clearly outweigh the effort costs described in Section 2, then inattention would not be rational. This section argues that, at least for many consumers, inattention to a substantial fraction of fuel cost variation creates small average losses from mistakes, which makes rational inattention plausible.[20]

The model shows that attention to lifetime fuel costs is more valuable to consumer choice as the uncertainty is greater—that is, as the variance of fuel costs around the unbiased belief gets larger. To study this empirically, I begin by summarizing the overall fuel cost variation in the automobile market. This is an upper bound on the relevant fuel cost variation for consumer choice, as it is the variation that would influence choice if consumers had no information about how fuel costs vary across models without exerting effort. I also summarize variation within vehicle class, which may be a closer approximation of the information that consumers have at low cost. It would be the relevant variation if consumers knew, for example, by how much compact cars differed from midsize cars but knew nothing about variation within those classes without exerting effort.

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 transaction price variation, which both serves to contextualize the fuel cost magnitudes and acts as a proxy for variation in all other attributes. Transaction data are from an industry source that directly samples a large, representative sample of dealers across the

---

[20] Sallee (2013) also includes an empirical analysis of home appliances. Inattention is probably less likely to be rational in the market for automobiles than in the market for many home appliances because fuel cost variation is greater for automobiles and because gasoline prices are arguably more salient than electricity prices. This paper focuses on automobiles because superior data are available for autos, but it should be kept in mind that this market may be a particularly unlikely place to find rational inattention.

Table 1
Median Transaction Prices and Fuel Costs across Vehicle Categories

	Median Price		Median Fuel Cost		Sample	
	$	SD	$	SD	N	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
Sport utility vehicle	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
Overall market	23,405	9,506	13,498	3,101	1,861,437	1,705

Note. Median transaction prices account 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 percent discount rate and a $2.50-per-gallon price of gasoline. The fuel economy value used is the Environmental Protection Agency's estimated combined fuel economy. VIN = vehicle identification number.

country.[21] 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 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. Model year 2006 was chosen as the most recent year for which data are available.

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 (U.S. Department of Transportation 2008). I use a 5 percent 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 with the transaction data and use the combined rating.

Table 1 shows statistics for the overall market and for the vehicle classes identified in the data—which delineates cars into compact, midsize, luxury, and sports cars and light-duty trucks into sport utility vehicles (SUVs), pickups, and vans. The table presents the number of unique transactions observed and the number of different models, identified by distinct vehicle identification number (VIN) stub. The VIN stub identifies the manufacturer, model name, model year, number of engine cylinders, engine displacement, drive type, body style, trim level, fuel type, transmission, and aspiration (for example, turbo charged) of the vehicle. For the analysis, each distinct VIN is treated as a different type of car. The mean and standard deviation of transaction price (net of incentives) vary within and across VINs; fuel cost varies only across VINs.

Fuel costs, which range from $10,000 to $18,000 on average across categories,

---

[21] The data contain information about the transaction price, trade-in allowance, cash rebates, and financing. The final prices are adjusted for incentives, including cash rebates and interest rate subsidies calculated relative to the Federal Reserve's survey of 48-month car loan interest rates from commercial banks. The same methodology is employed in Sallee (2011a), which provides additional detail.

are large—both in absolute value and relative to transaction prices. But it is the variance in these costs, not their average, that matters for rational inattention. Variation in average fuel costs across classes 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.

But consumers are likely to be drawn to vehicles with similar attributes, which means that they will often be choosing among vehicles in the same class. Table 1 shows that, even within class, fuel cost variation is large, ranging between $1,100 for vans and almost $2,500 for SUVs. This variation is substantial and on its face does not suggest that rational inattention is likely, but it is far smaller than the variation in transaction prices within a class. This suggests that other attributes may differ so much across cars that fuel costs are rarely pivotal.

To explore this further, I use second-choice data from the Vehicle Ownership and Alternatives Survey (VOAS), an Internet survey of about 3,000 consumers that is analyzed in Allcott (2011, 2013), as well as survey data from about 13,000 new-car buyers collected by a major market research firm that is used in Langer (2012). These surveys asked consumers what car they would have purchased had their first choice been unavailable. I use these data to estimate the distribution of fuel cost savings that consumers face between their first- and second-choice vehicles.[22]

To make results comparable to the fuel cost numbers in 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. 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 0. The standard deviation is $2,890, the 25th percentile is −$1,350, and the 75th percentile is $1,620. In Langer (2012), the mean of the difference is −$225, and the median is 0. The standard deviation is $3,112, the 25th percentile is −$1,472, and the 75th percentile is $984.

These standard deviations 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 6.[23] In the VOAS, around 18 percent 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 0. In Langer (2012), only 6 percent report their first- and second-choice vehicles as having identical

---

[22] The Vehicle Ownership and Alternatives Survey includes choices of 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 in which 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.

