NBER WORKING PAPER SERIES

THE LIGHTBULB PARADOX:
EVIDENCE FROM TWO RANDOMIZED EXPERIMENTS

Hunt Allcott
Dmitry Taubinsky

Working Paper 19713
<http://www.nber.org/papers/w19713>

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

We are grateful to Raj Chetty, Lucas Davis, Stefano DellaVigna, Mushfiq Mubarak, Sendhil Mullainathan, Emmanuel Saez, Josh Schwartzstein, and other colleagues, as well as seminar audiences at the ASSA Annual Meeting, the Behavioral Economics Annual Meeting, Berkeley, the European Summer Symposium for Economic Theory, Harvard, the NBER Public Economics Meetings, Resources for the Future, Stanford, the Stanford Institute for Theoretical Economics, UCLA, the University of California Energy Institute, and the University of Chicago for constructive feedback. We thank our research assistants - Jeremiah Hair, Nina Yang, and Tiffany Yee - as well as management at the partner company, for their work on the in-store experiment. Thanks to Stefan Subias, Benjamin DiPaola, and Poom Nukulkij at GfK for their work on the TESS experiment. We are grateful to the National Science Foundation and the Sloan Foundation for financial support. Files containing code for replication, the in-store experiment protocol, and audio for the TESS experiment can be downloaded from Hunt Allcott's website. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Hunt Allcott and Dmitry Taubinsky. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Lightbulb Paradox: Evidence from Two Randomized Experiments
Hunt Allcott and Dmitry Taubinsky
NBER Working Paper No. 19713
December 2013, Revised August 2014
JEL No. D03,D12,H21,H31,L94,Q41,Q48

ABSTRACT

Imperfect information and inattention to energy costs are important potential justifications for energy efficiency standards and subsidies. We evaluate these policies in the lightbulb market using a theoretical model and two randomized experiments. We derive welfare effects as functions of reduced-form sufficient statistics capturing economic and psychological parameters, which we estimate using a novel within-subject information disclosure experiment. In the context of the model, the main results suggest that moderate subsidies for energy efficient lightbulbs may increase welfare, but informational and attentional biases alone do not justify a ban on incandescent lightbulbs.

Hunt Allcott
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
hunt.allcott@nyu.edu

Dmitry Taubinsky
Harvard University
taubinsk@fas.harvard.edu

A fundamental assumption in traditional policy analysis is that people’s choices identify their true preferences. In practice, however, many policies are at least partially predicated on the idea that consumers’ choices may not maximize their own welfare. Examples include consumer financial protection, taxes and bans on drugs, alcohol, cigarettes, and unhealthy foods, and subsidies and mandates for energy efficient products. To evaluate such policies, it is necessary to extend traditional public finance analysis to allow for the possibility of consumer mistakes and to design empirical strategies that identify the necessary economic and psychological parameters. This paper carries this out in the context of energy efficiency policy.

Energy efficiency subsidies and standards are important examples of policies partially motivated by addressing consumer bias. It has long been suggested that consumers may be imperfectly informed about or inattentive to energy costs when they buy energy-using durables such as cars, air conditioners, and lightbulbs.¹ This suggestion is supported by recent empirical evidence that people are inattentive to other ancillary product costs such as sales taxes (Chetty, Looney, and Kroft 2009), shipping and handling charges (Hossein and Morgan 2006), and out-of-pocket insurance costs (Abaluck and Gruber 2011). Because energy is costly – American households spent \$325 billion on gasoline and another \$245 billion on electricity, natural gas, and heating oil in 2011 (BLS 2013) – even small inefficiencies can aggregate to substantial losses in the absence of corrective policies.

This paper focuses on the lightbulb market, a particularly compelling case study of what Jaffe and Stavins (1994) call the “Energy Paradox”: the low adoption of energy efficient technologies despite apparently large cost savings. Compared to standard incandescents, compact fluorescent lightbulbs (CFLs) last much longer and use four times less electricity, so a 60-Watt equivalent CFL saves about \$5 per year on average.² In 2010, however, only 28 percent of residential sockets that could hold CFLs actually had them (DOE 2010), and in that year, using incandescents instead of CFLs cost US households a total of \$15 billion.³ Of course, CFLs and incandescents are far from perfect substitutes, and many consumers dislike CFLs for various reasons. Is the CFL’s low market share simply an expression of well-informed preferences, or are consumers unaware of or inattentive to how much money they could save?

This question matters for policy. Electric utilities in the U.S. spent \$252 million promoting CFLs in 2010, largely through subsidies (DOE 2010). Furthermore, the Energy Independence and Security Act of 2007 sets minimum efficiency standards that ban traditional incandescent lightbulbs, and Argentina, Australia, Brazil, Canada, China, Cuba, the European Union, Israel, Malaysia, Russia,

¹Allcott (2014) includes an extended series of quotes from policymakers and policy analyses that document this argument. See also Anderson and Claxton (1982), Blumstein *et al.* (1980), Jaffe and Stavins (1994), Sanstad and Howarth (1994), Gillingham and Palmer (2013), and many others.

²The \$5 estimate reflects \$4.50 in electricity savings, based on an average usage of 1000 hours per year (DOE 2010) and a national average electricity price is \$0.10 per kilowatt-hour (DOE 2014), plus \$0.50 in bulb replacement savings at typical prices. Throughout the paper, we assume that incandescents and CFLs last an average of 1000 and 8000 hours, respectively. (To receive the Energy Star rating, a CFL model must last a median of 8000 hours in official tests.)

³This \$15 billion estimate is equal to 5.8 billion residential sockets (DOE 2012), times the 80 percent of sockets that can accommodate CFLs (DOE 2010) minus the actual “socket share” of 28 percent (DOE 2010), times \$5 per socket per year.

and Switzerland have similar bans. Although externalities and other market failures also play a role, many advocates argue that the incandescent lightbulb ban acts in consumers’ best interests by preventing them from buying a product with large “shrouded” costs. A rancorous debate has evolved in a void of relevant evidence, despite a simple testable hypothesis: fully informed and attentive consumers would have higher willingness-to-pay for a CFL.

We use two randomized experiments to answer two questions. First, *how much does information provision affect demand for CFLs?* Second, *if powerful information provision is costly or infeasible, does a CFL subsidy or a ban on incandescents increase welfare as a second best solution to imperfect information and inattention?* The first is a positive question that can be answered by the treatment effects of information, with no additional structure or assumptions. To answer the second question, we use an optimal policy framework to derive “sufficient statistic” formulas for welfare effects and carry out an experiment specifically designed to estimate those sufficient statistics.

Our optimal policy framework follows Allcott, Mullainathan, and Taubinsky (2014), Baicker, Mullainathan, and Schwartzstein (2013), DellaVigna (2009), Mullainathan, Schwartzstein, and Congdon (2012), and others,⁴ and provides a simple extension of classical optimal tax formulas. Just as Diamond (1973) shows that the optimal externality tax equals the average marginal externality, the optimal internality tax (or subsidy) equals the average marginal bias. The net welfare effect of a ban on incandescents is the loss of perceived surplus for consumers who had purchased the incandescent plus any gain from internality reduction. Two functions are sufficient statistics to evaluate a subsidy or ban: the market demand curve and the average marginal bias at each point on that demand curve.

Our first experiment is an “artefactual field experiment” (Harrison and List 2004) using a nationally-representative online platform called Time-Sharing Experiments for the Social Sciences (TESS). Two specific features allow it to identify the two sufficient statistics. First, it is a within-subject design: consumers make baseline choices between CFLs and incandescents at different relative prices using a multiple price list format, then there is a randomly-assigned information treatment, and then consumers make endline choices using another multiple price list. We thus observe the baseline market demand curve and the conditional average treatment effect (CATE) on willingness-to-pay (WTP) at each point on that curve. Second, the information treatment was specifically designed to provide only hard information, ensure comprehension, and minimize demand effects and other potential confounds. It is thus not unreasonable to assume that our information treatment is what we call a *pure nudge*: it informs all previously-uninformed consumers and draws full attention to energy costs, with no other effects. Under this pure nudge assumption, the CATEs on WTP from our information treatments equal the average marginal bias from imperfect information and inattention. Although this assumption is also made in Chetty, Looney, and Kroft (2009) and other work, it is perhaps the greatest weakness of this approach, and we view it only as a useful approximation. We evaluate it throughout the paper and provide robustness checks under

⁴Also closely related is Spinnewijn (2014), as well as earlier work by O’Donoghue and Rabin (2006) and Gruber and Koszegi (2004) on optimal sin taxes with present-biased preferences.

plausible alternative assumptions.

In the TESS experiment, information increases the CFL’s market share at market prices by about 12 percentage points. The treatment effects on willingness-to-pay for a 60-Watt equivalent CFL differ across points on the baseline demand curve, with an average treatment effect of \$2.30. Under the pure nudge assumption that the effects of information measure consumer bias, our optimal policy framework suggests that the optimal CFL subsidy is approximately \$3. This is slightly larger than typical CFL subsidies offered by many electric utilities in the U.S. However, a large group of consumers purchase incandescents at baseline and are still willing to pay substantially more for incandescents after the informational intervention. Banning incandescents imposes welfare losses on this population that outweigh the gains to uninformed or inattentive consumers. This implies that in our model, imperfect information and inattention do not justify a ban on traditional incandescents unless there are additional market failures. This qualitative conclusion holds in most, although not all, of our welfare analyses under alternative assumptions.

Our second experiment is a natural field experiment with a “2-by-2” design that randomly assigned subsidies and information provision across shoppers at a large home improvement retailer. It is a useful complement to the TESS experiment: while this in-store setting imposes design constraints that limit the parameters that can be identified, the results provide evidence from a more realistic shopping environment. In this experiment, information did not statistically significantly affect CFL market share, and we bound the effect at around five percentage points with 90 percent confidence. We discuss factors that could explain why the in-store market share effects were smaller than the TESS effects, including that there was additional information available to the control group in stores or that the more complex in-store environment attenuated effects on the treatment group. While we show formally that market share effects are not informative about the average marginal bias, the smaller in-store market share effects do suggest smaller bias. This would strengthen the TESS result that imperfect information and inattention do not justify the incandescent lightbulb ban. The two experiments are also qualitatively consistent in showing that meaningful shares of consumers still purchase incandescents even after substantial effort to inform and draw attention to energy costs.

The paper makes three main contributions. First, while we document in Section 1 that the internality rationale for energy efficiency policy is of great interest, there is confusion among policymakers and disagreement among economists about how to formalize and test it. This paper provides a simple model and theoretically-motivated empirical tests useful for energy-using durables and many other situations where consumer bias motivates policy intervention. The analysis adds to a broader empirical literature on whether durable good buyers “undervalue” energy costs relative to purchase prices.⁵ In this literature, our approach is innovative in that we test for undervaluation

⁵This literature includes including Allcott (2013), Allcott and Wozny (2012), Busse, Knittel, and Zettelmeyer (2013), Dubin and McFadden (1984), Goldberg (1998), Hassett and Metcalf (1995), Hausman (1979), Saltee, West, and Fan (2009), and many others. There are also several theoretical and simulation analyses of energy taxes, energy efficiency standards, or subsidies for energy efficient goods when consumers misoptimize, including Allcott, Mullainathan, and Taubinsky (2014), Heutel (2011), Fischer, Harrington, and Parry (2007), and Parry Evans, and

using randomized experiments instead of comparing demand responses to price vs. energy costs in market data. Our experiments provide two initial data points on whether information or inattention might justify incandescent lightbulb bans, which have become an important feature of energy policy in the U.S. and other countries.

Second, our two experiments are a proof of concept for how large randomized experiments can be used to measure the effects of energy use information on durable goods purchases.⁶ This dearth of evidence is especially remarkable given the large literature on information disclosure in other domains; see Dranove and Jin (2010). While there has been extensive work on other aspects of information and energy demand, including Jessoe and Rapson (2014) and many others, only experiments like ours that provide durable good energy cost information are directly relevant to the important policy debates around multi-billion dollar subsidies, standards, and information disclosure for energy-using durables.

Third, the average marginal bias is a key sufficient statistic not just for energy policy, but also for a broader set of questions in behavioral public finance. Our TESS experiment is the first, to our knowledge, to attempt to directly measure this object. Existing empirical analyses, including the influential work of Chetty, Looney, and Kroft (2009), focus on a parameter that we call the *equivalent price metric* (EPM), which equals the average marginal bias only under special homogeneity and linearity assumptions. The TESS experiment allows us to directly estimate the CATE on WTP and compare it to the EPM. We find that the EPM is a poor approximation in our setting, and we show theoretically why this is likely to be the case more generally.

Section 1 gives more background on lightbulbs and related policies. Section 2 lays out our theoretical framework and defines the sufficient statistics that must be estimated. Section 3 presents the TESS experiment, and Section 4 carries out welfare evaluation. Section 5 presents the in-store experiment, and Section 6 concludes.

Oates (2010).

⁶There are some related studies that differ from our experiments on one or more dimensions. Kallbekken, Saelen, and Hermansen (2013) study energy information disclosure at six retail stores in Norway, comparing purchases to a non-randomly selected control group. Anderson and Claxton (1982) study energy information labels with 12 stores assigned to treatment groups and six to control. There are a number of studies that randomly assign information disclosure across individual experimental subjects and study effects on stated preferences in hypothetical choices, including Newell and Siikamaki (2013) and Ward, Clark, Jensen, Yen, and Russell (2011). Deutsch (2010a, 2010b) studies information disclosure to online shoppers, measuring what products they click on and what products they put in online shopping carts, but he does not observe actual purchases. Houde (2012) uses quasi-experimental variation with a structural demand model to estimate how the Energy Star label affects consumer welfare, while Herberich, List, and Price (2011) study how prices and social norm information affect CFL purchases.

1 Background: Reasons for Subsidies and Standards

Why subsidize CFLs or ban⁷ traditional incandescents? One potential reason for subsidies and standards is unpriced externalities. But while the lack of a carbon price artificially depresses electricity prices, two other distortions that imply that electricity prices could actually be *above* social marginal cost. First, retailers typically include much of fixed distribution costs in marginal prices, as Borenstein and Davis (2012) and Davis and Muehlegger (2010) show for natural gas. Second, most residential customers are charged time-invariant prices instead of real-time market prices, which are lower at night and higher during the day. Thus, if lightbulbs are relatively more likely to be used at night, they use underpriced electricity. This suggests that if the primary distortion is mispriced residential electricity, it could actually be optimal to *subsidize* incandescents.⁸

Alternatively, subsidies for new or emerging products might help correct for uninternalized spillovers from research and development or consumer learning. However, the CFL is an established technology, and the vast majority of consumers already have experience with it: 70 percent of consumers report having at least one CFL in their home, compared to 80 percent who report having at least one incandescent (Sylvania 2012).

Asymmetric information in real estate markets could also justify subsidies and standards. For example, prospective renters cannot costlessly observe energy efficiency, which reduces the incentive of landlords to invest in energy efficient capital stock (Davis 2012, Gillingham, Harding, and Rapson 2012). Similarly, renters or owners who expect to move before the end of the investment life have reduced incentive to invest. Davis (2012) estimates that renters in the U.S. are five percent less likely to use CFLs, but this would explain only a small portion of the CFL’s smaller market share given that only a quarter of US households are renters.

A final set of inefficiencies are “internalities,” or choices that don’t maximize the consumer’s own welfare. Informational and attentional internalities play an important role in the policy debate. For example, the Regulatory Impact Statement for Australia’s ban on energy inefficient lightbulbs (DEWHA 2008, page vii) argues:

[Incandescent lightbulbs] continue to sell remarkably well because, if their energy costs are ignored, they appear cheap ... There are significant information failures and split

⁷The U.S. lighting efficiency standards do not ban all incandescents. Instead, they set a maximum energy use per unit of light output. Along with CFLs, light-emitting diodes (LEDs) and high-efficiency halogen bulbs also comply. We focus on the choice between CFLs and incandescents because these are *by far* the most important current technologies. In 2012, about 1.5 billion incandescents and 300 million CFLs were purchased, compared to only 23 million LEDs (Energy Star 2013). While our quantitative welfare calculations would certainly change in the future if LEDs become a relevant part of the choice set, the basic questions about imperfect information and inattention from CFLs might also apply to LEDs, given that LEDs also have high purchase prices, long lifetimes, and large energy cost savings relative to both incandescents and CFLs.

⁸California is a particularly stark example. Regulations encouraging low-carbon electricity generation mean that the carbon content of electricity consumed there is extremely low relative to other states, so the downward distortion to electricity prices from the lack of a carbon tax is particularly small. Meanwhile, residential electricity tariffs with sharply increasing block prices distort marginal prices upward. Despite the fact that these two forces significantly weaken or reverse the argument that underpriced electricity justifies energy efficiency policies, California implemented the federal lighting efficiency standards early.

*incentive problems in the market for energy efficient lamps. Energy bills are aggregated and periodic and therefore do not provide immediate feedback on the effectiveness of individual energy saving investments. Consumers must therefore gather information and perform a reasonably sophisticated calculation to compare the life-cycle costs of [incandescents] and CFLs. But many lack the skills. For others, the amounts saved are too small to justify the effort ...*⁹

An overview article by Gayer (2011) summarizes the argument. “Private net benefits represent the bulk of the benefits of the energy-efficiency standards,” according to the official cost-benefit analyses (CBAs). “Energy-efficiency regulations and fuel economy regulations are therefore justified by such CBAs only by presuming that consumers are unable to make market decisions that yield personal savings, that the regulator is able to identify these consumer mistakes, and that the regulator should correct economic harm that people do to themselves.”

In the absence of our results, this policy argument could be quite plausible, as empirical estimates from other contexts suggest large attentional biases. A CFL saves \$36 (undiscounted) over its expected life relative to an incandescent. This dwarfs the typical relative price difference, suggesting that a consumer who is inattentive to these savings would be much more likely to buy a CFL. If 20 percent of consumers don’t think about energy costs, which is a seemingly conservative estimate relative to estimates for other energy-using durables, sales taxes, and health insurance plans¹⁰, the average bias would be around \$7 and our informational interventions could have massive impacts on demand.

In summary, while there are other market failures that could justify lightbulb subsidies and standards, we focus on imperfect information and inattention because results from other literatures suggested that these two distortions could be large, while other market failures appear to be less relevant in this context.

⁹Analogously, the U.S. government’s economic impact analysis (DOE 2009) and advocates such as the NRDC (2011) argue for the ban by showing that consumers will enjoy billions of dollars in annual cost savings. Notice that these private cost savings could only reflect welfare gains in the presence of imperfect information, inattention, or some other consumer market failure. Opponents also focus on the behavioral argument, suggesting that the ban is “over-reaching government intrusion into our lives” (Formisano 2008). U.S. Senator Rand Paul says that he supports energy conservation but objects to the idea that Department of Energy regulators “know what’s best for me” (ABC 2011).

¹⁰Forty percent of Americans report that they “did not think about fuel costs at all” when buying their most recent vehicle (Allcott 2011). In their two empirical studies, Chetty, Looney, and Kroft (2009) estimate that consumers are only 35 percent and 6 percent as attentive to sales taxes as they are to product prices. Abaluck and Gruber (2011) find that consumers are five times more responsive to insurance plan premiums than to out-of-pocket costs.

2 Theoretical Framework

2.1 Consumer Choice and Optimal Policy

2.1.1 Consumer Choice

We consider consumers that make one of two choices, labeled E and I . In our empirical application, E represents the purchase of an energy efficient good (the CFL), while I is an energy inefficient good (the incandescent). We let p_j denote the price of good $j \in \{E, I\}$ and let $p = p_E - p_I$ denote the relative price of E .

We define v_j as the consumer's true utility from consuming product j and let $v = v_E - v_I$ denote the relative true utility from E . In our empirical application, v can depend on any and all differences between CFLs and incandescents, such as electricity costs, lifetimes, mercury content, brightness, and "warm glow" utility from reduced environmental impact.

