

Goal Setting and Energy Conservation

By

Matthew Harding

and

Alice Hsiaw

August 2014

COLLEGE OF THE HOLY CROSS, DEPARTMENT OF ECONOMICS
FACULTY RESEARCH SERIES, PAPER NO. 14-03*



Department of Economics
College of the Holy Cross
Box 45A
Worcester, Massachusetts 01610
(508) 793-3362 (phone)
(508) 793-3708 (fax)

<http://www.holycross.edu/departments/economics/website>

*All papers in the Holy Cross Working Paper Series should be considered draft versions subject to future revision. Comments and suggestions are welcome.

Goal Setting and Energy Conservation

By
Matthew Harding[†]
Stanford University

and

Alice Hsiaw^{††}
College of the Holy Cross

August 2014

Abstract

This paper develops a theoretical model of consumer demand for an energy conservation program that involves non-binding, self-set goals. We present evidence from a Northern Illinois goal-setting program, aimed at reducing residential electricity consumption, that is difficult to reconcile with standard preferences and is broadly consistent with a model of present-biased consumers with reference-dependent preferences. We find that the need for commitment is correlated with program adoption, higher pre-adoption consumption, and lower responsiveness to goals. Consumers choosing realistic goals persistently save substantially more, achieving savings of nearly 11%, than those choosing very low or unrealistically high goals..

Acknowledgements: We thank the Precourt Energy Efficiency Center for generous financial support. We are grateful to C3 staff for sharing the data and in particular we thank Andy Frank, Zeke Hausfather and Patrick McNamara for numerous conversations and assistance with data preparation. We thank Ana Meyer and Marcel Priebsch for excellent research assistance. We thank Alisdair McKay, Matt Kotchen, and seminar participants at U.C. Berkeley, Harvard, Stanford, College of the Holy Cross, and University of Tennessee, Knoxville for useful comments.

JEL Classification Codes:

Keywords:

[†]Matthew Harding, Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305; Phone: (650) 723-4116; Email: mch@stanford.edu

^{††}Alice Hsiaw, Department of Economics and Accounting, College of the Holy Cross, 1 College Street, Worcester, MA 01610; Phone: (508) 793-2687; Email: ahsiaaw@holycross.edu

1 Introduction

As the impact of climate change becomes increasingly stark, policymakers have renewed interest in implementing interventions that can improve energy conservation. In light of general federal inaction on implementing price-based interventions such as introducing a carbon tax or allowing emissions trading in the foreseeable future, economists have begun to advocate the use of behavioral “nudges,” which are non-price interventions grounded in psychology and behavioral economics, as a partial solution to attenuate the impact of climate change (Allcott and Mullainathan, 2010). The hope is that such approaches can generate economically meaningful energy savings at low cost, mirroring like interventions that have been documented in other areas of economics, such as microfinance and retirement savings (Bertrand et al., 2010; Madrian and Shea, 2001).

In recent years, utilities and private enterprises have deployed energy efficiency and conservation programs centered around behavioral nudges to millions of customers nationwide.¹ They typically provide consumers with information on energy saving actions and additional behavioral nudges, such as social comparisons. Lately, many programs have also offered consumers the option to choose a personal energy conservation goal. However, the overall success of such programs presents a puzzle, since there is no established economic framework that explains both individuals’ motivations to voluntarily participate in a program to save energy and the underlying mechanisms that affect consumer behavior in the absence of material incentives.

This is the first paper to address both of these questions, by studying consumers’ decisions to join an energy conservation program and their subsequent behavior upon joining. We present theory and evidence on adoption and post-adoption behavior in the context of a multi-part energy conservation program that included non-binding personal goal-setting, intended to reduce residential electricity use. We develop a model in which consumers sign up for the program because they are aware of their dynamic inconsistency, and thus their tendency to overconsume electricity, and respond to their self-set goals because they have reference-dependent preferences. We find substantial and persistent energy savings of 11% for informed customers who set realistic goals, and approximately 4.4% on average for all customers. This paper is unique, in that it models and finds empirical evidence that is broadly consistent with a precise behavioral mechanism that explains both the decision to opt into the program as well subsequent behavior once enrolled.

This paper contributes to a number of different literatures in both energy and behavioral economics. First, it contributes to the growing literature on the impact of behavioral, non-price interventions on energy consumption. Several studies suggest that utilities’ interventions that include social comparisons to others’ usage or social norms can meaningfully affect consumers’ behav-

¹One of the largest companies, OPower, works with over 80 utilities nationwide and claims over 2 Terawatt hours of electricity saved as a result of its behavioral programs.

ior (Allcott, 2011; Ferraro and Price, 2013; Costa and Kahn, 2010; Nolan et al., 2008; Bolsen et al., 2013). Harding and Rapson (2014) show that a utility's carbon offset program can lead to consumption reductions of 2-3% when the social cost of pollution is emphasized. By exploiting variation from the field, this study substantially enhances previous studies on the effects of personal goal-setting, which have found little effect of goals relative to energy-saving tips, due to rather unclear designs (Wilhite and Ling, 1995; van Houwelingen and van Raaij, 1989). Second, it contributes to the behavioral economics literature on non-binding goals. These have been explored extensively by psychologists (Locke and Latham, 2002; Heath et al., 1999), but they have only recently been studied by economists. Kőszegi and Rabin (2006; 2009) model reference dependence by endogenizing reference point formation through rational expectations, which has been supported by experimental laboratory (Abeler et al., 2011; Ericson and Fuster, 2011; Sprenger, 2010) and field (Fehr and Goette, 2007; Goette and Huffman, 2007) evidence. Hsiaw (2013) shows that goals, in the form of targets, must be sufficiently realistic for reference-dependent agents to counteract present-biasedness. This paper adds to the growing body of empirical evidence of reference-dependent preferences (Camerer et al., 1997; Crawford and Meng, 2011; Pope and Schweitzer, 2011; Card and Dahl, 2011) in a new context, with goals that act as endogenous and measured reference points. Moreover, to our knowledge, it is the first empirical study in which the *motivated* use of reference dependence is demonstrated, to counteract present bias.

The paper proceeds as follows. First, we describe the design and implementation of a multi-part program in Northern Illinois that allows utility customers to choose energy savings goals. It also provides them with information on how to implement energy-conserving actions to save energy, and constant feedback on their performance.

Second, we introduce a theoretical model of sign-up and post-adoption usage in the program. A consumer allocates consumption between an aggregate good and electricity, a good that yields immediate benefits but a delayed (environmental) cost. If she is present-biased, she may sign up for the program because she is aware that she will consume more electricity than she prefers *ex ante*. If she has reference-dependent preferences, she curbs her consumption in response to a goal, which acts as her reference point, because she derives utility directly from comparing her consumption against this goal. Thus, the consumer's present-biased preferences induce her to sign up for the goal-setting program, while her reference-dependent preferences lead her to curb her consumption in response to goals. We compare this behavioral model's predictions against those of a standard expected-utility model, and later discuss variants of them. Leading alternative mechanisms based on standard preferences to explain the observed savings and sign-up patterns, such as information-seeking and selection, are not consistent with the observed results.

Third, we explore consumers' decisions to opt into the program and the choice of goals. We find that the program is attractive to consumers who are present-biased and are aware of their

need for commitment to attenuate overconsumption. The choice of goals reveals substantial heterogeneity in consumers. While 73% of consumers choose realistic or optimistic goals, 15% of consumers choose the minimum goal possible. About 12% of consumers choose impossibly high goals. Overall, average savings over the year post-adoption are approximately 4.4%, but up to 8% savings are achieved in the first few months of the program. For well-calibrated consumers who choose realistic goals, we find substantial and persistent savings of nearly 11%, but no savings among those who set very low or unrealistically high goals. While it is not possible to derive a causal claim regarding the impact of goal setting on savings in the context of a non-experimental program implementation, our analysis shows that households who are aware of their present bias and desire to save electricity will select into the program and tend to save in accordance with their chosen goals when such goals are realistic.

We conclude by discussing leading alternative mechanisms, such as information seeking and selection, which are difficult to reconcile with our findings.

2 CUB Energy Saver Program

We now describe a concrete example of a behavioral program, where the primary user experience was focused on selecting and acting on an individually chosen energy conservation goal. We collected data from a pilot program which was designed by Efficiency 2.0 (recently acquired by C3), a provider of energy efficiency programs to utilities, and was funded by the Citizens Utility Board (CUB) of Illinois.² The program was restricted to residential users, all of whom were customers of ComEd, a large utility. It was advertised to a limited subset of ComEd customers through direct mail and community events. As a result, program adoption is heavily concentrated in the Chicago Metropolitan Area. We restrict our analysis to adopters during the first year of the program, starting with its introduction in June 2010, a period over which program implementation remained constant. The CUB Energy Saver program involves a website designed by C3, which can be accessed from the main CUB webpage. It provides an integrated user experience, separate from the utility website, with a focus on detailed information coupled with behavioral incentives.

Signing up for the program involves successfully completing three steps. First, a user must fill out basic contact information. Second, the user must provide utility account information, which allows the program to access past and future billing information. Third, the user is offered a menu of energy savings options corresponding to roughly 5%, 10% and 15% annual electricity savings, which are labeled as “No Cost,” “Low Cost” and “Home Investment” plans, reflecting the extent

²CUB was set up in 1983 by the Illinois General Assembly as a nonprofit, nonpartisan organisation designed to represent the interests of residential utility consumers.

to which a household may need to make further purchases (e.g. energy efficient appliances) to save energy. Each savings plan comes with a concrete list of energy savings recommendations, all of which the user can see *before* selecting a goal.³ Users can add additional energy savings actions from a long list of possible actions, effectively selecting a goal on a continuous savings range. Each action can be customized for a household’s specific technology base (for example, by entering information on exactly how many lightbulbs will be replaced with CFLs and the number of hours they will be used per day), providing a more accurate savings estimate.

Once a consumer has signed up for the website, she can easily monitor her monthly consumption, since the website obtains and displays all past monthly electricity bills directly from the utility. On average consumers log in to the website once every 2-3 months but a significant fraction of the consumers appears to be very involved, logging in at least once a month. The website also rewards customers with “points” computed by an algorithm that adjusts energy use for weather and seasonality. This algorithm is not available to the user, and points are *not* based on consumption relative to self-set goals. These points can then be redeemed for coupons offering discounts at local retail establishments. The point system was implemented by the company to promote engagement with the website, and our analysis indicates that these points are only weakly correlated to actual savings. Point redemption is also extremely low and a survey indicates that points are not a major incentive for program adoption. Thus, we believe that the unaltered billing data is the most useful information available to households in order to monitor their consumption goals.

3 Theoretical Framework

We model a consumer’s enrollment decision and her subsequent behavior as a two-stage optimization problem. In the first stage, she decides whether or not to sign up for the program. In the second stage, she chooses her monthly level of electricity consumption. If she has not signed up, then she receives no goal and consumes electricity accordingly. If she has signed up, then she receives a goal, in the form of a consumption target, and consumes electricity accordingly. Derivations and proofs are collected in the Appendix.

