NBER WORKING PAPER SERIES

DO CONSUMERS ACQUIRE INFORMATION OPTIMALLY? EXPERIMENTAL EVIDENCE FROM ENERGY EFFICIENCY

Andrea La Nauze
Erica Myers

Working Paper 31742
http://www.nber.org/papers/w31742

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2023

We acknowledge funding from the University of Pittsburgh Social Science Research Initiative. The authors have no material financial interests related to this research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Do Consumers Acquire Information Optimally? Experimental Evidence from Energy Efficiency
Andrea La Nauze and Erica Myers
NBER Working Paper No. 31742
September 2023
JEL No. D12,D83,D91,Q41

ABSTRACT

We use an experiment to test whether consumers optimally acquire information on energy costs in appliance markets where, like many contexts, consumers are poorly informed and make mistakes despite freely-available information. We find consumers acquire information suboptimally; there is little correlation between the revealed utility gain from improved decision making due to information and willingness to pay for information. We compare two behavioral interventions to address consumer mistakes: a conventional subsidy for energy-efficient products and a non-traditional subsidy paying consumers to view information on energy costs. We show that paying for attention can target welfare improvements more effectively.

Andrea La Nauze
University of Queensland
Level 6, Colin Clark Building
Brisbane, QLD 4072
Australia
a.lanauze@uq.edu.au

Erica Myers
Department of Economics
University of Calgary
414 Social Sciences Building
2500 University Drive NW
Calgary, AB T2N 1N4
Canada
and E2e
and also NBER
erica.myers@ucalgary.ca

A data appendix is available at
http://www.nber.org/data-appendix/w31742
A randomized controlled trials registry entry is available at
https://www.socialscienceregistry.org/trials/6656

1 Introduction

Despite readily available information, consumers are often poorly informed and as a result, make mistakes in a variety of decision-making settings including health care, financial investments, and energy efficiency (e.g., Hortaçsu and Syverson, 2004; Choi, Laibson and Madrian, 2010; Handel, 2013; Bronnenberg et al., 2015; Handel and Kolstad, 2015; Bhargava, Loewenstein and Sydnor, 2017; Choi, Laibson and Madrian, 2010; Houde and Myers, 2019). The failure of consumers to acquire valuable information could be a rational response to "frictions", in the form of costly time and attention required to acquire information. Or, it could result from "mental gaps", psychological distortions in how consumers gather, process, and attend to information (Handel and Schwartzstein, 2018).

In this paper, we go beyond identifying the consequences of imperfectly informed consumers to distinguish between these two possible causes. This endeavor is important for two main reasons. First, empirical evidence on whether frictions or mental gaps drive consumer information acquisition determines the suitability of different modeling approaches. Rational inattention models, for example, assume that individuals optimally trade off the costs and benefits of acquiring information (Sims, 2003; Woodford, 2012; Caplin and Dean, 2015; Gabaix, 2014; Maćkowiak, Matějka and Wiederholt, 2021). Alternative models incorporate mental gaps, for example by allowing consumers to have the wrong decision model or suffer from biased beliefs and overconfidence (Schwartzstein, 2014; Bénabou and Tirole, 2002). Second, whether frictions or mental gaps explain information acquisition matters for welfare and policy across innumerable settings where information is readily available. If consumers acquire information optimally in response to frictions, low-cost policies that reduce those frictions will improve welfare. However, in the presence of mental gaps, consumers will make mistakes even in the absence of frictions, and more heavy-handed policy, such as subsidizing information acquisition could better target welfare improvements (Sallee, 2014).

Our first contribution is to design and implement a sufficient-statistic test for whether consumers optimally acquire information. Our test builds on existing literature studying whether decision makers acquire more information when the stakes are higher (e.g., Bartoš et al., 2016; Hoopes, Reck and Slemrod, 2015; Fuster et al., 2018; Gabaix et al., 2006). Testing for *optimal* information acquisition is challenging because it requires three parameters that are typically unobserved: individual measures of willingness to pay for information, the utility value of information, and cognitive effort costs of processing information. We design a multi-stage experiment using incentive-compatible mechanisms to elicit these parameters. In our experiment, consumers choose between products and have the option to acquire information on a product characteristic, which they may be uncertain about. The product-characteristic information is valuable if it causes the consumer to make a better product choice that leads to higher utility. We quantify the utility gain due to improved choices from information using within-subject changes in product willingness to pay. Conditional on effort costs and under risk neutrality, optimal information acquisition implies the average realized value of information across consumers moves one-for-one with their willingness to pay for that information.

We implement the test of optimal information acquisition in the setting of appliance markets, where it has been established that many consumers make mistakes by under-valuing lifetime operating costs (e.g., Houde and Myers, 2019). We focus on the light bulb market, which has several advantages. First, light bulbs are a common appliance that consumers are likely to have experience with. Second, previous experiments demonstrate that lack of information about lifetime operating costs leads to imperfect light bulb choices (e.g., Allcott and Taubinsky, 2015; Andor, Gerster and Götte, 2019; Rodemeier and Löschel, 2022; Beattie, Ding and La Nauze, 2022). As a result, we can offer a well-defined information signal that, in expectation, will have value to consumers. For the purposes of external validity, the existing literature also allows for a direct comparison of consumer behavior and the effects of information in our experiment with other settings that may use different light bulb package offerings, have different experimental populations, different stakes, or provide information in alternative formats. To this end, we use a well-established technique for eliciting consumers' relative willingness to pay for a package containing more versus less efficient bulbs known as an incentive-compatible multiple-price list. Consumers are allocated a budget and asked to choose which package they prefer at a series of relative prices knowing that each choice has an equal probability of becoming their actual purchase. Like previous experiments, we elicit this before and after a consumer receives information on the lifetime operating costs. But, while other experiments are designed to measure sub optimal product choices, our experiment is designed to compare the revealed value of information with willingness to pay for information and test for sub optimal information acquisition.

In addition to the value of information, our test requires two new parameters that we experimentally recover: consumer willingness to pay for information on lifetime costs and a measure of effort costs. We elicit the former after the first multiple price list for the two bulb packages. We describe the well-defined information that consumers may view, which is electricity use and lifetime information based on values provided by the manufacturer on the back of the box in typical market settings. Then, consumers select whether to view the information at a series of positive and negative price points (i.e., reductions or increases in their budget). To make the choice incentive compatible, a price point is chosen at random, and the consumer's implied choice of whether to view the information at that price is implemented.

To elicit a measure of effort costs we mimic the task used to elicit willingness to pay for information, with one key difference. Instead of offering information that can improve product choices, we offer consumers similar information for obsolete light bulbs and elicit their willingness to accept to review this information and answer a question. As the information has no impact on product choices, it has a utility value of zero, and the elicited willingness to accept to undertake this task can be used to measure individual effort costs. During the experiment, we also measure each consumer's distribution of baseline and endline beliefs about the relative lifetime costs of each light bulb technology, elicit risk and time preferences, and collect demographic characteristics. We implemented the experiment online using a representative sample of 1902 adults in the United States.

We strongly reject consumers optimally acquire information, indicating errors in decision making are due to mental gaps and not purely the result of frictions. On average, across informed consumers, the value of information revealed through changes in product choices does not move one-for-one with their willingness to pay for information. In fact, the correlation is 0.04, the result of consumers both over- and under- valuing information.

Our next contribution is to understand whether competing decision models are consistent with observed behavior, and in particular to understand how biased beliefs affect information acquisition. Using an individual-level measure of bias, we also reject that otherwise rational consumers acquire information suboptimally due to biased beliefs. Instead, the probability of under-valuing information is correlated with initial preferences for energy efficiency and proxies for consumers' political leaning. This suggests behavior consistent with models of confirmation bias and motivated reasoning where individuals seek out information to confirm beliefs or a preferred state of the world may contribute to the valuation patterns we observe (Kunda, 1990; Epley and Gilovich, 2016; Charness and Dave, 2017).

Our final key contribution is to propose and analyze a non-traditional subsidy to address mental gaps in information acquisition that pays people to acquire and process information. We show that this subsidy is a function of reduced form sufficient statistics from our experiment (e.g., Mullainathan, Schwartzstein and Congdon, 2012; Allcott and Taubinsky, 2015; Farhi and Gabaix, 2020) and compute the optimal subsidy in our setting. A single public entity that can gather, organize, and disseminate information at low marginal cost may improve efficiency by reducing information frictions. However, if consumers irrationally ignore freely-available information, public provision of information may be less effective than anticipated. In this case, incentivizing consumers to view and attend to information could improve welfare. This type of intervention is increasingly relevant given the prevalence of online shopping and has been implemented in practice. For example, in 2023 the Victorian Government in Australia offered a \$250 incentive for consumers to compare the costs of rival electricity plans on their price comparison website.[1]

We compare the properties of this information subsidy to a more conventional product subsidy. The optimal information subsidy is equal to the average undervaluation of information by consumers whose information choice is marginal to the subsidy. The subsidy improves welfare because the value of improved product choice from information outweighs the effort cost of attending to information. The optimal product subsidy is equal to the average bias in product choice of consumers whose product choice is marginal to the subsidy. The product subsidy improves welfare because the gains from improved product choices outweigh the losses from any distortions in product choices the subsidy introduces.[2] Whether a product or information subsidy is preferred depends on the magnitude of the gains from improved decisions from each instrument versus the losses of mental effort (from the information subsidy) or distorted choices (from the product subsidy). Our results indicate that the optimal information subsidy is approximately \$9 on a \$20 package of LED bulbs while the optimal product subsidy is approximately \$12. In our setting, the information subsidy dominates, increasing welfare by 43 percent more than the product subsidy.

Our results contribute to several strands of literature. First, our study speaks to a growing body of work on inattention (see Gabaix (2019) for a recent review). We contribute to this literature by testing for optimality of information acquisition, which is important for welfare and policy in

[1] See https://compare.energy.vic.gov.au/, accessed March 16, 2023.