A consumer's utility from purchasing product j at price p_j is $v_j + (Z - p_j)$, where Z is the consumer's budget and $Z - p_j$ is utility from consuming the numeraire good. A fully optimizing consumer thus chooses E if and only if $v > p$. A misoptimizing consumer chooses E if and only if $v - b > p$, where b is a bias that affects choice but not true utility.

We let F denote the cumulative density function (CDF) of v , let $G(b|v, p)$ denote the CDF of b conditional on a true valuation v and at market price p , and let H denote the CDF of perceived valuations $\hat{v} = v - b$. We let $D_B(p) = 1 - H(p)$ and $D_N(p) = 1 - F(p)$, respectively, denote the demand curves corresponding to consumers' actual choices and to the choices consumers would make if they were unbiased. We assume that D_B and D_N are both smooth and downward sloping.

Our utility function is quasilinear, which is reasonable for purchases such as lightbulbs where p is small relative to Z . This simplifies the results, although the analysis can easily be generalized. We also assume that there are no externalities, although we will show in Section 2.2 that the welfare formulas are easily generalized to incorporate externalities.

2.1.2 Optimal Policy

The policymaker seeks to maximize social welfare and can set two policies: a subsidy of amount s for good E and a ban on either choice. We assume that the policymaker maintains a balanced budget through lump-sum recycling (taxes or transfers), and we let $Z(s)$ denote consumers' after-tax income when the policymaker sets a subsidy s . Under the lump-sum recycling and quasilinear utility assumptions, the subsidy does not distort other consumption and is thus purely corrective. Because of lump-sum recycling and no outside option (consumers choose either E or I), a subsidy for E is equivalent to a tax on I , and a ban on one choice is equivalent to a mandate for the other.

Goods E and I are produced in a competitive economy at constant marginal costs c_j , with relative cost $c = c_E - c_I$. Good E 's relative price after subsidy s is $p = c - s$.

Generalizing the classic analysis of Harberger (1964), we now derive a simple formula for the welfare impact of a subsidy. $W(s) = Z(s) - p_I + \int_{v-b \geq p} (v - p) dF dG$ denotes social welfare as a

function of the post-subsidy price $p = c - s$.

Proposition 1

$$W'(s) = (s - B(p))D'_B(p) \quad (1)$$

and

$$W(s - \Delta s) - W(s) \approx \overbrace{s\Delta s D'_B(p) + \frac{(\Delta s)^2}{2} D'_B(p)}^{\text{"Harberger distortion"}} + \overbrace{\underbrace{-\Delta s D'_B(p)}_{\text{Change in demand}} \underbrace{(E_H[B(x)|p - \Delta s \leq x \leq p])}_{\text{Average marginal bias}}}_{\text{"Externality reduction"}} \quad (2)$$

where $B(p) = E_G(b|v - b = p)$ is the average marginal bias at price $p = c - s$.¹¹

Equation (2) follows from Equation (1) by considering non-marginal changes in the subsidy. Both equations show that a subsidy has two effects. First, a subsidy distorts the market away from consumers' perceived private optimum. That is, it induces consumers to buy goods that they think they value at less than production cost. We call this the "Harberger distortion," and when the average marginal bias is zero, Equation (2) reduces to the standard Harberger formula. Second, when the average marginal bias is positive, the subsidy reduces externalities. That is, it induces consumers to buy products that they would value more if they did not misperceive true utility.

In our framework with no outside option, lump-sum taxation, and quasilinear utility, a ban on good I is equivalent to an infinite subsidy, so the welfare impact of a ban is simply $\int_0^\infty W'(s)ds = \int_0^\infty (s - B(c - s))D'_B(c - s)ds$. This can be approximated empirically by applying Equation (2) over increasing subsidy levels.

At the social optimum, Equation (1) must equal zero, which leads to a simple characterization of the optimal subsidy:

Corollary 1 *If s^* is an optimal subsidy, then $s^* = B(c - s^*)$.*

The Corollary is analogous to a result obtained by Allcott, Mullainathan, and Taubinsky (2014) in a richer framework in which consumers both choose a product (e.g., a car) and then choose how much to utilize it (e.g., miles driven). The Corollary extends the sin tax logic of O'Donoghue and Rabin (2006) to the case of general biases¹² and also extends Mullainathan, Schwartzstein, and Congdon (2012) to the case of arbitrarily heterogeneous biases.

In this analysis, there is a close analogy between externalities and externalities. Analogous to Proposition 1, the welfare impact of an externality tax can similarly be decomposed into (1) the negative impact of distorting the market away from the private optimum and (2) the positive

¹¹As usual, we subscript the expectation operator with the distribution over which the expectation is taken.

¹²When bias is non-negative for all consumers, but also positive for some (with at least some of those types on the margin), $B(p) > 0$ for all p , and thus the optimal subsidy is positive. This generalizes the O'Donoghue and Rabin (2006) result that the optimal sin tax must be positive when at least some consumers are present-biased.

impact of externality reduction. Analogous to Corollary 1, Diamond (1973) shows that the optimal externality tax equals the average marginal externality.

Proposition 1 and Corollary 1 show that the average marginal bias $B(p)$ and the market demand curve $D_B(p)$ are sufficient statistics for computing the welfare effects of a subsidy or ban. One powerful implication is that while different consumers might be biased for different reasons (for example, biased beliefs, inattention, or present bias), the underlying behavioral model of the bias does not matter conditional on $B(p)$. This is also important because some models have consumers that are either fully unbiased or fully biased with some probability, while other models might have all consumers with a partial bias, and it may be difficult to empirically distinguish between these models. Notice also that even if many consumers are biased, the standard Harberger (1964) formulas still hold exactly and the optimal corrective subsidy is still zero if the bias is not *systematic*, i.e. if the bias has mean zero at all values of $\hat{v}$.

2.2 Estimating the Average Marginal Bias

Chetty, Looney, and Kroft (2009), DellaVigna (2009), and Mullainathan, Schwartzstein, and Congdon (2012) categorize several approaches to estimating $B(p)$. One is to experimentally deliver what we call a *pure nudge*: a change in choice architecture whose only effect is to cause biased types to choose optimally. For example, active choice might eliminate default-related biases in retirement savings choices (Carroll *et al.* 2009), while a commitment device might address present bias. Assuming that a pure nudge is available, what strategies can identify $B(p)$ in our model?

One strategy is to identify the effect of the nudge at all levels of $\hat{v}$ using the following steps. First, elicit a baseline distribution of perceived valuations $H(\hat{v})$. Second, apply the pure nudge. Third, condition on $\hat{v}$ and estimate the average change in valuation induced by the nudge: $E_G[v - \hat{v} | \hat{v} = p] = E_G[b | \hat{v} = p] = B(p)$. The TESS experimental design derives from this strategy.

Such a strategy, however, requires within-subject identification that is difficult to implement as a natural field experiment. One potential alternative strategy might be to calibrate the change in price that has the same effect on market share as the nudge. Intuitively, if the effect of a nudge is twice the effect of a \$1 price change, then the nudge might be increasing valuations by approximately \$2, thus giving $B(p) = \$2$. Chetty, Looney, and Kroft (2009) implement this by estimating how labels with total tax-inclusive prices affect market shares and comparing this to the price elasticity of demand. We call this measure the *equivalent price metric*¹³:

$$EPM(p) = \frac{D_B(p) - D_N(p)}{D'_N(p)}. \quad (3)$$

The benefit of the EPM strategy is that it can be implemented with a much simpler “2-by-2”

¹³Like Mullainathan, Schwartzstein and Congdon (2012), Baicker, Mullainathan, and Schwartzstein (2013) and Chetty, Looney, and Kroft (2009), we divide by the slope of the unbiased demand curve, though in practice one could instead normalize by D'_B or the average of the slopes. We focus on this normalization because it approximates $B(p)$ under the broadest range of assumptions.

experimental design that varies nudges and prices, like our in-store experiment. Mullainathan, Schwartzstein, and Congdon (2012) show that the EPM approximates $B(p)$ under a restrictive homogeneous bias assumption which in our notation corresponds to when $G(\cdot|v, p)$ is degenerate for all v, p .

Unfortunately, $EPM(p) \neq B(p)$ in the general case with more realistic heterogeneity in bias. More broadly, any strategy such as the EPM that utilizes only the biased and unbiased demand curves D_B and D_N cannot identify $B(p)$, except under special conditions. For a simple example, suppose that $v \sim N(\mu, \sigma_v^2)$, $b \sim N(0, \sigma_b^2)$, and $Cov(v, b) = \sigma_{v,b}$. Using standard convolution and signal extraction formulas, $v - b \sim N(\mu, \sigma_v^2 + \sigma_b^2 - 2\sigma_{v,b})$ and $E_G(b|v - b = p) = \frac{-\sigma_b^2 + \sigma_{v,b}}{\sigma_v^2 + \sigma_b^2 - 2\sigma_{v,b}}(p - \mu)$. If we let $\sigma_{v,b} = \sigma_b^2/2$, then $v - b \sim N(\mu, \sigma_v^2)$, and thus $D_B(p) = D_N(p)$ for all p . However, $B(p) = E_G(b|v - b = p) = \frac{-\sigma_b^2/2}{\sigma_v^2}(p - \mu)$, which is positive for $p < \mu$ and negative for $p > \mu$. Thus, the average marginal bias can be arbitrarily large or small *even when the biased and unbiased demand curves are identical*, meaning that the nudge has no effect on market share at any price. Proposition 3 in Online Appendix D.2 extends this example and shows that even under strong restrictions including linear demand curves and tightly bounded support for the bias, it is still possible to have $B(p) \gg 0$ or $B(p) \ll 0$ when $EPM(p) = 0$.

To provide further intuition, Online Appendix D.2 includes an example with two bias types that illustrates the mechanisms causing the divergence between $EPM(p)$ and $B(p)$. A key statistic for understanding this divergence is each bias type’s “elasticity ratio”: the price elasticity in the biased state divided by the price elasticity in the unbiased state. We show that an approximate condition for $B(p) > EPM(p)$ is that the high type’s elasticity ratio is greater than the low type’s, and conversely for $B(p) < EPM(p)$.

Chetty, Looney, and Kroft (2009), DellaVigna (2009), and Mullainathan, Schwartzstein, and Congdon (2012) also discuss a second approach to estimating $B(p)$, which we call *comparing demand responses*. This exploits the fact that optimizing consumers should care only about a good’s total costs, so demand should be equally responsive to changes in purchase prices vs. changes in potentially less-salient add-on costs such as sales taxes, shipping and handling charges, or energy costs. A large literature uses this approach, including Abaluck and Gruber (2011), Allcott and Wozny (2012), Busse, Knittel, and Zettelmeyer (2013), the alcohol tax analysis in Chetty, Looney, and Kroft (2009), Hossein and Morgan (2006), and others. Using a general model that encompasses many settings including energy efficiency and tax salience, we show in Online Appendix D.3 that the comparing demand responses approach approximates $B(p)$ under an even stronger set of conditions than those that are required for the EPM. Thus, our theoretical and empirical results on how the EPM poorly approximates $B(p)$ also suggest that the comparing demand responses approach poorly approximates $B(p)$.

2.3 Biases Eliminated by Information Provision

In practice, a nudge addresses particular biases. For example, commitment contracts can address present bias, but they may not address inattention. In the context of lightbulbs, we are interested in identifying the effects of imperfect information and inattention. To do this, we use informational interventions that fully inform consumers about energy costs and bulb lifetimes and aggregate upfront and future costs into one total user cost.

The idea that information provision could identify bias is inspired by Chetty, Looney, and Kroft (2009), who identify inattention to sales taxes by informing consumers of tax-inclusive purchase prices in a supermarket. In justifying their approach, they write that “when tax-inclusive prices are posted, consumers presumably optimize relative to the tax-inclusive price.” Similarly, it seems reasonable to assume that consumers optimize relative to lightbulb lifetimes and energy costs after we provide them with information about these attributes. Providing information plausibly eliminates the following types of biases:

1. Biased beliefs, as tested by Allcott (2013), Attari *et al.* (2010), Bollinger, Leslie, and Sorensen (2011), and others. In our context, consumers may know that CFLs use less energy but misestimate the cost savings.
2. Exogenous inattention to energy as a “shrouded” add-on cost, related to Gabaix and Laibson (2006) or Heidhues, Koszegi, and Murooka (2013).
3. Costly information acquisition, as in Gabaix *et al.* (2006) and Sallee (2014). This includes many standard models of imperfect information in which the consumer incurs a cost to learn about energy efficiency and, in the absence of that information, assumes that different goods have the same energy efficiency.
4. “Noisy” and costly thinking models, as in Gabaix (2014), Sims (2004), Caplin and Dean (2014), and others. In these models, consumers might at first have only a noisy representation of the true value of energy efficiency, but thinking allows a more precise representation, subject to either a cognitive constraint or an explicit thinking cost.

Information interventions would not affect all biases that could affect lightbulb demand. For example, “bias toward concentration” (Koszegi and Szeidl 2013) could cause consumers to undervalue electricity costs because they occur in a stream of small future payments. Koszegi and Szeidl (2013) point out that re-framing the stream of payments as one net present value, as our interventions do, does not necessarily address this possible bias. Furthermore, consumers could be imperfectly informed about or inattentive to other attributes not discussed in our informational interventions.

Denoting $A(p)$ as the average marginal bias from other biases not addressed by information provision and $\phi(p)$ as the average marginal uninternalized externality, Equation (4) generalizes to

$$W'(s) = (s - B(p) - A(p) - \phi(p))D'_B(p). \quad (4)$$

The generalization of (2) would follow similarly. This equation illustrates that estimates of average marginal bias from imperfect information and inattention can be easily extended into a more comprehensive welfare analysis when combined with complementary estimates of $A(p)$ or $\phi(p)$. We illustrate this approach in Section 4.2.

This section has clarified that an experiment to identify the welfare effects of a subsidy or ban to address imperfect information and inattention must have two features. First, the treatment must plausibly approximate a “pure nudge”: it should provide clear information while minimizing demand effects and confounds. Second, the design must identify the sufficient statistics for welfare analysis: average marginal bias $B(p)$ and market demand curve $D_B(p)$.

3 TESS Experiment

3.1 Survey Platform and Population

We implemented the artefactual field experiment through Time-Sharing Experiments for the Social Sciences (TESS), which provides a nationwide sample of more than 50,000 consumers for computer-based experiments. Many economists have used this platform, including Allcott (2013), Fong and Luttmer (2009), Heiss, McFadden, and Winter (2007), Newell and Siikamaki (2013), and Rabin and Weizsacker (2009). One key feature of TESS is that the recruitment process generates a sample that is as close as practically possible to nationally representative on unobservable characteristics, which allows more credible generalization to the US population. Unrecruited volunteers are not allowed to opt in. Instead, potential TESS participants are randomly selected from the U.S. Postal Service Delivery Sequence File and recruited through an extensive series of mailings and telephone calls. About 10 percent of invitees actually become participants. Households without computers are given computers in order to complete the studies. We re-weight all TESS results to be nationally representative on observables.

Participants take an average of two studies per month, and no more than one per week. Of the qualified participants who began our survey, about 3/4 completed it, giving a final sample size of 1533. Per TESS rules, we could not force participants to answer all questions, although we successfully negotiated to require responses to the most important ones.

3.2 Experimental Design

3.2.1 Overview

Figure 1 gives a synopsis of the TESS experimental design. The study had four parts: baseline lightbulb choices, information provision screens, endline lightbulb choices, and a post-experiment survey. This design is both within-subject (we have both pre-information and post-information choices) and between-subject (consumers received different information screens).

Each consumer was randomly assigned to Treatment or Control, and within Treatment to

a matrix of four sub-treatments. These group assignments determined which two information screens the consumer would receive. As we discuss in more detail below, the “Positive” sub-treatment included information about the cost savings from CFLs, while the “Balanced” sub-treatment included information about cost savings *and* the the CFL’s negative attributes. The right column in the matrix of sub-treatments is the Endline-Only Treatment, in which consumers skipped the baseline choices and began directly with the information provision. Except when specified, we pool these four sub-treatments together and refer to them as the “Treatment” group; we show in Section 3.5 that effects of these four sub-treatments are not statistically distinguishable.

Choices were incentive compatible. Consumers were given a \$10 “shopping budget” that they could use to purchase packages of incandescents or CFLs at varying prices. Each consumer made 15 baseline choices and 15 endline choices via standard multiple price lists, and one of those 30 was randomly selected as the “official purchase.” TESS staff shipped consumers the lightbulb package they had chosen in that official purchase, charged the price of the package, and added the remainder of the \$10 to consumers’ TESS bonus accounts. Online Appendix A contains screen shots from the experiment.

3.2.2 Baseline and Endline Lightbulb Choices

Consumers chose between two lightbulb packages, one containing one Philips 60-Watt equivalent Compact Fluorescent Lightbulb, and the other containing four Philips 60-Watt incandescent lightbulbs. The two lightbulb packages were chosen to be as comparable as possible, except for the CFL vs. incandescent technology. While the choice screen had only pictures of the bulbs, consumers could click to a “Detailed Product Information” window, which included light output in Lumens, a quantitative measure of light color, energy use in Watts, and other information. About 19 percent of consumers opened this window. Both packages typically sell online for about \$4, so the market relative price is $p = 0$. We did not tell consumers these typical prices.

Half of consumers were randomly assigned to see the incandescent on the left, labeled as “Choice A,” while the other half were assigned to see the incandescent on the right, labeled as “Choice B.” The choice screen included 15 decisions in which the relative price of Choice A increased monotonically in decision number. For example, Decision Number 1 offered Choice A for free and Choice B for \$10, Decision Number 8 had equal prices of \$4, and Decision Number 15 offered Choice A for \$10 and Choice B for free. Consumers spent a median of three minutes on the baseline choice screen and one minute twenty seconds on the endline choice screen.

3.2.3 Information Provision

After the baseline lightbulb choices, each group received two information screens in random order. The screens were designed to closely parallel each other, to minimize the chance that idiosyncratic factors other than the information content could affect purchases. Each screen included about 10-15 lines of text, plus a graph to illustrate the key concept. The text was read verbatim on an audio

recording, which is available as part of the Online Supplementary Materials. At the bottom of the information screen, there was a “quiz” on a key fact.

Two design features help to ensure that consumers processed and understood the information. First, using multiple channels to convey information (text, graphical, and audio) means that people who learn in different ways had a higher chance of internalizing the information. Second, the quiz forced respondents to internalize the information if they had not done so initially.

The different groups received some combination of the following four screens:

1. **Treatment Information:** As described below, this screen compared electricity costs and replacement costs for CFLs and incandescents.
2. **Negative Information:** This screen was designed to present information about disposal and warm-up time, two ways in which CFLs might not be preferred to incandescents. It explained that *“because CFLs contain mercury, it is recommended that they be properly recycled instead of disposed of in regular household trash.”* It also explained that *“after the light switch is turned on, CFLs take longer to warm up than incandescents”* and graphed a typical CFL’s warm-up time.
3. **Number of Bulbs:** This screen presented information on the number of lightbulbs installed in residential, commercial, and industrial buildings in the U.S.
4. **Sales Trends:** This screen detailed trends in total US lightbulb sales between 2000 and 2009.

Control group consumers received the Number of Bulbs and Sales Trends screens. We designed these screens to have no impact on relative WTP, and neither screen mentioned energy costs or distinguished between CFLs and incandescents. The Positive Treatment group received the Treatment Information screen plus a randomly-selected one of the two Control screens. The Balanced Treatment group received the Treatment Information and Negative Information screens. We included the Balanced Treatment to both test whether consumers might be inattentive to or misinformed about product attributes other than energy costs and also help test for experimenter demand effects, as it is especially unlikely that this group would assume that the experimenter wanted them to purchase CFLs.

