

The Heterogeneous Effects of Eco-labels on Internalities and Externalities

Anshuman Sahoo*

Graduate School of Business

Steyer-Taylor Center for Energy Policy and Finance

Stanford University

and

Nik Sawe

Emmett Interdisciplinary Program in Environment and Resources

Stanford University

March 22, 2015

*Contact information:

A. Sahoo (corresponding author)

655 Knight Management Center, Graduate School of Business, Stanford University, Stanford, CA 94305.

Phone: +1 713 305 2204; e-mail: asahoo@stanford.edu.

N. Sawe

Yang and Yamazaki Building, Suite 226, Stanford University, Stanford, CA 94305. Phone: +1 650 814 4648;

e-mail: sawe@stanford.edu.

A previous version of this paper was circulated under the title “Negative Dividends: Internality Losses can Outweigh Externality Gains”. We gratefully acknowledge the financial support of the Precourt Energy Efficiency Center at Stanford University and John Beshears, Sebastien Houde, Brian Knutson, Sam McClure, and James Sweeney for their comments on the experimental procedure and analysis. We also thank Taurean Butler and Natalie Luu for their research assistance. All errors are our own.

The Heterogeneous Effects of Eco-labels on Internalities and Externalities

Abstract:

The Energy Star program labels energy efficient goods and has been credited with reducing the external costs of energy consumption. Its social value is nonetheless ambiguous if, in its absence, some consumers over-value while others under-value the energy consumption attribute of goods, relative to their economic preferences. The label must perform opposite tasks to guarantee an increase in consumer welfare: it must prompt some individuals to increase their valuation of the energy consumption attribute and others, to decrease it. Otherwise, the program could yield “negative dividends” by inducing losses in individual-level welfare that outweigh externality reductions. We develop a method to quantify the impact of the program on individual-level decision-making behavior and welfare. Using novel data from a stated choice experiment involving light bulbs, we illustrate the potential for negative dividends and that the value of programs such as the Energy Star depends on the choice set available to consumers.

Keywords: Energy Star, Energy efficiency, Internality, Stated choice, Light bulbs.

JEL codes: D12, H23, Q48

1 Introduction

Increases in energy efficiency are often considered among the most cost-effective methods to reduce energy demand and the environmental externalities associated with energy consumption. Since residential use accounts for about 22% of U.S. energy consumption, changes in appliance stock could meaningfully shift U.S. energy demand. However, consumers frequently choose not to purchase energy efficient alternatives that should provide them with net savings. In seminal work, Hausman (1979) showed that consumers choosing appliances do not buy goods that would have been optimal at a market interest rate but act instead as if they discount future savings on energy consumption at a much higher rate. Since then, a large body of literature has attempted to explain how consumers make such decisions (Gillingham and Palmer, 2014).

Embedded in several explanations is a notion that consumers do not pay enough attention to the energy consumption attribute of alternative appliances.¹ Standard economic theory assumes that all consumers make consumption decisions as if they were maximizing a utility function using all available information; the rational consumer would thus pay full attention to the energy consumption attribute. Nonetheless, inattentiveness may be pervasive in consumer choice. For example, Chetty, Looney, and Kroft (2009) show that consumers imperfectly consider taxes on common consumer goods. In the energy sector, Allcott (2011) shows that U.S. consumers do not pay adequate attention to fuel costs in making vehicle purchase decisions; Leard (2013) similarly suggests that approximately 30% of consumers shopping for a new vehicle completely ignore fuel costs. In the context of appliance choice, Houde (2014) presents evidence characterizing three types of appliance consumers, two of which effectively ignore electricity cost implications. If consumers are insufficiently attentive to the ramifications of energy consumption for themselves, they would certainly be insufficiently attentive from a social perspective, which adds the external costs of energy consumption. Thus, given an expected pattern of product use, the under-investment in energy efficiency by consumers implies a greater than socially optimal level of energy consumption.²

¹We focus on an inattentiveness hypothesis because this, in our opinion, most directly motivates the apparent goal of the Energy Star to increase the salience of the energy consumption attribute of goods. This also motivates the approach in Sallee (2014), in which the consumer must incur significant costs to determine the level of energy efficiency implied by alternative goods. Nonetheless, a large number of alternative explanations have been offered for the perceived low uptake of energy efficient goods.

²Throughout our discussion, we restrict attention to product choice decisions and assume that the pattern

A natural policy reaction is to increase the salience of the energy consumption attribute of alternatives. Common policy responses implement this reaction to differing degrees. One strategy, that of minimum efficiency standards, does not attempt to increase the salience but instead simply removes energy inefficient goods from the market. This approach entails a cost to certain consumers: those who plan to infrequently use the affected good may be compelled to invest in a higher level of efficiency than they would rationally choose. A strategy that can theoretically achieve the first-best outcome is that of attaching a Pigouvian tax on energy-consuming goods. Such goods exert external costs attributable to environmental damages.³ Standard theory prescribes that goods or services should be taxed at a level equal to the marginal external costs they produce. Thus, to achieve the first-best outcome, the policymaker would in general need to know each individual's rate of utilization of the affected energy-consuming good. Absent this information, we would expect a gap to exist between the first-best policy outcome and that actually obtained.

Two strategies attempt to affect only those consumers who do not fully consider the energy consumption attribute. One of these seeks to increase the salience of energy consumption information by explicitly providing it at the time of decision-making. Another class seeks to do so by using “nudges,” or changes in choice architecture, as described by Thaler and Sunstein (2008). By definition, these strategies preserve all alternatives in consumers' choice sets but change the manner in which they are presented. Critical to the appeal of such policies is the supposition that they should not affect the decisions of those consumers who act rationally in the absence of the intervention.

The Energy Star certification and labeling program administered by the U.S. Environmental Protection Agency (EPA) can be interpreted as one such policy. Manufacturers of goods that meet energy consumption benchmarks may voluntarily label compliant products with the Energy Star label. In general, such eco-labels have increased the uptake of

of use remains the same regardless of the product chosen. A study of the implications of the Energy Star for the misoptimization of product use would provide a useful contribution. Our statement about excessive energy consumption relative to a social optimum would be reversed only if consumers' product use increased so radically with goods of higher efficiency that this effect counteracted the lower expected consumption at a constant level of use. The increase in use that accompanies greater energy efficiency is known as the rebound effect, and most estimates of it suggest that our assertion holds (Gillingham and Palmer, 2014).

³They also exert external costs attributable to energy security, since societies maintain military forces partially to guarantee access to energy resources. Since this paper focuses on appliances in the U.S. context, where a very small portion of electricity is generated from resources that are sourced internationally, we do not focus on this second source of externalities.

energy efficient goods; Sammer and Wustenhagen (2006) provide evidence that consumers' willingness-to-pay for products with energy eco-labels exceeds that for products that are equally efficient but unlabeled. Further, Ward et al. (2011) estimate that customers are willing to pay an extra \$250 - \$350 for a refrigerator with an Energy Star label.

If the analyst's policy goal is solely to reduce environmental externalities, this evidence suggests that the Energy Star has been effective. The expression by consumers, on average, of a higher willingness-to-pay for labeled goods implies that they will more frequently select energy efficient alternatives. Nonetheless, this evidence does not comment on whether the program also makes individuals better off, as measured by the degree to which it encourages them to choose products in a manner consistent with utility maximization.

Indeed, while consumers may be willing to pay for eco-labels, it is unclear whether the latter increase or decrease the former's attention on the energy consumption attribute of alternatives. This bidirectionality may underpin recent work showing that consumers respond heterogeneously to eco-labels such as the Energy Star. Both Houde (2014) and Shen and Saijo (2009) suggest that consumers may (1) use labels as a substitute for conventional utility optimization and ignore electricity cost and consumption data, (2) use electricity cost and consumption data to optimize choice but ignore the eco-label, or (3) use neither electricity cost and consumption data nor eco-labels in making a choice. These observations imply a second level of uncertainty. After the Energy Star exerts an impact on the attention an individual pays to the energy consumption attribute, he or she may make choices that reflect a lower or higher valuation of savings on energy consumption, relative to that expressed in the absence of the label. An implication is that some individuals who considered energy consumption rationally in the absence of the label may use the attribute irrationally or not at all in the presence of the Energy Star. If such changes occur, one of the chief appeals of the label and similar behavioral instruments would be at risk.

More broadly, the possibility that the Energy Star could increase some consumers' valuation on energy consumption savings and decrease it for others raises the question that motivates this paper: what is the value of the Energy Star to consumers? We ask this question because the heterogeneous effects on consumers' valuation suggest that the Energy Star could sometimes influence consumers to make decisions that are less consistent with utility maximization. This paper presents results from a stated choice experiment on light bulb choice intended to answer our question. In the experiment, a cross-national sample of 1,550

individuals made hypothetical choices across compact fluorescent light (CFL) bulbs. Using the 46,500 choices we observe in these data, we estimate coefficients on a mixed logit specification stemming from an underlying random utility model. Importantly, the estimation procedure yields utility coefficients, conditional on the sequence of an individual's choice, on both the energy consumption attribute of alternatives and the interaction between it and the presence of the Energy Star. We thus estimate each individual's utility coefficient on energy consumption in the absence and presence of the Energy Star.

This allows us to include in the evaluation of the Energy Star program its impact not only on environmental externalities but also on consumers' internalities, as defined by Herrnstein et al. (1993) and as applied by, e.g., Allcott, Mullainathan, and Taubinsky (2014) and Leard (2013). We expand considerably on the concept of internalities subsequently, but they can be understood as losses in experienced utility, relative to a rational benchmark. Consider a given consumer at different times as different people. This is consistent with standard economic theory as long as the consumer assigns equal weights to all versions of herself in her decision-making process, as this implies that her choices reflect time-consistent preferences. However, if the consumer weighs different versions of herself unequally such that the experienced utility implications of a choice at different times are scaled unequally, her choices reflect time-inconsistent preferences. These choices could imply a loss in experienced utility, and we term such a loss as an "internality." If an under- or over-valuation of energy consumption savings influences a shift in the good selected away from that consistent with the consumer's true preferences, these tendencies will exert internalities on the consumer. If the Energy Star affects the valuation of energy consumption savings, it implies changes in the magnitude of internalities. Our central point in this paper is that these must be quantified in economic analyses of the Energy Star and, in general, of similar behavioral tools.

We facilitate this quantification by establishing a rational benchmark utility coefficient on energy consumption for each individual; this can be interpreted as reflecting time-consistent preferences for energy consumption savings. The benchmark follows from a discount rate elicited from each individual and assumptions about the value of savings on energy consumption. When the consumer acts as if she is applying this coefficient in decision-making, she chooses the good that maximizes her experienced utility. Since a wedge between the rational benchmark utility coefficient and that observed in either the absence or presence of the Energy Star could prompt the selection of a good that does not maximize experienced utility, it

could yield an internality. The larger this gap, the larger the potential for misoptimization of product choice and of losses in experienced utility to the consumer. Upon estimating the size of the gaps in the presence and absence of the Energy Star, we can determine whether the label mitigates or exacerbates such wedges at both the individual and population level.

The chief conceptual contribution of this paper is a method by which to value the impacts of behavioral policy instruments, including their implications for both externalities and internalities. The latter usually have not been formally included in such assessments. While our methodological approach is motivated and illustrated by the Energy Star, it can be applied to other contexts as long as the analyst is willing to make two assumptions. First, aside from the utility coefficient on the attribute(s) affected by the behavioral intervention, all estimated utility coefficients reflect time-consistent preferences. The second assumption is that of the nature of long-term costs and benefits. In contexts such as ours, in which individuals face an immediate cost and long-term benefits or an immediate benefit and long-term costs, the second assumption requires assumptions about the magnitude of these benefits or costs and the duration over which these obtain. The specification of these costs and benefits becomes much more difficult in scenarios where estimates of the experienced utility losses over time reflect much more heavily the analyst's judgment. Importantly, we highlight that the implications of behavioral instruments for internalities and externalities are functions of the choice set available to the consumer. Some choice sets may be characterized by a dominant alternative such that, even when wedges exist between the rational benchmark value of an utility coefficient and that actually used in decision-making, no changes in experienced utility (i.e., internalities) obtain. This is in contrast to those choice sets offering a rich set of trade-offs to the consumer. In such settings, the experienced utility realized is sensitive to the value of the utility coefficient.

Our main findings in the Energy Star context are as follows. In aggregate, the Energy Star label appears to do what it was designed to do: it increases consumers' valuation of savings on energy consumption. More often than not (for 53% of our participants), the effect of the Energy Star is aligned with the individual's own interest. The overall value of changes in experienced utility across the population and attributable to the program depends on the marginal utilities of wealth among the affected population. We examine the performance of the Energy Star on a choice set offered by a national retailer and find that the program yields modest increases in experienced utility and decreases in external costs. We observe

that the overall value of the Energy Star program derives principally from its impact on internalities instead of on externalities. While we do not present results from an alternate choice set, this positive result for both internalities and externalities is not a general result across all choice sets.⁴

The work in this paper contributes to two bodies of literature. One studies decision-making processes in appliance choice. Our work is particularly motivated by Houde (2014) in quantifying the implications of heterogeneous responses to the Energy Star that he identifies. It also extends the study of eco-labeling, which has hitherto focused mostly on measuring consumers' willingness-to-pay (Sammer and Wustenhagen, 2006; Shen and Saijo, 2009; Ward et al., 2011). The work by Newell and Siikamaki (2014) considers how the informational content of labels guide decision-making. While we overlap conceptually with them in using an objectively determined optimal behavior construct, we differ in focusing on individual differences in response to the Energy Star label rather than on differences in response to alternative labels. Like us, Min et al. (2014) focus on the Energy Star, but they present aggregate-level implications of eco-labeling on consumers' decision-making processes.

The other literature studies how the interactions of internalities and externalities change the choice and use of policy instruments. Our work follows Leard (2013) closely in its treatment of internalities but differs in both the context and the examination of a non-price policy instrument. More broadly, our work is related to Allcott, Mullainathan, and Taubinsky (2014), who consider how the presence of both externalities and internalities affect energy policy-making. They suggest that pecuniary instruments to offset externalities yield a double dividend by also reducing internalities. However, they suggest that if such taxes are poorly targeted, other policies, such as information provision and nudges may be more appealing. Their intuition is that these policies are designed to affect only those who do not act rationally. Our work suggests an important nuance: the heterogeneous effects of such behavioral instruments may imply a *negative dividend*.

Although our paper illustrates this nuance in the context of policy evaluation, analogous lessons apply to managers. While changes in choice architecture could increase profits, certain classes of consumers may preferentially gain while others lose. Though regulatory guidance on choice architecture is unclear, if the harm to the latter segment of consumers is sufficiently large, changes in choice architecture may prompt action by bodies with a mandate

⁴We provide a specific counter-example in Section 5.3.

to protect consumer welfare. More directly, managers should not assume that the average change in consumer response following a change in choice architecture will apply to an entire population of consumers. Our illustration in the eco-labeling context provides one example of the importance of individual-level differences in response behavior. While average changes in consumer behavior may imply an increase in profitability, the distribution of responses may suggest just the opposite. Managers should attempt to characterize individual differences in consumer behavior to ensure that they both avoid reprimands by consumer watchdogs and select profit maximizing strategies.

2 Experimental and Analytical Methods

Our analysis uses compact fluorescent light bulb (CFL) choice data from a stated choice experiment. In this section, we describe our experiment and economic model of choice.

2.1 Experimental Details

Figure 1 describes the experimental procedure. We contracted with a third-party survey administration firm to recruit a cross-national panel. Eligible participants were U.S. residents who were at least 18 years old, owned their primary residence, and had either purchased or remodeled their residences within five years. The last criterion sought to limit participants to those who had recently made decisions about energy-consuming goods.