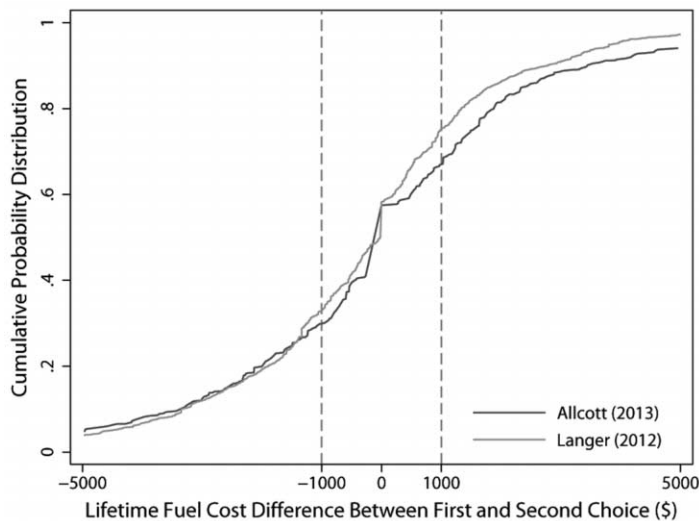[23] Data for Figure 6 are from Allcott (2013) and Langer (2012).

Figure 6. Cumulative distribution function of fuel cost differences

fuel economy, but the distributions are otherwise fairly similar. In particular, around 40 percent 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 are many others, however, choosing among vehicles with far greater cost variation. There are limitations to relying on second-choice data; the second-choice pairings may be the result of attention to fuel economy, and these comparisons 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.

4.1. 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 first obtain estimates of how much a random sample of consumers value each of the cars in their choice set. Then, using estimates of the value of fuel economy, I estimate how much consumers would value these same choices if they had an

incorrect perception (due to inattention) of each vehicle's fuel economy. These alternative values can be used to identify the vehicle that consumers would choose if they had incomplete information. If consumers would choose the same car in either case, then they are said to have no welfare loss from inattention (because the information involved was not pivotal). If the change in information does cause consumers to change their choice, then their welfare loss can be measured by comparing the difference between the valuations (under full information) of their true first choice and the vehicle they would mistakenly choose when information is incomplete. This procedure provides estimates of the welfare losses that consumers would experience if they were forced to make choices with incomplete information.

This procedure requires estimates from a discrete-choice model of the car market. Underlying any discrete-choice model is a random utility model, which supposes that the payoff to each consumer of each product that he or she might buy is a combination of the product's characteristics, how much the consumer values each characteristic, and a random error term that captures the purely idiosyncratic benefit of each consumer-product pair. The estimates of a discrete-choice model provide coefficients that indicate how much consumers value each characteristic, and it provides an estimate of the distribution of the idiosyncratic error term, but it does not include realizations of those errors. So for each vehicle-consumer pair, the model's results indicate the distribution of valuations. The simulation performed takes a large sample of random draws from these distributions to assemble a simulated data set that has a complete characterization of how much each consumer values each car and manipulates this data set to model inattention.

There are a variety of discrete-choice models of the vehicle market. Many use coarse categories—for example, Ford midsize car. These are not adequate for present purposes because they would hide the many cases in which inattention caused consumers to switch between very similar vehicles that fall within the same aggregated category. Discrete-choice models also differ in how flexibly they allow preferences to vary across individuals. The most restrictive assume common valuations of each attribute across consumers (for example, everyone values a 1-unit increase in horsepower at $10), whereas the most flexible allow the coefficients to vary randomly across consumers, according to a distribution the parameters of which are estimated along with the other coefficients. For present purposes, it is important to allow as much flexibility in preferences as possible, which is best done in a mixed logit model that estimates the shape of a heterogeneous distribution of taste for attributes, including fuel economy. For the mixed logit model, research has shown that having stated second-choice data as well as the 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 (for example, 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 an unbiased 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 car purchases and stated second choices. Langer estimates the utility function as follows:

$$U_{ij} = \delta_j + \alpha p_j z_i + \sum_k \beta_k x_{jk} z_i + \tilde{\alpha} \nu_{ip} p_j + \sum_k \tilde{\beta}_k \nu_{ik} x_{jk} z_i + \varepsilon_{ij}, \tag{6}$$

where

$$\delta_j = \bar{\alpha} p_j + \sum_k x_{jk} \bar{\beta}_k + \xi_j \qquad \text{for each } j = 1, 2, \ldots, J,$$

where i denotes a consumer, j denotes a vehicle, p_j is price, z_i is a set of consumer characteristics, and x_j is the set of attributes (denoted $k = 1, 2, \ldots, K$). 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_j and x_k (not interacted with individual characteristics or the random term ν). The term ξ is the value of unobserved attributes of the vehicle, which is analogous to a fixed effect for each model, and 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 (for example, men might value horsepower more than do women, on average) and by allowing individuals to have random variation in tastes around the mean taste (for example, some men have above-average valuations of horsepower, and others have below-average valuations). The random variation is assumed to be distributed normally with a mean of 0, and the variance is estimated in parameters $\tilde{\alpha}$ and $\tilde{\beta}$.[24]