The Treatment Information screen began by explaining that CFLs both last longer and use less electricity compared to incandescents, and it translated these differences into dollar amounts using simple calculations at typical prices. The bottom line was:

Thus, for eight years of light, the total costs to purchase bulbs and electricity would be:

- *\$56 for incandescents: \$8 for the bulbs plus \$48 for electricity.*
- *\$16 for a CFL: \$4 for the bulbs plus \$12 for electricity.*

The graph was a simple bar graph illustrating these bullets.

The quiz question at the bottom of the screen was: *For eight years of light, how much larger are the total costs (for bulbs plus electricity) for 60-Watt incandescents as compared to their CFL equivalents?* The correct answer could be inferred from the information on the screen: \$56 for incandescents - \$16 for CFLs = \$40. Sixty-four percent of consumers correctly typed \$40. Those who did not were prompted to try again. After this point, 73 percent of consumers had typed \$40. The remaining consumers were told that the correct answer was \$40. After this point, 89 percent of consumers had typed \$40. The 11 percent failure rate is higher than we expected. However, results in Online Appendix Table A.7 show that the ATE on WTP is only four percent higher and statistically indistinguishable when excluding consumers who failed, suggesting that the failures do not meaningfully affect our results.

Consumers spent a median of two minutes and 12 seconds reading the Cost Info Screen and completing the quiz question. This substantial time, along with the quiz, show that the vast majority of Treatment group consumers engaged with and understood the information.

3.3 Data

The multiple price list allows us to observe choices at relative prices of $p \in \{-10, -8, -6, -4, -3, -2, -1, 0, 1, 2, 3, 4, 6, 8, 10\}$. Consumers' relative willingness-to-pay (WTP) for the CFL, denoted w , must lie between the highest relative price when they choose the incandescent and the lowest relative price when they choose the CFL. We assume that consumers that switch choices between any two relative prices have relative WTP w equal to the mean of those prices. For example, consumers who choose CFLs at relative price $p = 0$ (i.e. when both packages cost \$4) but choose incandescents when incandescents are one dollar cheaper are assumed to have $w = \$0.50$.¹⁴

Some consumers had censored WTPs: they preferred either Choice A or Choice B at all relative prices. These consumers were asked to report a hypothetical relative price at which they would prefer the other choice. Across all censored consumers, the median absolute value of self-reported relative WTP was \$15, and we impute a relative WTP of \$15 and -\$15 for top-coded and bottom-coded consumers, respectively. We will demonstrate the sensitivity of the results to this assumed value.¹⁵

Online Appendix B presents sample characteristics and shows that the treatment groups are balanced. That appendix also reports correlations between baseline WTP and observable charac-

¹⁴Eight percent and 4.3 percent of consumers in the baseline and endline choices, respectively, did not choose monotonically: they chose Choice A at a higher relative price than another decision at which they chose Choice B. These consumers were prompted with the following message: *The Decision Numbers below are organized such that Choice A costs more and more relative to Choice B as you read from top to bottom. Thus, most people will be more likely to purchase Choice A for decisions at the top of the list, and Choice B for decisions at the bottom of the list. Feel free to review your choices and make any changes.* After this prompt, 5.3 percent and 3.6 percent of consumers still chose non-monotonically in baseline and endline choices, and we code their WTP as missing.

¹⁵For treatment effects and welfare analysis, we technically should impute using mean censored WTP, not median. However, the distribution of self-reports is skewed, with a small number of consumers reporting very large values. Because these self-reports were not incentive compatible, we wish to be cautious about using them in the primary analysis.

teristics. Men, democrats, environmentalists, consumers who had previously reported taking steps to conserve energy, and those with higher discount factors have higher demand for CFLs. (The discount factors are the δ parameter in a β, δ model of present bias, as calibrated from hypothetical intertemporal tradeoffs in the post-experiment survey.) These correlations conform to our intuition and build further confidence that the differences in WTP are meaningful. Renters and more present-biased (lower β) consumers do not have lower WTP for CFLs conditional on other observables. This provides no support for the hypotheses that real estate market failures or present bias affect the lightbulb market, which reinforces the importance of our focus on imperfect information and inattention.¹⁶

3.4 Results

The results in this section begin to answer our first research question: how does information affect demand? Unlike the policy analysis in Section 4, these results do not require the assumption that the treatment is a pure nudge that only eliminates biases.

3.4.1 Quantity Effects, Demand Slopes, and the Equivalent Price Metric

Figure 2 presents a histogram of the within-subject changes in WTP between baseline and endline. About 90 percent of Control group consumers either have exactly the same WTP or change by \$2 or less. In Treatment, there is a mass to the right of the figure, with 36 percent of people increasing WTP by between \$1 and \$10. This figure illustrates that the Control information screens were successful in the sense that they did not affect average WTP. It also shows that the Treatment information both increased average WTP and had very heterogeneous effects.

Figure 3 presents endline demand curves. If some Treatment group consumers want to be internally consistent between baseline and endline, endline choices would be biased towards the baseline compared to a design without baseline choices (Falk and Zimmermann 2012). The Endline-Only Treatment demand curve lies directly on top of the Baseline & Endline Treatment curve, illustrating that internal consistency does not bias the information effects.¹⁷ Both Treatment curves are shifted out relative to Control. At market prices ($p = 0$), Treatment group CFL market share is about 77 percent, a 12 percentage point increase relative to control. The market share effect differs substantially at different relative price levels. For example, at relative price $p = -1$ (i.e. after a \$1 CFL subsidy), the effect is 7 percentage points. A key result is that at the market

¹⁶This may not be surprising. Present bias over *cash flows* might cause consumers to buy an incandescent to reduce current expenditures, but as Andreoni and Sprenger (2012) and many others have pointed out, agents in most models would be present biased over *consumption*, and most American consumers have enough liquidity that paying the incremental few dollars for a CFL does not immediately affect consumption. Present bias could induce people to procrastinate in buying and installing CFLs, but this would play no role in the TESS experiment because we forced consumers to make an active choice.

¹⁷Formal tests confirm that Endline-Only demand is not statistically different than Baseline & Endline demand: the share of consumers with endline WTP $w^1 > w^+$ does not differ statistically between these two demand curves at any level of w^+ .

price, a meaningful share of consumers still prefer incandescents even after being informed about the CFL’s cost advantage.

Figure 3 shows that demand is highly price-responsive near market prices. For example, between relative prices of 0 and -1, Treatment group demand has slope of 10 percent per dollar. This is not just an idiosyncratic feature of the TESS experimental setting¹⁸: we estimate in Section 5 that demand is equally or perhaps even more price responsive at market prices in the in-store experiment. We return to this issue in Section 4.2.

If demand is fairly inelastic because many consumers have strong preferences, then it is more likely for information to have small effects on market share. We thus use the equivalent price metric to benchmark the effect of information against the effect of prices. Defining $\Delta p = p_h - p_l$, we approximate the average EPM over price interval $p \in [p_l, p_h]$ as:

$$EPM[p_l, p_h] \approx \frac{(\Delta Q(p_l) + \Delta Q(p_h)) \cdot \frac{1}{2}}{(D_N(p_l) - D_N(p_h)) / \Delta p} \quad (5)$$

This equation approximates the average EPM over an interval by the average of the quantity effects at the endpoints divided by the slope between the endpoints. The average EPM just below the market price, denoted $EPM[-1, 0]$, can be calculated using numbers in the past few paragraphs: $EPM[-1, 0] \approx \frac{(0.07+0.12)/2}{0.1} \approx \0.94 . On this interval, information affects CFL market share about as much as a \$0.94 price reduction.

3.4.2 Conditional Average Treatment Effects

Figure 4 presents the conditional average treatment effects on WTP. Each diamond on the figure represents the average treatment effect for consumers with baseline relative WTP w^0 in an interval $w^0 \in [p_l, p_h]$, where p_l and p_h are adjacent points on the multiple price list. There is a thinner density of consumers with outlying high or low values of w^0 , so for precision, we group all $w^0 < -\$3$ and all $w^0 > \$8$. Most of these CATEs are in the range of \$2-\$4, except at the highest baseline WTP, where the CATE is statistically zero. This is simply due to top-coding: consumers who start at the top WTP in the multiple price list cannot increase their WTP further. Because these inframarginal consumers are unaffected by the subsidy and the ban, this will not affect the welfare calculations. After excluding consumers with top-coded and bottom-coded baseline WTP, the CATEs are statistically significantly increasing in w^0 . Given that the effects at different price levels are important for our policy analysis, this slope highlights the importance of an experimental

¹⁸Two features of the TESS experiment could in theory have caused highly elastic demand. First, if lightbulbs were perishable and consumers did not immediately need one, consumers would buy the cheapest package instead of revealing the WTP they would have if they did need one. In practice, lightbulbs are easily stored, and we reminded consumers of this fact in the introductory text. Second, if it were costless to resell the experimental purchase and replace it with a different purchase outside the experiment, consumers who know that the typical retail prices are approximately equal would always buy the cheaper package. In order to avoid making this salient, the experiment did not include information about the bulbs’ typical retail prices. In practice, it seems unlikely that consumers resold the packages that they received.

design like this one that identifies a CATE function instead of approximating it with a single average treatment effect.

How closely does the EPM approximate the CATE on WTP in these data? Figure 4 shows that the CATE on the interval $p \in [-1, 0]$ is \$2.11, which is more than twice the \$0.94 EPM. In Section B.2 of the Online Appendix, we compare the EPM and CATE at all nine price intervals where both can be calculated. Four of the nine differ with more than 90 percent confidence, and on average, the EPM differs from the CATE by 49 percent. Given that the CATE on WTP is what we need for policy analysis, this highlights the importance of a design like the TESS experiment that directly identifies it, instead of a more standard design that can only approximate it with the EPM. Since the CATE-EPM divergence is driven by treatment effect heterogeneity and such heterogeneity is pervasive across contexts, TESS-style designs may be more broadly important in contexts other than energy efficiency.

3.5 Average Treatment Effects, Robustness Checks, and Alternative Estimates

To more formally assess robustness, we now calculate average treatment effects on relative WTP and discuss alternative estimates. Let T_i be an indicator for whether the consumer is in the Treatment group, denote X_i as consumer i 's vector of individual characteristics, and denote μ_i as a vector of indicator variables for each level of baseline WTP. We estimate the average treatment effects of information provision on endline WTP w_i^1 using OLS with robust standard errors:

$$w_i^1 = \tau T_i + \gamma X_i + \mu_i + \varepsilon_i \quad (6)$$

Table 1 presents the results. Because baseline WTP, and thus μ_i , are not available for the Endline-Only group, all columns except column 5 exclude that group. Column 1 presents the unconditional difference in means. Column 2 adds the μ_i controls, while column 3 further adds X_i to give the exact specification from Equation (6). The sample size decreases in column 3 because at least one characteristic is missing for 15 consumers. In column 3, information increased relative WTP for the CFL by an average of \$2.30, and the estimates in columns 1 and 2 are economically and statistically identical.

Top-coding and bottom-coding of WTP mechanically influence the treatment effect. Consumers with baseline WTP equal to the maximum could not reveal a post-treatment increase in WTP, and any consumers with baseline WTP equal to the minimum could not reveal a decrease in WTP. Because the treatment tends to increase WTP, the former effect should dominate, and the average treatment effect should be understated. This connects to the result in Figure 4 that consumers with the highest baseline WTP have statistically zero treatment effect. Column 4 excludes consumers with top-coded or bottom-coded baseline WTP of \$15 or -\$15, and the estimated effect increases to \$3.10.¹⁹

¹⁹Relatedly, the assumed mean censored value of \$15 caps the increase in WTP that any consumer can reveal. Since a larger share of endline WTP is top-coded in treatment relative to control (29 percent vs. 16 percent), increasing

Column 5 adds the Endline-Only Treatment group, while excluding the μ indicators. Estimates show that the Endline-Only group’s WTP is not statistically different from the Control group’s endline WTP. This confirms the graphical result in Figure 3 that internal consistency does not bias the estimates.

With any experiment other than a natural field experiment, demand effects might arise: participants might change their actions to comply with, or perhaps defy, the perceived intent of the study. We address demand effects in three ways. First, because the Balanced treatment disclosed both positive and negative information about CFLs, these consumers should be less likely to perceive that the experimenters intended to persuade them to purchase CFLs. If demand effects typically cause consumers to comply with the study’s perceived intent, the Positive treatment would have larger effects on WTP than the Balanced treatment. Column 6 of Table 1 includes an indicator for the Positive Treatment group, showing that the effects do not differ statistically, and the point estimates are similar. Because the effects do not differ between the Balanced and Positive Treatment groups, we have combined these groups in other parts of the analysis.

Second, demand effects are less likely if participants cannot identify the intent of the study. The post-experiment survey asked consumers what they thought the intent of the study was. Results available in Online Appendix B, Table A.4 show that there is substantial dispersion in perceived intent within groups, which suggests that there is no one clear way in which demand effects might act.

Third, if demand effects are present, they should differentially affect people who are more able to detect the intent of the study and are more willing to change their choices given the experimenter’s intent. We proxy for this ability using the Self-Monitoring Scale (Snyder 1974), and we find no evidence that self-monitoring ability moderates the treatment effect. Details are available in Online Appendix B.3.

3.6 Effects on Beliefs

How much did the information treatment affect choices through increased attention vs. updated beliefs? The post-experiment survey elicited beliefs over how much less it costs to buy electricity for a CFL vs. incandescents over the CFL’s 8000-hour rated life, at national average electricity prices. Figure 5 presents the cumulative density functions (CDFs) of responses in Treatment and Control. The figure has three key features. First, beliefs are highly dispersed. Second, the information treatment substantially reduces this dispersion, and about 30 percent of Treatment group consumers have beliefs that are “correct” in the sense that they correspond to lifetime cost savings provided in the Information Treatment screen.²⁰ Results in Online Appendix Table A.7

this assumed value should increase the treatment effect. Regressions in Online Appendix Table A.7 show that when we alternatively assume mean censored values of \$12 (\$20) instead of \$15, the ATE changes to \$1.98 (\$2.83).

²⁰According to information on the Information Treatment screen, the correct answer to this question was \$36 (\$48 for the incandescent minus \$12 for the CFL). While some Treatment group consumers put \$36, many others put \$40, apparently misreading the question and also including the \$4 in bulb replacement cost savings. Thirty percent of

show that these consumers have statistically significantly larger ATE on WTP, and the ATE is 34 percent higher when estimated only off of these consumers. We return to this group in alternative welfare analyses in Section 4.

Third, the treatment increases perceived savings at the median of the distribution and all percentiles below the 65th. These data suggest that the information treatment may act at least partially through belief updating - unlike in Chetty, Looney, and Kroft (2009), whose belief surveys suggests that sales tax information acts through increasing salience. Because of the wide dispersion, however, the average treatment effect on beliefs is statistically indistinguishable from zero and very imprecisely estimated, making it impossible to precisely compare it to ATEs on WTP.

4 Welfare Analysis of Subsidies and Bans

The theoretically ideal way to address imperfect information and inattention would be a powerful and costless nationwide information disclosure technology. Subsidies and standards have been proposed as second-best policies with the idea that practically feasible information disclosure programs either do not fully remove bias or are too costly to scale. We now combine the model in Section 2 with the TESS results in Section 3 to evaluate these policies in the lightbulb market.

Our base scenario assumes that bias from imperfect information and inattention is the only market distortion. We view this as a reasonable simplification given the discussion in Section 1, and Equation (4) shows how the welfare analysis can be easily generalized with credible estimates of additional bias $A(p)$ and uninternalized externality $\phi(p)$. Separately, our base scenario also assumes that the conditional average treatment effect on WTP equals the average marginal bias from imperfect information and inattention²¹. Formally, this is:

Assumption 1: $\tau(p) = B(p)$.

Assumption 1 holds if the information treatment is a pure nudge, although it is slightly weaker: by identifying off of the CATE, it allows the information treatment to have additional idiosyncratic effects on WTP that are mean-zero at each price level.

4.1 Base Scenario Results

Table 2 evaluates the welfare impacts of incremental increases in the CFL subsidy using Equation (2) with the TESS data. We begin at the market price $p = 0$ and then increment the subsidy by amounts corresponding to the price differences in the multiple price list. Column 2 contains the average relative WTP for the CFL for consumers marginal to each subsidy increment. Column 3 presents the CATE on WTP $\tau(p)$ for consumers at that level of baseline WTP, which is just the point estimates from Figure 4. Under Assumption 1, this equals the average marginal bias.

consumers' beliefs were between \$36 and \$40, inclusive.

²¹In the language of Bernheim and Rangel (2009), we define Control group choices as *provisionally suspect* due to the possibility of imperfect information processing. If choices differ between Treatment and Control, we delete Control group choices from the welfare-relevant domain.

Column 4 contains the change in market share from the subsidy increase. Column 5 presents the welfare effect of the subsidy increase, using Equation (2). This column is simply column 4 multiplied by the sum of columns 2 and 3. When the average marginal bias is larger than the absolute value of average marginal WTP, internality reduction outweighs the Harberger distortion, and the subsidy increment increases welfare. Once the subsidy is so large that average marginal WTP is highly negative, however, further increases in the subsidy decrease welfare. Column 6 presents the cumulative welfare effect of changing the subsidy from zero to the amount listed in that row. Columns 3-6 are measured with sampling error, which we consider below.

Figure 6 illustrates the calculations in Table 2. The dashed curve is demand from the baseline choices, $D_B(p)$. The vertical black lines are guides to illustrate the changes in market share from each increment in the subsidy. The shaded rectangles above the x-axis reflect the internality reduction in Equation (2). Their height is the average marginal bias at each level of baseline WTP. The (triangle and) trapezoids below the x-axis are the familiar Harberger (triangle and) trapezoids generated by the distortion to perceived utility. When the area of the internality reduction rectangles exceeds the area of the Harberger trapezoids, the subsidy increment increases welfare.

Table 2 shows that of the discrete values that we can assess given the TESS multiple price list, the globally optimal subsidy is \$3. Correspondingly, Figure 6 shows that once the subsidy exceeds \$3, the area of the incremental Harberger trapezoid exceeds the area of the incremental internality reduction. \$3 is slightly larger than the CFL subsidies offered by many electric utilities over the last decade, which were typically in the range of \$1-2 per bulb.

Under the assumptions of our model, the welfare effects of an incandescent ban equal the effects of an infinite CFL subsidy. In Figure 6, this is the sum of all internality reduction rectangles above the x-axis minus all Harberger trapezoids below the x-axis. As shown in Table 2, this sum is negative: in our model, the ban reduces welfare by \$0.44 per package sold. While a ban is mechanically weakly worse than the optimal subsidy under the model's assumptions, the empirical magnitude is remarkable: in absolute value, the losses from a ban are 65 percent larger than the gains from the optimal subsidy. Thus, while a ban improves welfare for some biased consumers, this is far outweighed by the losses to consumers who strongly prefer incandescents even after being informed of the CFL's benefits.

In practice, we view Assumption 1 as only an approximation: it is possible that the information treatment might have effects other than removing informational and attentional bias. Furthermore, there may be additional market distortions. Additional distortions or mis-measurement of imperfect information and inattention that increased our base case $B(p)$ function by a multiplicative factor of 1.69 or more would cause the ban to increase welfare in our model. Alternatively, any total homogeneous distortion larger than \$3.60 per package would cause the ban to increase welfare. We further explore these issues below.

4.2 Alternative Assumptions

Table 3 presents results under alternative assumptions. The first row restates the base estimates in Table 2, adding a third column that reports the welfare effects of a ban as a share of incandescent lightbulb buyers’ total perceived consumer surplus from having the incandescent available. Put differently, column 3 contains the net welfare effects divided by the total area of the shaded Harberger trapezoids.