[2] There will be some consumers who are induced to buy the more efficient product because of the subsidy but whose undistorted willingness to pay for the product is less than its marginal cost, generating a welfare loss.

a wide variety of contexts where information is easily accessible, though potentially underutilized. Our work is related to studies that test for optimality of other aspects of attention. For example, Howard (2022) finds that experts produce attention more efficiently than novices in online chess games. Bronchetti et al. (2022) find that people do not optimally acquire so-called "bandwidth enhancements" to manage future attention.[3] Our findings also add to the literature in economics and psychology on confirmation bias and motivated reasoning. Recent evidence suggests incentives are a powerful solution to overcome avoidance of moral information and belief in misinformation (Zimmermann, 2020; Serra-Garcia and Szech, 2022). The findings of our experiment speak directly to how incentives to acquire information overcome biased acquisition of information by consumers. The fact that consumers are not optimally acquiring information on light bulbs, a relatively common consumer purchase, suggests that there are likely other purchase and choice decisions where people are irrationally underinformed and incentives may be an effective intervention.

Our results also build on the nascent behavioral public economics literature describing optimal policy with behavioral agents (e.g., Bernheim and Rangel, 2009; Mullainathan, Schwartzstein and Congdon, 2012; Taubinsky and Rees-Jones, 2018; Farhi and Gabaix, 2020; List et al., 2022). Much of the previous work focuses on optimal allocation policies, such as quantity standards or taxes, to correct distorted product choice (e.g., Allcott, Mullainathan and Taubinsky, 2014; Allcott, Lockwood and Taubinsky, 2019; Houde and Myers, 2019), or, in some cases, on policies that reduce information frictions (e.g., Handel, Kolstad and Spinnewijn, 2019). We add to this

[3] Recent work also tested various predictions of rational inattention models in the laboratory (e.g., Caplin, Dean and Martin, 2011; Khaw, Stevens and Woodford, 2017; Dean and Neligh, 2023).

literature by considering a non-traditional policy designed to address mental gaps in information acquisition. We provide direct empirical evidence on the magnitude of distortions in consumer information acquisition. We then show how the welfare impacts of optimal information subsidies and optimal product subsidies can be compared.

Finally, our results speak to our understanding of the impact of information provision and labeling interventions. There is minimal evidence supporting the efficacy of government-mandated information disclosure programs across a wide range of contexts (see reviews in Winston, 2008; Loewenstein, Sunstein and Golman, 2014; Ho, Ashwood and Handan-Nader, 2019). Our results suggest consumers misperceive the value of attending to and incorporating information, so labeling and disclosure policies may have limited effects. This is of particular importance in the context we study. Appliances and building-related equipment account for almost all the energy used in buildings, which as a sector accounts for around 40% of U.S. energy consumption and associated greenhouse gas (GHG) emissions (Baldwin et al., 2015). Despite the fact that energy labeling policies for energy-using durables have been in place in the U.S. since the 1980s, a large share of consumers misperceive the associated costs (Houde and Myers, 2019). When people lack information or are inattentive to the operating costs of energy-using durables, they will sub optimally invest in efficiency even under carbon pricing policies. We demonstrate that mental gaps in information acquisition about operating costs may be an important driver of investment inefficiencies and show that a non-traditional policy instrument, information subsidies, can target welfare improvements more effectively than conventional product subsidies.

2 Test of Optimal Information Acquisition

We develop a straightforward test of optimal information acquisition comparing, for a group of consumers, ex ante willingness to pay (WTP) for information with the ex post value of that information in terms of improved decision making.

Suppose a consumer is purchasing an appliance and can choose to purchase information about the operating costs of different models (e.g., differences driven by the energy consumption and the appliance lifetime). We begin by assuming consumers only gain value from information through improving this product choice, for example, there is no external value of information outside the choice of which light bulb to buy. Let the utility of an informed, risk-neutral individual i for appliance j be $U_{ij} = \nu_{ij} - p_j - \kappa_{ij}$ where ν_{ij} includes the consumer's valuation of the appliance's characteristics that do not affect its operating costs (e.g., color), p_j is the price of appliance j and κ_{ij} is the consumer's lifetime operating cost of appliance j. When individuals are not perfectly informed about κ_{ij} they have an estimate k_{ij} where $\kappa_{ij} = k_{ij} + \epsilon_{ij}$ and ϵ_{ij} is an error term that is mean zero if consumers have unbiased beliefs. Therefore, an uninformed individual's experienced utility, which they actually receive from their choice, may differ from their choice utility, which they anticipate when they are uninformed about product costs. For an informed individual, the choice utility and the experienced utility are the same.

Ex post, the realized value of information is the gain in experienced utility from becoming informed. For the given utility function, information only changes experienced utility if it leads the consumer to choose a different appliance than they otherwise would have. Consider a consumer choosing between two appliance models: an efficient model 1, and an inefficient model 0. Let $l \in \{0, 1\}$ denote the preferred product of the informed consumer and $m \in \{0, 1\}$ denote the preferred product of the uninformed consumer at prices p_l and p_m. If $l = m$ (i.e., the preferred products of the informed and uninformed consumer are the same) then the realized value of information is zero. In this case, the consumer's experienced utility from the preferred product may be higher than their choice utility, but they would have realized this higher utility regardless of whether they become informed. If $l \neq m$, then the realized value of information is the difference in experienced utility between products: $U_{il} - U_{im}$. Define the informed consumer's relative WTP for product l: $W_{il} = (\nu_{il} - \kappa_{il}) - (\nu_{im} - \kappa_{im})$ (the utility derived from the non-price characteristics of purchasing product l versus product m) then the ex post value of information can be written:

$$V_i = W_{il} - (p_l - p_m), l \neq m \tag{1}$$

Ex ante, to choose optimally whether or not to acquire information, a rational consumer forms an expectation about the value of information $E_i[V_i]$ (where we use subscript i on the expectation operator E to denote this expectation is from the perspective of the consumer, taken over the distribution of ϵ). This reflects the difference between the expected value of the good maximizing experienced utility minus the utility of the good maximizing choice utility:

$$E_i[V_i] = E\left[\max_j(\nu_{ij} - p_j - \kappa_{ij})\right] - \max_j\left[\nu_{ij} - p_j - k_{ij}\right] \tag{2}$$

If consumers are rational, they have an accurate expectation of the value of information, meaning: they have well-defined priors over the long-run costs of each technology and a correctly-specified model of the distribution of signals they will receive from the information. Further, given risk neutrality, their WTP for information, W_i^I, would be the expected value of updating beliefs in response to the information signal net of any idiosyncratic cognitive effort cost e_i.

$$W_i^I = E_i[V_i] - e_i \tag{3}$$

Note that it may be rational for consumers to have a negative WTP for information, for example, if their expected value of information is low, and effort costs are high. Further, for any individual, the expected value of information may not equal the realized value of information given there is a random component to individuals' uninformed beliefs. However, across the population, if consumers are optimizing, the expectation of the revealed value of information should equal the mean of the consumers' ex ante expectations of its value $E[V_i] = E[E_i[V_i]]$ (where $E[]$ is the expectation over the distribution of consumers). Solving for $E_i[V_i]$ in 3, substituting, and taking the conditional expectation gives:

$$E[V_i|W_i^I, e_i] = W_i^I + e_i \tag{4}$$

Under risk neutrality, these equations suggest two possible avenues for a test of optimal information acquisition. First, if all the components of equation 4 are observable, then one can compare if, on average, the left-hand side is equal to the right-hand side. Second, if effort cost is not perfectly observable, a measure of consumers' realized value of information and their

WTP for that information are sufficient to perform the test as long as effort cost is not correlated with both W_i^I and V_i.[4] In this case, a necessary condition for optimal information acquisition is that the population expectation of the realized value moves one-for-one with WTP for information. While effort costs should determine WTP for information, they should not directly determine the expected value of information, which is a function of the likelihood the consumer changes their choice.

In the next section, we outline the experiment enabling us to implement this test across consumers by recovering revealed preference estimates of W_i^I and V_i and a proxy for effort costs. The effort cost proxy allows us to empirically test whether effort costs are uncorrelated with the value of information. Finally, we elicit other information to further test assumptions of our model, including measures of risk preference, and other potential drivers of the value of information.

3 Experimental Design and Implementation

Given that the parameters needed for our test of rational information acquisition are unobservable in natural market settings, we have to turn to an artefactual field experiment. However, we designed important aspects of the experimental environment to be as natural as possible. First, we chose a product that consumers are familiar with and are therefore likely to have well-defined priors over. Most consumers have some experience with light bulb purchases, and as we describe below, we elicit priors about the relative costs of the two types of bulbs as part of the experiment. Second,

[4] If effort cost is correlated with both W_i^I and V_i, it would create an omitted variables bias on the slope coefficient from regressing realized value of information on WTP for information.

the information we offer in the experiment is well-defined. Consumers can choose to view electricity consumption and lifetime metrics reported by manufacturers on appliance packaging; a replication of information found in typical market settings. This gives consumers the best shot, as good as in a natural environment, at having a correctly-specified model of the distribution of signals they will receive from the information.

In our online experiment, consumers purchase light bulbs and information on the lifetime costs of the bulbs. Consumers are offered two types of bulbs that are similar along many product dimensions but differ substantially in their lifetime energy costs and durability: light emitting diode (LED) bulbs and halogen (incandescent). LED bulbs are superior in terms of energy use and durability, but consumers may have high WTP for incandescents, for example, because of strong preferences for light quality. The key variables measured in the experiment are: the distribution of beliefs on the lifetime costs of bulbs before and after information; relative WTP for LED light bulbs before and after information; WTP for information; and willingness to accept (WTA) for an effort task. At the beginning of the experiment, participants were informed that one in ten people who complete the experiment would win a prize worth up to $100.[5] The value of the prize was determined by their choices and random price draws, and consisted of any light bulbs that were purchased as part of the incentive-compatible WTP elicitations and any remaining budget on an Amazon gift card. Participants were also paid for completing the experiment.[6]

[5]In expectation the stakes in our experiment are similar to the literature where consumers are certain to receive a light bulb package and have a budget of $10-13 such as Allcott and Taubinsky (2015); Andor, Gerster and Götte (2019).