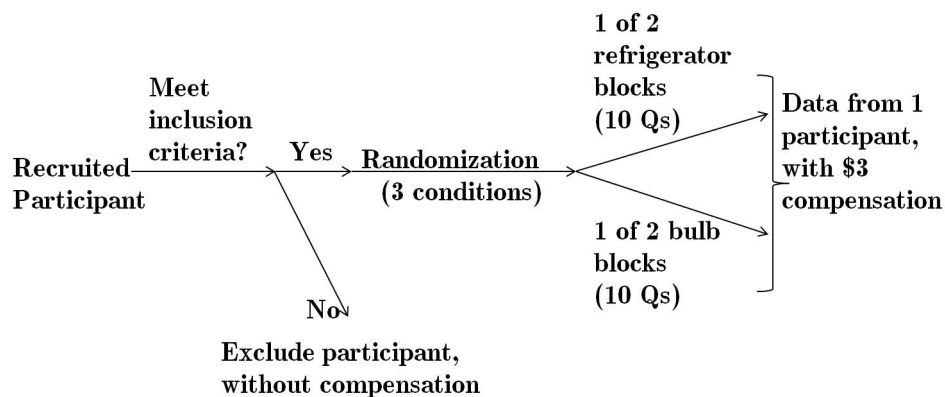


Figure 1: Overview of experimental flow

Those participants who did not meet the inclusion criteria or incorrectly answered one of two questions designed to test attentiveness were dismissed without compensation, while

those who did were randomized into one of three conditions.⁵ This variation was a precursor to an experiment that will be implemented separately. The three conditions differed in the denomination of the consumption information provided. We re-phrased the accompanying interpretive information to match the denominations. Condition 1 provided participants with information in only power (Watt (W)) units. In Condition 2, participants received information in terms of the estimated annual cost of operating the good. Condition 3 provided consumption information in both units. The Appendix provides the interpretive information accompanying each condition. We adjust the data to account for differences in error variances across conditions (see Section 3.1).

Participants performed a series of ten light bulb and refrigerator choice tasks. This paper focuses on the light bulb data to isolate our message about the implications of heterogeneous responses to the Energy Star for the value of the program. For each product, twenty choice sets were organized into two blocks of ten choice sets. Following the randomization step, individuals responded to randomly assigned light bulb and refrigerator choice blocks. We generated the constituent choice sets using Ngene (ChoiceMetrics, 2012). Our program applied priors on coefficient values from a pilot stated choice experiment to build attribute levels designed to optimize the efficiency of our estimation. The range of prices and energy consumption were such that all alternatives could have been Energy Star certified.




As Figure 2 illustrates, each participant selected one of three alternatives from each of ten choice sets. An alternative is characterized by three primary attributes: price, energy consumption, and the Energy Star certification. To increase the reality of the task, we labeled each bulb as providing light of one of four colors.⁶ Each participant received instructions about the task, information with which to interpret the attributes, and a set of assumptions to make about attributes, such as the lifetime of the light bulb, which we did not vary. The Appendix details this material. We collected 46,500 choice observations from our 1,550 participants.

Subsequently, each participant answered 144 demographic, psychological, and financial questions intended to characterize individuals.⁷ We consolidated many of these questions into standard metrics (e.g., 21 questions about time preferences implied an estimate of the

⁵As we discuss later, we did not find material differences among consumers' characteristics or choice behaviors across conditions.

⁶The color describes the warmth of the light as "warm white," "soft white," "neutral" or "cool."

⁷The Appendix provides details about these questions.

	Bulb 1	Bulb 2	Bulb 3
Price	\$7.69	\$8.72	\$2.45
Light color	Cool	Neutral	Cool
Energy Star certified?			
Energy demand (Watts)	8.94	12.60	13.02

Which of the above bulbs would you purchase?

- Bulb 1
- Bulb 2
- Bulb 3

Figure 2: Example bulb choice set

participant’s temporal discount rate). Following this consolidation, our data included 35 demographic, psychological and financial variables.

2.2 Economic Model

2.2.1 Base utility assumptions and choice model

We assume that participants choose the alternative yielding the highest decision utility and thereby model choices with the random utility model (RUM). Formally, individual i selects good j among J alternatives in set t of T choice situations to maximize her utility:

$$U_{ijt} = V_{ijt} + \varepsilon_{ijt} = \beta X_{ijt} + \varepsilon_{ijt} \quad (1)$$

X_{ijt} is a vector of explanatory variables that includes the attributes of alternatives, V_{ijt} or βX_{ijt} is the deterministic component of utility, and ε_{ijt} is a random error term. The probability, P_{ijt} , that individual i selects alternative j over k in choice set t is:

$$P_{ijt} = P(U_{ijt} > U_{ikt}) \forall k \neq j \in t \rightarrow P_{ijt} = P(\varepsilon_{ikt} - \varepsilon_{ijt} < \beta X_{ijt} - \beta X_{ikt}) \forall k \neq j \in t \quad (2)$$

A traditional assumption is that the error terms are identically and independently distributed (iid) according to an extreme value type I distribution. With this assumption, we obtain

the multinomial logit (MNL) model, which implies that the probability that alternative j is selected, conditional on β , is given by:

$$P_{ijt}(\beta) = \frac{e^{\sigma\beta X_{ijt}}}{\sum_{k=1}^K e^{\sigma\beta X_{ikt}}} \quad (3)$$

Above, σ is a scale parameter on utility that enters the model from our assumed joint distribution of ε_{ijt} . The distribution is characterized by both location and scale parameters. We return to the scale parameter in Section 3.1 but assume now that it has been normalized to 1.

The MNL model has three limitations that we overcome by estimating a mixed logit (ML) model. ML models allow for random taste variation, unrestricted substitution patterns by permitting any specification of the correlation pattern across alternatives, and a correlation in unobserved factors over choices made by the same individual (Revelt and Train, 1998).

2.2.2 Allowing preference heterogeneity through the mixed logit model

Mixed logit models “mix” MNL estimates by integrating over a joint density of parameters characterizing distributions of coefficients on particular attributes of the alternatives (Train, 2009). As such, we update the probability that alternative j is selected in choice set t to:

$$P_{ijt} = \int \frac{e^{\sigma\beta_i X_{ijt}}}{\sum_{k=1}^K e^{\sigma\beta_i X_{ikt}}} f(\beta | \Omega) d\beta \quad (4)$$

Equation 4 introduces a mixing distribution, $f(\beta | \Omega)$, over all random coefficients. The set Ω includes parameters of the distributions on random coefficients. If one specifies that marginal coefficient distributions are correlated, as we do, the parameter space also includes co-variance terms between correlated coefficients. Since an analytical solution for the likelihood function does not exist, we estimate coefficients by maximizing a simulated log likelihood function.⁸ The function is defined by the product of the probabilities that individual i purchases the product actually chosen, j^* , in choice situation t :

$$SLL = \sum_t \sum_i \ln P_{ij^*t} \quad (5)$$

⁸Train (2009) and Hensher and Greene (2003) provide additional details about simulated maximum likelihood.

Briefly, the procedure assumes parameter values for all elements of Ω , draws coefficient vectors, given these parameter values, calculates the probability of selection of the alternative actually selected, given the coefficient vector, repeats the first three steps many times, averages the probability implied by each of these repeats, and selects the parameter values Ω_{SMLE} that yield the highest simulated likelihood.⁹

2.2.3 From the mixed logit to “individual-level” parameters

Since we observe repeated choices from each participant, we can refine the mixed logit-based coefficient estimates by conditioning on the sequence of choices, \tilde{y}_i , made by participant i in the sequence of choice situations, \tilde{t}_i . We accordingly assign each individual a conditional parameter distribution, rather than the sample-wide parameter distribution. Following Train (2009), upon conditioning on \tilde{t}_i and \tilde{y}_i , we update the distribution of coefficients based on Ω_{SMLE} , denoted $g(\beta | \Omega_{SMLE})$, to $h(\beta | \tilde{y}_i, \tilde{t}_i, \Omega_{SMLE})$.

We do not provide the full intuition behind the distribution h , since it is available in Train (2009). However, h re-scales the distribution g , as shown in Equation 6:

$$h(\beta | \tilde{y}_i, \tilde{t}_i, \Omega_{SMLE}) = \frac{P(\tilde{y}_i | \tilde{t}_i, \beta)g(\beta | \Omega_{SMLE})}{P(\tilde{y}_i | \tilde{t}_i, \Omega_{SMLE})} \quad (6)$$

Equation 6 follows from an application of Bayes’ Rule and indicates that the conditional density of the coefficient vector, h , among those who choose the sequence \tilde{y}_i when choosing among \tilde{t}_i is proportional to the product of the unconditional density and the probability that \tilde{y}_i would be chosen if the coefficient vector were β .

We estimate the conditional distribution of coefficients for each individual by simulation. The simulation draws coefficient values β from the population density, $g(\beta | \Omega_{SMLE})$, calculates the probability of observing \tilde{y}_i given this draw, and takes the weighted average of the draws. The weights are determined by the ratio of the calculated probability for a particular draw to the sum of probabilities across all draws. (Train, 2009)

Since Ω_{SMLE} is itself estimated with sampling error, we add a second layer of simulation. This second layer draws Ω from $N(\Omega_{SMLE}, W_{SMLE})$, the estimated sampling distribution of Ω , using a Choleski decomposition of W_{SMLE} (Train, 2009). Thus, the overall simulation procedure entails multiple draws of both β and Ω . We accomplish the latter by taking 500 Halton draws and the former by taking 500 samples from the sampling distribution.

⁹The subscript “SMLE” denotes simulated maximum likelihood estimator.

Though we refer colloquially to the resulting β_i as individual-level coefficients, they are correctly interpreted as population coefficients, conditional on the choices made by individual i . Since the data generation process includes a fixed number of observations for each participant, the conditional mean coefficient vector, $\bar{\beta}_i$, is not a consistent estimate of β_i .

Finally, the individual-level coefficients imply individual-level WTP measures on attribute l of alternatives:

$$WTP_{il} = -\frac{\beta_{il}}{\eta_i} \quad (7)$$

Above, η_i is the marginal utility of wealth. The individual-level measures imply a sample-wide mean and standard deviation of the WTP.¹⁰

2.2.4 Specifying the observed component of utility

We have so far abstracted from the observable components of utility, which Equation 8 details:¹¹

$$V_{ij} = \eta_i \cdot p_j + \theta_i \cdot C_j + \lambda_i \cdot ES_j + \gamma_i \cdot (ES_j)(C_j) + \Upsilon_i^T Z_j \quad (8)$$

In the above model, p_j is the price of product j , θ_i is the marginal utility of an additional unit of power rating, C_j is the power rating of the good, λ_i is the marginal utility from the presence of the Energy Star logo, ES_j is a dummy variable equal to 1 if the product is Energy Star labeled, and γ_i is the incremental marginal utility of an additional unit of energy consumption when the product is Energy Star labeled. Finally, Υ_i is a vector of marginal utilities from the warmth of the light emitted, and Z_j is a vector of dummy variables on the warmth levels. Though we allow the utility to reflect the color of light produced by a bulb, we assume that the utility from the lighting service itself is uniform across all bulbs and irrelevant to our modeling of differences in utility. Given a particular pattern of product use, the energy consumption of the bulb is a monotone transformation of the power rating.

¹⁰We also compute parameters of the WTP distribution by performing 1,000,000 draws of the constituent mixed logit utility coefficients. Though we do not report the simulation-based WTP measures, the mean and median WTP measures are close when calculated by either approach. The standard deviations from the simulation-based method are much larger than those implied from the aggregation of the individual-level coefficients. This reflects the larger variance in the unconditional distribution relative to the conditional distribution.

¹¹As mentioned above, this model should be interpreted as describing decision utility instead of experienced utility.

Since our experiment entailed specific assumptions about the usage pattern, we refer to the attribute as energy consumption.¹² Note that this specification allows utility to increase or decrease by the presence of the Energy Star label independently of the energy consumption of the particular product. This affords us the ability to separately assess whether the Energy Star label leads people to place more or less weight on the actual energy consumption.

We allow four utility coefficients to vary randomly across individuals: η_i , θ_i , λ_i , and γ_i . Since theory suggests that utility decreases in prices and energy consumption, we constrain η_i and θ_i to be strictly negative by using a lognormal mixing distribution for both.¹³ Theory does not provide such guidance for the other parameters, and we use a normal mixing distribution for them. Since heterogeneity in preferences for light bulb warmth levels is uninteresting in this context, we set $\Upsilon_i = \Upsilon \forall i$. Finally, we allow the coefficient distributions to be correlated, since theory suggests, for example, that individuals with high marginal utilities of wealth may also have high marginal disutilities from energy consumption.

3 Aggregate Results

5,919 respondents entered our experiment, and we arrived at our sample of 1,550 upon removing those that either did not meet the inclusion criteria or correctly respond to the attentiveness questions. We present summary demographic, psychological, and financial data in Table 6 of the Appendix; the lack of significant differences across the three conditions suggests that our sample is balanced across them. In Section 3.1, we provide details about our estimation procedure, and in Section 3.2, we discuss our estimated ML coefficients.

3.1 Estimation Details

Three comments about our estimation procedure clarify the interpretation of the coefficients. First, regardless of the units of consumption with which a participant was presented, a deterministic link between bulb wattage and expected annual expenditure on electricity allows us to use a singly denominated vector for estimation. Second, the standard errors we use for inference are based on cluster robust variance estimators (CRVE), with observations clustered

¹²Moreover, consumption, not the rating, is the source of the internalities and externalities we subsequently consider.

¹³A lognormal distribution for price moreover assures finite moments for willingness-to-pay distributions (Daly, Hess, and Train, 2012).

at the individual level. This allows for heterogeneity in the error structure across individuals and for an arbitrary autocorrelation structure in the errors across choice situations faced by the same individual (Wooldridge, 2010). Finally, we relax the assumption that the scale parameter in Equation 3 equals one. This assumption generally allows one to estimate β even though it is not separately identified from σ by the data. Since we combine data from three experimental conditions, we need to ensure that our estimate of β is not influenced by differences in the error variance across conditions. To adjust for possible differences, we follow methods developed by Swait and Louviere (1993). If the hypothesis of equal scale factors across conditions is rejected, the method estimates condition-specific scale factors with which to adjust the co-variate matrix prior to estimation. The Appendix details our derivation of scale parameters on Conditions 1 and 2. The results we present here are based on estimations with condition-specific scale parameters, but they do not change appreciably if we reset them to equal 1.

3.2 Mixed Logit Coefficient Estimates

Table 1 presents the results of our mixed logit estimation.¹⁴ Since the ML estimates follow from a model with dummy terms on three of the four warmth levels, the outside option is a bulb of the remaining warmth level, of average price and consumption, and without Energy Star certification. Coefficients should thus be interpreted as conditional on bulb purchase. The baseline bulb can be of any of the four warmth levels, as the coefficients on price, energy consumption, Energy Star, and the interaction term average over all warmth levels. Column 3 provides the means from the average of conditional (individual-level) coefficients.¹⁵ The agreement between Columns 2 and 3 provides a check that the model is correctly specified

¹⁴It is common to comment on whether the signs matched expectations. Though the means of the coefficient distributions in Column 2 imply that the signs on price and consumption are as expected, we note that these simply follow the assumptions of our mixing distribution. However, we observe the same signs in a model with normal mixing distributions on these attributes. The notion of an ‘expected sign’ is ambiguous for bulb warmth levels, for the Energy Star label, and for the interaction between the Energy Star and consumption level. The positive coefficient on the Energy Star logo is expected to the extent that we assume consumers interpret this as some indicator of quality. We use the Stata ‘mixlogit’ command written by Hole (2007) for our estimation.