The top set of alternative results involve alternative assumptions for relative WTPs censored by the ends of the multiple price list. Top-coding and bottom-coding WTP has two opposing effects on the welfare calculation. First, the treatment causes many Treatment group consumers to be willing to pay the maximum for the CFL. Assuming a larger mean WTP for this top-coded group increases the treatment effect, implying a larger bias and thus larger welfare gains from corrective policies. Second, however, the welfare effects of a ban depend importantly on the lower tail of the WTP distribution: if some consumers very strongly prefer incandescents, banning them can cause large welfare losses. Recall that in the base case, we assume that the mean values of top-coded and bottom-coded WTPs are \$15 and -\$15, respectively. Rows 2 and 3 assume {\$12,-\$12} and {\$20,-\$20}, respectively, for top-coded and bottom coded WTPs. Row 4 assumes consumers’ self-reported hypothetical WTP, bounded at +/- 100.²²

The bottom set of results consider alternatives to Assumption 1. Rows 5 and 6 consider the possibility that not all consumers understood the information treatment, which could cause the CATEs to understate bias. Row 5 increases the base case $B(p)$ estimates by four percent to recognize that the 89 percent of consumers who passed the “quiz” on the Treatment Information screen have four percent larger ATEs than the Treatment group as a whole. Rows 6 analogously increases $B(p)$ by 34 percent, using the result that the 30 percent of consumers with “correct” CFL savings beliefs on the post-experiment survey had 34 percent larger ATE than the Treatment group as a whole.

Row 7 scales down the $B(p)$ function by seven percent, to reflect the estimate in Table 1 that the Balanced Treatment group had a seven percent smaller ATE than the Treatment group as a whole. The effects of the Balanced Treatment would be of primary interest if one thought that the Positive Treatment group were more affected by experimenter demand effects, or if one wanted to also incorporate potential imperfect information or inattention related to the disposal and warm-up attributes discussed on the Negative Information screen. As discussed in Section 3, the differences used to scale Rows 5 and 7 are not statistically significantly different from zero. Rows 8 and 9 replace the $B(p)$ function with the 10th and 90th percent confidence bounds of the CATEs, as graphed in Figure 4.

Equation (4) in Section 2 showed how the framework can be extended when there are additional

²²While it may seem unsatisfying to need to make these assumptions about censored WTPs, remember that the TESS experiment substantially improves over the standard approach to analyzing the removal of a product from the choice set, which is to assume a logit or otherwise parametric functional form for demand.

distortions other than imperfect information and inattention. In Section 3.4, we pointed out that demand curves are much more price-responsive around zero than they are away from zero. Visual inspection of the demand curves suggest that this “excess mass” of valuations in the Treatment group is primarily contained on the interval $w^1 \in [-2, 4]$. Such excess mass of WTP around zero would be unlikely if all consumers fully value a large lifetime cost savings from the CFL. To see this, use a model as in DellaVigna (2009) with $w^1 = \hat{e} + n$, where $\hat{e}$ is perceived total cost savings and n reflects preferences for other non-energy attributes. If $\hat{e}$ tends to be large and positive, then n must be symmetrically large and negative in a remarkably coincidental way in order to generate a mass of w^1 near zero.²³ Thus, it seems more likely that there is a mass of consumers for whom both $\hat{e}$ and n are close to zero.

One explanation for a large group of consumers with small $\hat{e}$ is that these consumers think that they might break or discard the lightbulbs before the end of their rated lifetimes. This would not be a distortion, as the true social value of lifetime cost savings is also small in this case. A second potential explanation is asymmetric information in rental markets. Additional (unreported) regressions show that renters are no more likely to have WTP in the interval $w^1 \in [-2, 4]$, which seems to rule this out. A third explanation is that this is a behavioral bias not addressed by information provision, such as the Koszegi and Szeidl (2013) bias toward concentration.²⁴ If this is true, how would this affect the welfare analysis?

Because the TESS experimental design does not directly identify average marginal bias functions for distortions not addressed by information provision, we provide an illustrative, back-of-the-envelope calculation using a three-step approach. First, we use an excess mass test inspired by Chetty *et al.* (2011) to show that about 38 percent of the Treatment group is excess mass on $w^1 \in [-2, 4]$ relative to a prediction based on the density on the rest of the demand curve.²⁵ This excess mass can also be visually approximated on Figure 3 by assessing the additional market share on the specific interval $w^1 \in [-2, 4]$ compared to a prediction based on the slope of demand outside that interval. Second, we calculate the share of consumers at each level of baseline WTP w^0 that

²³For a stark example, assume away any heterogeneity in $\hat{e}$ that could result from variation in electricity prices and discount rates. If $\hat{e} = \$40$ for all consumers and w is tightly distributed around zero, then n would need to be tightly distributed around $-\$40$. Then, the demand curves in Figure 3 would imply that if $\hat{e}$ were actually $\$35$ or $\$45$ instead of $\$40$, then the CFL market share would move substantially to 0.4 or 0.94, respectively.

²⁴We thank one of our referees for drawing our attention to the high elasticity around $p = 0$ and for suggesting that this may be the consequence of some additional bias.

²⁵To do this, we index endline WTP intervals by m , denote D_m as the sample-weighted share of Treatment group consumers with endline WTP in interval m , denote Δp_m as the width of interval m (either $\$1$ or $\$2$), denote ξ_m as an indicator for interval m , and run the following regression:

$$D_m = \omega_1 \Delta p_m + \omega_2 w^1 \Delta p_m + \sum_{m=[-2, -1]}^{[3, 4]} \xi_m + \epsilon_m$$

Assuming a quadratic approximation to the demand curve outside the interval $w^1 \in [-2, 4]$, the sum of the ξ_m coefficients identify the total excess mass on $w^1 \in [-2, 4]$. (A higher-order approximation is not merited given the limited number of data points, and following Chetty *et al.* (2011), the interval $w^1 \in [-2, 4]$ was chosen by visual inspection of the demand curve.) The total excess mass $\sum_{m=[-2, -1]}^{[3, 4]} \hat{\xi}_m$ represents about 38 percent of the Treatment group.

are part of the excess mass on $w^1 \in [-2, 4]$. Third, we compute the additional bias function $A(p)$ if consumers who are part of the excess mass on $w^1 \in [-2, 4]$ have average true utility $v = 7.66$, which is the Treatment group mean w^1 after re-weighting to eliminate the excess mass.

Row 10 of Table 3 presents results. The ban now increases welfare by 114 percent of incandescent lightbulb buyers’ total perceived consumer surplus from having the incandescent available. The optimal subsidy increases to \$8, although welfare is only slightly higher than at a subsidy of \$4. This is because, at least under our back-of-the-envelope assumptions, most of the consumers comprising the excess mass also have baseline WTP w^0 close to zero, and they are thus mostly inframarginal to larger subsidies.

In summary, the welfare losses from the incandescent lightbulb ban in most scenarios amount to 30-50 percent of incandescent buyers’ total perceived consumer surplus from having the incandescent available. In two scenarios, however, the ban is welfare-enhancing.

5 In-Store Experiment

5.1 Experimental Design

Would the effects of information provision be different in a more typical retail setting compared to the TESS platform? To answer this, we partnered with a large home improvement retailer to implement an in-store experiment. Between July and November 2011, three research assistants (RAs) worked in four large “big box” stores, one in Boston, two in New York, and one in Washington, D.C. The RAs approached customers in the stores’ “general purpose lighting” areas, which stock incandescents and CFLs that are substitutable for the same uses.²⁶ They told customers that they were from Harvard University and asked, “Are you interested in answering some quick research questions in exchange for a discount on any lighting you buy today?” Customers who consented were given a brief survey via iPad in which they were asked, among other questions, the most important factors in their lightbulb purchase decision, the wattage and number of bulbs they planned to buy, and the amount of time each day they expected these lightbulbs to be turned on each day. The survey did not mention electricity costs or discuss any differences between incandescents and CFLs.

In the taxonomy of Harrison and List (2004), this was a “natural field experiment.” Participants believed that they were answering a survey, but they did not know that they were in a randomized experiment or that their subsequent purchase behavior would be scrutinized. This experiment has complementary strengths and weaknesses to the TESS experiment: while we observe consumers naturally participating in a standard marketplace, we could not implement the multiple price lists and within-subject design that allow the TESS experiment to identify the average marginal bias. Instead, we randomize information and prices in a standard two-by-two design and focus on

²⁶This includes standard bulbs used for lamps and overhead room lights. Specialty bulbs like Christmas lights and other decorative bulbs, outdoor floodlights, and lights for vanity mirrors are sold in an adjacent aisle.

answering our positive research question about the effects of information on demand.

The iPad randomized customers into information Treatment and Control groups with equal probability. For the Treatment group, the iPad would display the annual energy costs for CFLs vs. incandescents, given the customer’s estimated daily usage, desired wattage, and desired number of bulbs. The treatment screen also displayed the energy costs and total user costs (energy plus bulbs) for CFLs vs. incandescents over the 8,000 hour rated life of a CFL. Online Appendix C presents the information treatment screen. The RAs would interpret and discuss the information with the customer, but they were trained to not advocate for a particular type of bulb and to avoid discussing any other issues unrelated to energy costs, such as mercury content or environmental externalities. The Control group did not receive this informational intervention, and the RAs did not discuss energy costs or compare CFLs and incandescents with these customers.

At the end of the survey and potential informational intervention, the RAs gave customers a coupon in appreciation for their time. The iPad randomized respondents into either the Standard Coupon group, which received a coupon for 10 percent off all lightbulbs purchased, or the Rebate Coupon group, which received the same 10 percent coupon plus a second coupon valid for 30 percent off all CFLs purchased. Thus, the Rebate Coupon group had an additional 20 percent discount on all CFLs. For a consumer buying a typical package of 60 Watt bulbs at a cost of \$3.16 per bulb, this maps to an average rebate of \$0.63 per bulb. The coupons had bar codes which were recorded in the retailer’s transaction data as the customers submitted them at the register, allowing us to match the iPad data to purchases. We do not observe the possible purchases of the 23 percent of consumers in the iPad whose coupon numbers do not appear in the transaction data; these consumers either purchased lightbulbs without submitting the coupon or did not purchase any lightbulbs.

After giving customers their coupons, the RAs would leave the immediate area in order to avoid any potential external pressure on customers’ decisions. The RAs would then record additional visually-observable information on the customer, including approximate age, gender, and ethnicity. The RAs also recorded this information for people who refused. Finally, the RA recorded the total duration of the interaction. The difference between Treatment and Control reflects the amount of time spent on the informational intervention. The difference in means (medians) is 3.17 (3.0) minutes, which suggests that Treatment group consumers did engage meaningfully with the information.

5.2 Data

Of the 1561 people who were approached, 459 refused, while 1102 began the iPad survey. Of these, 13 broke off after the first question, two broke off later, and 1087 were assigned to a treatment group and given a coupon. Column 1 of Table 4 presents descriptive statistics for the sample of customers who completed the survey and were given a coupon. Column 2 presents differences between the 474 people who refused or did not complete the survey and the 1087 who completed, using the

demographic characteristics recorded for those who refused. People whom the RAs thought were older, male, Asian, and Hispanic were more likely to refuse. Columns 3 and 4 present differences between the information Treatment and Control groups and between the Rebate Coupon and Standard Coupon groups. In one of the 18 t-tests, a characteristic is statistically different with 95 percent confidence: we have slightly fewer people coded as Asian in the information treatment group. F-tests fail to reject that the groups are balanced.

We restrict our regression sample to the set of consumers that purchase a “substitutable lightbulb,” by which we mean either a CFL or any incandescent or halogen that can be replaced with a CFL. The bottom panel of Table 4 shows that 77 percent of interview respondents purchased any lightbulb with a coupon, and 73 percent of survey respondents purchased a substitutable lightbulb. While information or rebates theoretically could affect whether or not customers purchase a substitutable lightbulb, t-tests show that in practice the percentages are not significantly different between the groups.

5.3 Results

We denote T_i and S_i as indicator variables for whether customer i is in the Treatment and Rebate Coupon groups, respectively. X_i is the vector of individual-level covariates. We estimate a linear probability model²⁷ with robust standard errors using the following equation:

$$1(\text{Purchase CFL})_i = \tau T_i + \eta S_i + \gamma X_i + \varepsilon_i \quad (7)$$

5.3.1 Quantity Effects, Demand Slopes, and the Equivalent Price Metric

Table 5 presents estimates of Equation (7). Column 1 excludes covariates X_i , while column 2 adds them. The estimates are statistically identical, and the point estimates are very similar. The rebate increased CFL market share by about ten percentage points. Column 3 shows that the interaction between information and rebates is statistically zero. Using column 3, the point estimate of Treatment group demand slope is $\frac{7.8+5.4}{\$0.63} \approx 21$ percentage points per dollar.

In column 2, the information treatment increased market share by 0.4 percentage points, which is statistically indistinguishable from zero. The standard errors rule out with 90 percent confidence that the information treatment increased (decreased) market share by more than 5.8 (5.0) percentage points. Using column 3, the point estimates of the information treatment effect are -2.2 percentage points with the Standard Coupon and $-2.2+5.4=3.2$ percentage points with the Rebate Coupon; both are statistically indistinguishable from zero.

Inelastic demand could cause this statistically small effect on market shares. For example, one might hypothesize that in-store demand would be less elastic because consumers might enter the

²⁷In typical cases like ours where the true probability model is not known, Angrist and Pischke (2009) advocate for using the LPM instead of an arbitrary non-linear model such as probit or logit, and we follow their recommendation. In any event, S and T are indicator variables, and the probit estimates are almost identical.

store having already decided what type of lightbulb to buy, perhaps as instructed by other family members. The fact that in-store demand responds strongly to prices rules out this hypothesis.

To more formally compare information vs. price effects, we again calculate the equivalent price metric. Inserting the above coefficients into Equation (5), we have the EPM between the Standard Coupon and Rebate Coupon price levels: $EPM[Rebate, Standard] \approx \frac{(0.032 + -0.022)/2}{0.21} \approx \0.02 . Using the Delta method, the 90 percent confidence interval is $[-\$0.24, \$0.28]$. In other words, the standard errors rule out with 90 percent confidence that information provision had more than the effect of a \$0.28 CFL subsidy.

5.4 Comparing and Generalizing from the Two Experiments

We can compare the three main parameters from the in-store experiment (information effects on market share, demand slope, and EPM) to their analogues from the TESS experiment. As Figure 3 illustrates, the effects of information on market share differ substantially by price level. We separately compare the effects near market prices ($p = 0$ in the TESS experiment and with the Standard Coupon in the in-store experiment) and at a small discount ($p = -1$ in TESS and with the \$0.63 average discount of the Rebate Coupon). Near market prices, the effects in the TESS and in-store experiments are statistically different with a p-value of 0.015. At a small discount, however, the 7 percentage point effect at $p = -1$ in TESS is statistically indistinguishable from the effect on Rebate Coupon recipients in the store (p-value=0.44). Point estimates suggest that in-store demand is even more price-responsive than in TESS, although the slopes are statistically indistinguishable (p-value=0.15). Finally, the TESS EPM on $p \in [-1, 0]$ is statistically larger than the in-store EPM on $p \in [Rebate, Standard]$, with a p-value of 0.012.

It is well understood that empirical results can differ across contexts. If the goal is to evaluate a nationwide policy, ideally one would estimate a nationwide parameter using many experiments with consumers and retailers across the country. When attempting to learn from a limited number of experiments, it is particularly useful if they differ on dimensions that could moderate effects out of sample. While we only have two experiments, this pair is relatively useful because they differ markedly on the three key dimensions: consumer populations, choice environments, and treatments.

First, consumer populations differ substantially: the TESS population is nationwide, while the in-store sample is drawn from four stores in three eastern states. There are some differences on observable characteristics, and there surely differences on unobservables as well. In combination with the substantial treatment effect heterogeneity suggested by Figure 2, these differences could generate different parameter estimates.²⁸ On this dimension, the TESS estimates are of greater interest because consumers are drawn from a wider geographic area and are weighted for national

²⁸Remarkably, the in-store data suggest that incandescent buyers do not buy more bulbs per trip than CFL buyers. Because incandescents have much shorter lives than CFLs, people who prefer incandescents will thus need to purchase bulbs more often and will appear with higher probability in the in-store sample than in the TESS sample. The TESS CATEs are smaller for consumers that have lower baseline WTP (and thus more strongly prefer incandescents), and this could partially explain the differences between the two experiments.

representativeness on observables.

Second, the choice environments also differ markedly. The TESS experiment has a deliberately simple and controlled choice environment with a small choice set and limited additional stimuli, while the in-store environment includes hundreds of different lightbulb packages and many other stimuli and purchasing needs competing for attention. These factors could make it more difficult for the in-store Treatment group consumers to internalize, recall, and apply the information when they actually choose a package. Furthermore, like most home improvement stores, the stores we worked in have displays in lightbulb aisles that provide information on different lightbulb technologies, including electricity use. If this existing information fully informed the Control group, incremental information could have no effect. If this is the case, the treatment effects are still the relevant parameters for policy analysis in the in-store environment: if existing information provision mechanisms are fully effective, then there is no remaining imperfect information and inattention to justify subsidies and standards.

The choice environments in both experiments are of interest: home improvement retailers are the most common retail channel through which households buy lightbulbs (DOE 2010), and our partner alone sells upwards of 50 million lightbulb packages each year, a non-trivial share of national sales. Notwithstanding, more than half of lightbulbs are sold at grocery stores, drugstores, and other retail channels that typically have less in-store energy cost information, and an increasing number of consumers buy online. The TESS Control group’s informational environment may be more representative of these channels.

Third, the treatments mechanically differ: the TESS treatments were online with recorded audio and graphs, while the in-store treatments were presented by a live person without graphs. While we assume for policy analysis that the information treatments were pure nudges and would thus have identical effects for a given consumer in a given choice environment, this could perhaps be violated in either of the experiments. For example, in-store information effects could be smaller if the in-store treatment were somehow more difficult to understand, or if consumers chose not to process that information in the absence of a quiz. However, our RAs for the in-store experiment report that most consumers did seem to engage with and understand the information.

The in-store experiment does not allow us to directly estimate the effects of information on WTP. We showed empirically in Section 3.4 that there is no clear relationship between the EPM and the CATE on WTP, and we showed theoretically in Section 2 that the average marginal bias can be large even when the EPM is zero. Furthermore, Online Appendix D.2 calibrates an example with demand parameters similar to the in-store experiment and shows that the CATE on WTP could easily be as large as \$2 when the effect on market shares is zero, even with fairly restrictive assumptions on the distribution of bias across consumers. We thus use the in-store results primarily to answer our positive research question about the effects of information on demand. Theoretically, we cannot reject the null hypothesis that the CATEs from the in-store experiment are equivalent to the CATEs in the TESS experiment. Notwithstanding, the fact that the EPM and the information

effects on market shares are generally smaller than in the TESS experiment is certainly consistent with the idea that the CATE on WTP would also be smaller. In this case, the in-store experiment results would strengthen the qualitative conclusion that information and inattention do not justify the incandescent lightbulb ban.

6 Conclusion

While imperfect information and inattention are commonly used to justify energy efficiency policies, the arguments are typically qualitative, without formal welfare analysis and relevant empirical tests. In this paper, we derived welfare effects of subsidies and standards in terms of two sufficient statistics, the baseline demand curve and the average marginal bias, and carried out a randomized experiment specifically designed to identify the two statistics.

Our main results suggest that moderate CFL subsidies may be optimal, but that imperfect information and inattention alone cannot justify a ban on incandescents. The approach requires the assumption that the effects of information on WTP equal the average marginal bias. While we carefully designed our treatments to make this as plausible as possible, we still view this assumption as only an approximation, and we explored plausible alternative assumptions in the welfare analysis. We showed how the analysis can be easily extended to incorporate externalities and other distortions not addressed by information provision, and we gave an example of this in studying what may be an excess mass of consumers with valuations near zero. In this alternative scenario, our model suggests that a ban does increase welfare.