The model has four key features: (1) the consumption problem is a constrained allocation decision between electricity and an aggregate good, (2) electricity is a good that yields immediate benefits and delayed costs, (3) the consumer is present-biased, so she tends to overconsume electricity in the absence of intervention, and (4) the consumer has reference-dependent preferences, so a goal influences her behavior by serving as a reference point. The first two assumptions describe the consumer’s basic problem and the key properties of electricity, in contrast to the aggregate

³The information provided is, to some extent, customized to each individual user, based on rough estimates derived from a statistical model of household energy consumption and appliance saturation.

good. Given billing structures, consumers derive immediate utility from electricity usage, but do not pay for it until the end of the month. Thus, electricity use is much like credit card borrowing, with which present bias has been correlated (Meier and Sprenger, 2010; Shui and Ausubel, 2005).⁴ An alternative, broader interpretation is that while the benefits of consumption are enjoyed in the present, the nontrivial negative consequences of current consumption will not be felt until the next generation(s) – extreme weather due to climate change, sufficiently polluted air and water – about which the consumer cares to some degree for selfish or altruistic reasons.⁵ In conjunction with these two features, the third assumption generates demand for a mechanism to attenuate future consumption. The fourth assumption explains how goal provision can serve as such a mechanism, by providing the consumer with an incentive to consume less electricity. Note that present-biasedness is *not* necessary to describe the consumer’s behavior in response to goals once she has signed up for the program. Rather, present-biasedness is necessary to explain why the consumer would sign up for the program in the first place—a time-consistent consumer would be able to consume less electricity without the aid of the program if that were optimal. In an alternative model in which a consumer has rationally decided not to reduce energy use (for example, if she finds it too costly to invest in a durable good that reduces energy use), revealed preferences imply that she would not select an intervention to change her behavior, especially if it is costly to sign up. Thus, the model is novel in that it explains both *how* goal-setting affects behavior through reference dependence, and *why* consumers demand such an intervention due to present bias.⁶

We also consider an alternative model in which the consumer is *neither* reference-dependent nor present-biased, and is driven entirely by material benefits. As a standard expected-utility maximizer, she signs up for the program to earn reward points (and perhaps to curb consumption); likewise, her post-adoption behavior is incentivized by reward points, not reference-dependent preferences. Because the expected utility model is a special case of the first model, we develop the former and then contrast their predictions. After comparing the results to these two models, we discuss and rule out other variants of expected utility, including information-seeking and selection.

The Consumer’s Problem. In the first stage, denoted period 0, the consumer decides whether or not to sign up for the goal-setting program. If she signs up, she incurs a one-time sign-up cost ($\epsilon \geq 0$) in period 0. We can think of this cost as the nuisance of filling out forms to sign up and link their billing records to the program, net of the sign-up bonus. She also sets her goal r for

⁴And if a consumer pays the electricity bill with a credit card, the cost of usage is delayed even further exactly by such borrowing.

⁵Here, “environmental concern” can be interpreted as an individual’s subjective beliefs about the discounted environmental costs from electricity usage, $\delta c(x)$.

⁶Hsiaw (2013) considers the interaction of present bias and reference dependence in an optimal stopping problem. Here, we model energy consumption as an intertemporal problem, which leads to dynamic inconsistency in the presence of present bias.

consumption in the next period, against which she will compare her actual consumption in the second stage.

In the second stage, denoted period 1, the consumer decides how much of her income m to allocate between consumption of an aggregate good y , which has a per-unit price normalized to 1, and consumption of electricity x , which has a per-unit price p . At the end of period 1, she derives utility from consuming x and y and from comparing her consumption x to her goal r . The consumer's benefits from consumption of electricity and the aggregate good are accrued in period 1 and are described by $u(x, y)$, where $u_x > 0$, $u_y > 0$, and $u_{xx} < 0$, $u_{yy} < 0$.

Consumption of electricity also leads to a future cost to the consumer, $c(x)$. For simplicity, we assume that this future cost is incurred in the period after consumption (i.e., period 2), where δ is the consumer's discount factor. We assume that the marginal cost of each additional unit is positive and weakly increasing in electricity consumption: $c(0) = 0$, $c'(x) > 0$ and $c''(x) \geq 0$. In contrast, there is no future cost associated with consumption of the aggregate good y . Thus, electricity is a good whose consumption yields immediate benefits and delayed costs.

Why might a consumer want to curb her electricity use? If she is aware that she has a self-control problem and tends to overconsume electricity, she may seek a device to counteract this tendency. The consumer has quasi-hyperbolic time preferences, so for all t ,

$$U^t(u_t, u_{t+1}, u_{t+2}, \dots) = u_t + \beta \sum_{\tau=t+1}^{\infty} \delta^\tau u_\tau, \quad (3.1)$$

where $0 < \beta \leq 1$ and $0 < \delta \leq 1$ (Phelps and Pollack, 1968; Laibson, 1997). Thus, the consumer's period-1 discounted utility from consuming electricity and the aggregate good is $u(x, y) - \beta\delta c(x)$. However, her ex-ante, period-0 discounted utility is equivalent to that of an exponential discounter, with $\beta = 1$. Thus, when $\beta < 1$, the consumer suffers from a time-consistency problem: she always prefers to consume more electricity in the present, but prefers that she consume less in the future.

How can a goal counteract this time-consistency problem? We assume that the consumer has additively separable, reference-dependent preferences. In addition to deriving utility from absolute consumption, she derives comparison utility from comparing her electricity consumption against a reference point r , which is the consumption goal or target that she faces under the goal-setting program, at the end of period 1. The goal is set during period 0, so it is taken as given and cannot be changed by the consumer at the time she makes the allocation decision in period 1.⁷ If she does not sign up for the goal-setting program, she does not have a goal against which to evaluate herself in period 1. There is extensive evidence that without a well-defined basis against which to make a

⁷The qualitative results would be weakened but still hold if the consumer were allowed to adjust his goal in the second stage, as long as the original goal from the first stage is somewhat "sticky." Hsiaw (2013) explores this possibility in the context of optimal stopping, but the qualitative results apply here.

comparison, vague goals have no effect on motivation and effort (Latham and Locke, 1991; Mento et al., 1992). In this case, we assume that her comparison utility equals zero.

Let the consumer's comparison utility be described by the function $f(r - x)$. The consumer derives no comparison utility from meeting her goal exactly, since she experiences neither a gain or a loss ($f(0) = 0$). She experiences positive comparison utility (i.e., a gain) when she consumes less electricity than her goal ($f(r - x) > 0$ if $r > x$). Likewise, she experiences a loss when she consumes more than her goal ($f(r - x) < 0$ if $r < x$). The function f is strictly increasing in the gap between the goal and consumption ($f'(r - x) > 0 \forall x$). Consistent with prospect theory (Kahneman and Tversky, 1979), she exhibits diminishing sensitivity to both gains and losses ($f''(r - x) < 0$ if $r > x$ and $f''(r - x) > 0$ if $r < x$).

The consumer derives utility $g(x_0 - x)$ from reward points. In contrast to comparison utility, points are awarded by the program when the consumer consumes less electricity relative to *last year's consumption in the corresponding period*, denoted x_0 , rather than the self-set goal r . Clearly, past consumption must be taken as given by the consumer and cannot be changed in the second stage, and she derives no utility from having no points ($g(0) = 0$). The number of awarded points increases with the amount of electricity saved relative to x_0 , but points are not subtracted when the consumer consumes *more* electricity than she did in the corresponding period last year ($g(x_0 - x) > 0$ if $x_0 > x$, $g(x_0 - x) = 0$ if $x_0 \leq x$). Receiving more points yields more, but diminishing marginal utility ($g'(x_0 - x) > 0$ and $g''(x_0 - x) < 0$ if $x_0 > x$, and $g'(x_0 - x) = g''(x_0 - x) = 0$ if $x \leq x$).

Thus, the consumer's problem in period 1 is

$$\max_{x,y} u(x,y) + 1_{\text{signup}}[f(r - x) + g(x_0 - x)] - \beta \delta c(x) \quad \text{subject to } m \geq px + y, \quad (3.2)$$

where $1_{\text{signup}} = 1$ if the consumer has signed up for the program and is zero otherwise. When deciding whether to sign up for the program in period 0, she takes into account her consumption behavior in response to the goal and the reward points she would earn.⁸

Characterizing Consumer Behavior. Let (x^*, y^*) be the consumer's consumption if she signs up for the program. When $\beta < 1$, the consumer is present-biased, undervaluing the future costs of electricity consumption relative to its immediate benefits. This leads her to consume more electricity than she prefers from an ex-ante perspective, when she would like to consume as though $\beta = 1$. The presence of a consumption goal r counteracts this tendency, because she is motivated to increase comparison utility by consuming less. When the goal is achievable ($r \geq x$), a more ambitious goal motivates her to consume less due to diminishing sensitivity to gains - the marginal

⁸In the Appendix, we account for period-1 costs of implementing goals in the model, for which qualitative results are unchanged.

gain from decreasing consumption is higher as r decreases in the gains region. But, because the consumer exhibits diminishing sensitivity to losses, a goal that is excessively ambitious ($r < x$) becomes *less* effective as it becomes even more ambitious - the marginal gain from decreasing consumption is lower as r decreases in the loss region. As long as the consumer cares about future utility ($\beta > 0$), she will decrease electricity usage as its discounted costs increase.

Proposition 1 *Given a consumption goal r , the consumer's electricity usage upon sign-up has the following properties:*

1. *Her usage increases with present-biasedness: $\frac{\partial x^*}{\partial \beta} < 0$.*
2. *When her consumption is lower than her goal, her usage decreases as her consumption goal decreases: $\frac{\partial x^*}{\partial r} > 0$ when $r \geq x^*$. When her consumption is higher than her goal, her usage increases as her consumption goal decreases: $\frac{\partial x^*}{\partial r} < 0$ when $r < x^*$.*
3. *Her usage decreases as the discounted environment cost increases: $\frac{\partial x^*}{\partial \delta} < 0$ when $\beta > 0$.*

In period 0, the consumer decides whether to sign up for the program, and sets a goal for the second stage if she does so. If she signs up, she incurs a one-time sign-up cost, denoted $\epsilon \geq 0$. Thus, the sign-up decision has an initial fixed cost (ϵ), and a delayed period-1 benefit that arises from consuming a level of electricity closer to the customer's ex-ante optimum.

At this point, the manner in which the goal is determined is crucial for the sign-up decision and characterizing subsequent behavior. We assume that the selected goal must be fully consistent with the consumer's beliefs about the outcome she will achieve. That is, she must set a goal that she believes she will actually meet, satisfying rational expectations. Under the interpretation that the goal acts as a reference point against which the consumer compares her consumption, this assumption is consistent with psychological findings, as well as prevailing theoretical work on reference-dependent preferences. Based on lab and field experiments, Latham and Locke (1991) conclude that goal choice is an integration of what one wants and what one believes is possible, suggesting that goals are realistic. In this vein, Kőszegi and Rabin (2006; 2009) assume that when making decisions and plans, reference-dependent individuals endogenously form reference points that must satisfy rational expectations.