In the heterogeneous portion of the utility function, Langer (2012) 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 term. 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. The first step is to take the demographic characteristics of each of the 13,454 consumers in Langer (2012) but discard the choice of car that they made. Their choice depends on the random error terms they experienced in reality, but these are unobservable. Thus, I ignore the choice

---

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

that they made in the original data and instead simulate their valuations of each car using the coefficients estimated by Langer and random draws of possible error terms. Note that all random error terms in equation (6) are normalized so that each ν_k is a standard normal term and ε is a type I extreme-value term. I draw the required number of ν_k (there are $i \times k$, that is, 13,454 × 9 such errors) and ε values for each consumer (there are $i \times j$, that is, 13,454 × 213 such errors). I combine these random draws with Langer's estimates of other parameters in the model (α, $\boldsymbol{\beta}$, $\boldsymbol{\delta}$, and ξ) and the data on consumer demographics ($\boldsymbol{z}_i$) and car attributes ($\boldsymbol{x}_j$ and p_j) to determine the valuation (consumer surplus) of each vehicle for each consumer. I identify the vehicle with the highest utility as the choice that consumers would make if they had full information. (The model does include an outside good, the value of which is normalized to 0, 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 assume that consumers observe true fuel costs c_j^*, I assume that they make decisions on the basis of 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 $\bar{\alpha}$.[25] 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 (using the true fuel consumption for both vehicles).[26]

I consider three scenarios that represent different levels of information and generate different perceived fuel consumption values. To represent total igno-

---

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

[26] Note that the Langer (2012) estimation procedure assumes that consumers are rational and informed about vehicle characteristics. As I discuss in Section 5, rational inattention could imply that Langer's coefficient estimates are biased. This raises an interesting methodological point that I do not attempt to resolve here, but note that the simulation can nevertheless be interpreted as an exercise under the null hypothesis that consumers are attentive and fully informed. Under that null hypothesis, Langer's coefficients are estimated correctly, and the simulation estimates the value of information. If information is valuable enough to offset requisite effort costs, then we would conclude that the data are consistent with the null hypothesis; otherwise we would reject it.

rance, 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 use the global mean of fuel consumption, I use the class mean. Classes in Langer's (2012) data are luxury cars, sports cars, all other cars, pickup trucks, vans, and small, medium, and large SUVs, 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 55 percent of the variation in fuel costs of a vehicle, so this scenario represents a considerable improvement over pure ignorance.

In the third scenario, I use a more sophisticated prediction of a vehicle's fuel economy that takes into account horsepower, weight, passenger capacity, 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 passenger capacity, and I assume that consumers' perception is the predicted value from this regression. The R^2-value on this regression is around .84, 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 passenger capacity 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 ($\bar{c}_j$ in the model's notation), 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 on the basis of the choice he or she would make given incomplete information but the utility he or she would derive given the vehicle's

Table 2
Estimated Welfare Impacts of Choice under Three Information Scenarios

	Global Mean	Class Mean	Attribute-Predicted Mean
Average welfare loss ($ per vehicle)	552	217	89
	(33)	(18)	(11)
Percentage who change choice of vehicle	19	12	7
	(.8)	(.7)	(.5)
Average welfare loss conditional on changing vehicle ($)	2,957	1,800	1,262
	(141)	(112)	(121)
SD of average welfare loss ($)	1,725	909	533
	(96)	(73)	(62)

Note. Statistics are averages over 500 simulations based on the data and coefficients from Langer (2012). Standard deviations of the test statistics over the 500 trials are in parentheses. 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, small sport utility vehicles [SUVs], medium SUVs, large SUVs, and vans). Class predicts 55 percent 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, passenger capacity, and a dummy for being domestically produced. These variables predict 84 percent of the variation in fuel consumption.

true characteristics. Consumers who choose the same vehicle under complete or incomplete information will have a utility loss of 0. 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. This simulation does not estimate the effort costs of attention. The welfare impacts reported here are decreases in consumer surplus from vehicle choice. These would need to be compared with the cost of paying attention to determine if inattention is rational.

Table 2 summarizes the results for all 500 trials.[27] The average welfare loss is an average of no utility loss for many consumers who do not change their choice of vehicle because of misperception and some utility losses from those who do change their choice of vehicle. The standard deviations across trials are fairly small because each trial has a large sample.

In the scenario assuming total ignorance, in which consumers assign the global mean fuel consumption to all vehicles and choose accordingly, the average welfare loss is $552. Even in this scenario, 81 percent of consumers will choose the same vehicle under full or 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 consumers' preferences are such that their first and second (and higher-order) choices are all similar cars, then the misperception will affect their views of all of those cars similarly, which preserves

[27] 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 in which the consumer chooses the outside good in both scenarios would drive down all of the estimated losses but also change the interpretation.

the ordering and does not change choice. When misperception does cause a consumer to switch vehicles, the welfare costs are large, an average of $2,957.

These welfare losses are reduced significantly when the consumer's perceived fuel economy is set to the class average. The average welfare loss from misperception when consumers take class into account is around half of the losses that arise from total ignorance.