¹⁵We modify Hole’s ‘mixlbeta’ command to estimate individual-level means and variance-covariance matrices.

and accurately estimated (Allenby and Rossi, 1998; Train, 2009).¹⁶

Attribute	<i>Dependent variable: Alternative j selected by person i (binary outcome)</i>					
	Mean	Mean (IL)	Variance	WTP, Mean	WTP, Med.	WTP, SD
Price	-0.684*** (0.045)	-1.277	1.825*** (0.129)	–	–	–
Energy Star	1.839*** (0.093)	1.837	2.403*** (0.229)	\$4.97	\$2.36	\$6.46
Consumption	-1.056*** (0.061)	-0.734	1.48*** (0.149)	-\$1.23	-\$0.71	\$1.65
ES*Cons.	-0.126*** (0.026)	-0.126	0.019*** (0.006)	-\$0.15	-\$0.13	\$0.31
Soft white	0.186*** (0.041)	–	–	\$0.51	\$0.35	\$0.48
Neutral	-0.189*** (0.054)	–	–	-\$0.51	-\$0.35	\$0.49
Cool	-0.743*** (0.060)	–	–	-\$2.02	-\$1.39	\$1.93
LogLik	-11914.764					
Obs.	46,500					
Correctly Predicted (%)	60.2					

Table 1: Estimated ML parameters for bulb choice. Columns 5 through 7 list summary statistics on WTP measures. Numbers in parentheses are standard errors. Note: ES*Cons. indicates the Energy Star and Consumption interaction term. Key to statistical significance – ***: ≤ 0.001 ; **: ≤ 0.01 ; *: ≤ 0.05 .

The variance terms in the fourth column validate our consideration of heterogeneous impacts of the Energy Star. All terms are statistically different from zero, and this confirms significant heterogeneity across all four coefficients. By considering this heterogeneity, the ML model outperforms the MNL in correctly predicting 60.2% of choices relative to the latter’s 51.6%.¹⁷

We subsequently consider the implications of heterogeneity, but our aggregate measures imply that participants express a positive WTP for Energy Star certification independently of actual energy consumption. To put the WTP for the Energy Star label in context, the

¹⁶The mean coefficients on price and consumption are parameters of the lognormal distribution; the corresponding parameters of the normal distribution are -1.257 and -0.728, respectively.

¹⁷For individual i and choice set t , the model correctly predicts choice if it assigns the highest probability of selection to j^* , the alternative actually selected.

mean price of bulbs in the study was \$5.56. Further, certification decreases the marginal utility from a 1W higher power rating; thus, Energy Star certification makes consumers more sensitive on average to the energy use of the bulb. In our upcoming discussion, we will refer interchangeably to the effect of the Energy Star on the marginal disutility of a 1W higher power rating and on the marginal utility of energy consumption savings implied. If the Energy Star increases the former, as on aggregate, it also increases the latter.

The estimate of the coefficient on the Energy Star label itself implies that the certification endows products with a “brand” that is almost always positively valued.¹⁸ The magnitude of the median willingness-to-pay for certification, \$2.36, is moreover greater than the overall incremental willingness-to-pay for a one watt lower power rating by a factor of approximately three.¹⁹ The label exerts a larger impact on light bulb choice by imparting a “brand name” than by affecting consumers’ sensitivity to energy consumption information. To the extent that the Energy Star label actually distinguishes efficient light bulbs in the market from inefficient ones, this energy-independent Energy Star effect yields additional reductions in the external costs of electricity generation and consumption.²⁰ However, as we make clear in the subsequent section, there are instances in which an individual is compelled by the Energy Star program to be overly attentive to energy consumption, relative to an individually rational level. In that case, the additional impact of this energy-independent Energy Star effect would make the individual even worse off than does the impact of the label on the utility coefficient on energy consumption. The social value of this effect is thus ambiguous. We comment more on the role of this coefficient in Section 5.2 and illustrate its effect in our example evaluation in Section 5.3.

Finding 1 summarizes our findings from the ML model:

Finding 1: *In aggregate and for our sample, the Energy Star increases consumers’ marginal utility from savings on energy consumption. It also increases consumers’ willingness-to-pay for a bulb, independent of the actual energy consumption. Mixed logit estimates confirm*

¹⁸60 of the 1550 individual-level coefficients on the presence of the Energy Star label are negative.

¹⁹The overall incremental WTP for a one watt lower power rating includes both the WTP stemming from the incremental Energy Star effect, \$0.13, and the baseline WTP of \$0.71. The “brand” effect of the Energy Star implies a median WTP that is roughly 18 times greater than the incremental WTP motivated by the Energy Star label for a one watt lower power rating.

²⁰We make this qualification because it is possible that a consumer’s choice set could include bulbs of equal efficiency with some bearing the label and others not. In this scenario, the energy-independent Energy Star effect does not influence reductions in external costs.

heterogeneity in the coefficients we allow to vary.

4 Heterogeneous Impacts of the Energy Star

The significant variance of the Energy Star and Consumption interaction term in Table 1 implies that the Energy Star could either increase or decrease a participant’s marginal utility coefficient on savings from energy consumption.²¹ Intuitively, increases would be valuable to those who undervalue these savings but harmful to those who do not or who overvalue them, relative to a rational benchmark. Section 4.1 develops this intuition by introducing the concept of internalities and demonstrating the implied economic costs to consumers.

4.1 The Origin and Cost of Internalities

In this subsection, we consider the consumer at different time points as different people to illustrate the origin and cost of internalities. We show how a consumer behaving as if she is applying time inconsistent preferences to a choice situation bears an economic cost, relative to the scenario in which she acts as if she applies the time consistent preferences that are central to standard economic theory.

In our discussion, we rely on two notions of utility: experienced utility (U_e) and decision utility (U_d). The experienced utility quantifies the benefits the consumer truly experiences from consumption. In contrast, the decision utility quantifies the benefits she believes she will experience from consumption at the time of decision-making. We illustrate these concepts with both an abstract choice setting and a numerical example. Both our abstract and numerical choice settings entail three time periods, with time progressing from $t = 0$ to $t = 2$. At time zero, the consumer becomes aware of the choice task, which she must complete at time one. Without loss of generality, the alternatives are such that they yield an immediate benefit enjoyed at time one and a delayed cost experienced at time two.

At time zero, the consumer studies the choices as a fully rational decision maker. She applies her preferences to evaluate the future costs and benefits of the alternatives, x , available in the choice set, S . She identifies the alternative that yields the largest stream of utility. Formally, she identifies, but does not yet choose, the alternative x^* that maximizes her experienced utility function: $x^* = \operatorname{argmax}(x \in S)U_e(x)$. The experienced utility function, U_e ,

²¹Figure 9 in the appendix presents the distribution of the coefficient estimates.

can be represented as in Equation 9:²²

$$U_e(u_t, u_{t+1}, \dots, u_T) = \sum_{\tau=t}^T \delta^\tau u_\tau \quad (9)$$

Above, U_e reflects an inter-temporal preference that considers both the current instantaneous utility, u_t , and future instantaneous utility levels. δ is a parameter that measures the degree to which utility experienced in a future period is equivalent to utility experienced in the current period; δ assumes values between 0 and 1, and the closer it is to zero, the lower the degree to which future utility is equivalent to current utility. We use x^* and the experienced utility it entails as a normative benchmark.

The consumer actually makes her consumption decision in time one, and internalities emerge from this decision-making context if her choice at time one is not x^* . To illustrate this, we follow a modeling strategy from behavioral economics (O'Donoghue and Rabin, 2000; Laibson, 1997). This introduces a parameter β_u to the utility representation in Equation 9 that applies to the utility streams in all but the current time period.²³ The utility function with this new parameter is the decision utility function, U_d :

$$U_d(u_t, u_{t+1}, \dots, u_T) = \delta^t u_t + \beta_u \sum_{\tau=t+1}^T \delta^\tau u_\tau \quad (10)$$

If $\beta_u \neq 1$, the consumer acts as if she is no longer applying time-consistent preferences at the time of decision-making. The alternative chosen at time one is x^d , where $x^d = \operatorname{argmax}(x \in S)U_d(x)$.

Traditionally, β_u has been restricted to the range $(0, 1]$, since Equation 10 has usually been applied to scenarios in which consumers are characterized by preferences for immediate gratification. In this range, when the consumer arrives at time one, she under-weights the utility streams in future periods. Such preferences are generally known as time-inconsistent preferences and the modeling strategy, as beta-delta discounting. Our treatment differs from the traditional application in requiring only that $\beta_u > 0$. When $\beta_u > 1$, the consumer acts as though she has a preference for delayed gratification.

Unless $x^d = x^*$, the optimization of the decision utility function instead of the experienced utility function implies a loss to the decision-maker. The experienced utility implied by x^d

²²The notation and discussion here are inspired by those in O'Donoghue and Rabin (2000).

²³We use the subscript 'u' to distinguish this parameter from the abstract coefficients of the ML model introduced in Section 2.2.

is given by $U_e(x^d)$. We refer to the gap between $U_e(x^d)$ and $U_e(x^*)$ as the internality: $Internality(x^d) = U_e(x^*) - U_e(x^d)$. The internality characterizes the loss in experienced utility due to the use of time-inconsistent preferences. If $\beta_u < 1$, the internality stems from an under-weighting of future utility streams. If we consider the individual at times one and two as different people, we observe that the loss in “social” value when $\beta_u < 1$ stems from the time-one consumer “overindulging” in the benefits of the consumption good such that it imposes extra costs on the time-two individual, relative to the rational benchmark.²⁴ Similarly, if $\beta_u > 1$, the internality stems from the time-two individual limiting too greatly the ability of the time-one individual to indulge in the benefits of the consumption good, relative to the rational benchmark.

The Appendix includes a numerical example that illustrates these utility concepts. The example demonstrates several points that help one understand the implications of our Energy Star experiment. Most importantly, the example highlights that a gap between the decision and experienced utility functions does not necessarily imply that the consumer’s choice will change. Consequently, such a gap does not immediately signal that the consumer bears a loss from time-inconsistency or other behavioral phenomena generating the wedge between the two utility functions.

4.2 Measuring the Utility Implications of the Energy Star

In this subsection, we develop an intuition about the effect of the Energy Star on losses in experienced utility and, equivalently, on internalities. We proceed here with three assumptions that we relax in our actual evaluation in the next section. First, we assume that the light bulb choice set is continuous, or that between the minimum and maximum price and power ratings, goods with all intermediary non-dominated combinations of prices and power ratings are available.²⁵ Second, we assume that consumers’ decision utility functions do not reflect an Energy Star “brand” effect; that is, we do not discuss here the role of the energy-independent Energy Star effect in directly influencing choice. Third, choices are known with certainty. Though we subsequently drop these assumptions in Section 5.2, they allow us to build a general perspective about the impact of a behavioral instrument on the

²⁴Our “society” here is the consumer comprised of both the time-one and time-two individuals.

²⁵A price and power rating combination would be dominated if, for a given power rating, a lower price option exists. Thus, for any given power rating, the infinite choice set would include only one option at a given price.

choices individuals make and on the implications for experienced utility.²⁶

Equation 8 implies an experienced utility function for person i who has selected good j :

$$U_{e_{ij}} = \eta_i \cdot p_j + \theta_i \cdot C_j + \Upsilon^T X_j \quad (11)$$

We note that the experienced utility formulation does not include the direct “brand” effect of the Energy Star label. This reflects a belief that it does not actually contribute to differences in experienced utility.²⁷ The time-consistent consumer who acts as if she were maximizing Equation 11 selects a bulb x^* and enjoys a utility $U_e(x^*)$. Time-consistency in the light bulb context relates to a full treatment of the future implications of energy consumption implied by the power rating of alternate bulbs. The utility coefficient on C_j thus determines the degree to which the individual is time-consistent. More specifically, if the utility coefficient on C_j implies that the consumer is willing-to-pay an amount equal to the present value of savings associated with a 1W lower power rating, we term her time-consistent. We derive this benchmark willingness-to-pay, WTP_i^R , for each consumer and a benchmark utility coefficient, θ_i^R , that would imply the individual would be willing to pay WTP_i^R . Note that when $\theta_i = \theta_i^R$, the consumer is acting in a manner analogous to that when $\beta_u = 1$ in the previous subsection.

Our calculation of WTP_i^R assumes an electricity price of \$0.1147/kWh, as shared with participants, and that a bulb is used 3 hours a day every day for eight years.²⁸ Formally,

$$WTP_i^R = \sum_{t=1}^{T=8} \frac{Savings_t}{(1+r_i)^t} \quad (12)$$

In Equation 12, r_i is the consumer’s temporal discount rate, and $Savings_t$ equals the annual savings that accrue to the consumer upon choosing a bulb with a 1W lower power rating. In general, we use an individual-level discount rate derived from an elicitation task embedded in the experiment and designed by Kirby and Marakovic (1996). To test the sensitivity of our results to the discount rate assumptions, we also use a common 6% discount rate.

²⁶It is especially helpful to temporarily assume away the direct effect of the Energy Star, since any given behavioral instrument may not necessarily generate an analogous effect.

²⁷This is consistent with the treatment by Houde (2013), and we echo his assertion that this is in reality open to debate. While we subsequently introduce it into our assumed decision utility function in Section 5.2, we do not discuss it here to focus instead on the impact of the Energy Star label on the coefficient on energy consumption.

²⁸We implicitly assume that the rational consumer believes in expectation that electricity rates over these eight years will remain constant.

WTP_i^R excludes any disutility from the environmental costs of consumption and should be interpreted as a “selfishly” rational measure. We exclude these costs because the costs incident on any one person should not make a first-order difference in WTP_i^R . We illustrate this with an extreme example in which all electricity consumed is generated from coal-fired facilities, which generally entail the largest external costs of electricity generation. In particular, Sundqvist (2004) suggests that the median external cost estimate in the literature is \$0.11/kWh, upon adjusting from 1998 to 2013 dollars. This reflects both localized costs, such as those from mercury and particulate matter emissions, and distributed costs, such as those associated with the emission of greenhouse gases. If we next assume that the \$0.11/kWh cost is exerted on only 100,000 people, any given person would include \$0.000001/kWh into his or her calculation of the value of energy consumption savings. While the total external cost of \$0.11/kWh should be considered in the evaluation of social welfare, the derivation of a given individual’s “selfishly” rational assessment will be essentially unchanged upon including his or her share of the environmental costs.²⁹ We incorporate the full external cost in Section 5 where we illustrate a comprehensive valuation of the Energy Star program.

The estimated marginal utility of wealth, η_i , points to a benchmark utility coefficient, θ_i^R , that would imply the individual would be willing to pay WTP_i^R :

$$\theta_i^R = WTP_i^R \cdot \eta_i \quad (13)$$

We define the wedge, $wedge_i$, as the difference between θ_i^R and θ_i , the estimated coefficient on energy consumption in the absence of the Energy Star:

$$wedge_i = \theta_i^R - \theta_i \quad (14)$$

With θ_i^R , we can update our expression for experienced utility to:

$$U_{e_{ij}} = \eta_i \cdot p_j + \theta_i^R \cdot C_j + \Upsilon^T X_j \quad (15)$$

Equation 15 allows us to determine the utility experienced by the time-consistent consumer with her utility-maximizing choice, x^* . This level of utility serves as the benchmark against which we compare the utility experienced by time-inconsistent decision-makers. When Equ-

²⁹Of course, the distribution of cost incidence may not be uniform, as we have assumed, but a very drastically different distribution would be required for the implication here to change.

tion 11 has $\theta_i \neq \theta_i^R$, it should be interpreted as the decision utility function. In the presence of the Energy Star, the decision utility function includes an additional term, $\gamma_i \cdot C_j$.