The consumer's sign-up decision and subsequent goal choice also depend on her beliefs about future behavior, particularly the degree to which she is aware of her time inconsistency. To capture the full spectrum of beliefs, we use O'Donoghue and Rabin's (2001) model of partial naivete, in which the consumer may be aware of her time-inconsistency but may underestimate its magnitude. The concept of partial naivete generalizes the polar cases of naive and sophistication (O'Donoghue and Rabin, 1999), and allows us to explore the effect of belief heterogeneity on observed behavior

in the actual program.⁹ Let $\hat{\beta}$ be her period-0 belief about parameter β , with which she actually makes the consumption decision in period 1. In period 0, she evaluates utility according to (the true) β , but she *believes* that in the future, she will make decisions based on $\hat{\beta}$, where $\hat{\beta} \in [\beta, 1]$. When $\hat{\beta} = \beta = 1$, the consumer is a standard exponential discounter, and is thus time-consistent. When $\hat{\beta} = \beta < 1$, the consumer is sophisticated, fully aware of her time inconsistency. When $\beta < \hat{\beta} = 1$, the consumer is naive, completely unaware of her time inconsistency and extremely optimistic about her ability to implement ex-ante plans.

When the goal is determined by expectations, the consumer *believes* that she will choose (x^*, y^*) according to $\hat{\beta}$ rather than β . Since her goal matches these expectations, then $r = x^*(\hat{\beta})$. Let $V(x^*(\hat{\beta}), y^*(\hat{\beta}) | r = x^*(\hat{\beta}))$ be her perceived indirect utility from signing up for the program, given goal $r = x^*(\hat{\beta})$. When $r = x^*(\hat{\beta})$, the consumer expects to receive zero comparison utility from meeting her goal. Let $V(\tilde{x}(\hat{\beta}), \tilde{y}(\hat{\beta}))$ be her perceived indirect utility from not signing up, where (\tilde{x}, \tilde{y}) is her consumption if she does not sign up:

$$V(x^*(\hat{\beta}), y^*(\hat{\beta}) | r = x^*(\hat{\beta})) = u(x^*(\hat{\beta}), y^*(\hat{\beta})) + g(x_0 - x^*(\hat{\beta})) - \delta c(x^*(\hat{\beta})) \quad (3.3)$$

$$V(\tilde{x}(\hat{\beta}), \tilde{y}(\hat{\beta})) = u(\tilde{x}(\hat{\beta}), \tilde{y}(\hat{\beta})) - \delta c(\tilde{x}(\hat{\beta})). \quad (3.4)$$

When she is not in the goal-setting program, she has no additional motivation to consume less electricity, so $\tilde{x}(\hat{\beta}) \geq x^*(\hat{\beta})$. Based on her beliefs, the consumer signs up for the program if its discounted net benefits outweigh its sign-up cost.

In the absence of reward points, a consumer who believes she will not overconsume in the future ($\hat{\beta} = 1$) would not sign up for the program even if $\varepsilon = 0$, since she believes that goals would distort optimal behavior. Thus, time-consistent and naive consumers will only sign up if reward points are sufficiently attractive. On the other hand, a consumer with a low (true) β is less inclined to sign up, since she does not value the future benefits from sign-up sufficiently relative to the immediate sign-up cost. Thus, consumers who recognize their need for self-control are more inclined to sign up for the goal-setting program ($\hat{\beta} < 1$). If reward points are not sufficiently attractive for time-consistent and naive consumers to sign up, they are certainly not sufficient to induce sign-up by sophisticated hyperbolic discounters, who anticipate earning fewer rewards due to the tendency to overconsume. In this case, program sign-up is driven by the desire to curb overconsumption due to present-biasedness.

Proposition 2 *When goals are determined by expectations, time-consistent ($\beta = \hat{\beta} = 1$) and naive ($\beta < \hat{\beta} = 1$) consumers only sign up for the goal-setting program if reward points are sufficiently attractive. Otherwise, a consumer must recognize her own time inconsistency to sign up for the*

⁹Ali (2011) derives the conditions under which partial naivete can arise when an individual can learn about the severity of her self-control problem through experimentation.

program ($\hat{\beta} < 1$).

When goals are determined by expectations, those who sign up and have more optimistic beliefs about their self-control (i.e., higher $\hat{\beta}$) will set more ambitious consumption goals (i.e., lower r), because they expect to achieve them. But those who are partially naive ($\beta < \hat{\beta}$) will fall short of them, since they are more present-biased than they had believed.¹⁰

Proposition 3 *When goals are determined by expectations, consumers who have signed up for the goal-setting program will either meet their goals (when fully sophisticated) or fall short of them (when partially naive).*

Freely Chosen Goals. We now consider the implications of an alternative model, where the consumer is an expected utility (EU) maximizer, neither reference-dependent nor present-biased. This is equivalent to assuming that $\hat{\beta} = \beta = 1$ and that $f(r - x) = 0$ for all x . The key difference between the standard EU maximizer and the reference-dependent consumer is that after sign-up, the former is not affected by the self-set goal at all. However, her behavior *is* affected by the presence of reward points, which affect her material outcome. She will be motivated by reward points to consume less electricity if they are negatively correlated with consumption.

Because the self-set goal (r) is completely divorced from the consumer's material payoff, her goal choice is completely irrelevant to her utility. Thus, there should be no correlation between her goal choice, r , and actual consumption. Here, the consumer has no self-control problem to mitigate, so she will sign up only if earning reward points outweighs their distortionary effect of reducing her consumption.

Proposition 4 *When the consumer is not reference-dependent, the selected goal is unrelated to actual consumption. Moreover, she will certainly consume less electricity after signing up for the program if she earns reward points from doing so.*

Thus, the two models of consumer preferences offer very different predictions about program sign-up and subsequent goal selection. When the consumer is reference-dependent and present-biased, she will only sign up if she desires commitment, and she will set realistic or overly-optimistic goals relative to her actual consumption. When the consumer is a standard EU maximizer, she will sign up purely to earn reward points. Because the accumulation of reward points is unrelated to the self-selected goal, her goal choice should be unrelated to actual consumption.

¹⁰The addition of uncertainty does not change the qualitative prediction when consumers have correct expectations about this uncertainty. Sophisticates will meet their goals on average, and those who are partially naive will fall short of their goals on average.

4 Model Evaluation

We conduct a comprehensive empirical analysis of the program data using a number of different econometric approaches.

Data. To conduct our empirical analysis, we combine data from various sources. C3 collects detailed information on program participants, including the time of sign-up and monthly electricity bills before and after adoption. For each consumer, we know the precise billing cycle, which typically spans days in two calendar months. We re-normalize each bill to correspond to a calendar month by assuming uniform consumption during the month. Thus, the estimated monthly bill for each consumer is constructed as the sum of the average daily consumption from two different bills. While this introduces some inevitable measurement error, it is preferable to ignoring the different billing cycles. We thus have an (unbalanced) panel dataset from January 2009 to December 2011.

Unfortunately, since the program was not run by the local utility ComEd, we do not have access to the full database of households that would have been eligible for the program. Since we wish to understand the decision to opt into the program in addition to the effect of adoption on consumers, we construct a “control” group of customers who could have opted in but did not. While it is impossible to determine exactly who was exposed to the varied marketing campaigns surrounding this program, we can look instead at customers who would have been eligible to opt in but did not, whatever the reason may be. To do so, we purchased voter registration files for the state of Illinois from a third party provider. We selected a random stratified sample of Illinois residents to reflect the proportion of non-adopters in the zip codes where adopters reside. Both groups of households were then linked to individual-level demographics from another third party provider. Some variables, such as age and income, are estimated at the household level. We do not know the precise age of the household head, but only whether someone of a certain age resides in the household. Consequently, households where more than one family reside together will appear twice in each respective age group. In addition to basic demographic information such as age, household income, education, presence of children and household size, we also acquired additional lifestyle information, which the third party provider obtains from past subscriptions or purchasing behavior.

To identify consumers who are motivated by environmental concerns, we acquired two variables: *environmental issues* and *green living*, which have been shown by Harding and Rapson (2014) to correlate with adoption into a carbon offsetting program in California. *Environmental issues* measures whether a household has expressed interest in environmental or wildlife issues through magazine subscriptions and/or mail response. *Green living* is a variable which aims to identify households that are living environmentally friendly, by buying green household cleaning

products, eating organic foods, donating funds to environmental causes, or driving hybrids. These variables were constructed without reference to the program analyzed.

We also acquired lifestyle variables that are potential indicators of self-control problems, i.e. difficulty behaving in a manner consistent with ex ante plans. Thus, we know whether the person signing up for the program is a smoker or regularly engages in lottery and casino gaming. We also know whether the consumer is or has recently made purchases of dieting and weight loss products or participates in weight loss programs. We captured two additional variables that may also be related to self-control when interpreted as conditional on income. One is a measure of loan-to-value ratio for the household, a commonly used measure of risk based on the ratio of the current mortgage relative to market values. A ratio of 80% or above typically identifies a risky borrower, so a high loan-to-value ratio, conditional on income, may indicate weaker self-control. The other financial variable we capture is whether or not the consumer has one or more retail store lines of credit (e.g. Old Navy), which does not include bank-issued credit cards. These variables are used extensively in consumer marketing and are associated with various measures of behavioral biases. For example, Meier and Sprenger (2010) show that experimentally measured present bias is highly correlated with credit card borrowing, while Ikeda et al. (2010) find that procrastination is positively correlated with body mass index.

Our sample consists of 2,487 households who signed up during the first year of the program and for which we have billing and demographic information. We also have information on 9,964 households derived from the random stratified sample of households in Illinois based on available voter registrations.¹¹ Table 1 shows the summary statistics for our data. Adopters are likely to be younger and more educated. They are much more likely to have gone to graduate school and live in smaller households. Adopters are also more likely to express environmental interest and to be engaged in environmentally friendly actions. Turning to our potential indicators of self-control problems, dieters are more likely to participate in the program, while smokers are less likely (though not statistically significant). Households are comparable in terms of gambling. Households with one or more credit lines are, however, less likely to participate, while households with a high loan to value ratio are more likely to sign up.¹²

Goal Choice. Goals are implicitly set by choosing a set of actions, each implying a certain level of electricity savings, which may vary depending on property characteristics. We quantify a consumer's selected goal levels as the implied electricity savings from implementing the chosen actions over the next year as a fraction of the total electricity usage in the 12 months prior to adoption.

¹¹We obtained information on 10,000 households chosen at random from the zip codes where adoption occurred, but 36 households were already registered for the program.

¹²Note that these statements are based purely on the summary statistics and do not consider the correlations between these different demographics as captured by our multivariate analysis.

Although users are actively encouraged to choose a goal of 5%, 10% or 15%, the actual distribution of chosen goals is continuous. Approximately 15% of the consumers choose a 0% savings goal (i.e. they choose to select no action which they plan to implement). Approximately 32% of the consumers choose a goal greater than 0% but less than 15%. We consider goals in this range, which the program deems achievable to consumers explicitly, to be “realistic.” About 41% choose savings goals in the range 15% to 50%, which we consider to be over-optimistic. In a recent consulting report, McKinsey (2009) suggests that savings of about 23% may be achievable. Meier (2009) reports that following an avalanche destroying the power transmission lines in the city of Juneau, Alaska, residents reduced energy use and improved energy efficiency by as much as 30% in the following 6 weeks. This is a fairly extreme case of what consumers can achieve under dire circumstances. Thus, goals in this range entail significant lifestyle changes and are rather over-optimistic for most consumers. The remaining 12% of consumers choose undoubtedly unreasonable goals that imply savings in excess of 50% of their previous year’s consumption.