If consumers make a more sophisticated prediction and accurately account for the impact of curb weight, horsepower, domestic production, and passenger capacity, then the average welfare loss falls to a modest $89 per vehicle. 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 class mean and the attribute-predicted estimates.[28]

The values reported in Table 2 indicate the reduction in consumer surplus that would occur when consumers choose vehicles with less than complete information about lifetime fuel costs. If effort costs exceed these losses, then inattention would be rational. I do not calculate estimates of these costs, but it seems plausible that they could be in the $100–$200 range, which makes inattention rational in the second and third scenarios. These scenarios, respectively, imply inattention to 45 percent and 16 percent of the variation in fuel costs across the entire car market. Note that it is not appropriate to interpret the consumer surplus reductions in Table 2 as net costs of inattention. If inattention is rational, then by definition the net costs of inattention are positive.

Inattention to attributes other than fuel costs might be rational. The simulation exercise can be adapted to other features of a vehicle. For illustration, I repeated the simulation assuming full information about fuel costs but incomplete information about vehicle horsepower, using the same three scenarios by replacing actual horsepower with a marketwide average, a vehicle class average, and a more sophisticated projection. This yields average consumer losses per vehicle of $288, $181, and $143 across the three scenarios. These are modestly smaller than the estimates for fuel consumption.

This does not necessarily imply that inattention to horsepower is more likely than inattention to fuel costs, however, because fuel costs are particularly difficult to perceive. Other attributes, including a vehicle's color, size, body type, passenger

---

[28] In response to the helpful suggestion of an anonymous referee, I also used this simulation to estimate the welfare loss implied by the EPA's revision to fuel economy ratings in 2008. This exercise interprets the revised numbers as accurate and simulates the loss in consumer surplus by randomly perturbing the perceived fuel costs of vehicles according to the degree of noise in the older ratings regime that was estimated in Section 2. This exercise suggests that 10 percent of consumers bought the wrong vehicle as a result of the inaccurate ratings, with an average welfare cost (including no loss for the 90 percent who were unaffected) of $124 per car purchased. Given the size of the U.S. car market, this aggregates to $1.8 billion per year.

capacity, brand, and sticker price are easy to perceive. Features such as the interior space, cargo volume, and comfort can be determined relatively easily on inspection. An 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 consumers 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, consumers have no way of seeing or experiencing an automobile's lifetime cost of refueling, and consequently they may struggle to value fuel economy. It is for this reason that I have focused on fuel costs, but it is nevertheless possible that inattention could extend to other attributes, in particular attributes like safety or reliability, that are also difficult to discern at the time of purchase.

Rational inattention is more likely as its welfare costs get smaller. But even small welfare costs add up across the market. In the past decade, the U.S. market for new automobiles 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 scenario 2 or 3, that would correspond to respective total welfare losses of $3.2 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 nontrivial. While $1–$3 billion is an important welfare impact, these values, which represent upper bounds on the welfare improvements possible from solving inattention-related inefficiencies, are dwarfed by the social costs of fuel consumption. For example, Davis (2014) calculates a global deadweight loss from the mispricing of gasoline and diesel of $92 billion per year.

4.2. Price Variation within Model

Until now, I have 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. Suppose that consumers have 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 with 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 VIN. 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

Table 3
Variations in Vehicle Transaction Prices and Fuel Costs: Standard Deviations

	All Prices (1)	Within-VIN Prices (2)	Fuel Costs (3)	Within-VIN Prices/ Fuel Costs (4)
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
Sport utility vehicle	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
Overall market	9,506	2,435	3,101	.8

Note. VIN = vehicle identification number.

in option 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 in which 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 1 standard deviation, conditional on having already selected a vehicle to purchase. There is substantially less variation within a VIN than across VINs, as would be expected, but the remaining variation is still quite large. In all categories, a 1-standard-deviation reduction in transaction price, within a 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 they choose to buy by 1 standard deviation than they would gain from making a 1-standard-deviation improvement in fuel economy, within their vehicle class, by changing vehicles.[29]

---

[29] 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 percent of transactions with the same vehicle identification number is $2,100 across all classes, and the sample that excludes interest-rate-subsidized vehicles has a standard deviation of $2,344, as compared with $2,435 for the full sample.
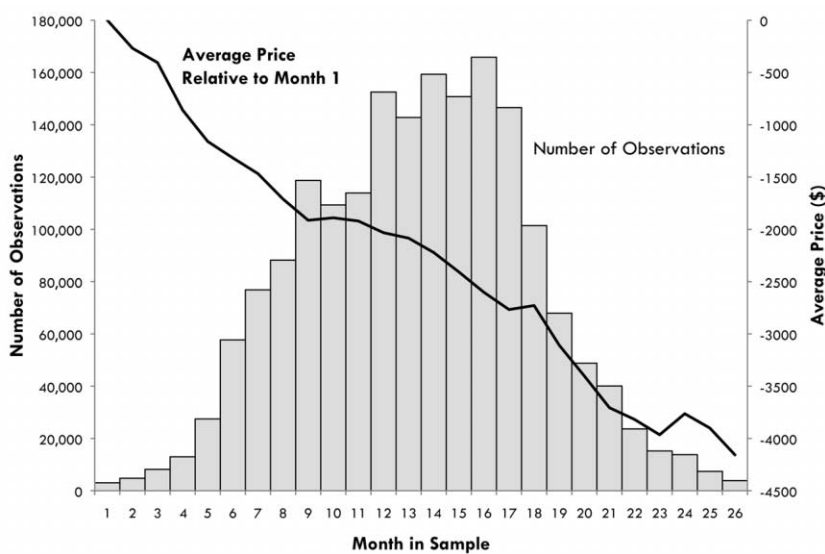
Figure 7. Vehicle prices over the model year cycle