[6]The full experimental transcript is here. We sent light bulbs and gift cards, panel providers contracted by Qualtrics made direct payments for survey completion.

3.1 Belief Elicitation

To test alternative theories of behavior we measure the full distribution of consumer beliefs about the relative lifetime costs of the bulbs before and after making this information available. To do so we asked consumers to consider the total costs of lighting one lamp for 10 years for 3 hours per day given a single bulb purchase price of $2. We then asked two follow-up questions. First, consumers were asked whether they thought the total cost would be higher using the LED or halogen bulbs. Then, consumers were asked to place 10 tokens in pre-determined ranges according to their expectations about the difference in total cost.[7] To assist consumers we asked them to suppose they were betting on the true costs using the 10 tokens and provided hints and examples. This latter task is a simplified version of the subjective elicitation task of Shrestha (2020). For each allocated token we compute the midpoint of the range and for each consumer, we compute the mean and standard deviation of beliefs at baseline and endline.

3.2 Light Bulb Choices

In our experiment, consumers chose between two packages of light bulbs: a package of 12 light-emitting diode (LED) bulbs and a package of 24 halogen (incandescent) light bulbs. The packages were the same brand (GE), as similar as possible on the characteristics of color and lumens, and the most common shape and fitting, but differed substantially in their energy consumption and durability. Basic information on the two packages, including wattage, which determines the electricity consumption of the bulbs, was

[7] Delavande and Kohler (2009) show probabilistic expectations elicited in surveys using similar tasks follow basic properties for probabilities, are correlated with qualitative measures of expectations, and are predictive of outcomes and behaviors.

provided to all consumers. To elicit WTP for the efficient bulb before and after information on lifetime costs, we used an incentive-compatible price list. Consumers selected their most preferred package at 17 price points for the LEDs relative to the halogens in the range -\$60 to \$60 where -\$60 corresponds to the cost of the LED package being 0 and the cost of the halogen package being \$60. The point at which consumers switch from one bulb type to another bounds their relative WTP for the package of LEDs. For example, if a consumer chose the package of halogens at a relative price of 0, and chose the package of LED bulbs at a relative price of -\$3.75 then the consumer's WTP for LEDs is between 0 and -\$3.75 (symmetrically, they have a WTP for the halogens of between 0 and \$3.75). For our analysis, we take the midpoint of this range, so that this consumer would have a WTP of -\$1.875 for the package of LEDs. If a consumer preferred one package at all prices between -\$60 to \$60, they were asked to provide the price at which they would switch to the alternative but were told this price would not impact their shopping budget, i.e., it was hypothetical.

Consumers were told that if they were a prize winner, one of the prices from the list would be picked by the computer (with each row having an equal chance of being picked) and they would be sent the package and any remaining shopping budget as an Amazon gift card. For those with WTP for the LED between -\$60 to \$60 this is an incentive-compatible elicitation, meaning it is in consumers' interests to reveal their true WTP. For those with WTP outside this range, it is in consumers' interests to reveal that their true value is outside the range and there is no strategic advantage to over- or under-stating WTP.

Prior to the baseline light bulb choice, the price list format was explained to consumers using the example of two cereal products. The example showed all possible monotonic answers to the prices given (referring to these answers as "logical"). To proceed through the experiment, participants had to correctly identify monotonic and non-monotonic answers. If consumers submitted non-monotonic responses in the remainder of the experiment they were prompted once to answer again. If they continued to answer non-monotonically they were not considered a quality response and were not eligible for the prize and we do not observe their response.

We use the baseline and endline light bulb choices to calculate the value of information for all consumers receiving information in the experiment. From Equation 1, at a given relative price, the revealed value of information is the endline relative WTP for the preferred product minus its relative price if the endline preferred product differs from the baseline. As we observe the preferred product at 17 different relative prices, each with an equal chance of being realized, we take the average of the value of information across price points to calculate V_i for the experiment.

3.3 Information Choice and Information Treatment

After their first set of light bulb choices, we elicited participants' WTP for information on the total cost of using the light bulbs using a "staircase" procedure. This incentive-compatible procedure is a streamlined version of a price list requiring fewer choices and providing comparable responses (Falk et al., 2016, 2018).[8] We first explained to consumers that information would be available to them before they made a second set of light bulb choices. We then used a series of scenarios to elicit their WTP for the

[8] Unlike the price list there is no possibility for consumers to submit non-monotonic responses in the staircase.

lifetime cost information. Appendix Figure A1 shows how the module started while Appendix Figure A2 outlines the full staircase. Depending on their response to the preceding questions, the consumer is offered an increase in their shopping budget for viewing/not viewing the information up to a value of \$10. Each consumer is asked up to four questions and their responses to these questions allow us to bound their WTP in the same way as a price list. As with the price list, we take the midpoint of the WTP range revealed by the staircase.[9] If consumers always answered they wanted information or did not what information, consumers were asked to state their WTP/WTA to receive the information.

Once the staircase was completed, each consumer's choices were summarized for them. They were then told the computer would randomly draw a payment amount and if their WTP/WTA was above that amount they would receive the information. Before the draw occurred, consumers were given the opportunity to revise their responses.

In the experiment, consumers were offered one of eleven payment amounts ranging from -\$10 to \$10 to view information. To ensure that our sample contained sufficient consumers with low WTP for information, we set the probability of the largest subsidy (\$10) to 0.6. All other payment amounts were offered with probability 0.04. Depending on their WTP for information, and their random price draw, some consumers were then provided with the lifetime cost information. Appendix Figure A3 shows the information provided, which is based on the mandatory product energy label provided by manufacturers. Over 10 years, the costs of using LED light

[9]For example, if a consumer selected information when no change in their shopping budget was offered, they were asked if they would prefer an increase in their shopping budget of \$5 and not to view the information. If they then selected no information, they were offered \$2.50 to not view the information. If at this point they chose information then their WTP for information would be the midpoint of \$3.75.

bulbs are expected to be $60 less than the costs of using halogen light bulbs. In order to proceed from the information screen to the rest of the experiment, consumers that were provided with the information had to correctly answer a multiple-choice question identifying the difference in operating costs of the bulbs over 10 years.

3.4 Consumer Characteristics

In addition to age, gender, income, and education, we asked consumers whether they pay electricity bills, the number of light bulb sockets in their home, how many people live in their home, and we measured their risk and time preferences using the staircase procedures of Falk et al. (2016, 2018).

3.5 Effort Task

Finally, we designed an incentive-compatible task to measure variation in effort costs, the additional costs of acquiring information that are idiosyncratic and unobservable. To do so, we asked consumers whether they would be willing to review information on the lifetime cost of obsolete light bulbs and answer an additional question in return for a series of payments presented according to a staircase.[10] As this information has no impact on light bulb purchasing decisions it has an expected utility value of zero in the experiment. From Equation 3 if the expected utility value of information is zero, then WTA to view the information is equal to the cost of effort required to process it. The cognitive task of processing the information on obsolete bulbs is almost identical to the cognitive task of processing information on light bulbs in the price list. Therefore, this proxy should

[10] As with the previous staircase, consumers are informed that one of the prices will be drawn at random and their choices will determine whether they undertake the task. Appendix Figure A4 outlines the full staircase.

capture across-subject differences in the determinants of effort costs, such as cognitive ability, needed for our test.

3.6 Data and Implementation

Our experiment was implemented online using the Qualtrics platform from October 2020 to February 2021.[11] The target sample were residents of the United States ages 18 and over, matched to census probabilities for age, gender, and education. Appendix Table B1 reports that the sample matches well on these target probabilities. The target sample size was 2000. Descriptive statistics of the sample are provided in Appendix Table B2. Our final sample consists of 1902 consumers, 1435 of these received the information treatment. Consumers took an average of 3 minutes 45 seconds to complete the baseline light bulb WTP elicitation and 1 minute 48 seconds to complete the endline light bulb elicitation. Those who viewed the information spent an average of 2 minutes on the information screen.

There are two well-known issues with price list elicitations. First, consumers may answer non-monotonically. In the implementation, Qualtrics removed consumers who answered the bulb choice price lists non-monotonically after being prompted to re-answer, hence we do not observe non-monotonic responses to the price lists.[12] Second, consumers with a WTP outside the price list report a stated WTP, which is hypothetical. To remove the influence of outliers from responses to hypothetical questions, we winsorize the following variables at the 5 percent level: standard deviation of beliefs, baseline, and endline WTP for bulbs and information, and WTA for

[11] We ran an initial pilot in early October. Details of the pilot are available on request.

[12] To minimize non-monotonic responses, consumers saw an example using cereal and had to correctly identify monotonic responses.

the effort task.[13] In Section 4.1 we demonstrate that results are robust to alternative approaches to dealing with outliers.

Before we turn to the primary hypothesis about optimal information acquisition, we present evidence our experiment measures our parameters of interest and that the behavior we observe closely aligns with previous literature. One concern with the experimental design may be that consumers maximize their gift card balance and their responses to the experimental procedures do not reveal true WTP. If this were the case, we would expect to see consumers in both baseline and endline light bulb elicitations always choose the cheapest light bulb package. In practice, at baseline 18% of consumers always choose this package, and at endline 15% of the uniformed consumers and 13% of informed consumers choose this package.[14]

Further, if consumers were maximizing their gift card balance, we would also expect to find that consumers have zero or negative WTP for information on the lifetime costs of bulbs, and that information, if received, has no effect. In contrast, we find that information increases consumers' WTP for the more efficient bulb by similar magnitudes as previous light bulb information experiments, which may have different experimental populations, light bulb technology, stakes, or information formatting.