Consumers have a strong propensity to overcommit by selecting too many goals. Most consumers go well beyond the three to five actions recommended on the website. The mean consumer chooses 10 actions, and 95% of consumers choose fewer than 25 actions. The mean consumer who chooses savings in excess of 50% commits to over 20 actions. Consumers tend to prefer options which do not require substantial financial investment, instead choosing options which rely on behavioral change.¹³ Low-cost or no-cost actions alone do not fully explain the degree of overconfidence observed. The median cost of implementing the savings actions is \$60, reflecting the large selection of low-cost actions, while the mean cost is just below \$1,500, indicating that some consumers commit to very expensive investment choices. In particular, consumers with savings commitments in excess of 50% of their past savings commit to mean investments of around \$3000.

Pre-adoption consumption. Our model predicts that before enrolling in the program, present-biased consumers are likely to consume more electricity than consumers who are not present-biased. In model A of Table 2, we regress monthly consumption before adoption on the household characteristics. The estimates suggest that variables associated with present-biasedness tend to have a (weak) positive effect on pre-adoption usage, indicating that present-biased consumers may in fact tend to consume more electricity. Specifically, a high loan-to-value ratio is significantly related to higher pre-adoption consumption; the other indicators of present bias are positively but not significantly related. Statistically, mean consumption is strongly driven by demographic fac-

¹³The most popular actions chosen are: Install CFLs in your lighting fixtures; Close your blinds during summer days; Wash only full loads of dishes; Use more natural lighting; Unplug your coffee maker when you’re done brewing; Clean the lint trap in your dryer before every load; Use a drying rack to dry your clothes; Turn up the temperature of your thermostat during the summer; Use a microwave oven instead of your oven for cooking; Hand clean your oven instead of using the auto clean.

tors such as household size and income.

Program adoption. A household is considered an adopter if she successfully “linked” her utility account to the CUB website. We assume that the households we observe from the voter registration lists had the opportunity to opt-in but did not. Thus, a natural way to investigate the characteristics of the adopters is to conduct a logit analysis on the decision to opt-in as function of the observable demographics, corresponding to a simple linear random utility choice model. Note that we are dealing with a non-random sample, since the data was constructed conditional on adoption status, a sampling framework that is usually referred to as “choice based sampling” or “retrospective sampling,” since it uses the ex-post outcomes as part of the sampling frame. It is well-known in this setting that estimation by maximum likelihood leads to inconsistent parameter estimates (Manski and Lerman, 1977). While several approaches are available to address this issue¹⁴, they typically require additional quantities that are not observed in the sample. An interesting feature of the logistic models is that consistent estimates of the slope coefficients and corresponding standard errors can be obtained even in the presence of choice-based sampling (Prentice and Pyke, 1979; Scott and Wild, 1986), while only the estimated intercept coefficient is inconsistently estimated. This imposes some restrictions, as it prevents us from computing marginal effects without imposing out of sample priors on these unobserved parameters.

The estimated coefficients are displayed in model B of Table 2. Column 2 shows that smaller, more educated households who are concerned for the environment are more likely to opt in. We find that smoking is a significant negative predictor of adoption, while dieting is a significant driver of adoption, suggesting that recognition of present bias is relevant for sign-up as predicted by the behavioral model. Dieters are more likely to be aware of time inconsistency and the need for commitment, since they seek additional intervention to curb caloric intake. Smoking in itself does not reliably predict the need for commitment, since smokers can include smokers who don’t want to quit.¹⁵ A high loan-to-value ratio is a positive and significant predictor of adoption, while having one or more store credit lines is a negative predictor. Both groups are likely to be present-biased, but given the recent collapse in the housing market and its severity, it is probable that the former is very much aware of its time inconsistency. Moreover, store credit lines alone are an incomplete measure of credit card borrowing and we do not have information on total credit card borrowing.

We can investigate the commitment behavior of households which have signed up for the program further by investigating their goal level and the number of actions to which they commit. Model C estimates an ordered logit for the level of commitment, while Model D estimates a neg-

¹⁴See Amemiya (1985) for a thorough review of this problem and associated classical solutions.

¹⁵Khwaja et al. (2007) find that smoking is not significantly correlated with measures of hyperbolic discounting, though it is related to more general measures of impulsivity.

ative binomial model for the number of actions. In order to estimate Model C we define the goal level of each household as an ordered choice corresponding to a fraction of committed savings to past usage of 0%, 0%-15%, 15%-50%, or 50%+. Both models focus on the decision of the households to commit to a certain level of savings, after they have successfully signed up for the program. As a result both models are estimated on the sample of 2,487 adopters.

Model C of Table 2 shows that smaller and less wealthier households are more likely to over-commit. We do not find a clear relationship between the commitment level and the included lifestyle variables. Model D of Table 2 attempts to explain the number of chosen actions that the consumers commit to. Education, especially a graduate degree, acts as a moderating force on the number of actions committed to by a consumer. Green living predicts a higher number of committed actions, which may indicate that some consumers may select actions they have already done. Wealthier households choose a higher number of actions, reflecting their greater savings opportunities. Interestingly, dieting indicates a propensity to set overly optimistic goals, in the number of actions. That they tend to be over-optimistic may be not surprising, in light of the fact that dieting itself is often an unsuccessful exercise. Since dieters are a group who exhibit a demand for commitment, their behavior is consistent with the behavioral model's prediction that those who recognize their self-control problem and their reference-dependence will set goals to counteract overconsumption. The EU model again predicts no systematic relationship between indicators of demand for self-control and adoption or goal choice.

Measuring post-adoption savings. We now present evidence that households who are aware of their present bias and have well-calibrated expectations tend to achieve their goals on average. Note that in the context of a large-scale real world program implementation we lack the ability to perfectly account for all selection factors.¹⁶ Thus, we do not claim that any households exposed to such a program would achieve post-adoption savings. Rather, we have shown that households who want to save energy and think they will respond to goals are signing up for the program and using the program as a mechanism to change their behavior and save electricity. Furthermore, we will show that some households do not have well-calibrated goals, and thus not all households will benefit from the program post-adoption.

In order to measure post-adoption savings we use a series of common econometric strategies from the labor economics literature. In particular, we use the variation in the timing of adoption in a difference-in-differences framework to measure the effect of adoption into the program on subsequent use, while also controlling for individual and time effects (Reber, 2005; Hoynes and

¹⁶In order to evaluate a stronger causal claim we would have required a randomized control trial where some households who sign up for the program are then denied service. In general such an approach is infeasible in the context of a heavily regulated industry, where considerations of customer equity are very important to firms and regulators alike.

Schanzenbach, 2009, 2012). This strategy assumes that conditional on time invariant household characteristics, the timing of adoption into the program is random and that we can compare the consumption of households who opt into the program with the consumption of households who have not yet opted in, but will do so later on.

We also conduct three additional robustness checks to validate the robustness of our approach to measuring post-adoption savings. First, we conduct a duration analysis of the timing of adoption using a Han and Hausman (1990) model with flexible baseline hazard and Gamma distributed unobserved heterogeneity, which reveals that while the shared baseline hazard is statistically significant, neither the observed covariates nor the unobserved heterogeneity can explain the timing of adoption (detailed results are available from the authors). This indicates that the timing of adoption depends solely on marketing exposure and does not relate to variables which may affect the trending behavior post-adoption. Second, our statistical model does not reveal the presence of any pre-existing savings trends before adoption and the savings trends after adoption show a consistent pattern across participants. Since each participant has a different start date, it is extremely implausible that the households were exposed to some other unobserved factors (e.g. community meetings), which happened to be temporally aligned with the timing of adoption into the program, and which were also aimed at inducing energy efficiency. Third, we show that, to the extent that unobservable factors are correlated with observable covariates, these cannot be explaining our savings by introducing a propensity score matching between individuals signing up for the program at different points in time using the method of Sianesi (2004). This adjusts for potential demographic differences between individuals signing up at different points in time and provides strong evidence of the robustness of our estimates.

To estimate average individual level savings while being as flexible as possible in terms of the econometric specification, we estimate the following “event study” equation:

$$\log(kWh)_{it} = \sum_{k=-\bar{m}}^{\bar{m}} \delta_k D_{it}^k + \alpha_t + \gamma_i + \varepsilon_{it} \quad (4.5)$$

where D_{it}^k are a set of indicator variables set equal to one if, in calendar month t , household i is k months away from its enrollment month in the CUB program. The model also accounts for individual effects γ_i and month-year effects α_t . The underlying assumption is that, conditional on time-invariant household characteristics and aggregate month-specific shocks, all households that are k months away from enrolling in the program are identical (in expectation). The specification implicitly models the response as a piecewise linear function of relative time to adoption, with no restrictions on the variation or pattern of the response over time. This provides a test of our main assumption, as we should not observe statistically significant trends pre-adoption.

We restrict the event study window such that $k \in [\underline{m}, \bar{m}]$, and normalize the coefficient of event

month prior to adoption to zero.¹⁷ We choose 12 months pre-adoption and 18 months post-adoption as the cut-off points for our event window. The relatively long post-adoption window allows us to evaluate the long run persistence of any observed behavioral change, but it prevents us from constructing a balanced panel, since data after December 2011 is not available to us. We thus remain cautious in our interpretation of the estimation results for the later months, since they may also reflect compositional changes resulting from a diminishing sample size and which will be reflected by increasing confidence bounds.

In Figure 1, we present the estimated coefficients on the indicator variables in event time. Before adoption, we see no statistically significant trends. Post adoption, however, the average estimated savings are 4.4%. Households realize very impressive saving of close to 8% by the second month after sign-up, but over the next six months, average savings diminish and become statistically indistinguishable from zero. Notice, however, that the standard errors also increase over time as the sample diminishes with the length of the event window. We can use propensity score matching across time periods to adjust the savings estimate for biases arising from the different selection into the program over time (Sianesi, 2004). The precise technical details of the matching algorithm are discussed in the Appendix. In Table 3 we report savings estimated from using both demographic variables and pre-adoption usage in two nearest-neighbor matching procedures. We find that estimated savings range from 2.9% to 4.4% depending on the matching set.

Goal level and savings. We estimate the effect of adoption on savings separately in relation to the chosen goals. In Figure 2 we report the estimated mean savings for the different groups of commitments (0%, 0%-15%, 15%-50%, 50%+). Consumers choosing realistic goals achieve the most substantial savings of 11% on average, indicating that their goals tend to be in line with actual consumption. In the figure, we also plot the 95% confidence bounds for the estimated mean savings for the group of consumers with realistic goals between 0% and 15%. The savings for these households are statistically significantly different from zero for the entire post-adoption time frame. By contrast, consumers choosing no goals save about 1.5% on average and those choosing over-optimistic goals save on average only 1% savings, but neither of these types of unrealistic goals results in statistically significant change in consumption behavior.¹⁸ Moreover, the yearly average savings achieved by the over-optimistic and unrealistic consumers are driven by their behavior in the first two months. These consumers quickly give up, realizing that they cannot

¹⁷Periods outside the event window are estimated at the same time as the bounds for the event window, i.e. the indicator variable for period $k = \underline{m}$ captures all periods $t \leq \underline{m}$ and the indicator variable for period $k = \overline{m}$ captures all periods $t \geq \overline{m}$.