To begin to address the question of whether the TESS results generalize to other populations and choice environments, we implemented a complementary in-store experiment. While the standard 2-by-2 design in this experiment does not allow a good approximation to the average marginal bias, the smaller effects on market shares might reinforce the qualitative conclusion that informational and attentional biases do not justify a ban.

The paper makes several important contributions. First, we provide provocative answers to our two research questions. While it was both common and plausible to believe that information provision could substantially affect the lightbulb market, both of our experiments show that large shares of consumers still prefer incandescents even after being powerfully informed. Similarly, while incandescent lightbulb bans have become important features of energy policy in many countries, our results suggest that more careful thought is needed about why these might increase welfare. Second, our experiments are a proof-of-concept for what we think will be an important effort to use RCTs to measure the effects of energy cost information provision on durable good purchases. Third, our basic approach is generally useful for studying behaviorally-motivated policies outside of the lightbulb market. In particular, we show that precisely identifying the parameters necessary for behavioral welfare analysis may require more complex empirical designs than had previously been anticipated.

References

- [1] Abaluck, Jason, and Jonathan Gruber (2011). “Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program.” *American Economic Review*, Vol. 101, No. 4 (June), pages 1180-1210.
- [2] ABC (2011). “Rand Paul’s Toilet Tirade.” Available from <http://abcnews.go.com/blogs/politics/2011/03/rand-pauls-toilet-tirade/>
- [3] Allcott, Hunt (2013). “The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market.” *American Economic Journal: Economic Policy*, Vol. 5, No. 3 (August), pages 30-66.
- [4] Allcott, Hunt (2014). “Paternalism and Energy Efficiency: An Overview.” In preparation for the *Annual Review of Economics*, Vol. 8.
- [5] Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky (2014). “Energy Policy with Externalities and Internalities.” *Journal of Public Economics*, Vol. 112 (April), pages 72-88.
- [6] Allcott, Hunt, and Nathan Wozny (2012). “Gasoline Prices, Fuel Economy, and the Energy Paradox.” *Review of Economics and Statistics*, forthcoming.
- [7] Anderson, Dennis, and John Claxton (1982). “Barriers to Consumer Choice of Energy Efficient Products.” *Journal of Consumer Research*, Vol. 9, No. 2 (September), pages 163-170.
- [8] Andreoni, James, and Charles Sprenger (2012). “Estimating Time Preferences from Convex Budgets.” *American Economic Review*, Vol. 12, No. 7 (December), pages 3333-3356.
- [9] Angrist, Joshua, and Jorn-Steffen Pischke (2009). Mostly Harmless Econometrics: An Empiricist’s Companion. Princeton, NJ: Princeton University Press.
- [10] Attari, Shahzeen, Michael DeKay, Cliff Davidson, and Wandu Bruine de Bruin (2010). “Public Perceptions of Energy Consumption and Savings.” *Proceedings of the National Academy of Sciences*, Vol. 107, No. 37 (September 14), pages 16054-16059.
- [11] Baicker, Kate, Sendhil Mullainathan, and Joshua Schwartzstein (2013). “Choice Hazard in Health Insurance.” Working Paper, Dartmouth College (October).
- [12] Bernheim, B. Douglas, and Antonio Rangel (2009). “Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics.” *Quarterly Journal of Economics*, Vol. 124, No. 1 (February), pages 51-104.
- [13] Blumstein, Carl, Betsy Krieg, Lee Schipper, and Carl York (1980). “Overcoming Social and Institutional Barriers to Energy Conservation.” *Energy*, Vol. 5, pages 355-371.
- [14] Borenstein, Severin, and Lucas Davis (2012). “The Equity and Efficiency of Two-Part Tariffs in U.S. Natural Gas Markets.” *Journal of Law and Economics*, Vol. 55, No. 1 (February).
- [15] Bollinger, Brian, Phillip Leslie, and Alan Sorensen (2011). “Calorie Posting in Chain Restaurants.” *American Economic Journal: Economic Policy*, Vol. 3, No. 1 (February), pages 91-128.
- [16] Bureau of Labor Statistics (BLS) (2013). “Consumer Expenditure Survey: Expenditure Tables.” <http://www.bls.gov/cex/csxstnd.htm>
- [17] Busse, Meghan, Christopher Knittel, and Florian Zettelmeyer (2013). “Are Consumers Myopic? Evidence from New and Used Car Purchases.” *American Economic Review*, Vol. 103, No. 1 (February), pages 220-256.
- [18] Caplin, Andrew, and Mark Dean (2014). “Revealed Preference, Rational Inattention, and Costly Information Acquisition.” Working Paper, New York University.
- [19] Carroll, Gabriel, James Choi, David Laibson, Brigitte Madrian, and Andrew Metrick (2009). “Optimal Defaults and Active Decisions.” *Quarterly Journal of Economics*, Vol. 124, pages 1639-1674.

- [20] Chetty, Raj, Adam Looney, and Kory Kroft (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review*, Vol. 99, No. 4 (September), pages 1145-1177.
- [21] Chetty, Raj, John Friedman, Tore Olsen, and Luigi Pistaferri (2011). "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *Quarterly Journal of Economics*, Vol. 126, No. 2, pages 749-804.
- [22] Davis, Lucas (2012). "Evaluating the Slow Adoption of Energy Efficient Investments: Are Renters Less Likely to Have Energy Efficient Appliances?" in *The Design and Implementation of U.S. Climate Policy*, Edited by Don Fullerton and Catherine Wolfram. Chicago, IL: University of Chicago Press.
- [23] Davis, Lucas, and Erich Muehlegger (2010). "Do Americans Consume too Little Natural Gas? An Empirical Test of Marginal Cost Pricing." *RAND Journal of Economics*, Vol. 41, No. 4 (Winter), pages 791-810.
- [24] Deutsch, Matthias (2010a). "Life Cycle Cost Disclosure, Consumer Behavior, and Business Implications: Evidence from an Online Field Experiment." *Journal of Industrial Ecology*, Vol. 14, No. 1, pages 103-120.
- [25] Deutsch, Matthias (2010b). "The Effect of Life-Cycle Cost Disclosure on Consumer Behavior: Evidence from a Field Experiment with Cooling Appliances." *Energy Efficiency*, Vol. 3, pages 303-315.
- [26] DEHWA (Australian Department of the Environment, Water, Heritage, and the Arts) (2008). "Regulatory Impact Statement: Proposal to Phase-Out Inefficient Incandescent Light Bulbs." http://www.energyrating.gov.au/wp-content/uploads/Energy_Rating_Documents/Library/Lighting/Incandescent_Lamps/200808-ris-phaseout.pdf
- [27] Diamond, Peter (1973). "Consumption Externalities and Imperfect Corrective Pricing." *Bell Journal of Economics and Management Science*, Vol. 4, No. 2 (Autumn), pages 526-538.
- [28] DOE (U.S. Department of Energy) (2009). "Technical Support Document: Impacts on the Nation of the Energy Independence and Security Act of 2007." http://www1.eere.energy.gov/buildings/appliance_standards/pdfs/en_masse_tsd_march_2009.pdf
- [29] DOE (U.S. Department of Energy) (2010). "Energy Star CFL Market Profile: Data Trends and Market Insights." http://www.energystar.gov/ia/products/downloads/CFL_Market_Profile.2010.pdf,
- [30] DOE (U.S. Department of Energy) (2014). "Electric Power Monthly." http://www.eia.gov/electricity/monthly/epm_table_grapher.cfm?t=epmt_5.3
- [31] Dranove, David, and Ginger Zhe Jin (2010). "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature*, Vol. 48, No. 4 (December), pages 935-963.
- [32] Dubin, Jeffrey, and Daniel McFadden (1984). "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica*, Vol. 52, No. 2 (March), pages 345-362.
- [33] Energy Star (2013). "ENERGY STAR Unit Shipment and Market Penetration Report Calendar Year 2012 Summary." http://www.energystar.gov/ia/partners/downloads/unit_shipment_data/2012_USD_Summary_Report.pdf 1224
- [34] Falk, Armin and Florian Zimmermann (2012). "A Taste for Consistency and Survey Response Behavior." *Cesifo Economic Studies*, Vol 59, No. 1, pages 181-193.
- [35] Fischer, Carolyn, Winston Harrington, and Ian Parry (2007). "Do Market Failures Justify Tightening Corporate Average Fuel Economy (CAFE) Standards?" *The Energy Journal*, Vol. 28, No. 4, pages 1-30.
- [36] Fong, Christina, and Erzo F.P. Luttmer (2009). "What Determines Giving to Hurricane Katrina Victims? Experimental Evidence on Racial Group Loyalty." *American Economic Journal: Applied Economics*, Vol. 1, No. 2 (April), pages 64-87.

- [37] Formisano, Bob (2008). “2007 Energy Bill - Are They Phasing Out or Making Incandescent Bulbs Illegal?” http://homerepair.about.com/od/electricalrepair/ss/2007_energybill.htm
- [38] Gabaix, Xavier (2014). “A Sparsity-Based Model of Bounded Rationality.” Working Paper, New York University (May).
- [39] Gabaix, Xavier, and David Laibson (2006). “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics*, Vol. 121, No. 2, pages 505-540.
- [40] Gayer, Ted (2011). “A Better Approach to Environmental Regulation: Getting the Costs and Benefits Right.” Hamilton Project Discussion Paper 2011-06 (May).
- [41] Gillingham, Kenneth, Matthew Harding, and David Rapson (2012). “Split Incentives and Household Energy Consumption.” *Energy Journal*, Vol. 33, No. 2, pages 37-62.
- [42] Gillingham, Kenneth, and Karen Palmer (2013). “Bridging the Energy Efficiency Gap.” Resources for the Future Discussion Paper 13-02.
- [43] Goldberg, Pinelopi (1998). “The Effects of the Corporate Average Fuel Economy Standards in the US.” *Journal of Industrial Economics*, Vol. 46, pages 1-33.
- [44] Gruber, Jonathan, and Botond Koszegi (2004). “Tax Incidence when Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes.” *Journal of Public Economics*, Vol. 88, pages 1959-1987.
- [45] Harrison, Glenn, and John List (2004). “Field Experiments.” *Journal of Economic Literature*, Vol. 42, No. 4 (December), pages 1009-1055.
- [46] Hassett, Kevin, and Gilbert Metcalf (1995). “Energy Tax Credits and Residential Conservation Investment: Evidence from Panel Data.” *Journal of Public Economics*, Vol. 57, pages 201-217.
- [47] Hausman, Jerry (1979). “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables.” *Bell Journal of Economics*, Vol. 10, No. 1, pages 33-54.
- [48] Heidhues, Paul, Botond Koszegi, and Takeshi Murooka (2014). “Inferior Products and Profitable Deception.” Working Paper.
- [49] Heiss, Florian, Daniel McFadden, and Joachim Winter (2007). “Mind the Gap! Consumer Perceptions and Choices of Medicare Part D Prescription Drug Plans.” NBER Working Paper 13627.
- [50] Herberich, David, John List, and Michael Price (2011). “How Many Economists Does it Take to Change a Light Bulb? A Natural Field Experiment on Technology Adoption.” Working Paper, University of Chicago (March).
- [51] Heutel, Garth (2011). “Optimal Policy Instruments forExternality-Producing Durable Goods under Time Inconsistency.” Working Paper, UNC Greensboro (May).
- [52] Hossain, Tanjim, and John Morgan (2006). “...Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay.” *Advances in Economic Analysis and Policy*, Vol. 6.
- [53] Houde, Sebastien (2012). “How Consumers Respond to Certification: A Welfare Analysis of the Energy Star Program.” Working Paper, Stanford University (October).
- [54] Jaffe, Adam, and Robert Stavins (1994). “The Energy Paradox and the Diffusion of Conservation Technology.” *Resource and Energy Economics*, Vol. 16, pages 91-122.
- [55] Jessoe, Katrina, and David Rapson (2014). “Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use.” *American Economic Review*, Vol. 104, No. 4 (April), pages 1417-1438.
- [56] Kallbekken, Steffen, Hakon Saelen, and Erlend Hermansen (2013). “Bridging the Energy Efficiency Gap: A Field Experiment on Lifetime Energy Costs and Household Appliances.” *Journal of Consumer Policy*, Vol. 36, pages 1-16.
- [57] Koszegi, Botond, and Adam Szeidl (2013). “A Model of Focusing in Economic Choice.” *Quarterly Journal of Economics*, Vol. 128, No. 1, pages 53-107.

- [58] Mullainathan, Sendhil, Joshua Schwartzstein, and William J. Congdon (2012). “A Reduced-Form Approach to Behavioral Public Finance.” *Annual Review of Economics*, Vol. 4, 17.1-17.30.
- [59] Newell, Richard, and Juha Siikamaki (2013). “Nudging Energy Efficiency Behavior: The Role of Information Labels.” Resources for the Future Discussion Paper 13-17 (July).
- [60] NRDC (Natural Resources Defense Council) 2011. “Better Light Bulbs Equal Consumer Savings in Every State.” www.nrdc.org/policy
- [61] O’Donoghue, Edward, and Matthew Rabin (2006). “Optimal Sin Taxes.” *Journal of Public Economics*, Vol. 90, pages 1825-1849.
- [62] Parry, Ian, David Evans, and Wallace Oates (2010). “Are Energy Efficiency Standards Justified?” Resources for the Future Discussion Paper 10-59 (November).
- [63] Rabin, Matthew, and Georg Weizsacker (2009). “Narrow Bracketing and Dominated Choices.” *American Economic Review*, Vol. 99, No. 4 (September), pages 1508-1543.
- [64] Saltee, James (2013). “Rational Inattention and Energy Efficiency.” *Journal of Law and Economics*, forthcoming.
- [65] Saltee, James, Sarah West, and Wei Fan (2009). “Consumer Valuation of Fuel Economy: A Microdata Approach.” Working Paper, University of Chicago (September). Available from <http://www.ntanet.org/images/stories/pdf/proceedings/09/032.pdf>.
- [66] Sanstad, Alan, and Richard Howarth (1994). “‘Normal’ Markets, Market Imperfections, and Energy Efficiency.” *Energy Policy*, Vol. 22, No. 10, pages 811-818.
- [67] Sims, Christopher (2003). “Implications of Rational Inattention.” *Journal of Monetary Economics*, Vol. 50, No. 3, pages 665-690.
- [68] Snyder, Mark (1974). “Self-Monitoring of Expressive Behavior.” *Journal of Personality and Social Psychology*, Vol. 30, pages 526-537.
- [69] Spinnewijn, Johannes (2014). “Heterogeneity, Demand for Insurance and Adverse Selection.” Working Paper, London School of Economics.
- [70] Sylvania (2012). “5th Annual Sylvania Socket Survey.” <https://www.sylvania.com/en-us/tools-and-resources/surveys/Pages/socket-survey.aspx>
- [71] Ward, David, Christopher Clark, Kimberly Jensen, Steven Yen, and Clifford Russell (2011). “Factors Influencing Willingness-to-Pay for the Energy Star Label.” *Energy Policy*, Vol. 39, No. 3 (March), page 1450 – 1458.

Tables

Table 1: **Effects of TESS Information Treatment**

	(1)	(2)	(3)	(4)	(5)	(6)
1(Treatment)	2.54 (0.55)***	2.28 (0.36)***	2.30 (0.37)***	3.16 (0.37)***	2.29 (0.54)***	2.14 (0.50)***
1(Endline-Only)					-0.44 (0.76)	
1(Positive Treatment)						0.35 (0.56)
R2	0.03	0.57	0.58	0.33	0.04	0.58
N	1,203	1,203	1,188	919	1,449	1,188
Baseline WTP Dummies μ	No	Yes	Yes	Yes	No	Yes
Individual Characteristics	No	No	Yes	Yes	Yes	Yes
Exclude Max./Min. Baseline WTP	No	No	No	Yes	No	No
Include Endline-Only Group	No	No	No	No	Yes	No

Notes: This table presents estimates of Equation (6). The outcome variable is endline willingness-to-pay for the CFL. 1(Treatment) pools all information sub-treatments. Robust standard errors in parenthesis. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively. Observations are weighted for national representativeness.

Table 2: **Welfare Analysis Using TESS Results**

(1)	(2)	(3)	(4)	(5)	(6)
	Average		Incremental		
	Relative WTP	Average	Demand	Incremental	Cumulative
CFL	of Marginal	Marginal	Change	Welfare	Welfare
Subsidy	Consumers	Bias	(Share of	Effect	Effect
(\$/package)	(\$/package)	(\$/package)	packages)	(\$/package)	(\$/package)
1	-0.5	2.11	0.126	0.20	0.20
2	-1.5	2.16	0.052	0.03	0.24
3	-2.5	3.41	0.028	0.03	0.26
4	-3.5	1.77	0.030	-0.05	0.21
6	-5	1.77	0.006	-0.02	0.19
8	-7	1.77	0.008	-0.04	0.15
10	-9	1.77	0.003	-0.02	0.13
∞	-15	1.77	0.043	-0.57	-0.44

Notes: This table uses the TESS experiment results to calculate the welfare effects at different levels of the CFL subsidy. Observations are weighted for national representativeness. Average Marginal Bias is the point estimates from Figure 4. Incremental Welfare Effect is from Equation (2).

Table 3: **Welfare Analysis Under Alternative Assumptions**

Row	Scenario	(1)	(2)	(3)
		Optimal Subsidy (\$/package)	Welfare Effect of Ban (\$/package)	Welfare Effect of Ban (Percent of Surplus)
1	Base	3	-0.44	-41
<u>Alternative Censoring Assumptions: If censored, assume ...</u>				
2	WTP={\$12,-\$12}	3	-0.34	-36
3	WTP={\$20,-\$20}	3	-0.60	-47
4	self-reported hypothetical WTP	3	-0.61	-43
<u>Alternatives to Assumption 1: Scale average marginal bias to match ...</u>				
5	consumers who pass Treatment Info screen “quiz”	3	-0.41	-38
6	consumers with “correct” post-experiment beliefs	3	-0.22	-21
7	Balanced Treatment group	3	-0.48	-45
8	10 percent confidence bound	1	-0.92	-86
9	90 percent confidence bound	(Ban)	0.05	4
<u>Additional Distortion</u>				
10	Excess mass consumers have $v = 7.66$	8	1.22	114

Notes: This table presents welfare results using the TESS data under alternative assumptions. Column 3 divides column 2 by incandescent lightbulb buyers’ total perceived consumer surplus from having the incandescent available. Observations are weighted for national representativeness.

Table 4: **Descriptive Statistics and Balance for In-Store Experiment**

	(1)	(2)	(3)	(4)
	Experimental Sample Mean	Refused - Sample Difference	Treatment - Control Difference	Rebate - Standard Difference
Individual Characteristics				
Energy an Important Factor in Purchase Decision	0.25 (0.43)		0.009 (0.026)	-0.024 (0.026)
Expected Usage (Minutes/Day)	333 (280)		12.8 (17.0)	2.7 (17.0)
Age	43.8 (11.4)	2.3 (0.6)***	0.7 (0.7)	-0.3 (0.7)
Male	0.66 (0.47)	0.06 (0.03)**	0.009 (0.029)	0.003 (0.029)
African American	0.16 (0.37)	-0.04 (0.02)**	-0.001 (0.022)	-0.008 (0.022)
Asian	0.06 (0.24)	0.04 (0.02)**	-0.030 (0.014)**	0.005 (0.015)
Caucasian	0.66 (0.47)	-0.07 (0.03)**	0.037 (0.029)	-0.005 (0.029)
Hispanic	0.07 (0.25)	0.06 (0.02)***	0.001 (0.015)	0.011 (0.015)
Middle Eastern	0.01 (0.12)	0.01 (0.01)	0.002 (0.013)	0.007 (0.007)
F-Test p-Value		0.00	0.742	0.896
Purchase Decisions				
Purchased Any Lightbulb	0.77 (0.42)		0.011 (0.025)	0.027 (0.025)
Purchased Substitutable Lightbulb	0.73 (0.44)		-0.008 (0.027)	0.011 (0.027)

Notes: Column 1 presents means of individual characteristics in the in-store experiment sample, with standard deviations in parenthesis. Column 2 presents differences in recorded demographic characteristics between those who refused or did not complete the survey and the experimental sample. Column 3 presents differences in means between Treatment and Control groups, while column 4 presents differences in means between the Rebate Coupon and Standard Coupon groups. Columns 2, 3, and 4 have robust standard errors in parenthesis. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively.