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:

$$p_{ijt} = \alpha + \sum_{s=2}^{26} 1(t = s)\delta_t + \sum_{v=2}^{1,705} 1(v = j)\gamma_j + u_{ijt}, \tag{7}$$

where p_{ijt} is the transaction price of observation i of VIN type j in month t, δ_t is a set of month dummies for the 26 months in the sample (with month 1 the omitted category), γ_j is a set of 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.

The solid line in Figure 7 plots these coefficients to show the average decline in price over the model year cycle for every month in the sample with VIN fixed effects. Month 1 is April 2005. Figure 7 also shows the sample size in each month to indicate the distribution of sales over the time period. The price decline is relatively smooth, with a slope of about −$156 per month. Thus, consumers can expect to gain more, on average, by waiting 1 month to buy an identical vehicle than they stand to lose by making the wrong choice because of inattention,

as estimated in the simulation in Section 4.1. The total decline in price over the cycle is around $4,000, which is 17 percent of the median vehicle price over the entire sample. Equation (7) can also be written with logged price on the left-hand side. Then the coefficients represent average percentage price declines relative to the first month in the sample. The slope of the coefficients from this regression has a very similar shape to the one shown in Figure 7. The coefficients indicate that prices decline by 13 percent over the 25 months in the sample, which is broadly similar to the average annual model year decline in price of 9.0 percent estimated by Copeland, Dunn, and Hall (2011).

5. Implications for Policy and Research

Rational inattention has important implications for both empirical research and public policy. First, in terms of research, a key implication of rational inattention is that it is possible for consumers to correctly value the energy efficiency of products that are available in the market, but there can simultaneously be an inefficient provision of energy efficiency. This runs counter to the logic of much of the literature on the energy paradox.

Before explaining further, it is useful to specify what is meant by an efficient provision of energy efficiency. I follow Allcott and Greenstone (2012) in using the term "investment inefficiency" to mean a situation in which a private (that is, before externalities are taken into account) cost-effective investment in energy efficiency is not made. One reason this might happen is if consumers (or firms) undervalue energy efficiency. Much of the literature on the energy paradox attempts to test for this undervaluation in revealed-preference settings. Typically these studies test how consumers trade off lifetime fuel costs against up-front purchase prices. If a consumer requires more than $1 in present discounted future savings in order to be willing to pay $1 more up front, then the consumer is said to undervalue energy efficiency.[30] Consumers who undervalue energy efficiency will underinvest in it, which creates investment inefficiency. The assumption implicit in much of this literature is that if, on the other hand, consumers do value energy efficiency fully, then there will be no investment inefficiency because firms will have the right incentive to adopt all cost-effective technologies.[31]

In the presence of inattention, however, rational valuation in a revealed-preference study need not imply investment efficiency. Suppose that some technologies are salient and easily understood by consumers, but others are not (that

---

[30] Empirical papers that take this approach include work on appliances (Hausman 1979; Dubin and McFadden 1984) and automobiles (Allcott and Wozny, forthcoming; Busse, Knittel, and Zettelmeyer 2013; Goldberg 1998; Kilian and Sims 2006; Sallee, West, and Fan 2009; Sawhill 2008).

[31] Allcott and Wozny (forthcoming) and Busse, Knittel, and Zettelmeyer (2013) employ the most credible research design on this topic. The estimates in those papers are consistent with modest undervaluation of energy efficiency, but they can rule out the extreme biases that would be necessary to make the most cost-effective technologies appealing.

is, they are shrouded in the sense of Gabaix and Laibson [2006]). If consumers are rationally inattentive to the nonsalient technologies, then producers will bring only the salient ones to market. Given the products on offer in equilibrium, consumers will rationally value the technologies, precisely because firms brought to market only those technologies that consumers would properly evaluate. There may simultaneously be cost-effective (nonsalient) technologies that are not adopted. By necessity, all revealed-preference empirical studies observe how consumers trade off the goods on offer in the market. Such studies will not be able to detect investment inefficiencies due to inattention, and inattention thus creates an important caveat to the dominant empirical methodology in the literature.

Those who believe that investment inefficiency is widespread tend to point to bottom-up engineering estimates, which frequently find that there are a host of cost-effective innovations that are not brought to market. Often these innovations are small in nature. For example, in the automobile market, the National Research Council (2011) concludes that the cost-effective technology improvements that appear to pay for themselves but are not undertaken 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. Related studies find similar results (see National Research Council 2002; National Highway Traffic Safety Administration 2010). Such facts are consistent with rational inattention.