¹⁸In order to explore the behavioral differences between consumers with different goal levels, we also estimate a model which measures the average post-adoption savings by group. We then conduct pairwise t-tests for the equality of the post-adoption savings between groups. We reject the hypotheses of equal savings in all cases except for the consumers with zero goals and those with overoptimistic goals.

achieve the projected savings. Together, these findings are consistent with the reference-dependent model, which predicts an approximately inverse U-shaped relationship between goal choice and savings. Those who set realistic goals achieve correspondingly increasing savings, while those who are partially naive will fall short of self-set goals and actually save less as the savings goal increases, due to diminishing sensitivity to relative losses. Since the goal is not payoff relevant, any relationship, or lack thereof, between goal choice and post-adoption savings would be consistent with the EU model.

Our reference-dependent, present-biased model predicts that, in addition to having higher post-adoption consumption, present-biased consumers also achieve lower post-adoption savings than those who are not. We restrict our attention to households who have chosen realistic goals, then further divide such households into subgroups as a function of the different characteristics and estimate the “event study” accounting for potential heterogeneity between the different subgroups. We find that on average, the group of consumers that is tagged as more likely to be present-biased saves less after adoption for variables for which we have sufficient observations to estimate the model by characteristic. In particular, we find that dieters and those with high loan-to-value ratios save significantly less than non-dieters and those with lower loan-to-value ratios, respectively.

For each of these variables, an analysis of households that includes both realistic and overly optimistic goals reveals an even wider gap in savings between subgroups, though average savings by both subgroups is unsurprisingly lower, since overly optimistic households save less on average than realistic ones. In contrast, the EU model, which precludes self-control problems by assumption, predicts no systematic relationship between indicators of demand for self-control and pre- and post-adoption consumption. The higher pre-adoption consumption and differential savings response to goals in the direction predicted by present bias suggest that our lifestyle variables are in fact indicators of present bias, rather than simply indicators of consistent preferences. For example, there is no parsimonious explanation for why those who simply prefer to lose weight and those who are borrow heavily on homes, perhaps with higher expectations of future home values, would both consume more electricity before adoption and be less responsive to goals after adoption than those who don't.

Feedback. In the CUB program, consumers receive monthly email feedback and reward points that are not directly related to their chosen goals. Rather, they depend on the consumer's consumption relative to her weather-adjusted usage in the previous year, where the adjustment algorithm is not known to the consumer. An analysis of the feedback and points awarded reveals that they are poorly correlated with consumption. Over the program duration, all customers received between 60 and 100 points per month. The mean number of points awarded for consumers who selected goals of 0%, 0-15%, 15-50% and 50%+ were 74, 81, 78 and 76 respectively, while the program-

reported mean monthly savings over the same month a year earlier were 91, 115, 90 and 78 kWh respectively. As we show in Section 4, actual savings by these groups was 1.5%, 11%, 1% and 0%, respectively. Thus, awarded points and email feedback are a very noisy and indirect signal of true savings relative to selected goals.

This suggests that if consumers were motivated to save solely through reward points, they would quickly give up after realizing the lack of correlation between their consumption and points earned. That is, reward points do *not* act as a material incentive for consumers to achieve their chosen goals. Therefore, the choice of goal level should be irrelevant to a standard expected-utility maximizing consumer. Nonetheless, we find positive savings in the program that persistent among a subset of consumers, and we observe a clear non-monotonic relationship between goal choice and savings.

Moreover, since consumers received similar rewards irrespective of the actual savings achieved or the goals chosen, consumers who sign up solely to earn reward points should quickly give up saving once they realize this. While the initial saving and subsequent drop-off in savings among those consumers who set 0% and overly ambitious savings goals is consistent with the EU model, the persistent and substantial savings among those who set realistic goals indicates that pure points-seeking is not a sufficient explanation for adoption and post-adoption behavior.

In sum, realistic goals imply persistently higher post-adoption savings than other goal levels despite ineffectual material incentives in the form of reward points. Indicators of demand for commitment predict higher pre-adoption and post-adoption consumption, as well as overly-optimistic goal choices, relative to their absence. Taken together, this evidence is consistent with the predictions of the behavioral model but cannot be easily reconciled with the expected utility model.

Alternative Mechanisms. Because the CUB program website includes reward points, personal goal-setting and energy-saving tips, the exact mechanism through which post-adoption savings occurs cannot be conclusively identified here. However, the observed pattern of post-adoption behavior is highly suggestive that information and reward points are not the sole drivers of savings behavior.

Table 4 summarizes our empirical findings and model predictions. We compare our findings for the predictions of the behavioral model with those of a standard EU model (described previously). We then consider several variants of the EU model, namely information-seeking, goals as indicators of preferences, and goals as indicators of expectations.

While we argue that behavior is driven by reference-dependent preferences, an alternative explanation is that non-binding personal goals are irrelevant to consumers, and that they are expected utility (EU) maximizers who sign up because they may want to reduce their electricity bills but have uncertainty about how to save effectively, which the energy-saving tips clarify. Thus, they

have no self-control problem and their goal choice is irrelevant for behavior. However, the evidence from the CUB program indicates that at least some proportion of consumers do not satisfy these predictions. First, consumers who exhibit demand for commitment in other domains (i.e., dieters or risky borrowers) are more likely to sign up than those who do not, and save significantly less than those who do not after signing up. Second, the relationship between self-set goals and savings behavior is not random. Moreover, at the action selection stage, all consumers are exposed to the full array of energy-saving tips, of varying investment cost and efficiency impact, *before* they select non-binding goals. Thus, it is not the case that consumers with higher goal levels are exposed to actions that are more difficult or costly to implement than those with lower goal levels. Consumers who select over-optimistic goals tend to select more actions that entail larger investments, which is unsurprising since these would generate the largest savings if actually implemented. However, there is no significant difference in the selection of low-cost behavioral actions across consumers with different goal levels. The degree of goal ambition is driven by the number of actions, rather than the type selected.

Another explanation is that consumers are EU maximizers who are information-seeking and interested in reward points, but their self-selected goals are reflective of underlying preferences, rather than random choices. But to explain the observed relationship between goals and consumption, this implies that those who set 0-15% savings goals wanted to save this amount (or earn the corresponding number of points), yet those who set above 15% savings goals actually wanted to save *less* than the 0-15% group (or earn fewer points). We find this quite unconvincing. If goals reflect underlying preferences, it is much more plausible that those who set less ambitious goals wish to save less, or are less motivated by points, than those who set more ambitious goals, who wish to save more or are more motivated by points. However, the non-monotonic relationship between goal level and savings behavior—in particular, that those who set goals above 15% actually save less than those who set goals under 15%, and yet those who set exactly 0% goals *also* save less than those who set positive but realistic goals—belies this, and is instead predicted by the model of reference-dependent preferences.

A third potential explanation is that those consumers who are the most responsive to goals are those who select realistic goals. That is, consumers who are most responsive to goals also have the most well-calibrated expectations, generating the non-monotonic relationship between goals and savings. We cannot necessarily rule out this explanation, which makes the same prediction as the assumption of diminishing sensitivity in the comparison utility function. Nonetheless, the fact that a group of consumers is (persistently) responding to non-binding goals at all still belies an expected utility-based model, which predicts that goals should have no effect, and is consistent with reference dependence, with or without diminishing sensitivity.

Alternatively, perhaps consumers are responding just to energy saving tips and not to goals, and

those who respond most to energy-saving tips happen to set realistic goals. That is, people who respond to information also have the most well-calibrated expectations, while those who don't respond to information happen to be either optimistic or pessimistic about future behavior. But it is unclear why responsiveness to information would be systemically correlated with miscalibrated expectations. Moreover, if very low initial goals reflect low expectations about the effectiveness of tips, then exceeding such expectations should lead to persistent savings among those with low initial goals. This explanation is consistent with the lack of persistence among those who set unreasonably optimistic goals, but cannot explain the lack of persistence among those who set very low goals.

The CUB program also provided consumers with information on the “ranking” of their community in terms of savings relative to other communities. As we mentioned before, we do not believe that this aspect of the program had a confounding effect. Evidence from the activity on the website shows that this feature of the program was the least utilized and rarely looked at.

In the Appendix, we extend our model to allow for implementation costs. The theoretical insights remain qualitatively very similar. Empirically, it is very challenging to consider ways of quantifying the costs to the consumer of implementing a specific savings plan, since many of the chosen actions are financially costless but may entail unobservable inconvenience or time costs. A team of C3 engineers created an estimated score of the cost to the user of implementing each action to supplement the observed financial cost. While this measure is clearly very noisy, it may still be informative. A careful analysis of this data does not support implementation costs as an explanation of the observed savings patterns in the data. In particular, it does not support the idea that households choosing more difficult actions are less likely to save or to save persistently.

Our analysis addresses a leading explanation in the existing environmental literature on energy conservation (e.g. Kotchen and Moore (2007)), that individuals are driven by altruistic concerns for the environment or social well-being, which is adversely affected by environmental factors. While this is entirely consistent with a desire to consume less electricity (as we acknowledged in describing and interpreting the model), altruism is neither necessary nor sufficient to explain why consumers would opt into the goal-setting program in the first place. If they were altruistic and time-consistent, they would have no difficulty implementing plans to consume less electricity. Thus, the only appeal of the program to such time-consistent altruists would be its informational value or reward points, if they are uncertain about how to save electricity and view the program as a source of information or if they want to earn points. However, we have already argued that neither of these motivations is likely to be sufficient to explain the results.

The evidence from the CUB program corroborates with the theory that a measurable, non-trivial proportion of consumers 1) have a tendency to overconsume electricity relative to their ex-ante preference, and 2) possess reference-dependent preferences, where the goal acts as a reference

point and is determined by expectations. These findings are quite difficult to reconcile with standard EU preferences, particularly seeking reward points, seeking information, or some selection effects. In our view, the most parsimonious explanation is that individuals seek the goal-setting program because they are aware that they have a tendency to overconsume due to present-biasedness. Because they are reference-dependent, they are aware that they will respond to even non-binding goals, and attenuate their consumption.

5 Conclusion

In recent years, major utility companies around the country have introduced energy conservation programs that rely on behavioral nudges to incentivize consumers to reduce energy consumption. One of the mechanisms employed is to encourage consumers to choose a personal energy conservation goal and provide them with the informational tools required for monitoring and achieving their goal. This paper provides insights into the popularity of such programs and the likely energy conservation response.

We develop a theory of consumer demand for energy conservation, in which consumers have present-biased preferences, which lead them to overconsume, and reference dependent preferences, where goals influence behavior by serving as reference points. We find evidence, based on the empirical analysis of an energy conservation program in Northern Illinois, that is broadly consistent with this theory, suggesting that goal setting can be an effective behavioral nudge for reducing residential electricity consumption. While on average consumers save 4%, savings are very heterogeneous. Consumers who set realistic goals persistently achieve substantially higher savings than those who do not. We consider and do not find support for several alternative explanations for the adoption decision and post-adoption behavior, including points-seeking and information-seeking. Thus, the evidence suggests that interest in energy-saving programs is driven by consumers' recognition of their present bias, and that goal setting can be quite effective at reducing energy consumption when goals are achievable.