The consumer who acts as if she were applying time-inconsistent preferences would display choice behavior consistent with the maximization of a utility function with a coefficient $\theta_i \neq \theta_i^R$. The *wedge_i* term allows us to quantify the degree to which the consumer is time-inconsistent and identify the nature of inconsistency; in particular, the consumer can be either present-biased ($\theta_i < \theta_i^R$) and under-value savings on energy consumption or future-biased ($\theta_i > \theta_i^R$) and over-value them. Continuing our assumption that the consumer faces a continuous choice set and thus changes her choices with any change in wedge value, as $|wedge_i|$ approaches zero, the choice of the time-inconsistent individual, x^d , becomes more similar to x^* . Consequently, the internality or loss in experienced utility, $U_e(x^*) - U_e(x^d)$, will become smaller. Thus, to evaluate whether the Energy Star provides value or harms a particular individual, we determine whether or not its effect is to shift her marginal utility coefficient closer to the rational benchmark.³⁰

Though we constrain the utility coefficients on energy consumption to be negative, for ease of interpretation we discuss the negative of the coefficients. Thus, individuals with positive wedges undervalue savings on energy consumption, and those with negative wedges overvalue savings on energy consumption.³¹ We discuss γ_i similarly; the higher it is, the greater is the impact of the Energy Star in increasing an individual's valuation of savings on energy consumption.

To gain an intuition for the effect of the Energy Star, we define the *mitigation*, mit_i , performed by the Energy Star logo as a link between the effect of the Energy Star on *wedge_i* and implications for experienced utility. We measure mitigation as the ratio of the coefficient

³⁰Note that our assumption of a continuous choice set in which choices change as wedges change is critical in our exposition here. To evaluate the Energy Star program more generally, we need to know if its impact on decision utility coefficients actually affects choices made and thus the experienced utility obtained by the consumer. We drop the simplifying assumption and present a more nuanced discussion of the effects of the Energy Star on experienced utility in Section 5.2.

³¹In the context of appliance choice, a natural retort to the concept of overvaluation of energy savings is that a large coefficient on energy consumption is also consistent with strong preferences for environmental protection. While that is true, even those individuals who have strong environmental preferences may overvalue savings on energy consumption. Since environmental quality can be affected by many actions, the individual can still pursue too much environmental conservation in appliance choice relative to other opportunities. Stated otherwise, the individual may have lower cost ways by which to achieve the same level of environmental conservation.

on the Energy Star and Consumption interaction term, γ_i , to $wedge_i$:

$$mit_i = \frac{\gamma_i}{wedge_i} = \frac{\gamma_i}{(\theta_i^R - \theta_i)} \quad (16)$$

Note that the mitigation measure is only defined when a wedge exists; it is undefined when $\theta_i^R = \theta_i$, since by definition no wedge exists to be mitigated. When an individual has $wedge_i = 0$ in the absence of a behavioral instrument like the Energy Star, any impact of the instrument on the apparent utility coefficients characterizing his choices implies a loss in experienced utility away from the normative utility level.³²




Figure 3: The mitigation scale is split into three areas: (1) wedge exacerbation, (2) wedge mitigation, and (3) an overshoot in mitigation. The first and second areas imply only losses and gains in experienced utility, respectively. In the third area, mitigation increases experienced utility in the range 1 to 2 but decreases it otherwise. The threshold value of 2 reflects a symmetric choice set.

Four features illustrate the connection between the impact of the Energy Star on wedges and experienced utility. First, when $mit_i < 0$, the Energy Star changes the utility coefficient on energy consumption in the opposite direction of that needed to set $wedge_i = 0$. This implies that the Energy Star yields a loss in experienced utility. The Energy Star delivers the largest gain in consumer utility when it exactly closes the wedge, or when $mit_i = 1$. When $0 < mit_i < 1$, the Energy Star partially closes the wedge and increases consumer utility. When $mit_i > 1$, the Energy Star overcompensates for the wedge; the sign of the wedge flips, with a wedge between θ_i^R and $\theta_i + \gamma_i$ remaining. As long as the new wedge does not imply a larger loss in experienced utility than implied by the old wedge, the Energy Star can still improve the consumer's utility. The threshold mitigation value will depend on the nature of

³²This follows from our assumption of a continuous choice set. Changes in the utility coefficients that do not imply changes in the alternative selected will not generate these losses.

the choice set, since that will drive the implications for experienced utility. Assuming that the choice set is symmetric, in the sense that deviations of equal magnitude from θ_i^R in either direction imply equal losses in experienced utility, the Energy Star will improve experienced utility as long as $1 < mit_i < 2$. If $mit_i > 2$, the overcompensation is sufficiently severe that the resultant wedge in the opposite direction implies a loss in experienced utility.³³

Figures 4 and 5 plot the wedges and mitigation levels observed in our sample. Analyzed in isolation, the means in both plots corroborate the motivation for the Energy Star program and imply that it accomplishes its goal of shifting individuals’ utility coefficients on energy consumption toward the “rational” level.³⁴ On average, participants’ utility coefficients on energy consumption are smaller than their rational benchmarks, and mitigation by the Energy Star is consistent with improvements in experienced utility. These average trends are in line with the estimates shared in Section 3.2.

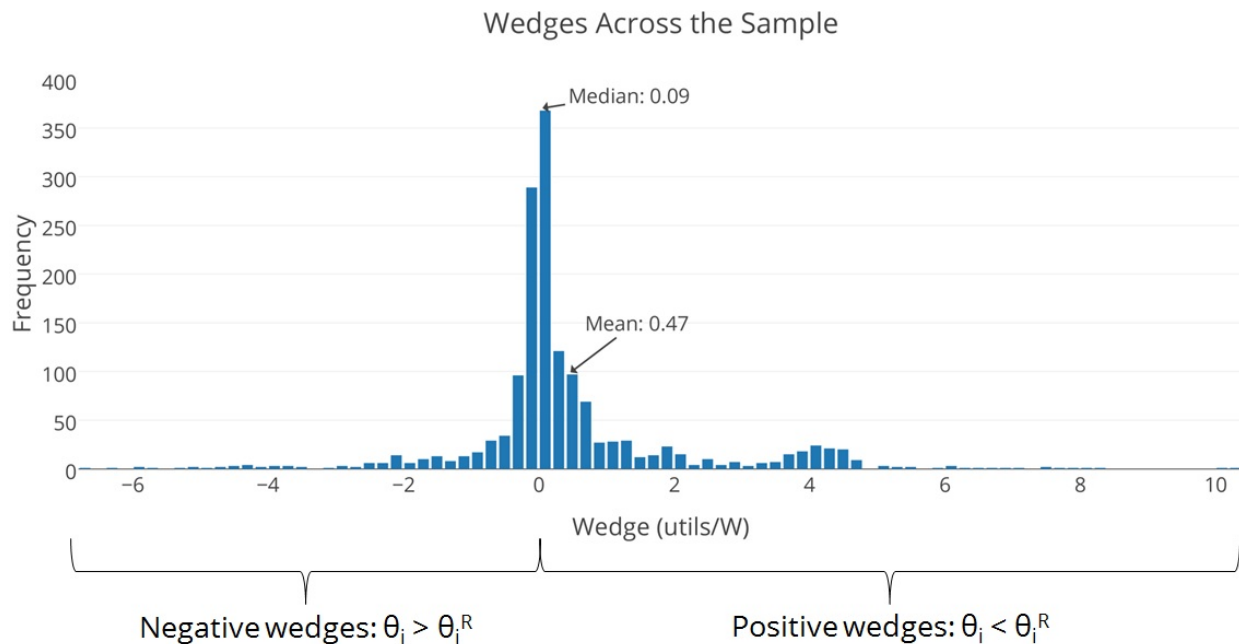


Figure 4: The observed wedges imply substantial heterogeneity in departures from the rational benchmark. Though the mean and median wedges equal 0.47 and 0.09, respectively, a substantial portion (37%) of the wedges are negative.

The distributions of wedge and mitigation values around the average levels indicate that

³³If the choice set available to the consumer implies an asymmetry around θ_i^R , the analyst should recalibrate this cutoff value.

³⁴Figure 5 is truncated at ± 5 . The mitigation levels range from -66.4 to 86.9.

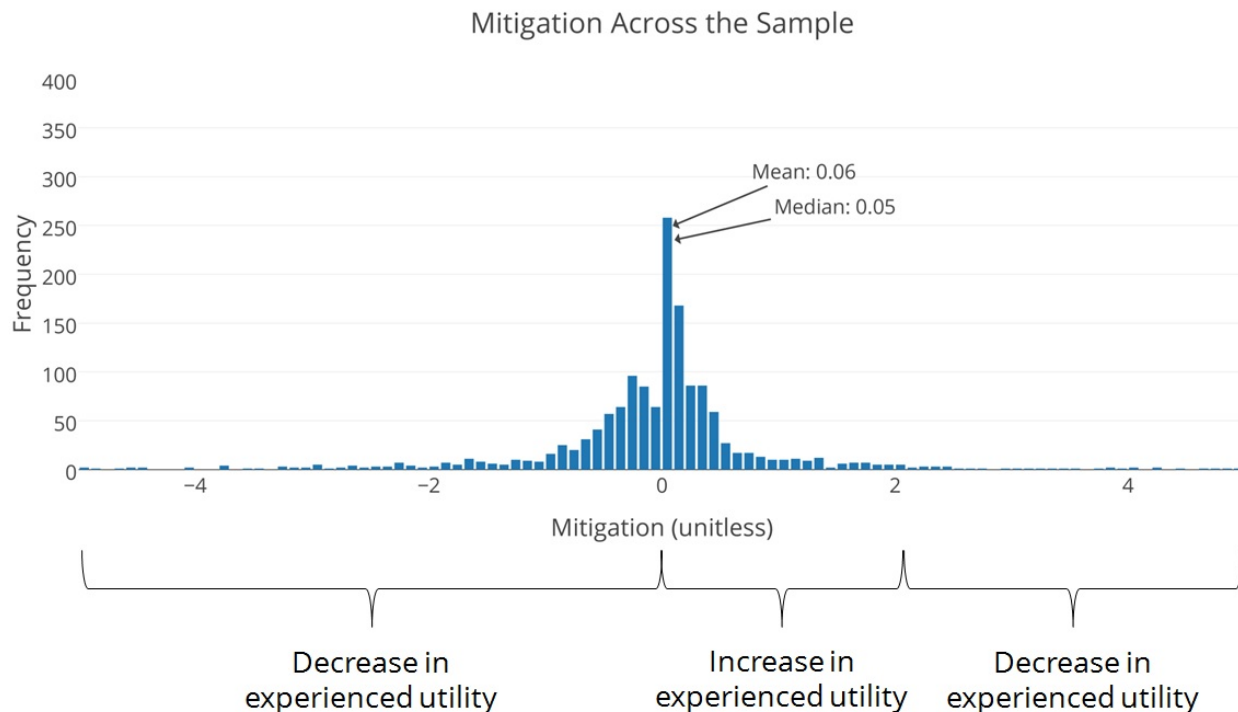


Figure 5: The observed mitigation levels imply substantial heterogeneity in consumers’ response to the Energy Star label. The plot is truncated at ± 5 , though the observed mitigation ranges from -66.4 to 86.9.

there are important differences in individual-level responses to the Energy Star. Our framework in Figure 6 categorizes the utility ramifications of the heterogeneous impacts of the Energy Star. The framework groups participants into one of four groups defined by the signs of $wedge_i$ and γ_i . These groups reflect the four possible effects of the Energy Star: people either under- or over-value savings on energy consumption, and the Energy Star either increases or decreases this tendency. Group membership reflects the experienced utility ramifications of these effects. Individuals in Groups A and D are those who, in the absence of the Energy Star, over-value savings on energy consumption; formally, $\theta_i > \theta_i^R$. From the perspective of wedge mitigation, the desired impact of the Energy Star would be to decrease θ_i closer to θ_i^R . Consequently, those for whom $\gamma_i < 0$ could realize a gain in experienced utility gain, while those for whom $\gamma_i > 0$ are subject to a loss in experienced utility. The tentative language for the former group reflects the possibility that γ_i could be sufficiently negative that it implies $mit_i > 2$ for the individual. As Figure 3 shows, this would lead to a loss in experienced utility. In contrast, participants in Groups B and C are those who, in the absence of the Energy Star, under-value savings on energy consumption, such that

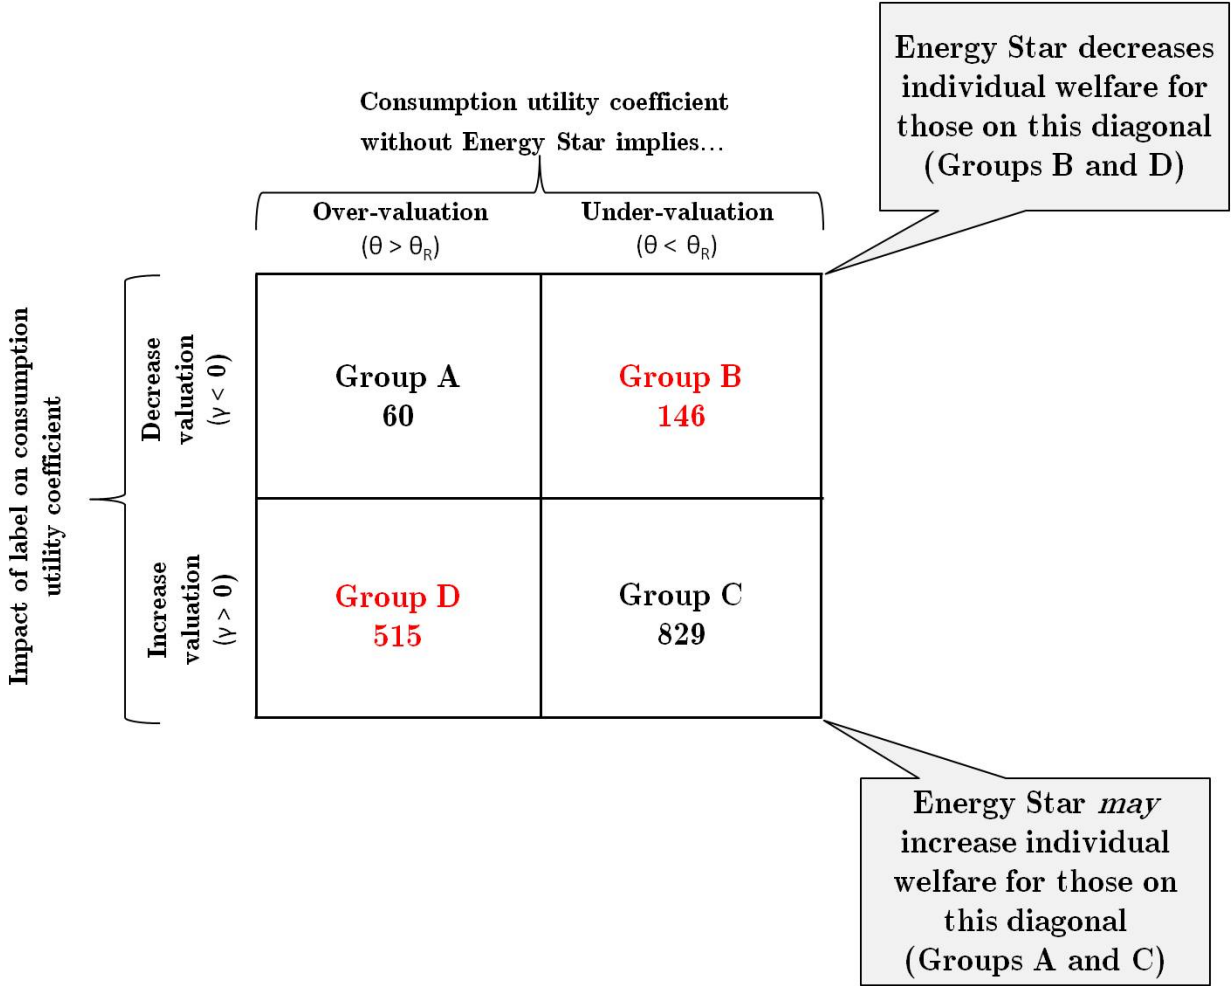


Figure 6: Participants are binned into four groups defined by $wedge_i$ and γ_i . The groups in red font are those for which the Energy Star implies a decrease in individual utility, since $mit_i < 0$. The numbers below the group labels indicate the number of participants in our sample of 1,550 binned into the group.