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 detect undervaluation. Busse, Knittel, and Zettelmeyer (2013), for example, relate 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 percent of vehicles with the worst fuel economy relative to the 25 percent of vehicles with the best fuel economy. They conclude that price variation is consistent with full valuation (though a null hypothesis of modest undervaluation cannot be ruled out), whereas Allcott and Wozny (forthcoming) use finer-grained variation and settle on estimates that suggest modest undervaluation.

Note also that rational inattention could create measurement bias in empirical studies. In the terminology of the model, rationally inattentive consumers use $\bar{c}_j$ to make choices, but the econometrician typically uses $c_j^\star$, which introduces mismeasurement. Mismeasurement could lead to attenuation bias, but it might also have more complicated biases in discrete-choice models.[32]

---

[32] See Bento, Li, and Roth (2012) for a closely related issue regarding how heterogeneity in valuation can bias discrete-choice models.

Second, in terms of policy design, rational inattention could cause investment inefficiency, but it is important to keep in mind that it implies that there are real economic costs involved in paying attention. For consumers, if making better choices requires costly effort, then a policy that coerces or incentivizes attention would create benefits (from improved choice) but also costs (from exerted effort). If inattention is rational, then consumers have demonstrated by revealed preference that the costs outweigh the benefits, and this balance cannot be overturned by a policy that simply induces attention. Such a policy could have a net positive social welfare cost if it causes a change in the set of products that producers provide to the marketplace, but in the absence of these product design effects, rational inattention implies that paying attention to energy efficiency is welfare decreasing.

This may shift the focus of policy away from corrective taxation and toward information provision. If policy can reduce the cost of attention (rather than incentivize it), then this will unambiguously improve welfare. The most obvious way of reducing attention costs is to direct more effort toward the effective 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 miles-per-gallon rating, though it has added the gallons-per-mile rating. 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. Fortunately, one recent study, Newell and Siikamäki (2013), takes up this issue using laboratory experiments.

A 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 increasing 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 Sallee (2013). These inefficiencies might be justified, however, if inattention causes consumers to ignore fine-grained information, whereas streamlined information, such as whether 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, which creates a justification for such policies. This possibility is explored in more detail in Houde (2013).

Rational inattention also suggests a rationale for minimum efficiency standards and technology mandates, especially for products like laptops or microwaves for which energy consumption is likely a secondary factor in choice. Consider again the situation in which some types of cost-effective technologies are not brought

into the marketplace because consumers are inattentive to them. Pareto improvements may be available because if all firms adopt the technology, their competitive balance could be unchanged, whereas all consumers would benefit (albeit unwittingly).[33] Note that there is a potential feedback loop between standards and inattention. As standards become more stringent, there is likely to be a compression in the distribution of energy efficiency across products. This lowers the variance of energy costs across products and makes inattention more appealing. A deeper examination of the relationship between minimum efficiency standards and rational inattention would be a promising area for future research.

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 that consumers themselves do not know—like how much they drive or idiosyncratic driving styles that affect fuel economy—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 hybrid vehicles was effective during a time period when there was excess demand for the model that dominated sales, the Toyota Prius, and therefore the subsidy had little to 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 hybrid vehicles), 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

---

[33] The optimality of standards is more complicated if some consumers are attentive and others are not. Parallels exist between this case and the model of Allcott, Mullainathan, and Taubinsky (2014).

of driving people's attention toward something are ambiguous. That is, reducing information acquisition costs is obviously beneficial, but causing consumers 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 developed here provides a framework for thinking about conditions under which rational inattention is likely to emerge. In particular, rational inattention is more likely when uncertainty around variation in energy costs is low, when variation in the differentiation of products in other dimensions is large, and when effort costs of paying attention are large. Empirically, this paper suggests that inattention might be rational in the automobile market.

The findings here suggest several directions for future research. I do not directly test for inattention or prove its existence; I demonstrate only that it is plausible that inattention is rational. Direct tests for inattention should be a key goal for future research. Houde (2014), which integrates costly attention into a discrete-choice model for refrigerators, is an important step in this direction. Future research should also aim to advance our understanding of how labels should be optimally designed to aid consumers' information processing. Newell and Sii-kamäki (2013), which uses stated-preference evidence on water heaters, makes a valuable contribution on this question, but additional work on other goods that uses revealed-preference data would also be valuable. Finally, formal models of the optimal design of policies—including information provision, technology mandates, and Pigouvian taxes—in the presence of inattention would give valuable guidance to policy makers.

Appendix

Variations in the Models

The baseline model in the paper assumes that the uncertain fuel costs are independent draws across products and that consumer beliefs are unbiased. This appendix first recasts the baseline model as one in which uncertainty about fuel costs owes to heterogeneity across consumers. Then it discusses the possibility of biased beliefs. In either case, it is not logical to assume that the fuel cost errors are independent of one another. Nevertheless, the same qualitative insights about when inattention would be rational emerge in either case.