Figure 2 shows the demand for the LED package measured as market share at endline and baseline for the group receiving information and not receiving information. For those receiving information, there is a clear increase in their demand for the LED package after receiving information. In

[13] Appendix Figure B1 shows the distribution of each variable winsorized at the 1st and 5th percentiles. In our sample, 3.7% and 17.5% were censored towards the halogen and LED at baseline. At endline, 3.1% and 31.2% were censored towards the halogen and the LED respectively. This is in line with previous literature (e.g., Allcott and Taubinsky, 2015; Beattie, Ding and La Nauze, 2022).

[14] We also ensure that dropping consumers who always choose this package doesn't change our conclusions.

contrast, demand for the LED package changes little between baseline and endline for those not receiving information. Appendix B.2 reports estimates of the treatment effect of information at each level of WTP for information and the average treatment effect computed by weighting these conditional treatment effects using the distribution of WTP for information.[15] We find that the average effect of information on WTP for the LED package is approximately \$23. As a percentage of baseline WTP for the efficient bulb, the treatment effect is comparable to previous literature including Allcott and Taubinsky (2015); Andor, Gerster and Götte (2019); Rodemeier and Löschel (2022); Beattie, Ding and La Nauze (2022), who use experiments to estimate the impact of light bulb lifetime cost information on demand for efficient bulbs in the United States, Europe, and China.

We also find, consistent with previous literature, that the mechanism for the effect of information appears to be that consumers are updating their beliefs in line with the information provided. Using a simple Bayesian learning model, we find the average treated consumer puts substantial weight (25%) on the information signal provided in the experiment relative to their priors (75%) in formulating their endline beliefs (See Appendix Table B8).[16] Among those receiving information, we also find that consumers who spend more time on the information screen or in reporting their beliefs appear to learn more from (i.e., put more weight on) the information provided (See Appendix Table B8).

[15] This approach avoids issues with the arbitrary weighting of treatment effects from a continuous instrument (the price draw) and binary treatment variable as in Heckman and Vytlacil (2005). Standard errors for the average treatment effect are bootstrapped.

[16] We estimate the following model:

$$\Delta MeanBeliefs_i = \alpha(Signal - Prior_i) \times ReceivedInformation_i + \beta(Signal - Prior_i) + \delta W_i^I + \epsilon_i$$

where α measures the "true" learning rate while β reflects spurious mean reversion. "Received Information" is random conditional on W_i^I.

A final concern about the experimental design is whether our effort task reveals meaningful variation in effort costs across participants. This task was designed to mimic the effort required to process light bulb cost information for another type of information with no value in the experiment. Therefore, the proxy should be highly correlated with individuals' effort costs in the information task. In Appendix Table B3 we show support that our measure does indeed appear to be a good proxy for effort cost. In particular, we find those with higher WTA in the effort task have lower WTP for information, are less patient, and spend less time on the information screen if they view it in the experiment, relationships that we would expect to find if our measure captures variation in effort costs.

4 Results

4.1 Testing Optimal Information Acquisition

We now turn to our test of optimal information acquisition, which compares the value of information with consumers' WTP for information. Figure 1 shows the distributions of the value of information and WTP for information in our sample. Panel (a) of Figure 1 shows the distribution of the winsorized value of information for the informed group. For around a third of consumers, the information has no value, meaning that it did not change their preferred bulb at any point in the price list. However, the mean value of information across all consumers was positive at \$8.36.

Panel (b) of Figure 1 shows the distribution of winsorized WTP for information for the group receiving information, which we use for our test of optimal information acquisition, and those not receiving information, which we do not include. Mean WTP for information among those receiving it was substantial, \$23.92. Of the 467 not receiving information, 68% had a WTP below the lowest price in the information price list (i.e., they were not marginal to the maximum subsidy we offered).

Recall from Equation 4, if consumers are risk neutral and their effort cost is not correlated with both W_i^I and V_i, then a test of optimal information acquisition is whether the revealed value of information increases 1 for 1 with WTP for it. Figure 3 plots this relationship for the subset of consumers that viewed information. From this figure it is clear that the line of best fit has a slope that is substantially lower than one. Individuals with low WTP for information appear to be undervaluing the impact of information on their product choice while individuals with high WTP for information appear to overvalue it.

Column (1) of Table 1 reports estimates from a formal test of the hypothesis. We regress the realized value of information on WTP for information.[17] Standard errors are robust and observations are weighted by the inverse probability of selection into treatment to account for the uneven probability of treatment conditional on WTP for information.[18] With this approach, we strongly reject that consumers are obtaining information optimally, i.e., a coefficient of one. The average value of information for consumers with zero WTP is \$7.50 while a \$1 increase in WTP for information leads to a \$0.04 increase in the value of information.

As mentioned, a condition for this to be a valid test of optimal information acquisition is that effort costs are not correlated with both the value

[17] We also consider a Poisson model specification. The results are consistent with the OLS presented here. See Appendix Table B9.

[18] For these regressions our sample is also truncated at a WTP for information of -\$8.75. Those with WTP less than -\$8.75 have a zero probability of receiving information. As this selection is a deterministic function of our independent variable our coefficients are not biased by this selection (Wooldridge, 2010).

of information and WTP for it. If so, it would create an omitted variables bias on the slope coefficient. To address this concern, we use consumers' revealed WTA to complete the effort task in the experiment as a proxy for effort costs in the information task. In Appendix Table B3 we show there is no correlation between WTA for the effort task and the revealed value of information[19], which suggests that effort costs are unlikely to be driving omitted variables bias in our context. Indeed, when we include the effort cost proxy as a control in our regression of value of information on WTP for it (Column (2) of Table 1), we find no significant or economically meaningful difference in the coefficient of interest.[20]

The test of optimality also requires an accurate representation of WTP, which can theoretically be influenced by choices in how we code responses to the price list. In columns (3) and (4) of Table 1 we consider whether WTP/WTA answers outside of the bounds of our price list (thus revealed in an open-ended, not-incentive-compatible way) could be driving the results. In column (3), we restrict the sample to consumers with an incentive-compatible WTP for information within the bounds of the information price list. Then, in column (4) we further restrict the sample to consumers with an incentive-compatible WTA within the bounds of the effort task price list. Neither of these restrictions significantly changes the coefficients, suggesting that responses outside the bounds of our experimentally defined price lists are not driving our findings. Further, our conclusions are not

[19] As noted earlier, the effort cost proxy is correlated with WTP for information, patience, and time spent on the information screen.

[20] Consumers have a one in ten chance of being a prize winner, receiving the value of information and paying for it. However, before prize winners are known, consumers receiving information invest the effort cost of processing information. For the purposes of this regression, re-scaling variables to account for the differences in uncertainty is unnecessary as the coefficient of interest is the slope of WTP for information, which has the same probability of being realized as the value of information.

sensitive to other choices in how we: (1) winsorize values outside of the price lists (Appendix Figure B3), (2) impose sample restrictions to account for the quality of responses (Appendix Table B6), or (3) assign levels of WTP for information within each range of values isolated by the price list–e.g., median, lower bound, upper bound of the price ranges (Appendix Table B7).

We also consider the role of risk tolerance. If consumers are risk-loving and optimally acquiring information the expected value of information would increase by less than one with increases in WTP for information. However, for risk tolerance to drive the weak relationship between WTP for information and its revealed value that we find, consumers would have to be extremely risk-loving. In contrast (consistent with the literature (e.g., Falk et al., 2016)), we find the vast majority of consumers in our experiment are somewhat risk averse (see Appendix Figure B2 for the distribution).

Finally, we consider the possibility that consumers may value the information for future light bulb purchases outside of the experiment, which could lead to the value of information increasing by less than one for a one-unit increase in WTP for it. If this were the case, with optimally behaving consumers, we would expect to find heterogeneity in the relationship between the value of information and WTP for it by a household's future demand for light bulbs. The difference in the expected annual lighting costs for LED versus halogen bulbs is primarily driven by the number of sockets in the home. For example, the expected annual savings of LED over halogen bulbs is approximately $600 for a home with 10 sockets and $1500 for a home with 25 sockets. However, we do not find any evidence consistent with homes with more sockets, and thus higher demand for lighting, having a weaker relationship between the value of information and WTP for it. Appendix Figure B4 displays this relationship for consumers with above and below the median number of sockets and there is no differential slope between the two groups. This suggests future demand for lighting is not a driving factor in the relationship between consumers' WTP for information and its value. Therefore, the relationship we see is not caused by consumers making optimal decisions about information that consider their future demand for light bulbs outside of the experiment.

4.2 Models of Information Acquisition

We have strong evidence that consumers are not acquiring information optimally. This suggests that consumers' lack of information about energy costs observed in other contexts is unlikely a rational response to the presence of "frictions", in the form of the time and attention required to acquire information (e.g., Maćkowiak, Matějka and Wiederholt, 2021). Rather, errors appear to be the result of "mental gaps", psychological distortions in how consumers make decisions including gathering, processing, and attending to information. In what follows we assess whether competing decision models are consistent with observed behavior.