References

- ABELER, J., A. FALK, L. GOETTE, AND D. HUFFMAN (2011): "Reference Points and Effort Provision," *American Economic Review*, 101, 470–492.
- ALI, S. N. (2011): "Learning Self-Control," *Quarterly Journal of Economics*, 126, 857–893.
- ALLCOTT, H. (2011): "Social Norms and Energy Conservation," *Journal of Public Economics*, 95, 1082–1095.

- ALLCOTT, H. AND S. MULLAINATHAN (2010): “Behavior and Energy Policy,” *Science*, 327, 1204–1205.
- AMEMIYA, T. (1985): *Advanced Econometrics*, Harvard University Press.
- BERTRAND, M., D. KARLAN, S. MULLAINATHAN, E. SHAFIR, AND J. ZINMAN (2010): “What’s Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment,” *Quarterly Journal of Economics*, 125, 263–306.
- BOLSEN, T., P. J. FERRARO, AND J. J. MIRANDA (2013): “Are Voters More Likely to Contribute to Other Public Goods? Evidence from a Large-Scale Randomized Policy Experiment,” *American Journal of Political Science*, forthcoming.
- CAMERER, C., L. BABCOCK, G. LOEWENSTEIN, AND R. THALER (1997): “Labor Supply of New York City Cabdrivers: One Day at a Time,” *The Quarterly Journal of Economics*, 112, 407–441.
- CARD, D. AND G. B. DAHL (2011): “Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior,” *Quarterly Journal of Economics*, 126, 103–143.
- COSTA, D. L. AND M. E. KAHN (2010): “Energy Conservation Nudges and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment,” NBER Working Papers 15939, National Bureau of Economic Research, Inc.
- CRAWFORD, V. P. AND J. MENG (2011): “New York City Cab Drivers’ Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income,” *American Economic Review*, 101, 1912–1932.
- ERICSON, K. M. AND A. FUSTER (2011): “Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments,” *Quarterly Journal of Economics*, 126, 1879–1907.
- FEHR, E. AND L. GOETTE (2007): “Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment,” *American Economic Review*, 97, 298–317.
- FERRARO, P. J. AND M. K. PRICE (2013): “Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large Scale Field Experiment,” *The Review of Economics and Statistics*, 95, 64–73.
- GOETTE, L. F. AND D. HUFFMAN (2007): “Affect and cognition as a source of motivation: a new model and evidence from natural experiments,” in *Do Emotions Help or Hurt Decision*

- Making?: A Hedgefoxian Perspective*, ed. by K. D. Vohs, R. F. Baumeister, and G. Loewenstein, New York: Russell Sage Foundation, 267–294.
- HAN, A. AND J. HAUSMAN (1990): “Flexible parametric estimation of duration and competing risk models,” *Journal of Applied Econometrics*, 5, 1–28.
- HARDING, M. AND D. RAPSON (2014): “Do Voluntary Carbon Offsets Induce Energy Rebound? A Conservationist’s Dilemma,” Mimeo.
- HEATH, C., R. P. LARRICK, AND G. WU (1999): “Goals as Reference Points,” *Cognitive Psychology*, 38, 79–109.
- HOYNES, H. AND D. SCHANZENBACH (2009): “Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program,” *American Economic Journal: Applied Economics*, 1, 109–139.
- (2012): “Work Incentives and the Food Stamp Program,” *Journal of Public Economics*, 96, 151–162.
- HSIAW, A. (2013): “Goal-Setting and Self-Control,” *Journal of Economic Theory*, 148, 601–626.
- IKEDA, S., M.-I. KANG, AND F. OHTAKE (2010): “Hyperbolic discounting, the sign effect, and the body mass index,” *Journal of Health Economics*, 29, 268–284.
- KAHNEMAN, D. AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47, 263–291.
- KŐSZEGI, B. AND M. RABIN (2006): “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 121, 1133–1165.
- (2009): “Reference-Dependent Consumption Plans,” *American Economic Review*, 99, 909–936.
- KHWAJA, A., D. SILVERMAN, AND F. SLOAN (2007): “Time preference, time discounting, and smoking decisions,” *Journal of Health Economics*, 26, 927–949.
- KOTCHEN, M. AND M. MOORE (2007): “Private Provision of Environmental Public Goods: Household Participation in Green-Electricity Programs,” *Journal of Environmental Economics and Management*, 53, 1–16.
- LAIBSON, D. (1997): “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 112, 443–477.

- LATHAM, G. P. AND E. A. LOCKE (1991): “Self-Regulation through Goal Setting,” *Organizational Behavior and Human Decision Processes*, 50, 212–247.
- LOCKE, E. A. AND G. P. LATHAM (2002): “Building a practically useful theory of goal setting and task motivation: A 35-year odyssey.” *American Psychologist*, 57, 705–717.
- MADRIAN, B. C. AND D. F. SHEA (2001): “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” *Quarterly Journal of Economics*, 116, 1149–1187.
- MANSKI, C. AND S. LERMAN (1977): “The Estimation of Choice Based Probabilities from Choice Based Samples,” *Econometrica*, 45, 1977–1988.
- MCKINSEY (2009): “Unlocking Energy Efficiency in the US Economy,” Company Report.
- MEIER, A. (2009): “How one city cut its electricity use over 30% in six weeks,” Tech. rep., Lawrence Berkeley National Laboratory.
- MEIER, S. AND C. SPRENGER (2010): “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal: Applied Economics*, 2, 193–210.
- MENTO, A. J., E. A. LOCKE, AND H. J. KLEIN (1992): “Relationship of Goal Level to Valence and Instrumentality,” *Journal of Applied Psychology*, 77, 395–405.
- NOLAN, J. M., P. W. SCHULTZ, R. B. CIALDINI, N. J. GOLDSTEIN, AND V. GRISKEVICIUS (2008): “Normative Social Influence is Underdetected,” *Personality and Social Psychology Bulletin*, 34, 913–923.
- O’DONOGHUE, T. AND M. RABIN (1999): “Doing It Now or Later,” *The American Economic Review*, 89, 103–124.
- (2001): “Choice and Procrastination,” *Quarterly Journal of Economics*, 116, 121–160.
- PHELPS, E. AND R. POLLACK (1968): “On Second-Best National Saving and Game-Equilibrium Growth,” *The Review of Economic Studies*, 35, 185–199.
- POPE, D. AND M. SCHWEITZER (2011): “Is Tiger Woods Loss Averse? Persistence Bias in the Face of Experience, Competition, and High Stakes,” *American Economic Review*, 101, 129–157.
- PRENTICE, R. AND R. PYKE (1979): “Logistic disease incidence models and case-control studies,” *Biometrika*, 66, 403–11.

- REBER, S. (2005): “Court-Ordered Desegregation: Success and Failures Integrating American Schools since Brown versus Board of Education,” *The Journal of Human Resources*, 40, 559–590.
- SCOTT, A. AND C. WILD (1986): “Fitting Logistic Models Under Case-Control or Choice Based Sampling,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 48, 170–182.
- SHUI, H. AND L. AUSUBEL (2005): “Time Inconsistency in the Credit Card Market,” Working paper.
- SIANESI, B. (2004): “An Evaluation of the Swedish System of Active Labor Market Programs in the 1990s,” *The Review of Economics and Statistics*, 86, 133–155.
- SPRENGER, C. (2010): “An Endowment Effect for Risk: Experimental Tests of Stochastic Reference Points,” Working paper.
- VAN HOUWELINGEN, J. H. AND W. F. VAN RAAIJ (1989): “The Effect of Goal-Setting and Daily Electronic Feedback on In-Home Energy Use,” *Journal of Consumer Research*, 16, 98–105.
- WILHITE, H. AND R. LING (1995): “Measured energy savings from a more informative energy bill,” *Energy and Buildings*, 22, 145–155.

Table 1: Summary statistics for CUB Adopters and a random control group of Northern Illinois households.

	Mean		Standard Deviation		p-val
	Control	Adopters	Control	Adopters	
<i>Demographic Variables</i>					
Age 35-65	0.761	0.762	0.426	0.426	0.930
Age 65+	0.403	0.277	0.490	0.448	0.000
HH income 75k-125k	0.289	0.318	0.453	0.466	0.005
HH Income 125k +	0.178	0.190	0.383	0.392	0.190
College	0.316	0.353	0.465	0.478	0.000
Graduate School	0.179	0.251	0.383	0.433	0.000
Children	0.373	0.355	0.484	0.479	0.090
Household Size	3.226	2.954	1.591	1.507	0.000
<i>Environmental Variables</i>					
Environmental Issues	0.176	0.224	0.381	0.417	0.000
Green Living	0.369	0.408	0.483	0.492	0.000
<i>Self-Control Variables</i>					
Smoker	0.057	0.051	0.231	0.221	0.309
Dieting	0.353	0.403	0.478	0.491	0.000
Gambling-Lottery	0.059	0.066	0.236	0.248	0.221
Gambling-Casino	0.137	0.141	0.344	0.348	0.675
<i>Financial Variables</i>					
Home loan to value 80% or +	0.215	0.265	0.411	0.441	0.000
Credit Line	0.751	0.703	0.432	0.457	0.000

Sample consists of 2,487 adopters and 9,964 control households.

Pairwise t-tests are conducted, and p-values corresponding to each t-test is reported in addition to the mean and standard deviation of each covariate by group.

Table 2: Empirical findings on adoption, goals and pre-adoption usage.

Model	(A)	(B)	(C)	(D)
Estimation procedure	OLS	Logit	Ordered Logit	Neg. Binomial
Dependent Variable	<i>log(kWh)</i>	<i>Adoption</i>	<i>Commitment levels</i>	<i># of actions</i>
Age 35-65	0.0472 (0.0309)	-0.258* (0.131)	-0.0110 (0.106)	0.0953 (0.0542)
Age 65+	-0.0736* (0.0291)	-0.700*** (0.0970)	0.0478 (0.111)	0.0190 (0.0537)
HH income 75k-125k	0.170*** (0.0251)	0.0871 (0.0557)	-0.0344 (0.0876)	0.100* (0.0478)
HH Income 125k +	0.248*** (0.0318)	0.0187 (0.0936)	-0.217** (0.0827)	-0.0254 (0.0582)
College	0.0126 (0.0265)	0.323*** (0.0598)	-0.0968 (0.0871)	-0.115* (0.0483)
Graduate School	-0.000237 (0.0293)	0.550*** (0.107)	-0.119 (0.0938)	-0.217*** (0.0546)
Children	-0.00780 (0.0284)	-0.0187 (0.0691)	0.138 (0.101)	0.0515 (0.0545)
Household Size	0.106*** (0.00945)	-0.106*** (0.0240)	-0.115*** (0.0272)	-0.0119 (0.0183)
Environmental Issues	-0.0473 (0.0328)	0.222** (0.0692)	0.0991 (0.112)	0.100 (0.0643)
Green Living	0.0555 (0.0310)	0.0992 (0.0608)	0.138 (0.107)	0.156** (0.0599)
Smoker	0.0684 (0.0487)	-0.265* (0.114)	-0.231 (0.157)	0.0214 (0.0958)
Dieting	0.0360 (0.0257)	0.273*** (0.0715)	0.152 (0.0817)	0.130* (0.0510)
Gambling-Lottery	0.00608 (0.0459)	0.106 (0.110)	-0.00834 (0.165)	-0.00478 (0.0946)
Gambling-Casino	0.0801* (0.0353)	-0.0406 (0.0802)	0.105 (0.115)	0.0163 (0.0699)
Home loan to value 80% or +	0.0780*** (0.0235)	0.234*** (0.0710)	-0.0818 (0.0846)	-0.0106 (0.0472)
Credit Line	0.0413 (0.0283)	-0.241** (0.0821)	-0.0980 (0.0972)	0.0527 (0.0498)
Observations	28090	12,498	2,487	2,487
R-squared	0.187	0.035		

Standard errors in parentheses, * p<0.05 ** p<0.01 *** p<0.001. Standard errors are clustered at the city level.