$\theta_i > \theta_i^R$. For these individuals, the “desired” impact of the Energy Star is that $\gamma_i > 0$. Again, if γ_i is too much greater than 0, individuals in Group C could bear a loss in experienced utility. However, those in Group B unambiguously face a decrease in experienced utility, as the impact of the Energy Star is in the direction opposite of that needed to close the wedge. Note that the diagonal including Groups A and C includes those for whom mit_i is greater than 0; for these individuals, the Energy Star could increase utility. The other diagonal with Groups B and D includes individuals for whom mit_i is smaller than 0 and who experience a loss in utility.

In our sample, 60, 146, 829, and 515 participants are in Groups A, B, C, and D, respec-

tively.³⁵ We highlight two points. First, we observe that more individuals have $mit_i > 0$ than the opposite sign. In particular, the Energy Star affects the utility coefficient on consumption in the desired direction among 57% of the sample, or 889 individuals. These counts would suggest an overall improvement in consumer welfare, but the actual implications will depend on the marginal utilities of wealth. Second, the Energy Star acts to increase the valuation of energy consumption savings for 87% of individuals. This is consistent with the mean coefficient on the Energy Star and Consumption interaction term in Table 1.

We summarize in Finding 2 the implications of the Energy Star for experienced utility.

Finding 2: *Assuming a continuous light bulb choice set, that choices are known with certainty, and that consumers’ decision utility functions do not include the “brand” effect of the Energy Star, the impact of the Energy Star is to exacerbate wedges among and decrease experienced utility for 661 (43%) of the 1550 individuals in our sample. It increases experienced utility by affecting wedges in the desired direction for at most 889 (57%) of the 1550 individuals in our sample.*

The overall value of the program will depend on the extent of mitigation among individuals, their marginal utilities of wealth, and the impacts the Energy Star has on experienced utility and environmental externality mitigation. We next turn to those aspects.

5 The Value of the Energy Star to Consumers

5.1 Evaluation Methodology

Our discussion in the previous section highlighted that the Energy Star label can influence a change in the good that the individual purchases. Since the alternative selected in the presence of the Energy Star program may imply experienced utility and external costs of electricity generation that are higher or lower than those obtained in its absence, we must value both changes. Equation 17 captures the two effects:³⁶

³⁵These counts are based on rational benchmarks defined by individual-level discount rates. If we use a common 6% discount rate, the group counts are 41, 165, 879, and 475, respectively.

³⁶While the equation will appear exactly the same as one that traditional public economics would prescribe, the consumer surplus terms reflect experienced utility levels stemming from *different* decision utility functions in the absence and presence of the instrument. In contrast, standard public economics would assume that surplus in both settings would reflect a stable experienced utility function.

$$Value_{ES} = (CS_{ES}^{exp} - CS_{NES}^{exp}) - (ECE_{ES} - ECE_{NES}) \quad (17)$$

In the expression above, CS^{exp} represents the consumer surplus based on the experienced utility from bulb purchase, ES the setting in which the Energy Star program exists, NES that in which it does not, and ECE the external costs of electricity generation and consumption. Changes in ECE enter negatively because consumers benefit from smaller external costs. Equation 17 reflects several assumptions. Our use of a traditional public economics valuation approach assumes that consumers incorrectly consider only the energy consumption attribute, such that the utility coefficients on the other attributes can be used to approximate the consumer's experienced utility. Moreover, notwithstanding the findings by Houde (2013) that firms respond strategically to the Energy Star certification, we assume a competitive supply side that sells the same products with and without the Energy Star. We further assume that firms do not modify the prices of their products in response to consumers' valuation of the Energy Star label itself.

To determine the change in the consumer surplus, we evaluate the expression below:

$$\Delta CS^{exp} = \sum_i -\Delta \frac{Internality_i}{\eta_i} = \sum_i \frac{1}{\eta_i} (U_{e_{ij^*}}^{ES} - U_{e_{ij^*}}^{NES}) \quad (18)$$

The division by η_i allows us to derive a dollar-denominated metric from the util-denominated measure of experienced utility, U_e . The j^* index denotes the good selected by the consumer.³⁷

To calculate the change in the external cost of electricity consumption, we must determine the external cost of each kilowatt hour of electricity consumed. We assign each product an external environmental cost, ECE_j . We derive this by first calculating the electricity consumed by a 1W bulb operating for 3 hours a day for 8 years and subsequently scale this figure by both the actual power rating of bulb j and ED , the environmental damage per kWh of electricity consumption. We discount future external costs using a 6% social discount rate, r_s .³⁸ Formally:

³⁷The negative sign preceding the Δ Internality term follows from the definition of the internality as the loss in experienced utility.

³⁸Since the marginal costs of greenhouse gas emissions reflect the aggregate stock of emissions, they increase at the rate of inflation. Strictly speaking, we should not discount these costs by the social discount rate. However, we use an aggregated measure of external costs; though a disaggregation would refine our valuation estimates, it does not add to the central point of the paper.

$$ECE_j = \sum_{t=1}^{T=8} \frac{C_j(3*365)}{1000} \frac{ED}{(1+r_s)^t} \quad (19)$$

We assume that ED equals \$0.068/kWh. This reflects a weighted average of damage estimates for electricity derived from coal, oil, natural gas, nuclear, hydro, wind, and solar sources, with weights determined by the share of U.S. electricity generation contributed by each source.³⁹

The evaluation methodology we have thus far outlined is based on the intuition developed in Section 4.2 for how the Energy Star would work under assumptions that we know the consumers’ choices with certainty, that the consumer accesses a continuous light bulb choice set, and that consumers’ decision utility functions do not include the direct “brand” effect of the Energy Star label. In the following subsection, we drop these assumptions and discuss the necessary modifications to our evaluation of the Energy Star program. The final subsection offers an evaluation for a particular choice set.

5.2 Three Complexities

Our evaluation methodology requires us to simulate consumer choice in both the absence and presence of the Energy Star program. As Figures 7 and 8 illustrate, the extent of internalities depends on the choice set available to the consumer and on the curvature of the experienced utility curve when plotted against the value of the utility coefficient, θ_i .⁴⁰ Figure 7 depicts a situation in which there are large potential losses from internalities. Assume that the consumer acts as if she applies a decision utility coefficient of θ_{i0} in the absence of the Energy Star. ΔU_{e1} describes the increase in experienced utility associated with an increase in the coefficient by the quantity γ_{i1} and towards θ_i^R . As the figure shows, the Energy Star

³⁹We scale the estimates from 1998 to 2013 dollars using the U.S. BLS Inflation Calculator. For each source, we use the median level of damages reviewed by Sundqvist (2004) and share of electricity generation in the U.S. from the U.S. Energy Information Agency (2013). In particular, we use data from a table of Net Generation by Energy Source. We allot generation from “petroleum liquids” to oil and remove the small amount of generation from petcoke, “other gas” and biomass. The median figures mask many underlying differences across studies, such as the inclusion or exclusion of CO_2 prices, the consideration of full fuel cycle effects, estimations by either abatement cost or damage cost approaches, and top-down versus bottom-up estimates (Sundqvist, 2004). However, the cost is in the range of others used; for example, Houde (2013) uses \$0.079/kWh, which reflects a \$67.1/t damage from CO_2 emissions.

⁴⁰The appendix includes a discussion that links the utility coefficient to the choice of products and ultimately to the experienced utility curve.

label can also imply losses in experienced utility. ΔU_{e2} denotes the lost experienced utility associated with an increase in the coefficient by the quantity γ_{i2} and away from θ_i^R . As the consumer acts as if she applies a θ_i that is farther away from her θ_i^R in her decision making, she chooses goods that imply a larger loss in experienced utility. This situation corresponds to one in which one light bulb does not dominate the others on the basis of both price and consumption.

Figure 8 depicts a scenario in which one alternative in the choice set dominates the others; namely, one bulb dominates the others on both price and consumption. Regardless of the θ_i that the consumer applies, her experienced utility remains the same, and the internality cost, zero. Thus, $\Delta U_{e1} = \Delta U_{e2} = 0$.




Figure 7: Implications for experienced utility when the choice set available to the consumer implies a highly non-linear experienced utility versus θ_i curve. This scenario arises when there is a rich trade-off between up-front prices and operating costs.

The difference between Figures 7 and 8 underscores a fundamental point: the greater the curvature in the experienced utility versus θ_i curve, the greater the loss from mis-optimization by consumers and the greater the potential gain from a ‘nudge’ by the Energy Star program

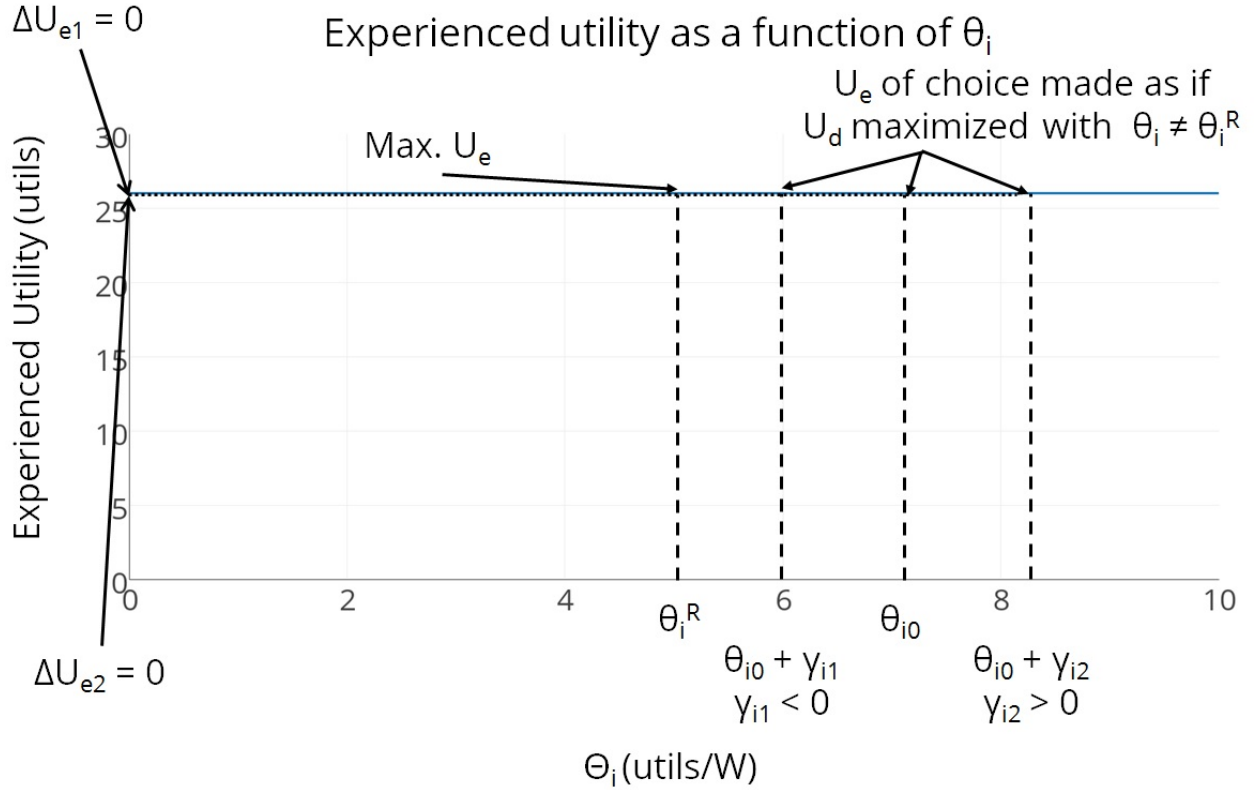


Figure 8: Implications for experienced utility when the choice set available to the consumer implies a horizontal experienced utility versus θ_i curve. This scenario arises when the choice set includes one alternative that dominates the others.

of θ_i in the right direction. The sensitivity of the implications of a wedge between θ_i and θ_i^R implies that the analyst can provide only an evaluation of the Energy Star program that is conditional on the assumed choice set. Since the consumer may evaluate alternatives unseen and unconsidered by the analyst, the latter is unable to observe the former's "true" $U_{e_{ij}}(\theta_i)$ curve. Of course, this possibility exists even when an analyst uses data on alternatives available in-market. In this case, the weak concavity of the experienced utility curve implies that estimates of the loss attributable to the Energy Star program can be interpreted as lower bounds.⁴¹

The logic connecting the nature of the choice set to implications for the change in experienced utility extends to implications for externalities as well. If the choice set includes one dominant bulb, the impact of the Energy Star program on θ_i will not yield any changes in the

⁴¹We cannot interpret estimates of the gain as upper bounds because it is possible that alternatives yielding a higher experienced utility exist within the consumer's choice set but were unobserved by the analyst.

alternative actually selected and, consequently, on the extent of externalities. In the setting in which the bulb selected is sensitive to the utility coefficient, however, we can identify a useful inequality to approximate whether the effect of the Energy Star on a given consumer implies a social gain or a social loss. When the effect of the Energy Star program is to increase θ_i closer to θ_i^R such that the consumer is more likely to purchase a bulb of a lower power rating, there is an unambiguous weak gain in social welfare, since the consumer's experienced utility weakly increases and the external costs of electricity consumption weakly decrease. Similarly, when the program decreases θ_i away from θ_i^R , there is an unambiguous weak loss in social welfare, as the consumer's experienced utility weakly decreases and the external costs of electricity consumption weakly increase. There are two situations in which there is some ambiguity about the social implications of the Energy Star program. These arise when the program either decreases θ_i toward the rational benchmark and thus encourages more consumption or increases θ_i away from the rational benchmark. In this case, we must compare the gain or loss in experienced utility with the gain or loss in environmental externalities. With "small" changes in the utility coefficient, the difference in experienced utility can be approximated by $\frac{(\theta_i + \gamma_i)}{\eta_i} \cdot C_{j^*}^{ES}(\theta_i + \gamma_i) - \frac{\theta_i}{\eta_i} \cdot C_{j^*}^{NES}(\theta_i)$ while the change in externality is given by $ECE_{j^*}^{ES}(\theta_i + \gamma_i) - ECE_{j^*}^{NES}(\theta_i)$. For small changes in the utility coefficient, the analyst can thus compare the approximated difference in experienced utility, which we have expressed in surplus terms to keep units consistent, to the difference in the external costs of electricity consumption.