4.2.1 Biased Beliefs and Rational Information Acquisition

We begin by assessing whether bias in prior beliefs leads otherwise rational consumers to make mistakes in information acquisition. Consider the utility function for uninformed individuals as presented in Section 2, where k_{ij} represents an estimate of true lifetime operating costs $\kappa_{ij} = k_{ij} + \epsilon_{ij}$. To derive testable predictions, we assume that ϵ_{ij} is a type 1 extreme value random variable with mean $\bar{\epsilon}_{ij}$ (representing bias in consumer beliefs) and scale parameter ϕ_i (representing the consumer's ex ante uncertainty). If m is the preferred appliance of the uninformed consumer and n is the alternative, then ex ante, the expected value of informing consumer i is:[21]

$$E[V_i] = \frac{1}{\phi_i} \ln\left(1 + \exp^{-\phi_i(\widehat{W}_{im}-(p_m-p_n)-(\bar{\epsilon}_{im}-\bar{\epsilon}_{in}))}\right) \tag{5}$$

where $\widehat{W}_{im}$ is the uninformed (or baseline) WTP for product m. According to this model, the expected value of information should be: decreasing with absolute relative WTP for the preferred bulb at baseline (smaller $|\widehat{W}_{i,m}|$), increasing with uncertainty over lifetime operating costs (lower ϕ_i), and increasing in the relative bias of the consumer *towards* their ex ante preferred product (larger $\bar{\epsilon}_{im} - \bar{\epsilon}_{in}$). However, consumers are unaware of their bias, and so make decisions as if $\bar{\epsilon}_{im} - \bar{\epsilon}_{in} = 0$. Biased beliefs could in part explain the failure to acquire information optimally if those who underweight the value of information are biased towards their ex ante preferred product and vice versa.[22]

Our experiment allows us to define individual-level bias in beliefs to test whether it can explain some of the observed distribution of mistakes. Because we elicit a distribution of beliefs from each individual, we can distinguish between bias and uncertainty. We define a consumer as having biased beliefs if the mean of their baseline belief distribution is statistically different at the 95% level to -\$62.50.[23] Figure B7 shows the distribution

[21] See Supplemental C.1 for derivation

[22] Information is valuable to consumers when it is pivotal - if individuals are biased against their ex ante preferred product, then correcting this bias will reinforce their choice. If individuals are biased towards their ex ante preferred product, then they will be more likely to change their mind in response to new information.

[23] According to our information treatment, the cost difference between the bulbs would be \$60. A consumer who was certain of this cost would enter all 10 tokens into the range -\$50 to -\$75, which has mid point -\$62.50. Alternative definitions of bias, such as a statistical difference in baseline and endline beliefs, provide similar results.

of mean ex ante beliefs by whether they are biased (53% of the sample) or unbiased (47% of the sample). We separate biased consumers into those who are optimistic about the LED (7% of the sample) and those who are pessimistic about the LED (46% of the sample).[24]

Equation 5 suggests that ex post, information should be more valuable to consumers who have beliefs biased towards their ex ante preferred product: optimists who initially prefer the LED and pessimists initially preferring the halogen. Further, these two groups of consumers should be more likely to undervalue information because the value of correcting their bias would not be reflected in their WTP. Panel (a) of Figure 4 plots the mean probability that a consumer undervalues information by whether they are optimistic, unbiased, or pessimistic about the costs of LEDs both for the group that initially prefers halogens and for the group that initially prefers LEDs. Error bars indicate 95% confidence intervals.

Interestingly, optimists, pessimists, and unbiased are equally likely to undervalue information among those who initially prefer LEDs (black bars). Similarly, the three groups are equally likely to undervalue information among those who initially prefer halogens (grey bars). Therefore being biased towards one's preferred product at baseline does not predict mistakes. However initial bulb preference does predict mistakes - consumers with an initial preference for the halogen are more likely to undervalue information. Appendix Table B10 reinforces the poor fit of the model of rational information acquisition with biased beliefs. In contrast to the predictions of this model, we find the WTP for information is increasing in the abso-

[24]These measures of bias have some predictive power. Compared to unbiased consumers: on average consumers who are optimistic about the LED are willing to pay $7 more for the LED package at baseline, while consumers who are pessimistic about the LED are willing to pay $6 less for the LED package at baseline.

lute value of baseline WTP for the LED, and decreasing in the consumer's uncertainty.[25] In Appendix Table B11 we also show that the rational information acquisition model is a better predictor of the revealed value of information than WTP for information.

4.2.2 Confirmation Bias and Motivated Reasoning

Having ruled out behavior predicted by rational individuals with biased beliefs, it is clear mental gaps are driving the results we see. While there may be a myriad of factors driving these mental gaps, we identify two candidate behavioral models that may help explain key features of behavior we observe: confirmation bias and motivated reasoning. These models and empirical evidence on belief formation suggest that information acquisition is affected in a non-Bayesian manner by preferences for information that confirms prior beliefs and/or a preferred state of the world (Epley and Gilovich, 2016; Charness and Dave, 2017; Zimmermann, 2020). The findings we observe could be consistent with some consumers engaging in motivated reasoning. For example, the fact that WTP for information is higher for those who are more certain in their prior beliefs is in line with individuals seeking out information if they are confident it will align with their beliefs. Further, motivated reasoning is also consistent with WTP for information being higher for those who initially prefer the LED bulb, if these consumers have a strong preference for a state of the world and are therefore more likely to seek out information to affirm their preference.

Panel (b) of Figure 4 provides further suggestive evidence on motivated

[25] In equation 5, ϕ_i may also capture consumers' anticipated attention to new information. Fuster et al. (2018) also find that WTP for information is higher among those who are ex ante more certain and show that these individuals also place more weight on information. Appendix Figure B6 shows that in our sample consumers who are more certain do not place more weight on information.

reasoning as a potential driver of mistakes in information acquisition. It shows the probability of undervaluing information for consumers based on the political orientation of their county using voting records in the 2020 Presidential Election (Election Data & Science Lab, 2022)[26]. In 2019, the Trump administration made headlines for removing regulations on the sale of inefficient light bulbs.[27] It is therefore plausible that some individuals evaluate information about the relative efficiency of light bulbs from a political standpoint. Panel (b) of Figure 4 shows the probability of undervaluing information in counties classified into "Democrat" (from a county in the top 20% of the Democratic vote share in the 2020 Presidential election), "Center" (from a county in the middle 20-80th percentiles of the Democratic vote share in the 2020 Presidential election), and "Republican" (from a county in the bottom 20% of the Democratic vote share in the 2020 Presidential election). We find consumers in strong Republican counties are more likely to undervalue information - the mean share of consumers who undervalue information in Republican counties is statistically different to the mean share of consumers in either Democrat or Center counties at the 5 per cent confidence level. This is consistent with some consumers engaging in politically motivated information acquisition that causes them to make costly mistakes.

4.3 Implications for Policy Design

In this section, we derive and compute optimal behavioral policies for addressing imperfect decision making in the presence of mental gaps as func-

[26]Matched using Department of Housing & Urban Development (2022).

[27]For media coverage of the decision see for example https://www.washingtonpost.com/climate-environment/2019/12/20/trump-administration-just-overturned-ban-old-fashioned-lightbulbs/.

tions of reduced-form sufficient statistics (e.g., Mullainathan, Schwartzstein and Congdon, 2012; Allcott and Taubinsky, 2015; Farhi and Gabaix, 2020). In the previous section, we showed suggestive evidence that motivated reasoning or confirmation bias are mental gaps that may partially drive the results we observe. However, the strength of the sufficient statistics approach is that to derive optimal policy and assess its impacts, it is sufficient to identify the impact of mental gaps on decision making (i.e., the internality), as our experiment is designed to do. It is not necessary to know the specific mechanisms driving mistakes in behavior.

4.3.1 Optimal Policy Interventions

In what follows, we consider two policy scenarios. First, we consider an approach where a planner chooses a single subsidy level to encourage consumers to view information when purchasing a product. Then we look at a more traditional approach where the planner offers a single subsidy level to encourage the purchase of the more efficient product. A subsidy for viewing information is increasingly applicable given that many purchasing decisions are being made online, making it possible to offer subsidies for viewing information screens and tutorials. As an example, in 2020 the Australian state of Victoria established a \$50 incentive for customers to view competing electricity contracts on their price comparison website. This was increased to \$250 in 2023.

Following the notation above, ν_{ij} is consumer i's valuation of the light bulb $j \in \{0, 1\}$, κ_{ij} is the lifetime operating costs for i of bulb j and p_j is the price of bulb j. For simplicity, we fix the value and lifetime operating costs of the inefficient bulb to be the same across consumers $\nu_{i0} = \nu_0$ and

$\kappa_{i0} = \kappa_0$ while allowing them to vary for the efficient bulb. For notational convenience, we also drop subscripts i such that the consumer's unbiased relative WTP for the efficient bulb is $W_1 = (\nu_1 - \kappa_1) - (\nu_0 - \kappa_0)$ and $\widehat{W_1}$ is their biased or uninformed relative WTP for the efficient bulb. We denote $F_X(x)$ the cumulative density function (CDF) of variable X[28] and $f_X(x)$ its probability density function (PDF). Then define the following: $Z(s^I)$ is the consumer's budget given subsidy s^I[29], $D_B(p^I) = 1 - F_{W^I}(p^I)$ is the biased demand curve of information, and $\tilde{V} = V - e$ is the value of information net of effort cost.

The welfare from subsidy s^I at price of information p^I can be written:

$$\mu(s^I) = Z(s^I) + \nu_0 - p_0 - \kappa_0 + \int_{\widehat{W_1} \geq (p_1 - p_0)} (W_1 - (p_1 - p_0)) f_{\widehat{W_1}, W_1}(\widehat{W_1}, W_1) d\widehat{W_1} dW_1$$
$$+ \int_{W^I > p^I - s^I} (\tilde{V} - p^I + s^I) f_{W^I, \tilde{V}}(W^I, \tilde{V}) dW^I d\tilde{V} \tag{6}$$

where $\nu_0 - p_0 - \kappa_0$ accounts for the utility of choosing the inefficient product, the first integral accounts for the incremental utility of choosing the efficient product when uninformed, and the second integral accounts for the utility of acquiring information. Taking the derivative of this expression with respect to the subsidy s^I gives first order condition:[30]

$$\mu'(s^I) = (s^I - A(s^I)) D'_B(s^I) \tag{7}$$

where $A(s^I) = E_{\tilde{V}|W^I}\left(B | \tilde{V} - B = p^I - s^I\right)$ is the average marginal bias in information acquisition at subsidy s^I, and $B = V - W^I - e$. Equation

[28] So, for example, $F_{W^I}(p^I)$ is the probability that the WTP for information W^I is less than the price of information p^I

[29] We assume the subsidy is revenue neutral and the government can levy non-distortionary lump sums.