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,

$\gamma_i^\star \bar{c}_j$. In the case of automobiles, an example would be if consumer i drives $\gamma_i^\star \bar{m}$ miles, where $\bar{m}$ is the mean number of miles driven, his or her fuel costs will be $\gamma_i^\star \bar{c}_j$. If consumers are certain of $\bar{c}_j$ but uncertain about $\gamma_i^\star$, then the model is very similar, and utility can be written as

$$U_{ij} = \beta' X_j - p_j - (\gamma_i + u_i)\bar{c}_j + \varepsilon_{ij} \equiv Y_{ij} - u_i \bar{c}_j. \tag{A1}$$

The assumption parallel to the model above would be to assume that they have an unbiased belief about their heterogeneity γ_i, such that $\gamma_i^\star = \gamma_i + u_i$ and $\mathbb{E}[\gamma_i] = \gamma_i^\star$. This is exactly the same as the original setup in equation (2), 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 of 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 applies.

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

$$\int_{-\infty}^{\infty} (Y_{ik} - u_i \bar{c}_k)dF(u_i), \tag{A2}$$

where $F(\cdot)$ is the distribution function of u_i. Define $u_i^\star$ as the value at which the consumer is indifferent between goods j and k; that is, $u_i^\star = (Y_{ik} - Y_{ij})/(\bar{c}_k - \bar{c}_j)$.[34] 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 $u_i < u_i^\star$ but will otherwise choose j. In this case, the expected payoff under full information is

$$\int_{-\infty}^{u_i^\star} (Y_{ik} - u_i \bar{c}_k)dF(u_i) + \int_{u_i^\star}^{\infty} (Y_{ij} - u_i \bar{c}_j)dF(u_i) - s. \tag{A3}$$

Taking the difference between expressions (A2) and (A3) and rearranging yields the following condition. Consumer i will exert effort if and only if

$$\int_{u_i^\star}^{\infty} u_i(\bar{c}_k - \bar{c}_j)dF(u_i) > [1 - F(u_i^\star)](Y_{ik} - Y_{ij}) + s. \tag{A4}$$

The left-hand side of expression (A4) is the expected gain from switching to

---

[34] Note that it makes sense to impose a restriction on the support of $u_i^\star$ 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 his or her mind. In such cases, the consumer will obviously not exert effort.

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.

Straightforward analysis of expression (A4) 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 increase the probability of effort because it will increase the truncated mean of u_i above $u_i^\star$.[35] 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 2.

The model can also be adapted to account for beliefs that are biased. 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 4.1, a limited number of automobile attributes account for over 80 percent 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 increase 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 choose between alternatives with similar attributes because those products will have correlated errors.

To see this, consider the simplest case in which a product has only a single attribute X_j, which determines utility and is correlated with fuel consumption. Let the true fuel consumption of product j be $c_j^* = \alpha X_j + c_j$, where c_j has a mean of 0 and is 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 ignores the attribute X_j of each model but is 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

---

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

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 no) effect on the consumer's rank ordering of the goods, which is all that matters for choice. That is, as β gets larger compared with α, failing to account for X in fuel cost beliefs will have less impact on choice.

For example, suppose that consumers have 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 increase the effort cost of becoming informed without necessarily increasing the benefits of information.

References

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

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

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

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

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. 2013. What Do Consumers Believe about the Future Price of Gasoline? *Journal of Environmental Economics and Management* 66:383–403.

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* 101:1375–1409.

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* 115:44–48.

Berry, Steven T. 1994. Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics* 25:242–62.

Bordalo, Peter, Nicola Gennaioli, and Andrei Shleifer. 2013. Competition for Attention. Working Paper No. 19076. National Bureau of Economic Research, Cambridge, Mass.

Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. Are Consumers Myopic? Evidence from New and Used Car Purchases. *American Economic Review* 103:220–56.

Copeland, Adam, Wendy Dunn, and George Hall. 2011. Inventories and the Automobile Market. *RAND Journal of Economics* 42:121–49.

Davis, Lucas W. 2014. The Economic Cost of Global Fuel Subsidies. *American Economic Review: Papers and Proceedings* 104:581–85.

Dubin, Jeffrey A., and Daniel L. McFadden. 1984. An Econometric Analysis of Residential Electric Appliance Holdings and Consumption. *Econometrica* 52:345–62.

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

Gabaix, Xavier. 2013. A Sparsity-Based Model of Bounded Rationality, Applied to Basic Consumer and Equilibrium Theory. Unpublished manuscript. New York University, Stern School of Business, New York.

Gabaix, Xavier, and David Laibson. 2005. Bounded Rationality and Directed Cognition. Unpublished manuscript. New York University, Stern School of Business, New York.

———. 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics* 121:505–40.

Gillingham, Kenneth, and Karen Palmer. 2013. Bridging the Energy Efficiency Gap. Discussion Paper No. 13-02. Resources for the Future, Washington, D.C.

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

Greene, David L. 2010. *How Consumers Value Fuel Economy: A Literature Review.* Technical Report No. EPA-420-R-10-008. Washington, D.C.: U.S. Environmental Protection Agency.