The dependent variable in model A corresponds to the log of total electricity consumption in the 12 months before signup

The dependent variable in model B consists of a binary indicator for adoption by household.

The dependent variables in model C consist of a variable which takes values between 1 and 4 corresponding to the different levels of commitment. (0%, 0%-15%, 15%-50%, >50%)

The dependent variable in model D consists of the number of selected actions at the point of adoption for each household.

Models don't report constant or auxiliary parameters.

Table 3: Propensity score matching of the savings ATT.

	(I)	(II)	(III)	(IV)	(V)
ATT	-0.044 (.009)	-0.029 (.033)	-0.044 (.037)	-0.031 (.023)	-0.031 (.027)
Matching variables					
Demographics		Y	Y	Y	Y
2009 usage		Y	-	Y	-
Matching method		Y	Y	-	-
Nearest Neighbor		-	-	Y	Y
Nearest 5 Neighbors					
Sample size	2,487	2,487	2,487	2,487	2,487

In all models the dependent variable is the household monthly log consumption.

Standard errors Models (II)-(V) were obtained using the bootstrap.

Details of the estimation procedure are provided in the Appendix.

Table 4: Summary of empirical findings in relation to model predictions.

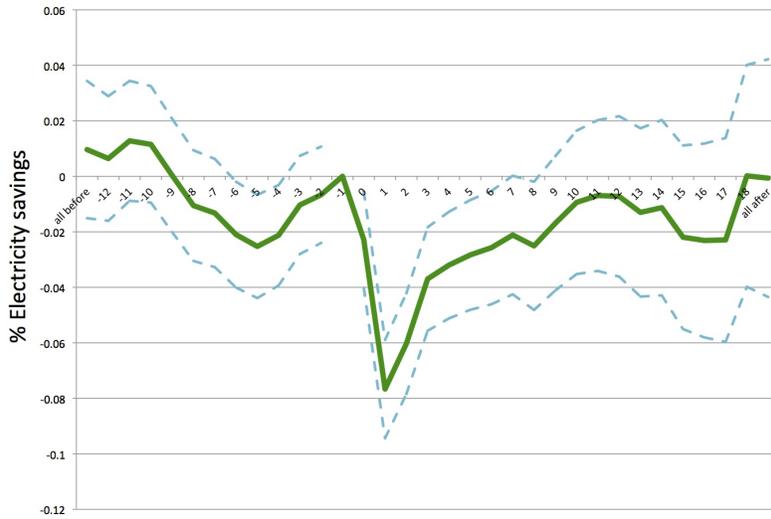
Feature	Finding	Finding Supported by Model Prediction?					
		EU1	EU2	EU3	EU4	PBRD1	PBRD2
Pre-adoption consumption	Hyperbolics consume more (weakly)	No ^A	No ^A	No ^A	No ^A	Yes	Yes
Program Adoption	Sophisticated hyperbolics more likely	No ^A	No ^A	No ^A	No ^A	Yes	Yes
Post-adoption savings	Initial saving then dropoff, for very low goals	Yes ^B	No ^C	Yes	No	Yes	Yes
	Persistently positive for realistic goals	No ^B	Yes ^C	Yes	Yes	Yes	Yes
	Initial saving then dropoff, for very high goals	Yes ^B	No ^C	No	Yes	Yes	Yes
	Hyperbolics save less	No ^A	No ^A	No ^A	No ^A	Yes	Yes
Model Key:	EU1 - standard expected utility (EU) maximizer						EU4 - info-seeking EU maximizer, most responsive to tips set realistic goals
	EU2 - info-seeking EU maximizer						PBRD1 - present-biased, reference dependent
	EU3 - info-seeking EU maximizer, goals reflect preferences						PBRD2 - present-biased, reference dependent, most responsive to goals set realistic goals

^ASince EU models assume no self-control problems, they predict no differences in behavior by variables consistent with present bias.

^BIf reward points are negatively correlated with electricity consumption, EU1 consumers will save after signing up for the program. But since reward points are uncorrelated with consumption, EU1 consumers' behavior will not ultimately change upon adoption.

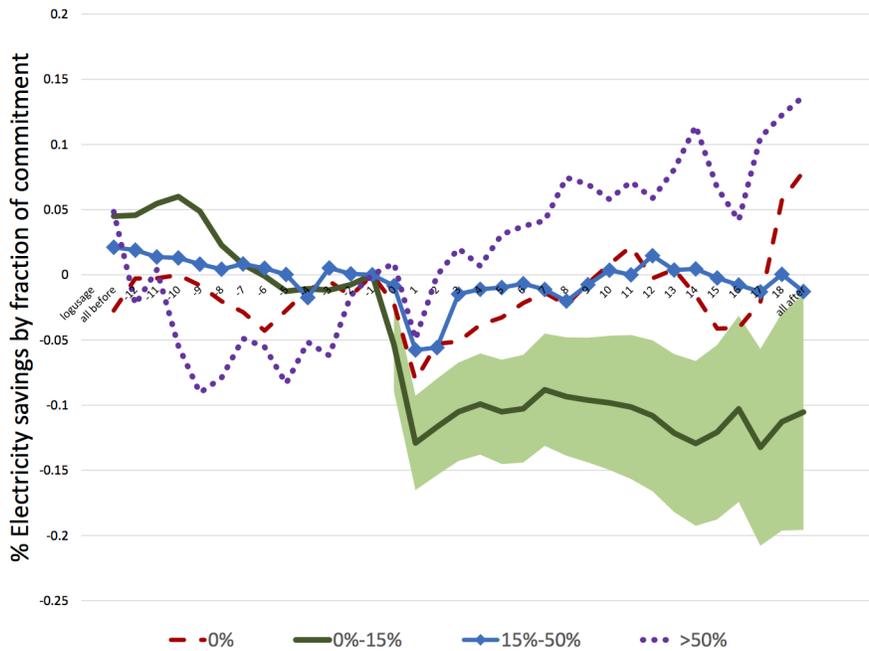
^CThe EU2 model predicts that we should observe the same savings across all goal levels, since consumers all receive the same savings tips. Savings should be persistent across all goal levels, or non-persistent across all goal level.

Figure 1: Estimated percent savings from program adoption in event time.



Month 0 denotes the month before sign up. 95% confidence bounds are plotted around the mean estimates.

Figure 2: Heterogeneity in electricity savings after program adoption in relation to the chosen goal.



Shaded area corresponds to the 95% confidence bounds for the estimated mean savings of the 0%-15% goals group.

Appendix

Proof of Proposition 1

In the second stage, the goal r is taken as given by the consumer. Assuming an interior solution, the consumer's optimal consumption given a goal, (x^*, y^*) , satisfies the first order condition when the consumer has signed up, where $px^* + y^* = m$:

$$\frac{\partial u}{\partial x}(x^*) - p\left(\frac{\partial u}{\partial y}(y^*)\right) - [f'(r - x^*) + g'(x_0 - x^*)] - \beta\delta c'(x^*) = 0. \quad (5.6)$$

If the consumer does not sign up for the program, then her optimal consumption (\tilde{x}, \tilde{y}) satisfies the following condition, where $p\tilde{x} + \tilde{y} = m$:

$$\frac{\partial u}{\partial x}(\tilde{x}) - p\left(\frac{\partial u}{\partial y}(\tilde{y})\right) - \beta\delta c'(\tilde{x}) = 0. \quad (5.7)$$

Applying the implicit function theorem to Equations 5.6 and 5.7 yields:

$$\frac{\partial x}{\partial \beta} = \frac{\delta c'(x)}{\frac{\partial^2 u}{\partial x^2} + 1_{\text{signup}}[f''(r-x) + g''(r-x)] + p^2\left(\frac{\partial^2 u}{\partial y^2}\right) - \beta\delta c''(x)} < 0 \quad (5.8)$$

$$\frac{\partial x}{\partial r} = \frac{f''(r-x)}{\frac{\partial^2 u}{\partial x^2} + f''(r-x) + g''(r-x) + p^2\left(\frac{\partial^2 u}{\partial y^2}\right) - \beta\delta c''(x)} \quad (5.9)$$

$$\frac{\partial x}{\partial \delta} = \begin{cases} \frac{\beta c'(x)}{\frac{\partial^2 u}{\partial x^2} + 1_{\text{signup}}[f''(r-x) + g''(r-x)] + p^2\left(\frac{\partial^2 u}{\partial y^2}\right) - \beta\delta c''(x)} < 0 & \text{if } \beta > 0 \\ 0 & \text{if } \beta = 0. \end{cases} \quad (5.10)$$

Note that $\frac{\partial x}{\partial r}$ is positive when $r > x$ and negative when $r < x$, due to diminishing sensitivity to gains and losses.

Proof of Proposition 2

If the consumer is time-consistent or naive ($\hat{\beta} = 1$), she only signs up if

$$\beta\delta[V(x^*(\hat{\beta}), y^*(\hat{\beta})|r = x^*(\hat{\beta})) - V(\tilde{x}(\hat{\beta}), \tilde{y}(\hat{\beta}))] - \varepsilon \geq 0.$$

When $\hat{\beta} = 1$, then $\tilde{x}(1)$ maximizes $u(x, y) - \delta c(x)$ so

$$u(x^*(1), y^*(1)) - \delta c(x^*(1)) < u(\tilde{x}(1), \tilde{y}(1)) - \delta c(\tilde{x}(1)).$$

Thus, she will only sign up if

$$g(x_0 - x^*(1)) \geq [u(\tilde{x}(1), \tilde{y}(1)) - \delta c(\tilde{x}(1))] - [u(x^*(1), y^*(1)) - \delta c(x^*(1))] + \varepsilon, \quad (5.11)$$

where the right-hand side is strictly positive. If Equation (5.11) is not satisfied, then a consumer with $\hat{\beta} = 1$ does not sign up for the program.

Suppose that Equation (5.11) is not satisfied, and that the consumer is aware of a self-control problem ($\beta < 1$). Since $\frac{\partial x^*}{\partial \beta} < 1$, then $g(x_0 - x^*(1)) > g(x_0 - x^*(\hat{\beta}))$, so rewards are less attractive when $\hat{\beta} < 1$ (because the consumer anticipates earning fewer rewards due to her tendency to over-consume). Thus, the consumer's motivation for sign-up is driven by the need for self-control when $\hat{\beta} < 1$.