[30] See Supplemental C.2 for derivation

7 is the standard first order condition in behavioral public finance that the subsidy should equal the average marginal bias, applied to the market for information. The information subsidy improves welfare if the average gain from improved product choices outweighs the average additional effort costs from the consumers whose information choice is marginal to the subsidy. The optimal subsidy balances this trade off at the margin such that the payment is equal to consumers' average marginal bias $s^{I*} = A(-s^{I*})$.

We compare this information subsidy to a more conventional product subsidy for encouraging take up of more efficient products. Analogous to the information subsidy, the optimal product subsidy equates the average marginal bias in product choice with the payment (Allcott and Taubinsky, 2015). The product subsidy improves welfare if the utility gain experienced by marginal consumers who truly value LEDs more than their marginal cost but are biased, outweighs any loss from distorting the product choice of other marginal consumers (the classic Harberger distortion (Harberger, 1964)). Whether a product or information subsidy is preferred depends on the benefit from improved product choices for each optimal instrument versus any losses from either distorted product choices or mental effort and therefore will be context specific.[31]

Our experiment is designed to recover the sufficient statistics necessary to calculate the optimal information and product subsidies and compare their welfare effects in the context of appliance purchasing behavior. We begin with the optimal information subsidy. As with similar studies, we initially assume that the information provision is a "pure nudge" so that consumers' informed demand revealed in our experiment is the "true" unbiased demand (e.g., Allcott and Taubinsky, 2015). We perform robustness

[31]In Appendix C.3 for the formal comparison between the two subsidies.

tests of our welfare ranking of the two policies when we relax this assumption.

To estimate the average marginal bias in information acquisition we need an estimate of individual effort costs. As previously discussed, consumers' revealed WTA for the effort task in the experiment (W^e) appears to be a good proxy for true effort cost (e). However, consumers have only a 1 in 10 chance of being a prize winner and receiving light bulbs but if they view information their effort costs are sunk. To rescale for the differences in probability between the effort and the prize, we assumed true effort costs $e = f(W^e)$, where $\frac{\partial W^I}{\partial f(W^e)} = -1$, i.e., an additional dollar of effort cost reduces WTP for information by exactly \$1. Assuming a linear function for f we can then compute effort costs e by estimating a linear regression of W^I on W^e and computing $\hat{e} = -\widehat{\beta}_1 \times W^e$.[32]

To compute the optimal subsidy we estimate the average marginal bias at each discrete level of W^I offered in the experiment as follows:

$$B = \sum_{n=1}^{N} \tau_n \mathbf{1}\left[W^I = w_n\right] + \rho \tag{8}$$

The parameters τ_n measure the average marginal bias for those who were willing to pay w_n for information and ρ is an idiosyncratic error term.

Figure 5 plots the estimated average marginal bias in information provision and 95% confidence intervals according to Equation (8) with effort costs estimated as above. At most levels of WTP for information the average marginal bias is positive (i.e., consumers underweight the value of information). The exception are the consumers stating a WTP above \$10

[32]This approach also rescales effort costs to account for the fact that the effort task is at the end of the experiment when effort may be more costly to consumers.

for information. This group of consumers overweight the value of information.

We compute the optimal information subsidy for those marginal to subsidies in our experiment by incrementing the subsidy according to the discrete levels of the information prices offered in the experiment, from an assumed initial price of zero–as if the information were already free.[33] Likewise, we assume that the baseline relative price of the bulb packages is zero (i.e., the packages have the same market price). If the baseline price of information is zero, a proportion of consumers with WTP for information ≥ 0 will be informed in the absence of a subsidy. With these assumptions, in the absence of subsidies, 86% of consumers are informed, and 88% of consumers purchase the LED.

Panel A of Table 2 reports the results of computing the change in welfare from incremental increases in the information subsidy from an initial price of zero. Column (1) reports the subsidy level corresponding to what was offered in the experiment. Column (2) reports the average marginal bias at each subsidy level. Column (3) reports the incremental increase in the proportion of consumers that become informed from a particular subsidy level relative to the previous level, or in the case of the first row, relative to zero. Column (4) reports the change in welfare due to the incremental change in the subsidy (relative to the previous row or zero) and column (5) reports the cumulative welfare effect of a particular subsidy, summing across the incremental changes in Column (4).[34] The subsidy should be

[33] This is consistent with the assumption that the social marginal cost of information provision is zero, or close to it.

[34] The change in welfare from a change in the subsidy can be computed as the change in market share from the subsidy increase (Column (3) of Table 2) multiplied by the sum of the average marginal bias (Column (2) of Table 2) and the average WTP for marginal consumers (Allcott and Taubinsky, 2015).

increased so long as the welfare effects of the incremental change are positive. We find an information subsidy of approximately \$8.75 maximizes welfare among these consumers.[35]

We now consider the welfare gain from a standard product subsidy to purchase more efficient LED bulbs. The columns of Panel B Table 2 mimic those of Panel A, with the exception of column (3) which reports the change in the share of consumers opting for the LED.[36] The average marginal bias in column (2) is the difference between the biased demand for LED, due to lack of information, and the unbiased demand at each subsidized price. Assuming the information acts as a "pure nudge," our estimate of bias is the treatment effect of information on demand for the LED at increments of baseline WTP for the LED. To do this, we use the random receipt of information conditional on WTP for information from the BDM mechanism to first estimate a treatment effect at each level of WTP for information (Berry, Fischer and Guiteras, 2020). We then compute the average treatment effect at each level of baseline WTP for the LED based on the distribution of WTP for information at each increment.[37] To mimic the setting above where the price of information is zero, we assume that the treatment effect for consumers with positive WTP for information is zero. We find a product subsidy of \$1.875 increases welfare and welfare is further increased at a product subsidy of \$12.1875 however further increases in the subsidy are not welfare improving (the average marginal bias at a subsidy of \$26.25 is \$3.068). The optimal product subsidy is therefore approximately

[35]This is a lower bound of the optimal subsidy which depends on the average marginal bias of those who were not marginal to the largest subsidy offered in the experiment.

[36]Due to limited sample size, we group consumers with relative WTP for the LED between -5.625 and -18.75 and assign them WTP -12.1875.

[37]As noted, this avoids issues with the arbitrary weighting of treatment effects from a continuous instrument and binary treatment as in Heckman and Vytlacil (2005).

\$12. For an LED package cost of \$20 this represents a 60% subsidy.

To evaluate whether information subsidies dominate product subsidies we then compare the cumulative change in welfare at each optimal subsidy (i.e., column (6) of Table 2). The optimal information subsidy increases welfare by 0.93, while the optimal product subsidy increases welfare by 0.651. In our setting the information subsidy therefore dominates the product subsidy. The change in welfare from the information subsidy is 35% of the aggregate welfare cost of biased information acquisition. The change in welfare from the product subsidy is 10% of the aggregate welfare cost of biased product choice. This suggests that the non-traditional policy instrument we propose, information subsidies, can potentially target welfare improvements more effectively than conventional product subsidies.

4.3.2 Sensitivity Analysis

We next explore the sensitivity of our policy results to our assumption that our information provision is a "pure nudge." It is possible that the intervention does not eliminate all biases in light bulb product choice. For example, it could be that information corrects beliefs but does not eliminate inattention. We, therefore, consider a scenario where, rather than a "pure nudge," the information intervention is a "nudge in the right direction." That is, the nudge pushes consumers closer to their true valuation, but does not fully debias them. To explore what this does to the welfare ranking of our two policies we assume that policy is designed based on the mismeasured estimates of bias, while actual welfare depends on true experienced utility that is not observed by the policy maker. We consider three levels of this experienced utility: low, i.e., for each consumer, the true change in WTP for the LED due to information is 10% higher than observed in our experiment, medium, i.e., the true change is 30% higher than we observe, and high, i.e., the true change is 50% higher. Appendix Table B12 shows that at all three levels of mismeasurement, the information subsidy continues to dominate the product subsidy for the "nudge in the right direction" and welfare is higher than the baseline.

We explore the sensitivity of the welfare ranking of the two policies to the demand curves used. It could be that bias would be different for another population of consumers. For this exercise, we go back to assuming that information serves as a "pure nudge," but consider what would happen if the change in WTP recovered by an experiment similar to ours were 10%, 30%, or 50% higher. To recalculate optimal subsidies, we retain consumers' baseline WTP for the LED, and their WTP for information and adjust their endline WTP for the LED (and therefore the revealed value of information) according to the three scenarios described above. We then recompute the average marginal biases in product and information space and recalculate optimal product and information subsidies.

Appendix Table B13 summarizes the optimal subsidy, product and information choices, and associated welfare from this exercise. In each case, we find that the level of optimal product and information subsidies are unchanged from the base case in Table 2, but the welfare impacts of the subsidies are larger reflecting the larger bias that these subsidies address.[38] We also find that the information subsidy dominates the product subsidy in all cases. Further, the relative welfare gain from the information sub-

[38]The optimal information subsidy already informed all consumers at a lower level of bias hence increasing the bias cannot result in a higher subsidy. The optimal product subsidy does not increase in part because the price increments at which we can measure bias in the experiment are relatively lumpy, and the fact that the average marginal bias in product choice is non monotonic and declines past a subsidy level of $12.

sidy (difference in welfare between the information and product subsidy) is higher than the base case for all three scenarios.

5 Conclusion

In a world saturated with information, consumers are often poorly informed and consequently routinely make mistakes. In this paper, we outlined an experiment designed to test whether remaining uninformed is the result of frictions in information acquisition, or mental gaps in decision making. The test is based on comparing ex ante WTP for information with the revealed value of this information ex post.

We implemented this experiment in the context of appliance choice where mental gaps in information acquisition about operating costs may be an important driver of investment inefficiencies. We found both under and over weighting of the value of information on energy efficiency and durability of the goods, evidence that consumers do not optimally trade off the costs and benefits of acquiring and processing information.