———. 2011. Uncertainty, Loss Aversion, and Markets for Energy Efficiency. *Energy Economics* 33:608–16.

Hausman, Jerry A. 1979. Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables. *Bell Journal of Economics* 10:33–54.

Heutel, Garth. 2011. Optimal Policy Instruments for Externality-Producing Durable Goods under Time Inconsistency. Working Paper No. 17083. National Bureau of Economic Research, Cambridge, Mass.

Houde, Sébastien. 2013. Bunching with the Stars: How Firms Respond to Environmental Certification. Unpublished manuscript. University of Maryland, Department of Agricultural and Resource Economics, College Park.

———. 2014. How Consumers Respond to Product Certification: A Welfare Analysis of the Energy Star Program. Working Paper No. 20019. National Bureau of Economic Research, Cambridge, Mass.

Howarth, Richard B., and Bo Andersson. 1993. Market Barriers to Energy Efficiency. *Energy Economics* 15:262–72.

Hutton, R. Bruce, and William L. Wilkie. 1980. Life Cycle Cost: A New Form of Consumer Information. *Journal of Consumer Research* 6:349–60.

Kilian, Lutz, and Eric R. Sims. 2006. The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data. Unpublished manuscript. University of Michigan, Department of Economics, Ann Arbor.

Langer, Ashley. 2012. Demographic Preferences and Price Discrimination in New Vehicle Sales. Unpublished manuscript. University of Michigan, Ford School of Public Policy, Ann Arbor.

Langer, Ashley, and Shaun McRae. 2013. Step on It: Evidence on the Variation in On-Road Fuel Economy. Unpublished manuscript. University of Arizona, Department of Economics, Tucson.

Larrick, Richard P., and Jack B. Soll. 2008. The MPG Illusion. *Science* 320:1593–94.

Li, Shanjun, Christopher Timmins, and Roger H. von Haefen. 2009. How Do Gasoline

Prices Affect Fleet Fuel Economy? *American Economic Journal: Economic Policy* 1:113–37.

Lu, S. 2006. *Vehicle Survivability and Travel Mileage Schedules.* Technical Report No. DOT HS 809 952. Washington, D.C.: National Highway Traffic Safety Administration.

Matějka, Filip, and Alisdair McKay. 2013. Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. Unpublished manuscript. Boston University, Department of Economics, Boston.

McNeill, Dennis L., and William L. Wilkie. 1979. Public Policy and Consumer Information: Impact of the New Energy Labels. *Journal of Consumer Research* 6:1–11.

Metcalf, Gilbert E. 1994. Economics and Rational Conservation Policy. *Energy Policy* 22: 819–25.

National Highway Traffic Safety Administration. 2010. *Final Regulatory Impact Analysis: Corporate Average Fuel Economy for MY 2012–MY 2016 Passenger Cars and Light Trucks.* Washington, D.C.: National Highway Traffic Safety Administration.

National Research Council. 2002. *Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards.* Washington, D.C.: National Academies Press.

———. 2011. *Assessment of Fuel Economy Technologies for Light-Duty Vehicles.* Washington, D.C.: National Academies Press.

Newell, Richard G., and Juha V. Siikamäki. 2013. Nudging Energy Efficiency Behavior: The Role of Information Labels. Working Paper No. 19224. National Bureau of Economic Research, Cambridge, Mass.

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

Sallee, James M. 2011a. The Surprising Incidence of Tax Credits for the Toyota Prius. *American Economic Journal: Economic Policy* 3:189–219.

———. 2011b. The Taxation of Fuel Economy. *Tax Policy and the Economy* 25:1–37.

———. 2013. Rational Inattention and Energy Efficiency. Working Paper No. 19545. National Bureau of Economic Research, Cambridge, Mass.

Sallee, James M., and Joel Slemrod. 2012. Car Notches: Strategic Automaker Responses to Fuel Economy Policy. *Journal of Public Economics* 96:981–99.

Sallee, James M., Sarah E. West, and Wei Fan. 2009. Consumer Valuation of Fuel Economy: A Microdata Approach. *Proceedings of the 102nd Annual National Tax Association Conference on Taxation,* pp. 254–59.

Sawhill, James. 2008. Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the U.S. Automobile Market. Working paper. University of California, Haas School of Business, Berkeley.

Sims, Christopher A. 2003. Implications of Rational Inattention. *Journal of Monetary Economics* 50:665–90.

Stigler, George J. 1961. The Economics of Information. *Journal of Political Economy* 69: 213–25.

Train, Kenneth E., and Clifford Winston. 2007. Vehicle Choice Behavior and the Declining Market Share of U.S. Automakers. *International Economic Review* 48:1469–96.

Turrentine, Thomas S., and Kenneth S. Kurani. 2007. Car Buyers and Fuel Economy? *Energy Policy* 35:1213–23.

U.S. Department of Transportation. 2008. *Corporate Average Fuel Economy for MY 2011–2015 Passenger Cars and Light Trucks: Preliminary Regulatory Impact Analysis.* Washington, D.C.: U.S. Department of Transportation.