Table 5: **Effects of In-Store Information Treatment**

	(1)	(2)	(3)
1(Treatment)	-0.002 (0.035)	0.004 (0.033)	-0.022 (0.045)
1(Rebate)	0.094 (0.035)***	0.105 (0.033)***	0.078 (0.047)*
1(Rebate and Treatment)			0.054 (0.066)
R2	0.01	0.16	0.16
<i>N</i>	794	793	793
Individual Characteristics	No	Yes	Yes

Notes: This table presents estimates of Equation (7), a linear probability model with outcome variable 1(Purchased CFL). The dependent variable has mean 0.38. Robust standard errors in parenthesis. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively.

Figures

Figure 1: **TESS Experimental Design**
Groups and Shares of Population

Process

1. **Baseline choices** (multiple price list)
2. **Information provision** (two screens, content varies by group)
3. **Endline choices** (multiple price list)
4. **Post-experiment survey** (beliefs, time preferences, etc.)

Figure 2: **Histogram of Relative WTP Changes**

Note: This figure plots the histogram of changes from baseline to endline in relative willingness-to-pay for the information Treatment and Control groups. Treatment pools both Positive and Balanced Treatment groups, although the Endline-Only Treatment group is excluded because there is no baseline WTP from which to calculate a change. Observations are weighted for national representativeness.

Figure 3: Endline CFL Demand Curves

Notes: This figure plots the endline demand curves from the TESS experiment. Observations are weighted for national representativeness.

Figure 4: Conditional Average Treatment Effects by Level of Baseline WTP

Note: This graph presents the conditional average treatment effects of information provision for consumers at each level of baseline relative WTP. Due to limited sample size, baseline WTPs less than -\$3 are grouped together, as are baseline WTPs greater than \$8. Dotted lines are 90 percent confidence intervals. Observations are weighted for national representativeness.

Figure 5: Cumulative Density Function of Cost Savings Beliefs

Note: This figure plots the cumulative density of beliefs about the electricity cost savings from CFLs compared to incandescents for 8000 hours of light, from the TESS post-experiment survey. The Treatment group pools all information sub-treatments. Observations are weighted for national representativeness.

Figure 6: Welfare Calculation Using TESS Experiment

Notes: This figure illustrates the welfare effects of increases in the CFL subsidy using the TESS experiment results. Internality Reduction and Harberger Distortion are as defined in Equation (2). Observations are weighted for national representativeness.

Online Appendix: Not For Publication

The Lightbulb Paradox: Evidence from Two Randomized Experiments

Hunt Allcott and Dmitry Taubinsky

A Details of TESS Experiment

Introductory Screen

In appreciation for your participation in this study, we are giving you a \$10 shopping budget. During the study, we will ask you to make 30 decisions between pairs of light bulbs using that \$10 shopping budget. There will be a first set of 15 decisions, then a break, and then a second set of 15 decisions.

After you finish with all 30 decisions, one of them will be randomly selected as your “official purchase.” In approximately four to six weeks, GfK will send you the light bulbs you chose in that official purchase. After your official purchase has been paid for from the \$10 shopping budget we are giving you, any money left over will be provided to you in the form of **bonus points awarded to your account**. This means that after the study is completed, you will receive 1) the light bulbs you selected in the decision that is randomly selected to be your “official purchase” and 2) an amount between zero and 10000 bonus points, corresponding to whatever money is left in your shopping budget after the purchase.

Light bulbs are frequently shipped in the mail. There is not much risk of breakage, but if anything does happen, GfK will just ship you a replacement. Even if you don't need light bulbs right now, remember that you can store them and use them in the future.

Since each of your decisions has a chance of being your official purchase, you should think about each decision carefully.

Next

Baseline Choices (Top of Screen)

We have given you a \$10 shopping budget to purchase a package of light bulbs. Your first 15 purchase decisions will concern the two packages of light bulbs shown below.

Choice A
Philips 60-Watt-Equivalent
Compact Fluorescent Light Bulb, 1-Pack

Choice B
Philips 60-Watt Incandescent
Light Bulbs, 4-Pack

[Click for detailed product information](#)

Between the 15 decisions, the only thing that varies is the price. Each of these decisions has a chance of being the one choice (out of 30) that will become your official purchase, so you should think about each purchase carefully. Whatever money you do not spend on the light bulbs, you get to keep: any remaining money will be provided to you as cash-equivalent bonus points. Please think about each decision carefully.

Here is an example of how this might work. After you make all your decisions, suppose that Decision Number 6 from the set below were selected as your official purchase.

- If you had chosen Choice A, you would pay \$2 from your \$10 shopping budget. You would receive the Choice A light bulb package in the mail within 4-6 weeks, as well as the remaining $\$10 - \$2 = \$8$ in your shopping budget (You would receive that \$8 in the form of 8000 bonus points credited to your account.)
- If you had chosen Choice B, you would pay \$4 from your \$10 shopping budget. You would receive the Choice B light bulb package in the mail within 4-6 weeks, as well as the remaining $\$10 - \$4 = \$6$ in your shopping budget. (You would receive that \$6 in the form of 6000 bonus points credited to your account.)

Now please make your decisions for each of the 15 choices below.

Detailed Product Information Window

<i>Detailed Product Information</i>		
Choice:	A	B
Manufacturer:	Philips	Philips
Type:	Compact Fluorescent (CFL)	Incandescent
Number of Bulbs:	1	4
Light Output:	60 Watt-equivalent	60 Watts
Light Output:	900 Lumens	840 Lumens
Color Temperature:	2700K	2700K
Energy Use:	13 Watts	60 Watts
Manufacturer's Home Country:	USA	USA

Back to Questionnaire

Baseline Choices (Bottom of Screen)

Now please make your decisions for each of the 15 choices below.

Decision Number	Choice A 60-Watt-Equivalent Compact Fluorescent Light Bulb, 1-Pack	Choice B 60-Watt Incandescent Light Bulbs, 4-Pack
	Purchase Choice A for free	Purchase Choice B for \$10
1)	☐	☐
	Purchase Choice A for free	Purchase Choice B for \$8
2)	☐	☐
	Purchase Choice A for free	Purchase Choice B for \$6
3)	☐	☐
	Purchase Choice A for free	Purchase Choice B for \$4
4)	☐	☐
	Purchase Choice A for \$1	Purchase Choice B for \$4
5)	☐	☐
	Purchase Choice A for \$2	Purchase Choice B for \$4
6)	☐	☐
	Purchase Choice A for \$3	Purchase Choice B for \$4
7)	☐	☐
	Purchase Choice A for \$4	Purchase Choice B for \$4
8)	☐	☐
	Purchase Choice A for \$4	Purchase Choice B for \$2

Note: This does not show all of the 15 Decision Numbers.

Balanced Treatment Introductory Screen

For this next part of the study, you will have the opportunity to learn more about compact fluorescent light bulbs (CFLs) and incandescent light bulbs. We will focus on the following two issues:

- Total Costs
- Disposal and Warm-Up Time

The discussion of each issue will be followed by a one-question quiz. Please pay close attention to the discussion so that you can correctly answer the quiz question.

Next

Treatment Information Screen

CFLs last longer than incandescents. At average usage:

- Incandescents burn out and have to be replaced every year.
- CFLs burn out and have to be replaced every eight years.

If one incandescent bulb costs \$1 and one CFL costs \$4, this means that the total purchase prices for eight years of light are:

- \$8 for incandescents
- \$4 for CFLs

Also, CFLs use less electricity than incandescents. At national average usage and electricity prices:

- A standard (60-Watt) incandescent uses \$6 in electricity each year.
- An equivalent CFL uses \$1.50 in electricity each year.

Thus, for eight years of light, the total costs to purchase bulbs and electricity would be:

- \$56 for incandescents: \$8 for the bulbs plus \$48 for electricity
- \$16 for a CFL: \$4 for the bulbs plus \$12 for electricity

The graph below illustrates this:

Negative Information Screen

After they burn out, CFLs need proper disposal:

- Because CFLs contain mercury, it is recommended that they be properly recycled, and not simply disposed of in regular household trash. CFLs can be recycled through:
 - Local waste collection sites
 - Mail-back services that you can find online
 - Local retailers, including Ace Hardware, IKEA, Home Depot, and Lowe's, as well as other retailers.
- No special precautions need to be taken to dispose of an incandescent light bulb. Incandescents can be disposed of in regular household trash.

After the light switch is turned on, CFLs take longer to warm up than incandescents:

- An incandescent reaches full brightness immediately.
- A typical CFL can take 60 to 90 seconds to reach its full brightness.

The graph below illustrates this:

Question: About how much longer does it take a typical CFL to reach full brightness, as compared to an incandescent?

Type your answer below.

to seconds

Control Introductory Screen

For this next part of the study, you will have the opportunity to learn more about light bulbs. We will focus on the following two issues:

1. Number of Bulbs by Sector
2. Sales Trends

The discussion of each issue will be followed by a one-question quiz. Please pay close attention to the discussion so that you can correctly answer the quiz question.

Next

Number of Bulbs Screen

According to official estimates, there are slightly more than eight billion light bulbs installed in the United States.

The US economy can be divided into three major sectors: residential, commercial, and industrial. Each sector has a different number of light bulbs:

- There are about 5.8 billion light bulbs installed in residential buildings in the U.S.
- There are about 2.1 billion light bulbs installed in commercial buildings in the U.S.
- There are about 0.14 billion light bulbs installed in industrial buildings in the U.S.

The graph below illustrates this:

Question: About how many more light bulbs are installed in residential buildings compared to commercial buildings in the U.S.?

To answer this question, you can enter whole numbers and/or decimals.

Type your answer below.

billion

Sales Trends Screen

According to official sales data, sales of light bulbs in the United States have had the following trend:

- Sales increased in each year between 2000 and 2007.
- Sales decreased slightly in 2008 and 2009.

Total light bulb sales were different at the end of the decade compared to the beginning:

- Sales in 2000 were just over 1.7 billion bulbs.
- Sales in 2009 were just under 1.8 billion bulbs.

The graph below illustrates this:

Question: About how many light bulbs were sold in the United States in 2009?

To answer this question, you can enter whole numbers and/or decimals.

Type your answer below.

billion

Endline Choices (Top of Screen)

Remember, we have given you a \$10 shopping budget to purchase a package of light bulbs. Your second 15 purchase decisions will concern the two packages of light bulbs shown below.

Choice A
Philips 60-Watt-Equivalent
Compact Fluorescent Light Bulb, 1-Pack

Choice B
Philips 60-Watt Incandescent
Light Bulbs, 4-Pack

[Click for detailed product information](#)

Between the 15 decisions, the only thing that varies is the price. Each of these decisions has a chance of being the one choice (out of 30) that will become your official purchase, so you should think about each purchase carefully. Whatever money you do not spend on the light bulbs, you get to keep: any remaining money will be provided to you as cash-equivalent bonus points. Please think about each decision carefully.

Here is an example of how this might work. After you make all your decisions, suppose that Decision Number 21 from the set below were selected as your official purchase.

- If you had chosen Choice A, you would pay \$2 from your \$10 shopping budget. You would receive the Choice A light bulb package in the mail within 4-6 weeks, as well as the remaining \$10-\$2=\$8 in your shopping budget (You would receive that \$8 in the form of 8000 bonus points credited to your account)
- If you had chosen Choice B, you would pay \$4 from your \$10 shopping budget. You would receive the Choice B light bulb package in the mail within 4-6 weeks, as well as the remaining \$10-\$4=\$6 in your shopping budget (You would receive that \$6 in the form of 6000 bonus points credited to your account.)

Now please make your decisions for each of the 15 choices below.

Endline Choices (Bottom of Screen)

Now please make your decisions for each of the 15 choices below.

Decision Number	Choice A 60-Watt-Equivalent Compact Fluorescent Light Bulb, 1-Pack	Choice B 60-Watt Incandescent Light Bulbs, 4-Pack
	Purchase Choice A for free	Purchase Choice B for \$10
16)	☐	☐
	Purchase Choice A for free	Purchase Choice B for \$8
17)	☐	☐
	Purchase Choice A for free	Purchase Choice B for \$6
18)	☐	☐
	Purchase Choice A for free	Purchase Choice B for \$4
19)	☐	☐
	Purchase Choice A for \$1	Purchase Choice B for \$4
20)	☐	☐
	Purchase Choice A for \$2	Purchase Choice B for \$4
21)	☐	☐
	Purchase Choice A for \$3	Purchase Choice B for \$4
22)	☐	☐
	Purchase Choice A for \$4	Purchase Choice B for \$4
23)	☐	☐
	Purchase Choice A for \$4	Purchase Choice B for \$3

Note: This does not show all of the 15 Decision Numbers.

Post-Experiment Survey Questions

Question 1. *How important were the following factors in your purchase decision? [Rate from 1-10]*

1. *Energy use*
2. *Time required for the bulb to reach full brightness after it is turned on*
3. *Bulb lifetime*
4. *Mercury content and protocols for proper disposal*
5. *Purchase Price*

Question 2. *Do you think that the intent of the study was to ...*

Select all answers that apply

1. *Understand the effect of price changes on purchasing patterns*
2. *Measure whether people make consistent purchases in similar situations*
3. *Understand why people buy incandescents vs. CFLs*
4. *Test how well people are able to quantify energy costs*
5. *Test whether ability to quantify energy costs affects purchases of incandescents vs. CFLs*
6. *Test whether the number of bulbs in a package affects purchasing patterns*
7. *Test whether consumer education affects purchases of incandescents vs. CFLs*
8. *Understand what features of lightbulbs are most important to people*
9. *Predict the future popularity of incandescents vs. CFLs*
10. *None of the above*

Question 3. Part A: *The typical CFL lasts 8000 hours, or about eight years at typical usage rates. Do you think it costs more or less to buy electricity for that 8000 hours of light from compact fluorescent light bulbs (CFLs) compared to incandescent light bulbs?*

- More
- Less

Part B: *At national average electricity prices, how much [more/less] does it cost to buy electricity for that 8000 hours of light from compact fluorescent light bulbs (CFLs) compared to incandescent light bulbs? Just give your best guess.*

Question 4. *Some states and local areas have rebates, low-interest loans, or other incentives available for energy efficiency. These might include rebates for Energy Star appliances or energy efficient light bulbs, low-interest loans for energy-saving home improvements, government-funded weatherization, and other programs. Are any such programs available in your area?*

1. *Yes*
2. *I think so, but I'm not sure*
3. *I'm not sure at all*
4. *I think not, but I'm not sure*
5. *No*

Question 5. *This question is about hypothetical choices and does not affect your earnings in this study.*
Suppose that you could get the amount under “Option A” (i.e. \$100), or the amount under “Option B” a year later. Assume it’s no more work for you to receive the money under Option A than under Option B, and that you would receive the money for sure, regardless of when you choose to receive it. Which would you prefer?

This question is about hypothetical choices and does not affect your earnings in this study. Suppose that you could get the amount under “Option A” (i.e. \$100), or the amount under “Option B” a year later. Assume it’s no more work for you to receive the money under Option A than under Option B, and that you would receive the money for sure, regardless of when you choose to receive it. Which would you prefer?

For each row, please choose whether you would prefer the payment under “Option A” or under “Option B”.

	Option A	Option B
	\$100 today	\$50 in one year
1)	☐	☐
	\$100 today	\$90 in one year
2)	☐	☐
	\$100 today	\$100 in one year
3)	☐	☐
	\$100 today	\$110 in one year
4)	☐	☐
	\$100 today	\$130 in one year
5)	☐	☐
	\$100 today	\$150 in one year
6)	☐	☐
	\$100 today	\$170 in one year
7)	☐	☐
	\$100 today	\$200 in one year
8)	☐	☐
	\$100 today	\$250 in one year
9)	☐	☐
	\$100 in one year	\$50 in two years
10)	☐	☐
	\$100 in one year	\$90 in two years
11)	☐	☐

Notes: This does not show all of the 18 choices. Participants were randomly assigned to receive either this table or another table that was identical except that the bottom half and top half were switched, so that the one year vs. two year tradeoffs were presented first.

Question 6. Please indicate how much you agree or disagree with the following statements:

Select one answer from each row in the grid

[Strongly Agree Agree Neutral Disagree Strongly Disagree]

1. *It's important to me to fit in with the group I'm with.*
2. *My behavior often depends on how I feel others wish me to behave.*
3. *My powers of intuition are quite good when it comes to understanding others' emotions and motives.*
4. *My behavior is usually an expression of my true inner feelings, attitudes, and beliefs.*
5. *Once I know what the situation calls for, it's easy for me to regulate my actions accordingly.*
6. *I would NOT change my opinions (or the way I do things) in order to please someone else or win their favor.*

B Additional TESS Results

B.1 Descriptive Statistics and Baseline Willingness-to-Pay

Column 1 of Table A.1 presents descriptive statistics. All statistics for baseline relative WTP necessarily exclude the Endline-Only group. Liberal is self-reported political ideology, originally on a seven-point scale, normalized to mean zero and standard deviation one, with larger numbers indicating more liberal. Party is self-reported political affiliation, similarly normalized from an original seven-point scale, with larger numbers indicating more strongly Democratic. Environmentalist measures the consumer’s answer to the question, “Would you describe yourself as an environmentalist?” Conserve Energy is an indicator for whether the consumer reports having taken steps to conserve energy in the past twelve months. Homeowner is a binary indicator variable for whether the consumer owns his or her home instead of rents. Except for baseline WTP, these variables were recorded when the consumer first entered the TESS panel, not as part of our experiment.

Column 2 presents the difference in means between all Treatment groups vs. Control. Column 3 presents the difference in means between the Positive and Balanced Treatment groups. All 20 t-tests fail to reject equality, as do the joint F-tests of all characteristics.

Table A.2 shows the association between baseline WTP and the individual characteristics in Table A.1. Column 1 shows that men, Democrats, and those who report having taken steps to conserve energy have higher demand for CFLs. Columns 2-5 separately test individual variables of environmentalism and political ideology which are correlated, providing additional evidence that liberals tend to have higher WTP. These correlations conform to our priors and build further confidence that the WTP measurements are meaningful.

The table also provides suggestive evidence on two distortions other than imperfect information and inattention which might justify subsidies and standards. The first is that renters might have lower CFL demand because they might leave the CFLs in the house’s light sockets when they move and be unable to capitalize on their investment. Lacking random or quasi-random assignment in renter vs. homeowner status, Davis (2012) and Gillingham, Harding, and Rapson (2012) correlate durable good ownership with homeowner status conditional on observables. Columns 1 and 6 replicate their approach in the TESS data, showing that homeowners do not have higher WTP for CFLs. However, additional (unreported) regressions with market share at market prices as the dependent variable show that we cannot reject the Davis (2012) result that homeowners are five percent more likely to prefer CFLs.

The second potential distortion considered in Table A.2 is present bias. In the post-experiment survey, we estimate the β and δ of a quasi-hyperbolic model through a menu of hypothetical intertemporal choices at two different time horizons: \$100 now vs. $\$m^1$ in one year, and \$100 in one year vs. $\$m^2$ in two years. Denoting $\bar{m}_i^1$ and $\bar{m}_i^2$ as the midpoint between the values at which participant i switches from preferring money sooner to later, the long run discount factor is $\delta_i = 100/\bar{m}_i^2$, and the present bias parameter is $\beta_i = \bar{m}_i^2/\bar{m}_i^1$. We dropped non-monotonic responses and top-coded and bottom-coded $\bar{m}_i^1$ and $\bar{m}_i^2$ at \$300 and \$40, respectively. The median δ is 5/7, meaning that the median consumer prefers \$100 in one year to \$130 in two years but prefers \$150 in two years to \$100 in one year. A slight majority of consumers (52 percent) have $\beta = 1$, meaning that they are not present or future biased by this measure, and the median β is also 1.