We consider whether competing decision models are consistent with the observed behavior in our experiment. First, we show that consumers' behavior is inconsistent with the rational model's prediction that consumers' WTP for information is correlated with the degree of uncertainty and the strength of their initial preference. Then, using an individual-level measure of bias, we demonstrate behavior is also inconsistent with otherwise rational consumers acquiring information suboptimally due to biased beliefs. Rather, the observed behavior is consistent with mental gaps. In particular, behavior in line with models that incorporate confirmation bias and motivated reasoning may contribute to the valuation patterns we observe.

If consumers are not optimizing subject to frictions, paying decision makers to acquire information may improve welfare. We derive the optimal information subsidy in the presence of mental gaps and compute it using our experimental data. In our setting, it is optimal to pay consumers close to \$9 to acquire information. We then compare welfare under this subsidy with welfare under the optimal product subsidy of around \$12. We find that the information subsidy dominates the product subsidy in our context. More generally, we show that the choice between policy interventions depends on the benefit from improved product choices for each optimal instrument versus any losses from either distorted product choices (in the case of the product subsidy) or mental effort (in the case of the information subsidy).

Governments frequently intervene to make information more readily accessible to decision makers. For example, governments mandate that appliance manufacturers provide energy cost information to consumers in a standardized format. Governments also frequently intervene in product markets by providing direct subsidies. We demonstrate that a non-traditional policy instrument, information subsidies, can target welfare improvements more effectively than conventional product subsidies. In a world where decisions are taken and information is delivered online, paying for attention could therefore result in substantial welfare gains in many settings.

References

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating behaviorally motivated policy: experimental evidence from the lightbulb market." *American Economic Review*, 105(8): 2501–2538.

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. 2019. "Regressive sin taxes, with an application to the optimal soda tax." *Quarterly Journal of Economics*, 134(3): 1557–1626.

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

Andor, Mark A, Andreas Gerster, and Lorenz Götte. 2019. "How effective is the European Union energy label? Evidence from a real-stakes experiment." *Environmental Research Letters*, 14(4): 044001.

Baldwin, S, G Bindewald, A Brown, C Chen, K Cheung, C Clark, J Cresko, M Crozat, J Daniels, J Edmonds, et al. 2015. "Quadrennial technology review: An assessment of energy technologies and research Opportunities." Technical Report.

Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka. 2016. "Attention discrimination: Theory and field experiments with monitoring information acquisition." *American Economic Review*, 106(6): 1437–75.

Beattie, Graham, Iza Ding, and Andrea La Nauze. 2022. "Is there an energy efficiency gap in China? Evidence from an information experiment." *Journal of Environmental Economics and Management*, 115: 102713.

Bénabou, Roland, and Jean Tirole. 2002. "Self-confidence and personal motivation." *Quarterly Journal of Economics*, 117(3): 871–915.

Bernheim, B Douglas, and Antonio Rangel. 2009. "Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics." *Quarterly Journal of Economics*, 124(1): 51–104.

Berry, James, Greg Fischer, and Raymond Guiteras. 2020. "Eliciting and utilizing willingness to pay: Evidence from field trials in Northern Ghana." *Journal of Political Economy*, 128(4): 1436–1473.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to lose: Health plan choices from a menu with dominated option." *Quarterly Journal of Economics*, 132(3): 1319–1372.

Bronchetti, Erin T, Judd B. Kessler, Ellen B. Magenheim, Dmitry Taubinsky, and Eric Zwick. 2022. "Is Attention Produced Optimally? Theory and Evidence from Experiments with Bandwidth Enhancements." *Forthcoming, Econometrica.*

Bronnenberg, Bart J, Jean-Pierre Dubé, Matthew Gentzkow, and Jesse M Shapiro. 2015. "Do pharmacists buy Bayer? Informed shoppers and the brand premium." *Quarterly Journal of Economics*, 130(4): 1669–1726.

Caplin, Andrew, and Mark Dean. 2015. "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review*, 105(7): 2183–2203.

Caplin, Andrew, Mark Dean, and Daniel Martin. 2011. "Search and satisficing." *American Economic Review*, 101(7): 2899–2922.

Charness, Gary, and Chetan Dave. 2017. "Confirmation bias with motivated beliefs." *Games and Economic Behavior*, 104: 1–23.

Choi, James J, David Laibson, and Brigitte C Madrian. 2010. "Why does the law of one price fail? An experiment on index mutual funds." *Review of Financial Studies*, 23(4): 1405–1432.

Dean, Mark, and Nate Leigh Neligh. 2023. "Experimental tests of rational inattention." *Forthcoming, Journal of Political Economy.*

Delavande, Adeline, and Hans-Peter Kohler. 2009. "Subjective expectations in the context of HIV/AIDS in Malawi." *Demographic Research*, 20: 817.

Department of Housing & Urban Development, HUD. 2022. "HUD USPS ZIP CODE CROSSWALK FILES."

Election Data & Science Lab, MIT. 2022. "County Presidential Election Returns 2000-2020."

Epley, Nicholas, and Thomas Gilovich. 2016. "The mechanics of motivated reasoning." *Journal of Economic Perspectives*, 30(3): 133–40.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. "Global evidence on economic preferences." *Quarterly Journal of Economics*, 133(4): 1645–1692.

Falk, Armin, Anke Becker, Thomas J Dohmen, David Huffman, and Uwe Sunde. 2016. "The preference survey module: A validated instrument for measuring risk, time, and social preferences." *IZA Discussion Paper No. 9674.*

Farhi, Emmanuel, and Xavier Gabaix. 2020. "Optimal Taxation with Behavioral Agents." *American Economic Review*, 110(1): 298–336.

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2018. "Expectations with endogenous information acquisition: An experimental investigation." *Review of Economics and Statistics*, 1–54.

Gabaix, Xavier. 2014. "A sparsity-based model of bounded rationality." *Quarterly Journal of Economics*, 129(4): 1661–1710.

Gabaix, Xavier. 2019. "Behavioral inattention." In *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 2, 261–343. Elsevier.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly information acquisition: Experimental analysis of a boundedly rational model." *American Economic Review*, 96(4): 1043–1068.

Handel, Benjamin, and Joshua Schwartzstein. 2018. "Frictions or mental gaps: what's behind the information we (don't) use and when do we care?" *Journal of Economic Perspectives*, 32(1): 155–78.

Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review*, 103(7): 2643–82.

Handel, Benjamin R, and Jonathan T Kolstad. 2015. "Health insurance for "humans": Information frictions, plan choice, and consumer welfare." *American Economic Review*, 105(8): 2449–2500.

Handel, Benjamin R, Jonathan T Kolstad, and Johannes Spinnewijn. 2019. "Information frictions and adverse selection: Policy interventions in health insurance markets." *Review of Economics and Statistics*, 101(2): 326–340.

Harberger, Arnold C. 1964. "The measurement of waste." *American Economic Review*, 54(3): 58–76.

Heckman, James J, and Edward Vytlacil. 2005. "Structural equations, treatment effects, and econometric policy evaluation 1." *Econometrica*, 73(3): 669–738.

Ho, Daniel E, Zoe C Ashwood, and Cassandra Handan-Nader. 2019. "New Evidence on Information Disclosure through Restaurant Hygiene Grading: Reply." *American Economic Journal: Economic Policy*, 11(4): 404–428.

Hoopes, Jeffrey L, Daniel H Reck, and Joel Slemrod. 2015. "Taxpayer search for information: Implications for rational attention." *American Economic Journal: Economic Policy*, 7(3): 177–208.

Hortaçsu, Ali, and Chad Syverson. 2004. "Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds." *Quarterly Journal of Economics*, 119(2): 403–456.

Houde, Sébastien, and Erica Myers. 2019. "Heterogeneous (mis-) perceptions of energy costs: Implications for measurement and policy design." National Bureau of Economic Research.

Howard, Greg. 2022. "A Check for Rational Inattention." *working paper.*

Khaw, Mel Win, Luminita Stevens, and Michael Woodford. 2017. "Discrete adjustment to a changing environment: Experimental evidence." *Journal of Monetary Economics*, 91: 88–103.

Kunda, Ziva. 1990. "The case for motivated reasoning." *Psychological Bulletin*, 108(3): 480.

List, John, Matthias Rodemeier, Sutanuka Roy, Gregory Sun, et al. 2022. "Judging Nudging: Toward an Understanding of the Welfare Effects of Nudges Versus Taxes." The Field Experiments Website.

Loewenstein, George, Cass R Sunstein, and Russell Golman. 2014. "Disclosure: Psychology changes everything." *Annual Review of Economics*, 6(1): 391–419.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2021. "Rational inattention: A review."

Mullainathan, Sendhil, Joshua Schwartzstein, and William J Congdon. 2012. "A reduced-form approach to behavioral public finance." *Annu. Rev. Econ.*, 4(1): 511–540.

Rodemeier, Matthias, and Andreas Löschel. 2022. "Information Nudges, Subsidies, and Crowding Out of Attention: Field Evidence from Energy Efficiency Investments." *Working Paper.*

Sallee, James M. 2014. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics*, 57.

Schwartzstein, Joshua. 2014. "Selective attention and learning." *Journal of the European Economic Association*, 12(6): 1423–1452.

Serra-Garcia, Marta, and Nora Szech. 2022. "The (in) elasticity of moral ignorance." *Management Science*, 68(7): 4815–4834.

Shrestha, Maheshwor. 2020. "Get rich or die tryin': Perceived earnings, perceived mortality rates, and migration decisions of potential work migrants from Nepal." *World Bank Economic Review*, 34(1): 1–27.

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

Taubinsky, Dmitry, and Alex Rees-Jones. 2018. "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment." *Review of Economic Studies*, 85(4): 2462–2496.

Winston, Clifford. 2008. "The Efficacy of information policy: A review of Archon Fung, Mary Graham, and David Weil's full disclosure: The perils and promise of transparency." *Journal of Economic Literature*, 46(3): 704–17.

Woodford, Michael. 2012. "Inattentive valuation and reference-dependent choice."

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data.* MIT press.

Zimmermann, Florian. 2020. "The dynamics of motivated beliefs." *American Economic Review*, 110(2): 337–61.