A second complexity arises because we do not know with certainty the choices that the consumer will make. We must assume a deterministic setting for Equation 17 to hold. However, our utility parameters were derived only with an assumption about the nature of unobserved utility. Consequently, and as Small and Rosen (1981) and Leard (2013) highlight, Equation 17 must be evaluated in expectation. We thus re-express that equation as follows:

$$\mathbb{E}[Value_{ES}] = \mathbb{E}[(CS_{ES}^{exp} - CS_{NES}^{exp})] - \mathbb{E}[(ECE_{ES} - ECE_{NES})]$$

To evaluate the above expression, we evaluate both Equation 20, which computes the expected change in the external costs of electricity generation, and Equation 21:

$$\mathbb{E}[\Delta ECE] = \sum_i \sum_j (P_{ij}^{ES} - P_{ij}^{NES}) \cdot ECE_j \quad (20)$$

$$\mathbb{E}[\Delta CS^{exp}] = \sum_i \sum_j (P_{ij}^{ES} - P_{ij}^{NES}) \cdot \frac{U_{e_{ij}}}{\eta_i} \quad (21)$$

In the equations above, P_{ij}^{ES} and P_{ij}^{NES} are the simulated probabilities that participant i selects alternative j when the Energy Star program exists and when it does not, respectively. This addition of the probability terms implies that the Energy Star imposes costs and benefits in expectation. Said otherwise, it can change the probability with which the consumer chooses a given alternative in her choice set. For example, if a consumer's optimal choice entails the lowest consumption but highest priced good available in a choice set and the Energy Star acts to decrease the probability with which she selects this good by reducing the magnitude of the coefficient on energy consumption in her decision utility function, it exerts a cost in expectation.

The third addition to our evaluation analysis is the effect of the Energy Star on decision utility, independent of its effects on θ_i . This implies an update to our decision utility function in the presence of the Energy Star:

$$U_{d_{ij}} = \eta_i \cdot p_j + (\theta_i + \gamma_i) \cdot C_j + \Upsilon^T X_j + \lambda_i \cdot ES_j \quad (22)$$

The impact of this consumption-independent effect of the Energy Star is ambiguous and must be empirically evaluated for a specific choice set. Since only 60 of the 1550 individuals in our sample had a value of $\lambda_i < 0$, we can characterize the modal impact of this Energy Star effect. If the utility-maximizing alternative is Energy Star certified, this Energy Star effect will simply act to increase the probability with which the optimal good is selected. If, however, the utility-maximizing alternative is not certified, this effect will have the opposite implications and imply a lower expected experienced utility. To the extent that this Energy Star effect compels individuals to select goods at least as energy efficient as they would have in its absence, it contributes to a weak reduction in environmental externalities. In the subsequent subsection, we quantify the incremental impact of the consumption-independent effect of the Energy Star on consumer welfare.

5.3 The Value of the Energy Star

To evaluate the expected change in value, we apply the individual-level utility coefficients to simulate consumer choice in the absence and presence of the Energy Star. The simulation

yields choice probabilities for both contexts. We use a hypothetical choice set containing CFLs available for purchase from a large national retailer in early 2014.⁴² The bulbs have a narrow band of power ratings (13 – 16W), since we restricted the lumen range of the choice set to remain between 800 and 900 lumens.

We perform three versions of the evaluation procedure. In the first version, we isolate the impact of the Energy Star on consumers’ utility coefficients on energy consumption. In the second, we include the energy-independent effect of the Energy Star in consumers’ decision utility. The final version restricts the effect of the Energy Star on the utility coefficient on energy consumption to bulbs that are Energy Star certified. This final version is most consistent with our econometric estimation. The other versions require an assumption that in an actual market, the presence of the Energy Star label prompts changes in consumers’ apparent decision utility functions across both labeled and unlabeled goods.

Table 2 summarizes the value of the Energy Star for the three versions described above. The top block of rows details the expected gain in experienced utility, and the middle block, the expected reduction in external cost. The “Total Increase in Value” is the sum of the two “All groups” totals listed below in each block of rows. At the bottom of the table, we scale the value of the Energy Star to our participants to the entire population of U.S. households. This row estimates the value per year across owner-occupied households in the U.S., assuming they would respond as the participants of our experiment did and purchase four bulbs each in a given year.⁴³

We make three comments based on Table 2. First, the increases in experienced utility do not match the theoretical guidance provided by Figure 6. In particular, we would have expected utility losses among those in Groups B and D and gains only among those in Groups A and C. The discrepancy between expected and actual results stems from the fact that we are evaluating changes *in expectation*. In the choice set we use for the simulation, one alternative dominates the others by offering the lowest consumption and lowest price option. This good is the experienced utility maximizing alternative for all individuals in

⁴²In particular, we derived this choice set from the larger choice set available to consumers at www.lowes.com, the online storefront for the retailer. We exclude certain bulbs with warmth levels beyond those considered in the choice experiment, since we are unable to estimate the marginal utility of this level. We also exclude decorative bulbs, or those whose light is a color other than white.

⁴³We scale the value of the instrument by the product of (1) the ratio of the total number of U.S households (approximately 115 million) to the 1550 households in our sample and (2) the share of households that are owner-occupied. We use 65.2% for the latter, per data from the U.S. Census Bureau (2014).

	Version 1: Indirect ES effect only	Version 2: Both ES effects	Version 3: Both ES effects; Indirect effect on labeled bulbs only
Exp. Gain in Experienced Utility			
Group A: ES reduces positive wedges	-\$2	\$16	\$25
Group B: ES increases negative wedges	-\$8	\$23	\$54
Group C: ES reduces negative wedges	\$11	\$96	\$42
Group D: ES increases positive wedges	\$13	\$178	\$69
All groups	\$13	\$313	\$189
	\$0.01	\$0.20	\$0.12
Exp. Reduction in External Cost			
Group A	-\$1	\$1	\$3
Group B	-\$3	\$1	\$6
Group C	\$3	\$12	\$6
Group D	\$4	\$16	\$11
All groups	\$4	\$30	\$25
	\$0.00	\$0.02	\$0.02
Total Increase in Value	\$17	\$343	\$215
	\$0.01	\$0.22	\$0.14
Scaled to all U.S. Households	\$4M	\$60M	\$36M
Assumes four bulbs purchased per year			

Table 2: In the “Total” rows, the top number quantifies the impact across the entire population and the bottom number, the per-capita impact. The final row scales the impact across the 1,550 households in the study to the population of U.S. owner-occupied households. Differences between totals and the sum of group-specific amounts reflect rounding errors. Note: “Exp.” denotes “Expected,” “ES,” “Energy Star,” and “M,” million. Recall that the indirect effect of the Energy Star is that on the utility coefficient on energy consumption and that the direct effect is the “brand” effect that is independent of energy implications.

our sample.⁴⁴ The increase in value observed in all three versions reflects the fact that the

⁴⁴While it may be surprising that a commercially available choice set would include both a dominant bulb and others, we note that the dominance follows upon considering only price, consumption, and bulb warmth. Our specifications of the utility functions abstract away from brand considerations that may be relevant to the consumer. Moreover, our example evaluation retains the assumption of our stated choice experiment that the lifetime of all bulbs is equal. In the choice set used for simulation, the bulbs’ lifetime ranged from 7 to 11 years. Thus, the bulb we consider dominant may not be, given a more comprehensive description of

Energy Star program more often increases the probability with which consumers select this experienced-utility maximizing good than it decreases it. Consumers in Groups A and B, who act as if their decision utility coefficient on energy consumption decreases in the presence of the Energy Star, experience a loss in expected experienced utility in Version 1 because the probability with which they select this alternative decreases. In contrast, those in Groups C and D act as if their decision utility coefficient on the attribute has increased, and the relevant probability increases.

A comparison of Version 2 with Version 1 reveals the powerful impact of the energy-independent effect of the Energy Star label. Since the dominant good in the choice set is labeled with the Energy Star logo, we observe that individuals in all four groups experience a larger gain in experienced utility in Version 2 than in Version 1. The size of the energy-independent Energy Star effect is sufficiently large that its impact on the probabilities of selecting the alternatives more than compensates for the decrease in probability of selecting the experienced utility maximizing good among Groups A and B in Version 1. Version 3 retains the assumption that the Energy Star exerts this effect but limits changes in the decision utility coefficient on energy consumption to those bulbs that are Energy Star labeled. Though the overall effect of this restriction is a reduction in the total increase in value in Version 3 relative to Version 2, those in Groups A and B obtain an increase in their expected gain in experienced utility and those in Groups C and D, a decrease. In Version 3 and for those in Groups A and B, the impact of the Energy Star on the energy consumption utility coefficient is equivalent to imparting a decrease in the perceived consumption level of labeled goods, relative to those without the label. The Energy Star exerts the opposite effect on individuals in Groups C and D.

We note finally that the expected impact on experienced utility is larger than that on external costs. Though instruments such as the Energy Star may be motivated by their potential to reduce external costs, their impact on private costs may be bigger. On the one hand, this reflects our decision to restrict the simulation bulb choice set to bulbs of 800 – 1000 lumens. On the other hand, this restriction may reflect how consumers make their bulb purchase decisions: bulbs are installed in a variety of locations throughout the house, and the search for bulbs may be limited to those within a particular lumen range and therefore a specific range of brightness. Since the variation in power requirements for a given range is

individuals' apparent utility functions.

relatively small, the reductions in external cost are unlikely to be large. Finally, though we observe an expected increase in value from the Energy Star, conditional on the simulation choice set we used, our numerical example does not imply that the Energy Star will perform this way with any choice set. An example choice set in which we would calculate a negative value is that with all bulbs of consumption between 5W and 16W in 0.25W increments associated with all prices from \$11.50 to \$0.50, respectively, in -\$0.25 increments.

The Energy Star appears to be a viable instrument if the goal of policymakers is solely to reduce externalities. Nonetheless, the instrument has the potential to exert large costs on consumers that are attributable to changes in their experienced utility. More broadly, in contexts where societal goals of mitigating market failures could be in conflict with individuals' experienced utility, the analyst should consider both effects in determining the attractiveness of a behavioral instrument. Finding 3 summarizes our findings:

Finding 3: *Using a choice set in which one alternative dominates the others, we find that the Energy Star improves consumer welfare in expectation. This expected improvement stems more from the impacts of the instrument on experienced utility than those on external costs. We emphasize that the unconditional value of the Energy Star is ambiguous.*

6 Conclusion

Given patterns of under-investment in energy efficient goods, the analyst with a goal of reducing environmental externalities could conclude that the more often the Energy Star implies an increase in the valuation of savings on energy consumption, the more valuable the program. However, this would ignore the effect of the Energy Star on individual welfare. For any given individual, the effect of the Energy Star is ambiguous. At one level, it is unclear if the Energy Star will actually increase the attention paid to energy consumption information, since the label could instead provide a license to ignore the attribute. This implies a second level of indeterminacy: the impact of the Energy Star may be to prompt the individual to make choices that reflect a lower or higher valuation of savings on energy consumption, relative to that expressed in the absence of the label.

Ideally, the Energy Star program would not only address environmental externalities but also encourage consumers to make choices that strike a balance between the upfront price and long-run costs of alternate goods that is more aligned with utility maximization.

To do so, it would address tendencies among consumers to under- or over-value savings on energy consumption. Such over- or under-valuation could exert externalities on consumers if it compels them to select goods that imply changes in their experienced utility. The reduction of these externalities would increase individual welfare. Since an under- or over-valuation of energy consumption savings can exert externalities on the consumer, any impact of the Energy Star on the valuation of energy consumption savings and expected changes in experienced utility must be quantified in policy analysis.

This realization implies the question that motivated this paper: given its effects on both externalities and internalities, what is the value of the Energy Star to consumers? This paper presented and analyzed data from a stated choice experiment intended to answer this question. Our chief contribution was an analytic framework that uses these data to quantify the impact of the Energy Star on externalities and to value the program. The fundamental strategy in our framework is to establish a benchmark for experienced utility coefficients that reflect standard time-consistent preferences. Specifically, we used our data to estimate utility coefficients in the absence and presence of the Energy Star. Assuming a particular price of electricity and lifetime of the bulb products considered by our participants, we derived a rational benchmark for the utility coefficient on energy consumption. The benchmark coefficient is analogous to $\beta_u = 1$ in the beta-delta conceptualization presented in the paper, since it is the parameter consistent with the maximization of an experienced utility function by the consumer. We quantified the impact of the Energy Star on externality losses by comparing the difference in expected experienced utility in the absence and presence of the Energy Star, relative to that predicted with the rational benchmark. Our discussion emphasized that wedges between the rational benchmark coefficient on energy consumption and those observed in the presence and absence of the Energy Star are not themselves sufficient to imply a loss to the consumer. Such losses occur only if the wedge compels the consumer to select a good other than the one that maximizes experienced utility. The same commentary applies to changes in externalities. Given this, our methodology can value externality and externality effects, *conditional* on a given choice set.

This framework can be used in other contexts upon making the two assumptions we have made. Namely, the analyst must first assume that the utility coefficients on the attributes other than those affected by the behavioral instrument reflect standard preferences. The second assumption establishes a rational benchmark experienced utility against which

internality losses in the absence and presence of the Energy Star can be measured. The framework will be more difficult to apply in contexts in which it is not as straightforward to measure the experienced utility implications of choosing a given alternative.

Our results suggest that the impact of the Energy Star on internalities is larger than that on environmental externalities. We applied our evaluation methodology to a choice set of light bulbs offered by a national retailer. When evaluating expected changes in internalities and externalities, we find that the Energy Star delivers modest reductions in the expected cost of both. This result does not generalize to all choice sets subject to the influence of the Energy Star. If the choice set entails a richer trade-off between up-front prices and costs of operation, a poor performance of the Energy Star would stem from the situations in which it increases and decreases the magnitude of consumers' utility coefficients on energy consumption among those who, in its absence, already overvalue or undervalue consumption savings, respectively. For these individuals, the Energy Star would significantly increase internality losses when selecting from such a choice set.

The potential for non-pecuniary measures such as the Energy Star to impose internality costs implies that such instruments can yield negative dividends. That is, the instruments may deliver greater internality losses than externality gains. This adds a nuance to the assertion by (Allcott, Mullainathan, and Taubinsky, 2014) that taxes aimed at the reduction of externalities could yield a "double dividend" by also addressing internalities. Non-price instruments are receiving increasing attention in light of popular discussions of strategies to "nudge" consumers. Our finding here should encourage a fuller evaluation of such proposals with a careful accounting of their impacts on internalities and consumer welfare.

7 Appendix

Distribution of individual-level coefficients on the ML interaction term




Figure 9: The mean individual-level utility coefficient on the Energy Star and energy consumption interaction term, -0.126 , implies that the Energy Star increases the weight consumers place on energy consumption. However, the distribution reveals that the label both increases and decreases this weight. The label decreases the weight for the 13% of individuals with an estimated coefficient on the interaction term greater than zero.

Interpretive information provided in the stated choice experiment

Description of the task

We are interested in understanding how you purchase light bulbs, and we will ask you a series of questions in which you will pick your preferred bulb among three alternate light bulbs.

As you answer these questions, please assume that all light bulbs:

1. Are compatible with standard fixtures (for example, ceiling or lamp fixtures)
2. Are dimmable; that is, the intensity of their output can be adjusted by using, for example, a light switch that allows one to choose a level of light intensity
3. Have an equal delay between when you turn them on and when they reach full brightness
4. Have an equal brightness
5. Have the same expected lifetime. All bulbs in this study are compact fluorescent light bulbs (CFLs). CFLs last approximately 8,000 hours, while standard incandescent bulbs typically last 750 hours

The light bulbs we will ask you to compare will differ along some or all of the following dimensions:

Price The price of the light bulb is the price you would pay in a retail store.