If there is a distribution of β and δ , consumers with higher β and δ should be more likely to purchase CFLs. Column 1 shows that there is a conditional correlation between δ and baseline WTP, suggesting that people who are more patient may be more likely to purchase CFLs. However, there is no statistically signif-

icant correlation between β and WTP. Column 7 repeats the estimates without any conditioning variables, and the coefficients are comparable. The results in column 1 rule out with 90 percent confidence that a one standard deviation increase in β increases WTP for the CFL by more than \$0.47.²⁹

B.2 Estimating the Equivalent Price Metric in the TESS Experiment

The TESS experiment allows us to directly estimate the conditional average treatment effect on WTP for consumers marginal between points on the multiple price list. Figure 4 shows the 11 intervals over which we calculate CATEs, which are bounded by $-\infty, -3, -2, -1, 0, 1, 2, 3, 4, 6, 8, \infty$. In this section, we calculate EPMs over the same intervals and compare them to the CATEs. We exclude the highest and lowest intervals, because it is not possible to calculate a demand slope on an interval bounded by an infinite price.

To estimate $EPM[p_l, p_h]$, we reshape the TESS data so that there are two purchase observations per consumer, one at p_l and one at p_h . Denoting S_p as an indicator for whether this observation is at the lower price, we then estimate the following equation in a linear probability model:

$$1(\text{Purchase CFL})_{ip} = \tau T_i + \eta S_p + \alpha T_i S_p + \varepsilon_{ip} \quad (8)$$

Because there are multiple observations per consumer, we cluster standard errors by consumer.

Coefficients from this regression can be inserted into Equation (5) to approximate $EPM[p_l, p_h]$:

$$EPM[p_l, p_h] \approx \frac{(\Delta Q(p_l) + \Delta Q(p_h)) \cdot \frac{1}{2}}{(D_N(p_l) - D_N(p_h)) / \Delta p} = \frac{\hat{\tau} + \hat{\alpha}/2}{(\hat{\eta} + \hat{\alpha})/\Delta p} \quad (9)$$

Standard errors are calculated using the Delta method.

Table A.3 presents results for each of the nine relative price intervals. Column 1 gives the numerator of Equation (9), column 2 gives the denominator, and column 3 gives the ratio of columns 1 and 2. Column 4 presents the CATE estimates, which are from Figure 4. Column 5 presents the p-value of the difference between the EPM and CATE on WTP, while column 6 presents the absolute value of the difference divided by the CATE. In four of the nine price intervals, the EPM and CATE differ with greater than 90 percent confidence, and on average, the two quantities differ by 49 percent.

B.3 Self-Monitoring Scale

If demand effects are present, they should differentially affect people who are more able to detect the intent of the study and are more willing to change their choices given the experimenter’s intent. One existing measure of these issues is the Self-Monitoring Scale, a battery of personality questions developed by Snyder (1974). Snyder writes that the scale is designed to identify individuals who “tend to express what they think and feel, rather than mold and tailor their self-presentations and social behavior to fit the situation.”

From the set of standard Self-Monitoring Scale statements, we took the most relevant six:

- It’s important to me to fit in with the group I’m with.

²⁹This may not be surprising. Present bias over *cash flows* might cause consumers to buy an incandescent to reduce current expenditures, but as Andreoni and Sprenger (2012) and many others have pointed out, agents in most models would be present biased over *consumption*, and most American consumers have enough liquidity that paying the incremental few dollars for a CFL does not immediately affect consumption. Present bias could induce people to procrastinate in buying and installing CFLs, but this would play no role in the TESS experiment because we forced consumers to make an active choice.

- My behavior often depends on how I feel others wish me to behave.
- My powers of intuition are quite good when it comes to understanding others' emotions and motives.
- My behavior is usually an expression of my true inner feelings, attitudes, and beliefs.
- Once I know what the situation calls for, it's easy for me to regulate my actions accordingly.
- I would NOT change my opinions (or the way I do things) in order to please someone else or win their favor.

At the very end of the post-experiment survey, we asked consumers to respond to each of these six statements on a five-point Likert scale from "Agree" to "Disagree." We normalize responses to each question to mean zero, standard deviation one, and interact each with the treatment indicator while also controlling for lower-order interactions. Table A.5 presents results. While the six Self-Monitoring Scale variables are correlated with each other, none is correlated with endline CFL demand or with the treatment effect, nor is a composite of the six.

B.4 Effects on Purchase Priorities

The post-experiment survey also asks consumers to rate on a scale of 1-10 the importance of price, energy use, bulb lifetime, warm-up time, and mercury and disposal in their purchase decisions. Table A.6 presents how the treatments affected these ratings. Both Positive and Balanced treatments decreased the stated importance of purchase prices, consistent with consumers re-orienting away from purchase price as a measure of total cost. Point estimates suggest that both the Positive and Balanced treatments increased the importance of energy use and that the Positive treatment also increased the importance of bulb lifetimes. These are the only estimates in the entire analysis whose significance level is affected by the weighting: they are not significant in Table A.6, but (unreported) regressions show that they are statistically significant when weighting all observations equally. The Positive Treatment group and Control group do not differ on the importance of warm-up time or mercury and disposal, which is to be expected because neither group received information on these two issues. Interestingly, the Balanced treatment decreased the importance of warm-up time. One potential explanation is that consumers had previously believed that CFL warm-up times were longer, and the treatment reduced the importance of this difference between CFLs and incandescents.

Appendix B Tables

Table A.1: **Descriptive Statistics and Balance for the TESS Experiment**

	(1)	(2)	(3)
	Population	Treatment - Control	Positive - Balanced
Individual Characteristics	Mean	Difference	Treatment Difference
Baseline Relative Willingness-to-Pay for CFL (\$)	2.9 (7.1)	0.20 (0.52)	-0.25 (0.70)
Household Income (\$000s)	70.9 (51.8)	-2.86 (3.92)	-3.79 (3.86)
Education (Years)	13.8 (2.5)	-0.04 (0.18)	0.18 (0.21)
Age	46.7 (16.9)	0.26 (1.26)	0.22 (1.45)
Male	0.48 (0.50)	-0.007 (0.035)	-0.009 (0.040)
Liberal	0.00 (1.00)	0.056 (0.076)	-0.005 (0.084)
Party	0.00 (1.00)	0.080 (0.072)	0.078 (0.082)
Environmentalism	0.30 (0.32)	-0.024 (0.023)	0.019 (0.026)
Conserve Energy	0.55 (0.50)	0.008 (0.036)	0.032 (0.041)
Homeowner	0.70 (0.46)	0.022 (0.035)	-0.012 (0.038)
F-Test p-Value		0.848	0.983

Notes: Column 1 presents means of individual characteristics in the TESS experiment population, with standard deviations in parenthesis. Column 2 presents differences in means between the Treatment and Control groups. Column 3 presents differences in means between Positive and Balanced treatment groups. Comparisons of baseline WTP and F-tests for all covariates necessarily exclude the Endline-Only group. Columns 2 and 3 have robust standard errors in parenthesis. Observations are weighted for national representativeness.

Table A.2: **Association Between Individual Characteristics and Baseline CFL Demand**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income (000s)	0.005 (0.006)						
Education (Years)	0.010 (0.148)						
Age	0.003 (0.018)						
Male	0.931 (0.533)*						
Liberal	0.091 (0.389)	0.374 (0.285)					
Party	0.573 (0.344)*		0.562 (0.266)**				
Environmentalism	0.682 (0.791)			1.429 (0.804)*			
Conserve Energy	0.970 (0.525)*				0.863 (0.545)		
Homeowner	0.047 (0.716)					0.116 (0.616)	
Present Bias β	0.281 (0.298)						0.144 (0.284)
Discount Factor δ	1.215 (0.620)*						0.962 (0.605)
R ²	0.03	0.00	0.01	0.00	0.00	0.00	0.01
N	1,163	1,226	1,229	1,221	1,219	1,229	1,178

Notes: The left-hand-side variable is baseline relative WTP for the CFL. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively. Observations are weighted for national representativeness.

Table A.3: **Equivalent Price Approach vs. Conditional Average Treatment Effects on WTP**

	(1)	(2)	(3)	(4)	(5)	(6)
Range of Relative WTP (\$/package)	Treatment Effect on Market Share	Demand Slope (Market Share/\$)	Equivalent Price Metric (\$/package)	Conditional Average Treatment Effect (\$/package)	p-value of Difference	Absolute Value of Difference (Percent of CATE)
6 to 8	0.17	0.02	7.03 (1.79)	5.35 (1.36)	0.46	31
4 to 6	0.19	0.02	7.96 (1.9)	4.05 (1.03)	0.07	96
3 to 4	0.22	0.11	1.97 (0.36)	4.11 (1.06)	0.06	52
2 to 3	0.24	0.08	3.13 (0.64)	5.14 (1.12)	0.12	39
1 to 2	0.21	0.10	2.02 (0.41)	3.82 (0.91)	0.07	47
0 to 1	0.14	0.05	2.64 (0.81)	2.39 (0.91)	0.83	11
-1 to 0	0.10	0.10	0.94 (0.33)	2.11 (0.51)	0.05	56
-2 to -1	0.05	0.03	1.46 (0.82)	2.16 (1.4)	0.66	33
-3 to -2	0.01	0.01	0.75 (1.67)	3.41 (1.77)	0.28	78

Notes: This table calculates the equivalent price metric for comparison to the conditional average treatment effect on WTP on every bounded interval in the TESS multiple price list. Standard errors are in parentheses. Standard errors in column (3) are calculated by applying the Delta method to Equation (9) using the covariance matrix from estimates of Equation (8). See Appendix B.2 for details.

Table A.4: **Perceived Intent of TESS Study**

	(1)	(2)	(3)
	Control	Balanced Treatment	Positive Treatment
<i>Do you think that the intent of the study was to . . .</i>			
Understand the effect of price changes on purchasing patterns	0.44	0.34	0.37
Measure whether people make consistent purchases in similar situations	0.31	0.25	0.26
Understand why people buy incandescents vs. CFLs	0.31	0.48	0.47
Test how well people are able to quantify energy costs	0.27	0.38	0.46
Test whether ability to quantify energy costs affects purchases of incandescents vs. CFLs	0.33	0.50	0.54
Test whether the number of bulbs in a package affects purchasing patterns	0.37	0.22	0.26
Test whether consumer education affects purchases of incandescents vs. CFLs	0.41	0.60	0.64
Understand what features of lightbulbs are most important to people	0.30	0.41	0.34
Predict the future popularity of incandescents vs. CFLs	0.30	0.34	0.34
None of the above	0.05	0.08	0.05
Number of Respondents	461	545	519

Notes: This table presents the share of consumers in each group who responded that the intent of the study was as listed in the leftmost column. Observations are weighted for national representativeness.

Table A.5: **Association of Treatment Effects with Self-Monitoring Scale**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Important to fit in	0.209 (0.396)						
Behave as others wish		0.413 (0.391)					
Good intuition for others' motives			0.312 (0.306)				
(-1)*Behavior expresses true feelings				-0.544 (0.341)			
Regulate my actions					-0.268 (0.322)		
(-1)*NOT change opinions to please someone						-0.212 (0.378)	
Self-Monitoring Mean							-0.009 (0.008)
R2	0.57	0.58	0.57	0.57	0.57	0.57	0.58
N	1,185	1,184	1,184	1,184	1,184	1,184	1,188

Notes: This table presents estimates of Equation (6) with the addition of Self-Monitoring Scale variables and the interaction of these variables with the Treatment indicator. The coefficients presented are these interaction terms. The outcome variable is endline willingness-to-pay for the CFL. Robust standard errors in parenthesis. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively. Observations are weighted for national representativeness.

Table A.6: **Effects on Important Factors in Purchase Decision**

	(1)	(2)	(3)	(4)	(5)
	Price	Energy Use	Bulb Lifetime	Warm-Up Time	Mercury and Disposal
1(Balanced Treatment)	-0.86 (0.21)***	0.15 (0.21)	0.02 (0.20)	-0.94 (0.24)***	-0.29 (0.25)
1(Positive Treatment)	-0.55 (0.22)**	0.20 (0.21)	0.25 (0.19)	0.04 (0.23)	-0.09 (0.24)
Constant	7.75 (0.13)***	7.43 (0.15)***	7.76 (0.13)***	5.41 (0.17)***	6.03 (0.18)***
R2	0.02	0.00	0.00	0.02	0.00
N	1,533	1,478	1,512	1,506	1,518

Notes: This table reports treatment effects on self-reported importance of different factors in purchase decisions. Robust standard errors in parenthesis. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively. Observations are weighted for national representativeness.

Table A.7: **Additional Estimates of TESS Average Treatment Effects**

	(1)	(2)	(3)	(4)
1(Treatment)	2.39 (0.38)***	3.07 (0.50)***	1.98 (0.30)***	2.83 (0.48)***
1(Treatment & Fail Info Quiz)	-1.06 (1.24)			
1(Treatment & Incorrect Survey Beliefs)		-1.17 (0.56)**		
R2	0.58	0.58	0.58	0.57
N	1,188	1,188	1,188	1,188
Assumed Censored Mean WTP	15	15	12	20

Notes: This table presents alternative estimates of Equation (6). The outcome variable is endline willingness-to-pay for the CFL. 1(Treatment) pools all information sub-treatments. Robust standard errors in parenthesis. *, **, ***: Statistically significant with 90, 95, and 99 percent confidence, respectively. Observations are weighted for national representativeness.

C iPad Total Cost Comparison Screen

Notes: This is the information screen presented to Treatment group consumers in the in-store experiment. Numbers in this screen shot represent a consumer buying one CFL at typical purchase prices and national average electricity prices.

D Appendix to Theoretical Framework

D.1 Proof of Proposition 1

By definition, the average true value of consumers marginal at price $p = c - s$ is given by

$$E_G(v|v - b = p) = E_G(v - b|v - b = p) + E_G(b|v - b = p) = p + B(p).$$

The social cost of transferring an additional unit of product E to a consumer is c . Thus the social efficiency gain from inducing all marginal consumers to purchase E instead of I is given by $p + B(p) - c = B(p) - s$. The total change in demand is $-D'_B(p)$ by definition, and thus the total change in welfare is $W'(s) = (s - B(p))D'_B(p)$.

Now as in Harberger (1964), we have that under the approximation $D''_B(p) \approx 0$,

$$\int_{s-\Delta s}^s x D'_B(c-x) dx \approx (\Delta s) s D'_B(p) + \frac{(\Delta s)^2}{2} D'_B(p),$$

which yields equation (2) by integrating $W'(x)$ from $s - \Delta s$ to s .

D.2 Comparing the equivalent price metric to the average marginal bias

D.2.1 A two-type example

Suppose that conditional on a perceived value of $\hat{v}$, there are only two possible values of bias: a *high* value $\tau_H(\hat{v})$ and a *low* value $\tau_L(\hat{v}) \leq \tau_H(\hat{v})$. That is, if a consumer's perceived valuation is $\hat{v}$ then his true valuation is either $\hat{v} + \tau_L(\hat{v})$ or $\hat{v} + \tau_H(\hat{v})$. Assume that τ_L and τ_H are differentiable and let $D_{B,L}$ and $D_{B,H}$ correspond to the demand curves of agents corresponding to the low and high values, respectively. Let $D_{N,L}$ and $D_{N,H}$ correspond to the demand curves that would be obtained if these consumers were debiased. That is, $D_{N,H}$ is the demand curve that would be obtained if all type H consumers were debiased. The relationship between D_N and D_B is now given simply by $D_{B,k}(p) = D_{N,k}(p + \tau_k(p))$ for $k = L, H$.

By definition,

$$B(p) = \frac{D'_{B,L}(p)\tau_L(p) + D'_{B,H}(p)\tau_H(p)}{D'_B(p)},$$

while to a first order approximation,

$$\begin{aligned} EPM(p) &= \frac{D_B(p) - D_N(p)}{D'_N(p)} = \frac{D_{B,L}(p) - D_{N,L}(p) + D_{B,H}(p) - D_{N,H}(p)}{D'_N(p)} \\ &= \frac{D_{N,L}(p + \tau_L(p)) - D_{N,L}(p) + D_{N,H}(p + \tau_H(p)) - D_{N,H}(p)}{D'_N(p)} \\ &\approx \frac{D'_{N,L}(p)\tau_L(p) + D'_{N,H}(p)\tau_H(p)}{D'_N(p)}. \end{aligned}$$

It thus follows that the EPM correctly approximates $B(p)$ if and only if

$$\frac{D'_{B,L}(p)}{D'_{N,L}(p)} = \frac{D'_{B,H}(p)}{D'_{N,H}(p)} \tag{10}$$

Note, however, that $D'_{B,j}(p) = (1 + \tau'_j(p))D'_{N,k}(p)$, and so it is not generally true that the condition in Equation (10) holds. The condition holds if $\tau'_L(p) = \tau'_H(p)$.³⁰ When consumers with higher bias are relatively more likely to be on the margin in their biased state ($D'_{B,H}/D'_{N,H} > D'_{B,L}/D'_{N,L}$ because $\tau'_H > \tau'_L$), the EPM will underestimate $B(p)$. Conversely, when consumers with higher bias are relatively less likely to be on the margin, the EPM will overestimate $B(p)$. This same principle is exactly what underlies our example with normal distributions in Section 2: consumers with high b are relatively less likely to be on the margin at high prices p , and relatively more likely to be on the margin at low prices p .

Notice also that having D'_B instead of D'_N in the denominator of the EPM does not expand the set of cases in which the EPM is a first-order approximation to $B(p)$. In fact, having D'_B instead of D'_N in the denominator narrows the set of cases in which the EPM is a first-order approximation to $B(p)$. When $\tau_L(p) = \tau_H(p)$ for all p , so that bias is homogeneous, the EPM with D'_N in the denominator constitutes a first-order approximation to $B(p)$. However, moving $D'_B \neq D'_N$ in the denominator will introduce a bias in the approximation.

D.2.2 Conditions for exact equivalence

With a slight abuse of notation, we now let $F(\cdot|b)$ denote the CDF of v conditional on a value of b and we let $G(\cdot)$ denote the unconditional CDF of b . To ease notation and exposition, we will restrict here to the case where $F(\cdot|b)$ and G continuously differentiable, with respective density functions $f(\cdot|b)$ and g . The argument for finite or mixture distributions would follow almost identically.

Proposition 2 *Assume that $G(b|v, p)$ does not depend on p . Then for all generic distributions G , $EPM(p) = B(p)$ if and only if $\frac{f(p+b|p)}{D'_B} = \frac{F(p+b|b) - F(p|b)}{bD'_N}$.*

The condition in the proof is slightly more general than the condition that $f(v|b)$ is linear on $v \in [p, p+b]$. And the linearity condition on $f(v|b)$ is roughly equivalent to requiring that D_N is linear and that b is independent of v in a neighborhood of $v = p + b$.

Proof.

By definition,

$$B(p) = \frac{\int b f(p+b|b) g(b) db}{\int f(p+b|b) g(b) db}$$

and

$$EPM(p) = \frac{(F(p+b|p) - F(p|b))g(b)}{\int f(p|b)g(b)db}.$$

Now $D'_B = \int f(p+b|b)g(b)db$ and $D'_N = \int f(p|b)g(b)db$ and thus comparing the equations for B and EPM shows that the only way these two equations will hold for all generic density weights $g(b)$ is if the condition in the Proposition holds for all b .