Implementation Costs

Suppose that the consumer also incurs a cost of implementing goals that result in reductions in energy usage relative to not signing up. Many of the suggested actions for saving energy, such as installing CFLs into lighting fixtures or regularly unplugging appliances after use, require additional effort or inconvenience to implement them. Due to diminishing marginal returns, the cost of saving energy is increasing and convex. The function $h(\tilde{x} - x)$ describes the cost of implementing energy-saving measures, where \tilde{x} is the consumer's energy consumption in the absence of program sign-up:

$$h(0) = 0, \quad (5.12)$$

$$h(\tilde{x} - x) \begin{cases} > 0 & \text{if } \tilde{x} \geq x \\ = 0 & \text{if } \tilde{x} < x, \end{cases} \quad (5.13)$$

$$h'(\tilde{x} - x) \begin{cases} > 0 & \text{if } \tilde{x} > x \\ = 0 & \text{if } \tilde{x} < x, \end{cases} \quad (5.14)$$

$$h''(\tilde{x} - x) \begin{cases} > 0 & \text{if } \tilde{x} > x \\ = 0 & \text{if } \tilde{x} < x. \end{cases} \quad (5.15)$$

We do not model $h(\cdot)$ as a direct function of the goal r , since the consumer would not incur any additional implementation cost if she were to set an unambitious goal that she would exceed without any additional effort (e.g., if $r > \tilde{x} > x$). Rather, implementation costs only are incurred when the agent decides to use less electricity than she would have had she not signed up for the program. The goal indirectly affects implementation costs through its effect on the consumer's

consumption decision.

Thus, the consumer's problem in the second stage is

$$\max_{x,y} u(x,y) + 1_{\text{signup}}[f(r-x) + g(x_0-x) - h(\tilde{x}-x)] - \beta\delta c(x) \quad \text{subject to } m \geq px + y, \quad (5.16)$$

where $1_{\text{signup}} = 1$ if the consumer signs up for the program and is zero otherwise. When deciding whether to sign up for the program, she takes into account her consumption behavior in response to the goal, her anticipated costs of implementing them, and the reward points she would earn.

With the inclusion of the goal implementation cost, the consumer's energy consumption in the absence of program sign-up, (\tilde{x}, \tilde{y}) , is still described by Equation (5.7). Her optimal consumption upon sign-up now satisfies Equations (5.17) and (5.18):

$$\frac{\partial u}{\partial x}(x^*) - p\left(\frac{\partial u}{\partial y}(y^*)\right) - [f'(r-x^*) + g'(x_0-x^*) - h'(\tilde{x}-x^*)] - \beta\delta c'(x^*) = 0, \quad (5.17)$$

$$px^* + y^* = m. \quad (5.18)$$

The consumer would never sign up for the program if the cost of goal implementation were so high that she would end up using more electricity after sign-up. Thus, $f'(r-x^*) + g'(x_0-x^*) - h'(\tilde{x}-x^*) \geq 0$ and $x^* \leq \tilde{x}$. We now re-derive the main propositions for the model with implementation costs.

Proof that $x^* \leq \tilde{x}$ Suppose that $x^* > \tilde{x}$, or equivalently that $\tilde{x} - x^* < 0$. Comparing Equations (5.7) and (5.6), $x^* > \tilde{x}$ only if $f'(r-x^*) + g'(x_0-x^*) - h'(\tilde{x}-x^*) < 0$. But if $\tilde{x} - x^* < 0$, then $h'(\tilde{x}-x^*) = 0$. Since $f'(r-x^*) > 0$ and $g'(x_0-x^*) \geq 0$, then it must be that $x^* \leq \tilde{x}$ from Equations (5.7) and (5.6). Thus, $x^* \leq \tilde{x}$.

Proof of Proposition 1: Applying the implicit function theorem to Equations 5.6 and 5.7 yields:

$$\frac{\partial \tilde{x}}{\partial \beta} = \frac{\delta c'(\tilde{x})}{\frac{\partial^2 u}{\partial x^2}(\tilde{x}, \tilde{y}) + p^2 \left(\frac{\partial^2 u}{\partial y^2}(\tilde{x}, \tilde{y}) \right) - \beta \delta c''(\tilde{x})} < 0 \quad (5.19)$$

$$\frac{\partial x^*}{\partial \beta} = \frac{\delta c'(x^*) - h'(\tilde{x} - x^*) \left(\frac{\partial \tilde{x}}{\partial \beta} \right)}{\frac{\partial^2 u}{\partial x^2}(x^*, y^*) + f''(r - x^*) + g''(r - x^*) - h''(\tilde{x} - x^*) + p^2 \left(\frac{\partial^2 u}{\partial y^2}(x^*, y^*) \right) - \beta \delta c''(x^*)} < 0 \quad (5.20)$$

$$\frac{\partial x^*}{\partial r} = \frac{f''(r - x^*)}{\frac{\partial^2 u}{\partial x^2}(x^*, y^*) + f''(r - x^*) + g''(r - x^*) - h''(\tilde{x} - x^*) + p^2 \left(\frac{\partial^2 u}{\partial y^2}(x^*, y^*) \right) - \beta \delta c''(x^*)} \quad (5.21)$$

$$\frac{\partial \tilde{x}}{\partial \delta} = \begin{cases} \frac{\beta c'(\tilde{x})}{\frac{\partial^2 u}{\partial x^2}(\tilde{x}, \tilde{y}) + p^2 \left(\frac{\partial^2 u}{\partial y^2}(\tilde{x}, \tilde{y}) \right) - \beta \delta c''(\tilde{x})} < 0 & \text{if } \beta > 0 \\ 0 & \text{if } \beta = 0. \end{cases} \quad (5.22)$$

$$\frac{\partial x^*}{\partial \delta} = \begin{cases} \frac{\beta c'(x^*) - h'(\tilde{x} - x^*) \left(\frac{\partial \tilde{x}}{\partial \delta} \right)}{\frac{\partial^2 u}{\partial x^2}(x^*, y^*) + f''(r - x^*) + g''(r - x^*) - h''(\tilde{x} - x^*) + p^2 \left(\frac{\partial^2 u}{\partial y^2}(x^*, y^*) \right) - \beta \delta c''(x^*)} < 0 & \text{if } \beta > 0 \\ 0 & \text{if } \beta = 0. \end{cases} \quad (5.23)$$

Note that $\frac{\partial x}{\partial r}$ is positive when $r > x$ and negative when $r < x$, due to diminishing sensitivity to gains and losses.

Proof of Proposition 2: If the consumer is time-consistent or naive ($\hat{\beta} = 1$), she only signs up if

$$\beta \delta [V(x^*(\hat{\beta}), y^*(\hat{\beta}) | r = x^*(\hat{\beta})) - V(\tilde{x}(\hat{\beta}), \tilde{y}(\hat{\beta}))] - \varepsilon \geq 0.$$

When $\hat{\beta} = 1$, then $\tilde{x}(1)$ maximizes $u(x, y) - \delta c(x)$ so

$$u(x^*(1), y^*(1)) - \delta c(x^*(1)) < u(\tilde{x}(1), \tilde{y}(1)) - \delta c(\tilde{x}(1)).$$

Thus, she will only sign up if

$$g(x_0 - x^*(1)) \geq [u(\tilde{x}(1), \tilde{y}(1)) - \delta c(\tilde{x}(1))] - [u(x^*(1), y^*(1)) - \delta c(x^*(1))] + h(\tilde{x}(1) - x^*(1)) + \varepsilon, \quad (5.24)$$

where the right-hand side is strictly positive. If Equation (5.11) is not satisfied, then a consumer with $\hat{\beta} = 1$ does not sign up for the program.

Suppose that Equation (5.24) is not satisfied, and that the consumer is aware of a self-control problem ($\beta < 1$). Since $\frac{\partial x^*}{\partial \beta} < 1$, then $g(x_0 - x^*(1)) > g(x_0 - x^*(\hat{\beta}))$, so rewards are less attractive when $\hat{\beta} < 1$ (because the consumer anticipates earning fewer rewards due to her tendency to over-

consume). Thus, the consumer's motivation for sign-up is driven by the need for self-control when $\hat{\beta} < 1$.

Robustness of Savings Behavior Using Matching

Our estimate of the extent to which households save as a result of this program may depend on the extent to which individuals with different characteristics opt into the program at different points in time. To evaluate the robustness of these savings estimates, we use propensity score matching to account for the potentially varying demographics of the adopters over time as they select into the program. We follow the method for computing the average treatment effect on the treated (ATT) introduced in Sianesi (2004), by comparing those individuals who adopt at time t with those who have not adopted up until t (they adopt in future periods), while correcting for discrepancies in observables.

First, we determine the treatment and control groups. Each period u has its own treatment and control group. The treatment is starting the program in month u . The control group at time u is composed of all those that have not joined the program yet. Formally, let D_{it} be equal to 1 if the individual i had join the program by period t . Then the treatment and the control groups at time u , T^u and C^u respectively, are defined as follows:

$$i \in T^u \Leftrightarrow D_{iu} = 1 \ \& \ D_{it} = 0 \ \forall t < u$$

$$i \in C^u \Leftrightarrow D_{it} = 0 \ \forall t \leq u.$$

In our case, we consider 19 months (June 2010-December 2011) when the individuals could have adopted, i.e. $u = 1, 2, \dots, 19$. Hence, we have 19 control and treatment groups.

Second, for each period u we are interested in obtaining the average impact at time t in log electricity usage, for those adopting the program in their u^{th} month, Δ_t^u . Δ_t^u is defined as:

$$\Delta_t^u \equiv E(Y_t^{1(u)} - Y_t^{0(u)} | D^{(u)} = 1); \ t \geq u. \quad (5.25)$$

Each Δ_t^u is estimated using propensity score matching. For each of the 19 treatment and control groups, we perform propensity score matching to obtain the ATT on log usage at time t . We match on the available demographics and on pre-adoption monthly usage, and also employ several varieties of the nearest neighbor method for matching.

Finally, we are interested in the synthetic overview of adoption's effect at time t . That is, we want to know the effect on electricity usage at period t after adoption, independently of when the individual adopted. Hence, we compute the average Δ_t^u , by weighting the different ATT's by the

proportion of treated that join in period u :

$$E_U(\Delta_t^u | D = 1) = \sum_{u \leq t} E(Y_t^{1(u)} - Y_t^{0(u)} | D(u) = 1) \times P(D^{(u)} = 1 | D = 1), \quad (5.26)$$

where $P(D^{(u)} = 1 | D = 1)$ is estimated as the number of individuals that joined in period u over the total individuals that joined up to period t . We compute the standard errors by 500 bootstrap repetitions.

For presentational simplicity, we summarize the estimation results in Table 3 by reporting the average of the estimated savings over the post-adoption period (instead of reporting period by period estimates). Model (I) estimates the difference-in-differences equation (Equation 4.5) using one post-adoption indicator. The average estimated post-adoption savings for consumers who opt in is 4.4%. The different matching approaches produce estimates ranging from 2.9% to 4.4%.¹⁹

¹⁹Note that the standard errors increase as a result of the matching algorithm utilizing only a small subset of the observations at each step. This is unavoidable given the limited data available.