Light color The light color is a measure of the color temperature of the light bulb. In residential settings, most light bulbs produce warm white, soft white, neutral or cool light. The color temperature profiles of the light bulbs we will describe shortly are described below.

Energy Star Certification Energy Star certified light bulbs use about 75% less energy than traditional incandescent bulbs and last at least 6 times longer. The Energy Star program is a voluntary certification program, and manufacturers may elect to apply labels such as the one below to products that meet energy consumption criteria set by the U.S. Environmental Protection Agency (EPA). Since the program is voluntary, some light bulbs that meet the criteria do not carry the label. An example label appears below.

We will use the following label for light bulbs that are not Energy Star labeled; remember that a bulb may be unlabeled either because (1) it does not meet the federal standard for

Energy Star certification or (2) although it meets the federal standard for Energy Star certification, the manufacturer has opted to not certify it.

[The instructions displayed the Energy Star label and a blank blue label used in the experiment to denote a non-certified bulb. Below, we include the three versions of details provided to participants in experimental conditions 1, 2, and 3, respectively.]

Consumption: Energy Demand (W) The energy demand refers to the ‘watt’ rating of the bulb. Watts are a standard unit of power and measure the rate at which a light bulb consumes energy. For a given amount of light outputted, different light bulb technologies are characterized by different watt ratings.

Example: a low watt CFL light bulb can produce an equivalent light intensity as a high watt traditional incandescent light bulb.

Consumption: Energy Demand (Estimated Annual Cost) The estimated annual energy cost refers to the monetary cost of energy consumed by a light bulb in one year. The actual energy cost for a bulb will vary depending on how often it is used and the price of electricity.

Information is provided about the annual cost of electricity associated with a given bulb. To derive this annual cost, we assume that the light bulb is used 3 hours a day and use an average U.S. residential electricity rate of \$0.1147/kWh.

Consumption: Energy Demand (Watts and Estimated Annual Cost)

The energy demand refers to the ‘watt’ rating of the bulb. Watts are a standard unit of power and measure the rate at which a light bulb consumes energy. For a given amount of light outputted, different light bulb technologies are characterized by different watt ratings.

Example: a low watt CFL light bulb can produce an equivalent light intensity as a high watt traditional incandescent light bulb.

Information is also provided about the estimated annual cost of electricity associated with a given bulb. To derive this annual cost, we assume that the light bulb is used 3 hours a day and use an average U.S. residential electricity rate of \$0.1147/kWh.

Details of the scale analysis

Our application of the Swait and Louviere analysis begins by testing whether one can reject the hypothesis of equal scale factors across the three experimental conditions. Though the scale factor cannot be identified, the ratio of the scale factor within the data from one experimental condition to that within another can be identified. We extend the Swait

and Louviere hypothesis testing procedure to three experimental conditions, wherein the overall hypothesis is one of coefficient and scale parameter equality: $H1 : \beta_1 = \beta_2 = \beta_3$ and $\sigma_1 = \sigma_2 = \sigma_3$. To test the hypothesis, we first estimate separate parameter vectors β_3 , $\sigma_2 \cdot \beta_2$, $\sigma_1 \cdot \beta_1$, which provide log likelihoods of L_1 , L_2 , and L_3 , respectively.⁴⁵ In a second step, we impose a constraint that $\beta_1 = \beta_2 = \beta_3$ but allow σ_1 and σ_2 to vary.⁴⁶ This second estimation generates a single log likelihood estimate, L_{step2} , and we form a test statistic $\lambda_1 = -2 [L_{step2} - (L_1 + L_2 + L_3)]$, which is asymptotically distributed with $K + 2$ degrees of freedom, where K is equal to 7. Table 3 summarizes our log likelihood estimates and informs our rejection of $H1$.

L_3	L_1	L_2	L_{Step2}	λ_1	p-value	σ_1	σ_2
-8996.128	-8909.386	-8850.386	-26850.037	188.270	<0.001	0.960	0.935

Table 3: Summary output from the test of coefficient and scale parameter equality

Given the rejection of $H1$, we follow the suggestion by Swait and Louviere to scale our data by the σ_1 and σ_2 values that yielded the maximum log likelihood in the step 2 estimate. Table 3 details these scale factors in the last two columns.

An example to illustrate the origin of internalities

Consider an infrequent caffeine consumer who benefits from an immediate large gain in mental acuity when she consumes caffeine but suffers, with delay, from considerable caffeine-induced anxiety. Facing a work deadline that has limited her sleep recently, she is deciding whether to consume an espresso drink with 1 or 2 shots of espresso. The outside option, which is associated with a baseline utility of zero utils, is to skip the drink altogether; we refer to this as the zero shot option. Table 4 presents the net benefit, in units of utility experienced, relevant to her decision. The consumer instantaneously experiences the benefits of increased alertness in time one and the disutility of the caffeine-induced anxiety at time two. The implication is that the latter costs are delayed while the benefit of consumption is immediately realized. This allows the costs to be only partially salient to the consumer making a choice at time one. Recall that the time-zero decision-maker, in contrast, is assumed

⁴⁵Condition 3 serves as our reference condition because it includes both descriptions of consumption information.

⁴⁶We search in a grid from 0.01 to 2.25 with a step size of 0.025.

to fully consider both the future benefits and costs of the alternatives available to her.

Shots of espresso	Time = 1	Time = 2
Zero	0	0
One	5	-3
Two	7	-6

Table 4: Utility schedule, u_t , for the infrequent caffeine consumer

Holding without loss of generality an assumption that $\delta = 1$, we can determine what the caffeine drinker will do. At time zero, the consumer begins to consider the choice but is not yet at the point of purchase. She evaluates the alternatives as if she were maximizing U_e . Since U_e equals 0, 2, and 1 upon consuming 0, 1, and 2 shots of espresso, respectively, the time-zero individual would want the time-one consumer to opt for one shot of espresso. If the time-one consumer applies the same time-consistent preferences (i.e., acts as if her $\beta_u = 1$), her choice will match her desired choice at time zero.

Let us now allow the time-one consumer to have $\beta_u \neq 1$. Consider what happens when $\beta_u = \frac{1}{2}$ and the consumer pays too little attention to the negative utility in time two. The decision utility function implies that 0, 1, and 2 shots of espresso will deliver 0, 3.5, and 4 utils of utility, respectively. Since $x^d = 2$, the consumer experiences a utility equal to 1 (i.e., $U_e(x^d = 2) = 1$) and an internality equal to 1, since $U_e(x^* = 1) - U_e(x^d = 2) = 1$. A similar economic loss can occur if the time-one decision maker acts as if $\beta_u = 2$ and the consumer pays too much attention to the negative utility in time two. The decision utility function equals 0, -1, and -5 for 0, 1, and 2 shots, respectively. She now opts for the zero-shot option. The consumer experiences a utility equal to 0 (i.e., $U_e(x^d = 0) = 0$) and an internality equal to 2, since $U_e(x^* = 1) - U_e(x^d = 0) = 2$. These examples illustrate that a gap between experienced and decision utility can impose economic costs on the consumer. These economic costs stem from the wedge between a rational parameter value (i.e., $\beta_u = 1$) and that actually used by the consumer in her decision utility function. As demonstrated in these examples, the internality equals the difference between the maximum level of experienced utility attainable by a time-consistent consumer and the experienced utility stemming from the choice made by the time-inconsistent consumer.

Before concluding our example, we highlight an implication of the choice set available to the consumer. The choice set we considered is discrete, since the consumer cannot choose a countably infinite number of shots between 0 and 2. This feature implies that for some values

β_u	x^d	$U_e(x^d)$	$U_e(x^*)$	<i>Internality</i>
1	One shot	2	2	0
$\frac{3}{4}$	One shot	2	2	0
$\frac{1}{2}$	Two shots	1	2	1
$\frac{3}{2}$	One shot	2	2	0
2	Zero shots	0	2	2

Table 5: Utility measures associated with the choice of zero, one, or two shots of espresso for a time-one consumer.

of β_u sufficiently close to 1, $x^d = x^*$ and the consumer does not experience the economic costs of internalities. Consider the utility outcomes when $\beta_u = \frac{3}{4}$ or $\beta_u = \frac{3}{2}$. Table 5 summarizes the experienced and decision utility outcomes and internalities associated with these and the above β_u values. We observe that when $\beta_u = \frac{3}{4}$ or $\beta_u = \frac{3}{2}$, the consumer opts for 1 shot of espresso. Though wedges between the decision and experienced utility functions exist in each case, they do not imply a loss in consumer welfare.

This example of espresso decision-making can be immediately extended to the context of the Energy Star program and appliance choice. When the consumer acts as if she is optimizing a decision utility function with a β_u equal to something other than one, the consumer makes decisions that she would not make if she were to “rationally” consider the entire stream of utilities stemming from her decision. In the context of the Energy Star program, the purchase of lower-priced but less efficient goods delivers immediate benefits, while the stream of implied energy consumption implies future costs. Consumers under-valuing the costs of energy consumption can be described as acting with a $\beta_u < 1$ and those over-valuing, with a $\beta_u > 1$. Though β_u has usually reflected impatience among consumers, $\beta_u < 1$ could equally well follow from inattentiveness and $\beta_u > 1$ from a tendency to be overattentive.

Moreover, the caffeine example illustrates several concepts that help understand our subsequent analysis of the Energy Star label and its effects. First, in our assessment of the Energy Star program, we calculate a normative benchmark for experienced utility that is analogous to the setting in which $\beta_u = 1$. Second, larger gaps between the time-consistent parameter value and that applied by the time-one decision maker imply a greater potential for the individual to make choices that imply a loss in welfare. Third, because the choice set is discrete, there are instances in which the use of a time-inconsistent parameter does

not imply a loss in welfare to the consumer. Fourth, when such losses occur, they imply real economic costs to the consumer.

As evident in Section 4.2, our analysis of the Energy Star experiment makes two strong assumptions. First, we assume that aside from the utility coefficient on the energy consumption attribute, all estimated coefficients reflect standard and thus time consistent preferences. Since we model utility as additively separable, this is consistent with beta-delta discounting. In our context, the β_u parameter applies only to that component of utility stemming from energy consumption, and any gaps between decision and experienced utility are attributable to either an over- or under-weighting of the implications of the future costs of energy consumption.

The second assumption is that of the magnitude and duration of energy savings that obtain for each individual. We were able to quantify internalities in the simple example only because we could apply a known experienced utility schedule, but such valuation is frequently quite difficult. This difficulty may at times preclude an economic evaluation with all relevant factors considered. Nonetheless, for a policy analysis that focuses only on externalities to be correct, it must make either of two strong assumptions itself. One assumption is that the Energy Star has no impact on the weights consumers assign to the energy consumption attribute of alternatives and thus on internalities. The alternative states that though the Energy Star exerts some influence on internalities, changes in internalities would be the same no matter which instrument is used. Assumption 2 seems intuitively unreasonable. The significant Energy Star and Consumption interaction term in Table 1 suggests Assumption 1 is also unjustified.

We note that our discussion of beta-delta preferences provides just one way of conceptualizing mistakes in optimization by consumers. Our simple model here should not be interpreted as a structural representation, as many phenomena could explain an under-investment in energy efficient goods. For example, Tsvetanov and Segerson (2013) model decisions over energy-consuming goods as an internal conflict between a fully rational planner and a less-than-rational doer, in the tradition of Gul and Pesendorfer (2001). The essential element of any modeling representation is that it admits the possibility for non-optimal decision-making by consumers. It is this non-optimality that leads to the experienced utility gap we term the internality.

Additional intuition for why policy analysis must condition on choice sets

Figure 10 provides an example of a choice set that may be available to the consumer. The top curve depicts the total cost of ownership of a bulb, and this is equal to the vertical sum of the other two curves, which plot the purchase price and lifetime costs of electricity consumption. A consumer who is over- or under-attentive to costs of electricity consumption would act as if the lifetime cost curve were shifted upward or downward. Consumer behavior consistent with the use of a $\theta_i \neq \theta_i^R$ in the decision utility function could lead the consumer to select a bulb with a power rating other than the one that is her experienced utility-maximizing alternative.

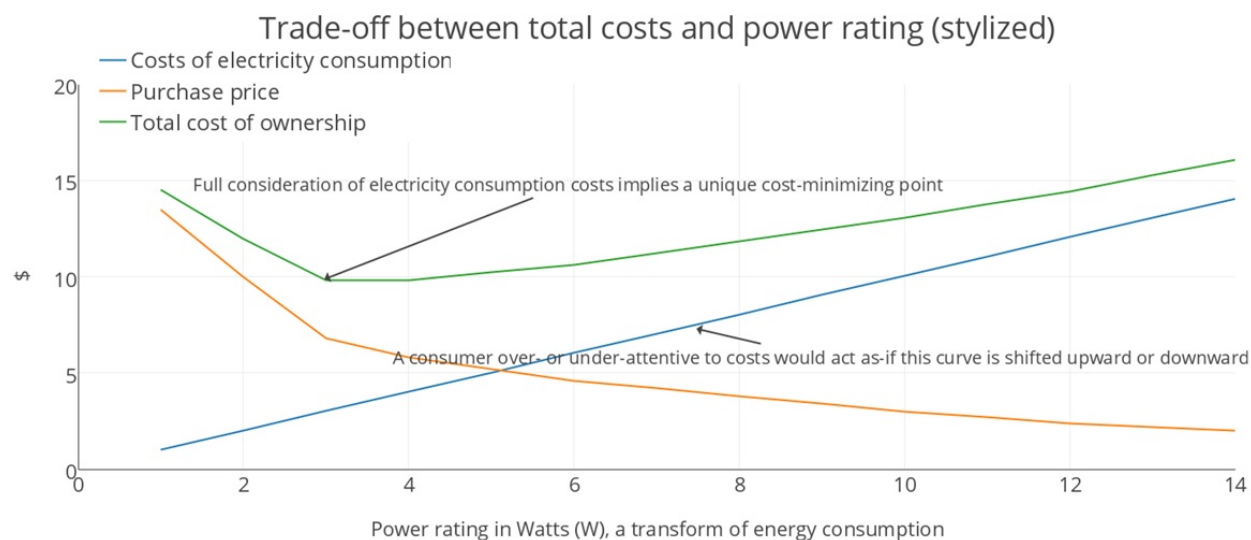


Figure 10: An over- or under-valuation of the costs of electricity consumption is equivalent to shifting upward or downward the blue curve (the bottom-most curve on the left side of the graph). This shift would alter the “perceived” total cost of ownership curve such that the consumer identifies a different power rating as minimizing costs.

By definition, the selection of a such a bulb would imply a loss in experienced utility. Though we depict in the main text curves relating θ_i and U_e directly, the relationship is indirect in the sense that when $\theta_i \neq \theta_i^R$ the bulb selected may not be the one that maximizes experienced utility. Any difference in the experienced utility function stems from subsequent changes in the bulb selected.

This last point is critically important. If Figure 10 instead included cost and price curves that implied a total cost of ownership that sharply increased from the true cost-minimizing

option, the use of a θ_i value different from θ_i^R would be unlikely to imply a difference in the utility experienced by the consumer. Even if the consumer were to apply $\theta_i \neq \theta_i^R$, the selected bulb is quite likely to be the same as the one that maximizes experienced utility. Thus, the degree of curvature in the choice set as embedded in the total cost curve is important in determining the likelihood with which reductions in experienced utility will obtain.

Details on measures collected and sample size

Measures We collected measures that were relevant to either appliance or financial decision-making, given the similarities in the theorized decision-making processes involved in both contexts. To ensure that the questions relevant to these measures did not influence participants' choice behavior, our queries followed the stated choice exercise. Our measures can be grouped into four sets: (1) environmental concern and financial status, (2) cognitive ability, (3) demographics, and (4) psychometrics. We describe the measures included in each of these buckets below.