6 Figures

Figure 1: Distributions of Value and Willingness to Pay (WTP) for Information

(a) Value of Information

(b) Willingness to Pay

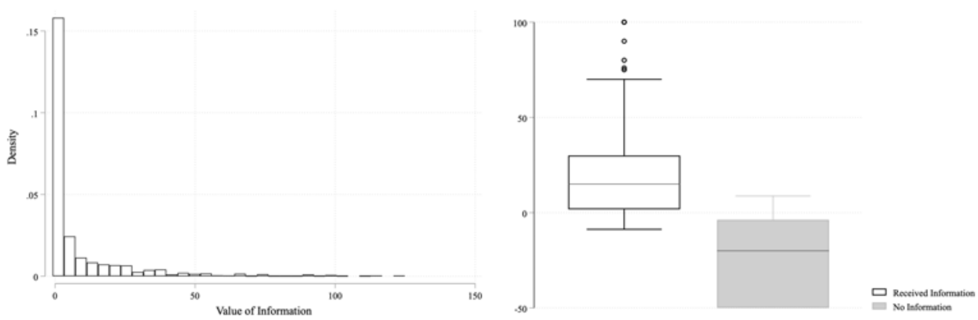

Notes: Panel (a) is a histogram of the value of information revealed by changes in product choice for those receiving information in the experiment. Panel (b) is a box plot of WTP for information on lifetime costs of bulbs for those receiving and not receiving information in the experiment.

Figure 2: Demand for LED at Baseline and Endline by Information Group

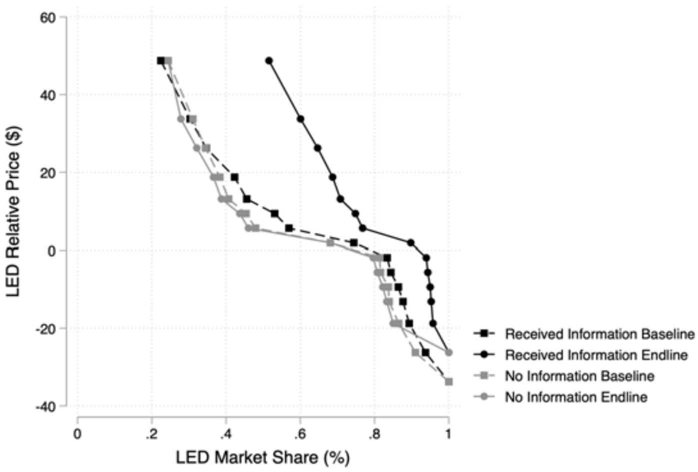

Notes: Demand for LED package measured as market share. Baseline is before information on lifetime costs is made available, Endline is after information is available. “Received Information” received information in the experiment. “No Information” did not receive information in the experiment.

Figure 3: Willingness to Pay (WTP) for Information and Value of Information

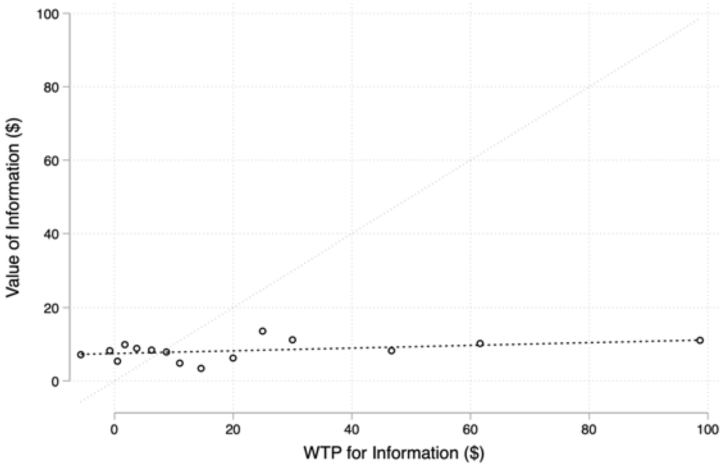

Notes: Figure shows a binned scatter plot of revealed value of information vs WTP for information.

Figure 4: Biased Beliefs, Political Motivations and Information Acquisition

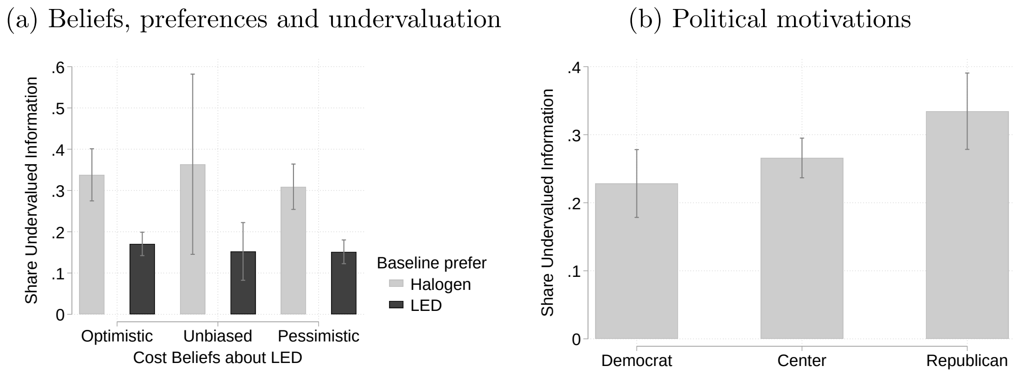

Notes: Panel (a) plots the probability a consumer undervalues information. Bias is evaluated for each consumer based on a test that the mean of the distribution of their beliefs is equal to -\$62.50 (the information provided in the experiment). Optimists are biased and have mean belief $< -\$62.50$, pessimists are biased and have mean belief $> -\$62.50$. Negative values indicate belief that the LED would cost less. Panel (b) plots the probability a consumer undervalues information in the following categories: "Democrat" (from a county in the top 20% of the Democratic vote share in the 2020 Presidential election), "Center" (from a county in the middle 21-79th percentiles of the Democratic vote share in the 2020 Presidential election), and "Republican" (from a county in the bottom 20% of the Democratic vote share in the 2020 Presidential election).

Figure 5: Average Marginal Bias in Information Acquisition

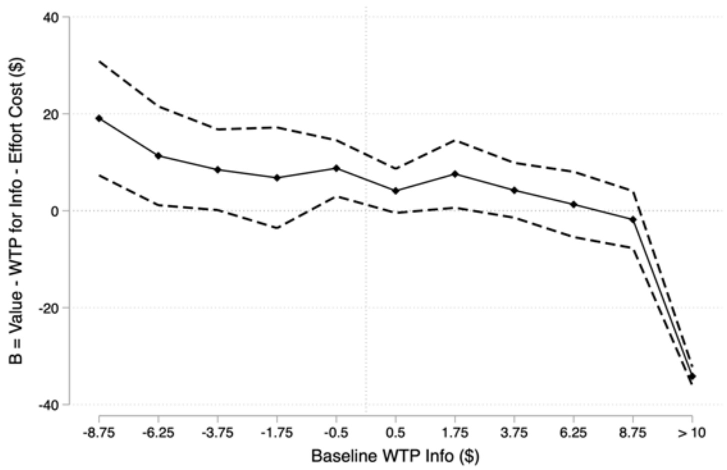

Notes: Figure plots point estimates and 95% confidence intervals for average marginal bias in information acquisition at each level of baseline WTP for information. Baseline is prior to information choice.

7 Tables

Table 1: Information Acquisition

	Dependent Variable = Value of Information			
Samples:	Info	Info	Marginal (a)	Marginal (a)+(b)
	(1)	(2)	(3)	(4)
(a) WTP Info	0.037**	0.039**	0.059	0.062
	(0.019)	(0.019)	(0.156)	(0.214)
(b) WTA Effort		-0.011	-0.007	0.310
		(0.016)	(0.022)	(0.429)
Constant	7.469***	7.629***	7.677***	7.707***
	(0.571)	(0.635)	(0.810)	(1.093)
Observations	1435	1435	644	440
Test optimal acquisition				
F-statistic (a) = 1	2697.41	2683.27	36.31	19.14

Notes: Table shows results of regressing the revealed value of information (Value of Information) on WTP for information (WTP Info) and WTA to undertake the experimental effort task (WTA Effort). Info Sample is the sample that viewed information. Marginal Samples are samples that were marginal to prices in the price lists (i.e., with WTP/WTA within the range of prices offered in the experiment to view information or undertake the effort task). Observations are weighted using probability weights to account for the probability of treatment. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.001$. F statistic (a) reports the test statistic for the null hypothesis that the coefficient on WTP Info (row (a)) is equal to one (necessary for optimal information acquisition).

Table 2: Welfare from Information and Product Subsidies

Panel A: Information Subsidy

	Subsidy	Bias	Δ Info(%)	Δ Welfare	Cum Δ Welfare
(1)	1.75	6.795	.079	.47	.47
(2)	3.75	8.458	.029	.167	.637
(3)	6.25	11.33	.02	.124	.761
(4)	8.75	19.069	.015	.169	.93

Panel B: Product Subsidy

	Subsidy	Bias	Δ LED(%)	Δ Welfare	Cum Δ Welfare
(1)	1.875	12.015	.044	.486	.486
(2)	12.1875	13.408	.026	.164	.651
(3)	26.25	3.068	.04	-.642	.009
(4)	33.75	18.021	.01	-.117	-.108

Notes: Table reports welfare analysis for information and product subsidies. In Panel A Bias is the Average Marginal Bias in information acquisition. In Panel B Bias is the Average Marginal Bias in product choice. In Panel A Δ Info (%) is the change in the share of consumers that are informed at the information subsidy in column (1). In both panels, Δ Welfare is the net welfare gain from incrementing the subsidy and Cum Δ Welfare is the running sum of Δ Welfare. In Panel B Δ LED (%) is the change in the share of consumers choosing the LED given the subsidy to the LED package in column (1). All calculations assume consumers with WTP for information above zero are informed without any subsidy and that the unsubsidized relative price of bulb packages is zero.