■

D.2.3 Quantifying possible deviations between the equivalent price metric and the average marginal bias

Here we will restrict to a simpler scenario in which we can partition consumers into finitely many types θ such that a consumer with WTP $\hat{v}$ has a true value of $\hat{v} + \tau_\theta(\hat{v})$. Clearly, the divergence between $B(p)$

³⁰This is a special condition generalizing the homogeneous bias assumption under which Mullainathan, Schwartzstein, and Congdon (2012) show that the EPM correctly approximates $B(p)$.

and $EPM(p)$ can only be higher in the slightly more general set up. We now show that even under fairly restrictive regularity conditions, the difference between $B(p)$ and $EPM(p)$ can be quite large.

Proposition 3 *Suppose that $\tau_\theta(p) \in [\underline{\tau}, \bar{\tau}]$. For $p_1 < p_2$, suppose $D_N(p_1)$, $D_N(p_2)$, $D_B(p_1)$, $D_B(p_2)$ are all measured, and that the following additional restrictions are known to apply:*

1. *For each pair θ_1, θ_2 , $D_{N,\theta_1}(p)/D_{N,\theta_2}(p)$ is constant on $[p_1, p_2]$.*
2. *$D_{B,\theta}$ and $D_{N,\theta}$ are linear on $[p_1, p_2]$ for all θ .*

Then $((B(p_1) + B(p_2))/2 - (EPM(p_1) + EPM(p_2))/2)$, the difference between the average B and the average EPM over the range $[p_1, p_2]$, can be as high as:

$$\frac{D'_N}{2D'_B} \min [\bar{\tau} - EPM(p_2), -D_B(p_2)/D'_N - EPM(p_2)] \cdot \max \left[\frac{D'_B}{D'_N}, \frac{EPM(p_1) - \underline{\tau}}{p_2 - p_1} \right] \quad (11)$$

and as low as:

$$-\frac{D'_N}{2D'_B} (EPM(p_1) - \underline{\tau}) \max \left[\frac{D'_B}{D'_N}, \frac{\bar{\tau} - EPM(p_1)}{p_2 - p_1} \right] \quad (12)$$

Proposition 3 shows that $B(p)$ and $EPM(p)$ can differ significantly even under the following restrictions: demand curves are linear, treatment effects are uncorrelated with the slopes of the unbiased demand curves $D_{N,\theta}$, and treatment effects are restricted to lie in a reasonably narrow range. A consequence of Proposition 3 is that even when information has no effect on demand, the conditional average treatment effect on WTP can still be substantial. Suppose, for example, that 40 percent of consumers purchase E at baseline prices, that a \$1 subsidy move demand by 10 percentage points, that information provision has no effect on demand at subsidized and unsubsidized prices. This roughly corresponds to the estimates in the in-store experiment. Finally, suppose that the treatment effect on any one consumer's WTP cannot be any greater than \$5 or smaller than -\$5. Proposition 3 then shows that despite there being absolutely no treatment effect on quantity purchased, and despite the very restrictive regularity conditions, the CATE on WTP for consumers who are marginal to the \$1 subsidy can still be as high as \$2 and as low as -\$2.50. This generalizes our normal distributions example in Section 2 to the case of biases with quite narrow support and more restrictive distributional assumptions.

It is important to note that Proposition 3 does not compute the maximum possible divergence between $B(p)$ and the EPM. Rather, it only provides a sense of how much the two can diverge. As of now, we have not yet been able to compute the maximum possible divergence between these two measures, though we conjecture that the quantities in Proposition 3 do, indeed, correspond to the maximum possible divergence.

D.2.4 Proof of Proposition 3

Consider two types, $\theta = 1, 2$ and let $\gamma = D'_{N,1}(p)/D'_{N,2}(p)$. Note that restriction 2 guarantees that $D'_{N,1}(p)/D'_{N,2}(p)$ is constant in p .

What we want to do is maximize

$$B(p) = \frac{1}{D'_B(p)} (\gamma \tau_1(p) D'_{B,1}(p) + (1 - \gamma) \tau_2(p) D'_{B,2}(p)) \quad (13)$$

$$= \frac{D'_N(p)}{D'_B(p)} (\gamma \tau'_1(p) \tau_1(p) + (1 - \gamma) \tau'_2(p) \tau_2(p)) \quad (14)$$

where the second equality is also an implication of restrictions 1 and 2 in the proposition.

Part 1: Maximizing $B(p) - EPM(p)$

Now set $\tau_1(p_1) = \tau_2(p_1) = EPM(p_1)$ and let $\tau'_1 = m_1$ and $\tau'_2 = m_2$, with $m_1 > m_2$. By definition, we must have

$$\gamma\tau(p) + (1 - \gamma)\tau_2(p) = EPM(p)$$

from which it follows that

$$\gamma m_1 + (1 - \gamma)m_2 = \frac{D'_B - D'_N}{D'_N} \equiv \tilde{m}$$

and thus

$$\gamma = \frac{\tilde{m} - m_2}{m_1 - m_2}.$$

We now have

$$\begin{aligned} B(p) - EPM(p) &= [\gamma m_1^2 + (1 - \gamma)m_2^2 - \tilde{m}^2] (p - p_1) \\ &= [\gamma(m_1^2 - m_2^2) + m_2^2 - \tilde{m}^2] (p - p_1) \\ &= [(m_1 + m_2)(\tilde{m} - m_2) + m_2^2 - \tilde{m}^2] (p - p_1) \\ &= [(m_1 + m_2)\tilde{m} - m_1 m_2 - \tilde{m}^2] (p - p_1) \\ &= (m_1 - \tilde{m})(\tilde{m} - m_2)(p - p_1). \end{aligned} \tag{15}$$

The last equation above shows that to maximize $B(p) - EPM(p)$ optimal to have m_1 as high as possible and m_2 as low as possible.³¹ So the final step is to determine these bounds for m_1 and m_2 .

To this end, note that we must have $EPM(p_1) + m_1(p_2 - p_1) \leq \bar{\tau}$. This implies that $m_1 \leq \frac{\bar{\tau} - EPM(p_1)}{p_2 - p_1}$.

Additionally, since we are requiring linearity of demand curves in the region of interest, we must have $D_N(p_2 + \tau_1(p_2)) \geq 0$; or $D_N(p_2) + D'_N \tau_1(p_2) \geq 0$. From which it follows that $\tau_1(p_2) \leq -D_N(p_2)/D'_N$. This similarly implies that $m_1 \leq \frac{-D_N(p_2)/D'_N - EPM(p_1)}{p_2 - p_1}$.

Similarly, $EPM(p_1) + m_2(p_2 - p_1) \geq \underline{\tau}$, and thus $m_2 \geq \frac{\underline{\tau} - EPM(p_1)}{p_2 - p_1}$. Additionally, demand must be downward sloping, which implies that $m_2 > -1$.

Altogether, we thus want to set

$$\begin{aligned} m_1 &= \frac{\min(\bar{\tau}, -D_N(p_2)/D'_N) - EPM(p_1)}{p_2 - p_1} \\ m_2 &= \max(-1, \frac{\underline{\tau} - EPM(p_1)}{p_2 - p_1}) \end{aligned}$$

But now $EPM(p_2) = \tilde{m}(p_2 - p_1) + EPM(p_1)$, from which it follows that

$$m_1 - \tilde{m} = \frac{\min(\bar{\tau}, -D_N(p_2)/D'_N) - EPM(p_2)}{p_2 - p_1}$$

The first part of the proposition now follows by combining (15) with the bounds we computed for m_1 and m_2 .

Part 2: Minimizing $B(p) - EPM(p)$

This other part follows analogously. Set $\tau_1(p_2) = \tau_2(p) = EPM(p_2)$. Then as in (15) we analogously get

³¹Just consider a perturbation where m_1 is increased by ϵ and m_2 is decreased by $\beta/(1 - \beta)\epsilon$

that

$$B(p) - EPM(p) = -(m_1 - \tilde{m})(\tilde{m} - m_2)(p - p_1). \quad (16)$$

Again, for $m_1 > m_2$ we similarly want m_1 as high as possible and m_2 as low as possible; that is, we want the low bias types to be the ones who are most elastic.

Analogously to the preceding computations, we must have $EPM(p_2) - m_2(p_2 - p_1) \leq \bar{\tau}$ from which it follows that $m_2 \geq -\frac{\bar{\tau} - EPM(p_2)}{p_2 - p_1}$. And as before, $m_2 > -1$ to generate downward-sloping demands.

Similarly, $EPM(p_2) - m_1(p_2 - p_1) \geq \underline{\tau}$ from which it follows that $m_1 \leq \frac{EPM(p_2) - \underline{\tau}}{p_2 - p_1}$.

D.3 A more general framework and the comparing demand responses approach

To analyze other strategies for quantifying consumer bias, we now propose a more general framework that will allow us to compare product subsidies to energy taxes. The framework will formally encompass an extension of our energy efficiency model, as well as the salience models of Chetty, Looney, and Kroft (2009) and Goldin and Homonoff (2013). The framework is also applicable to other situations in which there may be opaque attributes, such as the work by Hossain and Morgan (2006) on shipping charges, the work by Abaluck and Gruber (2011) on out-of-pocket insurance costs, and the work by Lacetera, Pope, and Sydnor (2013) on left-digit bias.

As before, we will continue working with the somewhat simpler setup in which we can partition consumers into finitely many types θ , where each type is a correspondence between true and perceived valuations. True demand curves are given by $D_{N,\theta}(p)$, while biased demand curves are given by $D_{B,\theta}(p) = D_{N,\theta}(p + \tau_\theta(t_1, t_2))$ where $p = c + t_1 + t_2$ is the total price after taxes t_1 and t_2 . We can think of t_1 as the tax on product prices and t_2 as the tax on energy costs, or we can think of t_1 as the tax on producers that passes through to price and t_2 as a tax on consumers.

In this more general framework, the welfare impact of increasing t_i depends on the marginal bias with respect to t_i :

$$B_i(t_1, t_2) = \sum \zeta_{D,\theta}^{t_i}(t_1, t_2) \tau_\theta(t_1, t_2)$$

where, $\zeta_{D,\theta}^{t_i} = \frac{D_{B,\theta}^{t_i}(t_1, t_2)}{\sum_\theta D_{B,\theta}^{t_i}(t_1, t_2)}$ is the portion of consumers who are type θ out of all those consumers who respond to a marginal increase in the tax t_i ; and $D_{B,\theta}^{t_i}(t_1, t_2)$ denotes the derivative of $D_{B,\theta}$ with respect to t_i evaluated at (t_1, t_2) . As in Proposition 1,

$$\frac{d}{dt_i} W(t_1, t_2) = (p - c + B_i(p)) D_B^{t_i}(t_1, t_2)$$

where $D_B^{t_i}$ is the derivative of D_B with respect to t_i .

Note that when consumers are debiased, changes in t_1 and t_2 should have identical impacts on demand. Let D'_N denote the derivative of D_N with respect to either t_1 or t_2 , and analogously to Section 2, define

$$EPM(t_1, t_2) = \frac{D_B(t_1, t_2) - D_N(t_1, t_2)}{D'_N}.$$

D.3.1 The EPM and the average marginal bias in a more general setting

As before, the EPM and the average marginal bias will typically not be the same. Notice also that as long as $B_1 \neq B_2$, it will always be true that the EPM is an imperfect approximation to either B_1 or B_2 (and possibly

both). And as we briefly show below, $B_1 \neq B_2$ for very simple and natural examples of this framework.³² Fleshing out this fact, we now show how the EPM can fail to provide the necessary sufficient statistic even for very simple models of bias, as long as bias is heterogeneous.

First, consider the simple case in which type θ consumers underweight energy costs by θ , and consider their choice of the relatively less efficient appliance I . Suppose that $c = c_1 + c_2$, where c_1 is the relative cost of producing the energy using appliance and c_2 is the relative energy cost associated with utilization. Then $\tau_\theta(t_1, t_2) = -(1 - \theta)(c_2 + t_2)$. Notice that in this framework, $D_{N,\theta}^{t_1}(t_1, t_2) = D_{B,\theta}^{t_1}(t_1, t_2)$ and therefore $EPM \approx B_1$.

At the same time, $D_{B,\theta}^{t_2}(t_1, t_2) = \theta D_{N,\theta}^{t_2}(t_1, t_2)$, and thus $EPM \neq B_2$. To see this concretely, note that to a first order approximation,

$$EPM \approx \frac{\sum_{\theta} D'_{N,\theta} \tau_{\theta}}{\sum_{\theta} D'_{N,\theta}} \quad (17)$$

whereas

$$B_2 = \frac{\sum_{\theta} D_{B,\theta}^{t_2} \tau_{\theta}}{\sum_{\theta} D_{B,\theta}^{t_2}} = \frac{\sum_{\theta} \theta D'_{N,\theta} \tau_{\theta}}{\sum_{\theta} D'_{N,\theta}} \quad (18)$$

In fact, it is easy to see that as long as θ is heterogeneous, $|EPM| > |B_2|$ because the most biased types are also the least elastic to the energy tax.³³ Assuming homogenous bias and then using the EPM to calculate a sufficient statistic for the optimal energy tax would thus produce a number that is too high.

The implications of the EPM for the tax salience frameworks of Chetty, Looney, and Kroft (2009) and Goldin and Homonoff (2013) are identical. To obtain these frameworks, let t_2 be the tax on consumers and let $\tau_{\theta}(t_1, t_2) = (1 - \theta)t_2$. Then proceed as above for the energy cost salience framework. Again, the conclusion here will be that in the presence of heterogeneity, $|EPM| > |B_2|$. Thus assuming a representative agent framework and using the EPM as a sufficient statistic for B_2 would lead one to underestimate the excess burden of increasing t_2 .

D.3.2 The comparing demand responses approach with homogeneous consumer bias

Another common approach for measuring bias is the “comparing demand responses” approach, as summarized by DellaVigna (2009), and used by Allcott and Wozny (2012) and Busse, Knittel, and Zettelmeyer (2013) to study undervaluation of energy costs and by Chetty, Looney, and Kroft (2009) to study tax salience. Variations of this approach have also been used by Hossain and Morgan (2006), Abaluck and Gruber (2011), and Lacetera, Pope, and Sydnor (2014).

The idea of this approach is as follows: Suppose consumers underweight future energy costs. Then they should also react less to changes in energy costs than to changes in the upfront prices of energy-using appliances. The approach then is to use the ratio $\frac{D_B^{t_2}(t_1, t_2)}{D_B^{t_1}(t_1, t_2)}$ as a measure of bias.

³²This issue of $B_1 \neq B_2$ is also central to Allcott, Mullainathan and Taubinsky’s (2014) derivation of the optimal combination of product and energy taxes. The crucial quantity for determining the optimal subsidy in that framework is $B_1 - B_2$.

³³Alternative versions of the EPM that place $D_B^{t_1}$ or $D_B^{t_2}$ in the denominator do not do any better here either. The simply scale up the EPM by the ratio $\frac{D'_N}{D'_B}$, which does not fix the fact that (17) and (18) are two very different weighted averages.

With a homogeneous consumer who underweights energy costs by θ , the comparing elasticity approach is simple. Because $D_B^{t_1}(t_1, t_2) = \theta D_N^{t_2}(t_1, t_2)$, we have that $\theta = \frac{D_B^{t_2}(t_1, t_2)}{D_B^{t_1}(t_1, t_2)}$.

Notice, however, that the success of this derivation depends on the fact that the bias $b = -(1 - \theta)(c_2 + t_2)$ is linear in t_2 and constant in t_1 . More generally, $\frac{D_B^{t_2}(t_1, t_2)}{D_B^{t_1}(t_1, t_2)} = \frac{1 + \frac{d}{dt_2}\tau}{1 + \frac{d}{dt_1}\tau}$, a quantity that without a number of additional structural assumptions is not generically related to the representative consumer's bias b . Depending on how b changes with t_1 or t_2 , the comparing demand responses approach will either over- or under-estimate the level of the bias. This is in contrast to the EPM, which always provides an accurate first order approximation when bias is homogenous.

In practice, $\frac{d}{dt_1}\tau$ is likely to be non-zero and $\frac{d}{dt_1}\tau$ may not be constant for several reasons. The first reason is that attention is likely to be endogenous. For example, the higher the energy tax, the more likely are consumers to pay attention it. Or a higher product tax might make a consumer more likely to consider the alternative seriously. The second reason is that bias is likely to depend on true valuations. Depending on the model of attention, the bias could be increasing or decreasing with the true valuation. Consumers who value the energy efficient product the most, for example, are the most likely to be “Green” consumers who are very attentive to energy costs.

D.3.3 The comparing demand responses approach with heterogeneous consumer bias

So far, we've shown that when consumer bias is homogeneous, the success of the comparing demand responses approach depends critically on the elasticity of consumer bias with respect to t_1 and t_2 . This is contrast to the EPM which (under the full debiasing assumption) always provides an accurate first-order approximation to consumer bias. How does the comparing demand responses approach compare to the EPM or $B(p)$ under heterogeneity?

Consider again the simple setting in which consumers underweight energy-related costs by θ , so that $\tau_\theta(t_1, t_2) = -(1 - \theta)(c_2 + t_2)$. Then the comparing demand responses approach gives

$$\begin{aligned} \rho := \frac{D_B^{t_2}}{D_B^{t_1}} &= \frac{\sum_\theta D_{B,\theta}^{t_2}}{\sum_\theta D_{B,\theta}^{t_1}} \\ &= \frac{\sum_\theta \theta D_{B,\theta}^{t_1}}{\sum_\theta D_{B,\theta}^{t_1}} \\ &= \frac{\sum_\theta D_{N,\theta}^{t_2} \tau_\theta}{\sum_\theta D_{N,\theta}^{t_2}} + (c_2 + t_2) \\ &\approx EPM + c_2 + t_2 \end{aligned}$$

where the last line is a consequence of (17). Thus to a first order approximation, the comparing demand responses approach produces the same quantity as the EPM. The analysis in section D.2 thus shows that $\rho = B_1 \neq B_2$; that is, just like the EPM, the comparing demand responses approach can produce a consistent estimate of B_1 but not of B_2 .

However, our TESS experiment suggests that the conditions under which the EPM produces an accurate estimate of B_1 do not hold, which suggests that $\frac{d}{dt_1}\tau_\theta \neq 0$. Under these kinds of more general conditions, the comparing demand responses approach cannot provide an accurate estimate of B_1 either. Indeed, more generally we have that

$$\begin{aligned}
\rho := \frac{D_B^{t_2}(p)}{D_B^{t_1}(p)} &= \frac{\sum_{\theta} D_{B,\theta}^{t_2}(p)}{\sum_{\theta} D_{B,\theta}^{t_1}(p)} \\
&= \frac{\sum_{\theta} (1 + \frac{d}{dt_2} \tau_{\theta}) D_{N,\theta}^{t_2}(p + \tau_{\theta})}{\sum_{\theta} (1 + \frac{d}{dt_1} \tau_{\theta}) D_{N,\theta}^{t_1}(p + \tau_{\theta})}
\end{aligned}$$

Under these more general conditions, ρ is the ratio of two weighted averages of “bias response” functions $1 + \frac{d}{dt_i} \tau_{\theta}$. Without various special assumptions, this difficult to interpret quantity is not closely related to the EPM, B_1 , or B_2 . Again, $\frac{d}{dt_i} \tau_{\theta}$ is likely to be non-zero and not even constant because of endogenous attention or because bias depends on true valuations.

References

- [1] Goldin, Jacob and Tatiana Homonoff (2013). “Smoke Gets In Your Eyes: Cigarette Tax Salience And Regressivity.” *American Economic Journal: Economic Policy*, Vol. 5, No. 1 pages 302-336.
- [2] Lacetera, Nicola, Devin Pope, and Justin Sydnor (2013). “Heuristic Thinking and Limited Attention in the Car Market.” *American Economic Review*, Vol. 102, No. 5 pages 2206-2236.