Environmental concern and financial status: Consumers may differ in the degree to which environmental and financial aspects of energy consumption are salient in their decision-making processes and thus in the extent to which these concerns are reflected in their utility coefficient on energy consumption. We measure environmental concern with scores on the New Ecological Paradigm (NEP) scale (Dunlap et al., 2000). We complemented this measure with questions about environmental self-efficacy, or the degree to which the participant believed his decisions could impact environmental outcomes. We do not use the latter here because the measure uses questions primarily relevant to refrigerator choice.

To characterize financial status, we elicited temporal discount rates, household assets and liabilities, and liquidity constraints. Our estimates of discount rates use the Kirby and Marakovic (1996) methodology.⁴⁷ We used these discount rates to establish the rational benchmark coefficient on energy consumption. Asset and liability measures allow us to gauge

⁴⁷Each participant answered 21 questions in which they were asked to choose either a smaller immediate payoff or a larger delayed payoff. When ordered, the questions imply 22 exponential and hyperbolic discount rates that predict a switch from the immediate payoff to the delayed payoff. The original paper includes 21 values, but we add two more bounds, 0 and 1, that allow participants to always select the immediate or delayed payoff. For each participant, we calculate the predicted choices based on each of the 22 time preferences and assign to the subject the discount rate that maximizes the match between actual and predicted choices. We treated accuracy ties between rates as in Kirby and Marakovic (1996).

the wealth of each household, and liquidity constraints, to understand the degree to which this one factor contributes to price sensitivity in consumers.

We include two other measures that have provided a richer description of financial decision making. A metric of self-continuity describes the degree to which one identifies with one's future selves. Our measure is based on that used by Ersner-Hershfield et al. (2009), who show that differences in self-continuity can explain differences in saving rates. This corresponds to the β_u parameter from Section 4.1 and, in the context of the utility coefficient on energy consumption, to over- or under-valuation of savings on energy consumption. The metric could thus explain differences in the valuation of savings in energy consumption which accrue to future selves. The other metric measures the emotional states of our participants. Emotional affect has implications for belief formation and financial decision-making: positive emotions such as excitement induce risk taking and confidence, while negative emotions trigger opposite reactions (Kuhnen and Knutson, 2011). Decision-making involving energy consumption entails risks about future electricity prices and environmental damages associated with energy consumption. We use responses from the Positive and Negative Affect Schedule (PANAS) (Watson, Clark, and Tellegan, 1988) to measure the degree to which individuals are predisposed toward positive or negative affect.

Eight questions were intended to examine risk preferences among participants. One set asked how individuals would allocate a sudden windfall of \$10,000 among stocks, bonds, and cash. Another asked how their current assets were allocated. A final measure sought to measure a gap between consumers' willingness to pay to take and to avoid a gamble. We do not use the first two measures, which are part of a survey in development but not definitively linked to risk preferences. We do not use the questions about willingness to pay because the answers we received indicated a misunderstanding of the questions.

Cognitive ability: Behavior consistent with the maximization of a time-consistent experienced utility function requires the ability to focus on all relevant information and translate this into utility implications. We include two measures to characterize participants' abilities to perform these steps. One is the cognitive reflection test (CRT) that determines the degree to which the individual performs slow and deliberate calculations instead of quick and reflexive ones, and we use a three-item measure developed by Frederick (2005). A separate numeracy measure is based on questions developed by Peters et al. (2007). We supplement the measure with an additional question that tests' individuals understanding of conditional

probabilities.

Demographics: Our data include place of birth and residence, age, household size, household income, energy bill burden, home ownership status, subjective socioeconomic status, education, intra-household income distribution, and expected length of stay in current residence. We also asked whether the respondent was responsible for financial and appliance decisions. Our education metric includes four buckets. The first includes those with at most a 2 year college degree, the second, those with a 4 year college degree, the third, those with a professional degree (e.g., an MBA or JD), and the fourth, those with a non-professional masters, doctoral degree, or medical degree.

Psychometrics: We measure the degree to which each participant is characterized by each of the Big-Five personality domains, namely, extraversion, agreeableness, conscientiousness, emotional stability, and openness to new experiences. Our inclusion of the measures is driven largely by an expectation that differences in conscientiousness could help explain attentiveness to energy consumption information and the utilization of the Energy Star logo. We use a ten item inventory developed by Gosling, Rentfrow, and Swann (2003).

Sample size We based our target sample size of 1,500 on the results of a pilot phase of data collection. For our pilot, we recruited participants via the Amazon Mechanical Turk platform and gathered data using the Qualtrics survey tool. We randomized participants into one of 8 blocks of 16 choice sets. Each participant selected one of three alternate light bulbs in 16 light bulb choice sets. We used Ngene (ChoiceMetrics, 2012) to generate the 128 choice sets included in the experiment. Of the 110 submitted responses, we deemed 103 to be complete and attentive by checking whether participants responded correctly to questions about their online identity in the middle and at the end of the task. We used a subset of 92 individuals in the upper 90 percentiles of task completion time to estimate coefficients. All participants were aged 18 or over, U.S. residents, and registered users of the Amazon Mechanical Turk system. We provided each participant \$7. We derived priors for our main choice experiment by estimating coefficients of a multinomial logit model of the pilot choice data. Among the coefficients we estimated was one on a binary variable that coded whether an alternative light bulb was labeled with the Energy Star.

Our initial objective for our main experiment was to understand whether the presence of the Energy Star was a significant predictor of choice and whether the magnitude of the Energy Star’s effect differed by the three experimental conditions. We used the estimated

coefficient and standard error on the binary variable that coded for the presence of the Energy Star label to calculate the sample size requirement.

Our main experiment reflected several changes to the pilot data set. First, our pilot data set included two types of light bulb technologies (i.e., CFL and light emitting diodes (LED)), while our main experiment used just CFL bulbs. Second, our utility specification in the main experiment was richer than that in the pilot setting. Third, and finally, we planned to use a mixed logit formulation to model the choice data from the main experiment instead of the fixed coefficient multinomial logit approach we used with the pilot data. Given all of these changes, we scaled up our sample size as permitted by the study's budgetary constraints.

Attribute	Energy Units		Dollar + Energy Dollar Units			Total	Diff		Diff (2/3)
	Cond. 1	Cond. 2	Cond. 2	Cond. 3	(1/2)		(1/3)		
Participants	514	507	507	529	1550	7	15	22	
Age	50.6	49.9	49.9	50.3	50.3	0.7	0.3	0.4	
% Women	49.4%	48.3%	48.3%	52.2%	50.0%	1.1%	2.8%	3.9%	
Household Size	2.84	2.77	2.77	2.79	2.76	0.07	0.05	0.02	
Education (of 4)	2.30	2.36	2.36	2.30	2.32	0.06	0.00	0.06	
Primary Earner, Wage Share (%)	78.4	79.8	79.8	80.2	79.4	1.4	1.8	0.4	
Hyperbolic Discount Rate	0.027	0.031	0.031	0.022	0.027	0.004	0.005	0.009	
Exponential Discount Rate	0.021	0.023	0.023	0.017	0.020	0.002	0.004	0.006	
CRT Score (Number correct of 3)	0.72	0.79	0.79	0.74	0.75	0.07	0.02	0.05	
Numeracy (Number correct of 15)	10.34	10.26	10.26	10.21	10.27	0.08	0.13	0.05	
Household Income (Median)	\$60 - 80K	\$60 - 80K	\$60 - 80K	\$60 - 80K	\$60 - 80K	-	-	-	
Socio-economic Status (Median)	5	5	5	5	5	-	-	-	
NEP (min: 15; max: 75)	50.7	51.3	51.3	50.9	51.0	0.6	0.2	0.4	
PANAS-Negative (min: 10; max: 50)	18.7	17.9	17.9	17.8	18.2	0.8	0.9	0.1	
PANAS-Positive (min: 10; max: 50)	33.4	32.6	32.6	32.5	32.8	0.8	0.9	0.1	
Big Five-Extraversion (min: 2; max: 14)	8.46	8.08	8.08	8.24	8.26	0.22*	0.38	0.16	
Big Five-Agreeable (min: 2; max: 14)	10.49	10.47	10.47	10.51	10.49	0.02	0.02	0.04	
Big Five-Conscientious (min: 2; max: 14)	11.14	11.42	11.42	11.31	11.29	0.28	0.17	0.11	
Big Five-Emotion Stability (min: 2; max: 14)	10.16	10.14	10.14	10.12	10.14	0.02	0.04	0.02	
Big Five-Openness to Exp. (min: 2; max: 14)	9.82	9.91	9.91	9.80	9.84	0.09	0.02	0.11	
Future Self-Similarity (min: 1; max: 7)	4.88	4.94	4.94	4.94	4.92	0.06	0.06	0.00	

Table 6: Overview of demographic, psychological and financial metrics across the sample. Significant differences at the 0.05 significance level are marked by an asterisk. Some variables are not summarized above, but these are factors that are distributed similarly across experimental conditions; e.g., the distributions of financial and appliance decision-making are similar across the three conditions.

References

- Allcott, H. (2011), “Consumers’ perceptions and misperceptions of energy costs,” *The American Economic Review*, 101(3), 98–104.
- Allcott, H., S. Mullainathan, and D. Taubinsky (2014), “Energy policy with externalities and internalities,” *Journal of Public Economics*, 112, 72–88.
- Allenby, G. M., and P. E. Rossi (1998), “Marketing models of consumer heterogeneity,” *Journal of Econometrics*, 89(1), 57–78.
- Chetty, R., A. Looney, and K. Kroft (2009), “Salience and Taxation: Theory and Evidence,” *The American Economic Review*, 99(4), 1145 – 1177.
- ChoiceMetrics (2012), “Ngene,” Version 1.1.1.
- Daly, A., S. Hess, and K. Train (2012), “Assuring finite moments for willingness to pay in random coefficient models,” *Transportation*, 39(1), 19–31.
- Dunlap, R., K. V. Liere, A. Mertig, and R. Jones (2000), “Measuring Endorsement of the New Ecological Paradigm: A Revised NEP Scale,” *Journal of Social Issues*, 56(3), 425 – 442.
- Ersner-Hershfield, H., M. Garton, K. Ballard, G. Samanez-Larkin, and B. Knutson (2009), “Don’t stop thinking about tomorrow: Individual differences in future self-continuity account for saving,” *Judgment and Decision Making*, 4(4), 280 – 286.
- Frederick, S. (2005), “Cognitive Reflection and Decision Making,” *Journal of Economic Perspectives*, 19(4), 25 – 42.
- Gillingham, K., and K. Palmer (2014), “Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence,” *Review of Environmental Economics and Policy*, 8(1), 18–38.
- Gosling, S., P. Rentfrow, and W. Swann (2003), “A very brief measure of the Big-Five personality domains,” *Journal of Research in Personality*, 37(6), 504 – 528.
- Gul, F., and W. Pesendorfer (2001), “Temptation and self-control,” *Econometrica*, 69(6), 1403–1435.

- Hausman, J. A. (1979), “Individual discount rates and the purchase and utilization of energy-using durables,” *Bell Journal of Economics*, 10(1), 33–54.
- Hensher, D., and W. Greene (2003), “The Mixed Logit Model: The State of Practice,” *Transportation*, 30(2), 133 – 176.
- Herrnstein, R. J., G. F. Loewenstein, D. Prelec, and W. Vaughan (1993), “Utility maximization and melioration: Internalities in individual choice,” *Journal of Behavioral Decision Making*, 6(3), 149–185.
- Hole, A. R. (2007), “Estimating mixed logit models using maximum simulated likelihood,” *Stata Journal*, 7(3), 388–401.
- Houde, S. (2013), “Bunching with the Stars: How Firms Respond to Environmental Certification,” Mimeo, University of Maryland.
- Houde, S. (2014), “How Consumers Respond to Environmental Certification and the Value of Energy Information,” *NBER WP No. 20019*.
- Kirby, K., and N. Marakovic (1996), “Delay-discounting probabilistic rewards: Rates decrease as amounts increase,” *Psychonomic Bulletin and Review*, 3(1), 100 – 104.
- Kuhnen, C., and B. Knutson (2011), “The influence of Affect on Beliefs, Preferences, and Financial Decisions,” *Journal of Financial and Quantitative Analysis*, 46(3), 605 – 626.
- Laibson, D. (1997), “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 112(2), 443 – 478.
- Leard, B. (2013), “Consumer Heterogeneity and the Energy Paradox,” Mimeo, Cornell University.
- Min, J., I. L. Azevedo, J. Michalek, and W. B. de Bruin (2014), “Labeling energy cost on light bulbs lowers implicit discount rates,” *Ecological Economics*, 97, 42–50.
- Newell, R. G., and J. Siikamaki (2014), “Nudging Energy Efficiency Behavior: The Role of Information Labels,” *Journal of the Association of Environmental and Resource Economists*, 1(4), 555–598.

- O'Donoghue, T., and M. Rabin (2000), "The economics of immediate gratification," *Journal of Behavioral Decision Making*, 13(2), 233–250.
- Peters, E., N. Dieckmann, A. Dixon, J. H. Hibbard, and C. Mertz (2007), "Less is more in presenting quality information to consumers," *Medical Care Research and Review*, 64(2), 169–190.
- Revelt, D., and K. Train (1998), "Mixed logit with repeated choices: households' choices of appliance efficiency level," *Review of Economics and Statistics*, 80(4), 647–657.
- Sallee, J. M. (2014), "Rational Inattention and Energy Efficiency," *Journal of Law and Economics*, 57(3), 781–820.
- Sammer, K., and R. Wustenhagen (2006), "The influence of eco-labelling on consumer behaviour Results of a discrete choice analysis for washing machines," *Business Strategy and the Environment*, 15(3), 185 – 199.
- Shen, J., and T. Saijo (2009), "Does an energy efficiency label alter consumers' purchasing decisions? A latent class approach based on a stated choice experiment in Shanghai," *Journal of Environmental Management*, 90(11), 3561 – 3573.
- Small, K. A., and H. S. Rosen (1981), "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 105–130.
- Sundqvist, T. (2004), "What causes the disparity of electricity externality estimates," *Energy Policy*, 1753–1766.
- Swait, J., and J. Louviere (1993), "The role of the scale parameter in the estimation and comparison of multinomial logit models," *Journal of Marketing Research*, 305–314.
- Thaler, R., and C. Sunstein (2008), *Nudge: Improving Decisions About Health, Wealth, and Happiness*, Penguin.
- Train, K. (2009), *Discrete Choice Methods with Simulation*, Cambridge University Press.
- Tsvetanov, T., and K. Segerson (2013), "Re-evaluating the role of energy efficiency standards: A behavioral economics approach," *Journal of Environmental Economics and Management*, 66(2), 347–363.

- U.S. Census Bureau (2014), “Residential Vacancies and Homeownership in the Fourth Quarter 2013,” <http://www.census.gov/housing/hvs/files/qtr413/q413press.pdf>.
- Ward, D., C. Clark, K. Jensen, S. Yen, and C. Russell (2011), “Factors influencing willingness-to-pay for the ENERGY STAR label,” *Energy Policy*, 39(3), 1450 – 1458.
- Watson, D., L. Clark, and A. Tellegan (1988), “Development and validation of brief measures of positive and negative affect: The PANAS scales,” *Journal of Personality and Social Psychology*, 54(6), 1063 – 1070.
- Wooldridge, J. M. (2010), *Econometric Analysis of Cross Section and Panel Data*, MIT press.