NBER WORKING PAPER SERIES

DO THE EFFECTS OF SOCIAL NUDGES PERSIST? THEORY AND EVIDENCE FROM 38 NATURAL FIELD EXPERIMENTS

Alec Brandon
Paul J. Ferraro
John A. List
Robert D. Metcalfe
Michael K. Price
Florian Rundhammer

Working Paper 23277
http://www.nber.org/papers/w23277

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2017

We thank Arhan Gunel and Nancy Hersch for sharing the data. Hunt Allcott, Michael Greenstone, Jim Kapsis, Marc Laitin, David Novgorodsky, Steve Puller, and David Rapson offered helpful comments. We would also like to thank participants of the American Economic Association annual conference, the Muenster Energy Conference, the Association of Environmental and Resource Economists annual conference, Camp Resources XXII, the University of Chicago, and Georgia State University. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Do The Effects of Social Nudges Persist? Theory and Evidence from 38 Natural Field Experiments
Alec Brandon, Paul J. Ferraro, John A. List, Robert D. Metcalfe, Michael K. Price, and Florian Rundhammer
NBER Working Paper No. 23277
March 2017
JEL No. C93,D01,D03,Q4

ABSTRACT

This study examines the mechanisms underlying long-run reductions in energy consumption caused by a widely studied social nudge. Our investigation considers two channels: physical capital in the home and habit formation in the household. Using data from 38 natural field experiments, we isolate the role of physical capital by comparing treatment and control homes after the original household moves, which ends treatment. We find 35 to 55 percent of the reductions persist once treatment ends and show this is consonant with the physical capital channel. Methodologically, our findings have important implications for the design and assessment of behavioral interventions.

Alec Brandon
Department of Economics
University of Chicago
5757 S. University Ave.
Chicago, IL 60637
alec.brandon@gmail.com

Paul J. Ferraro
Bloomberg School of Public Health
Carey School of Business
Whiting School of Engineering
Johns Hopkins University
100 International Dr.
Baltimore, MD 21202
pferraro@jhu.edu

John A. List
Department of Economics
University of Chicago
1126 East 59th
Chicago, IL 60637
and NBER
jlist@uchicago.edu

Robert D. Metcalfe
University of Chicago
Saieh Hall for Economics
5757 S. University Avenue
Chicago IL, 60637
rdmetc@gmail.com

Michael K. Price
Experimental Economics Center
Andrew Young School of Policy Studies
Georgia State University
P.O. Box 3992
Atlanta, GA 30302-3992
and NBER
mprice25@gsu.edu

Florian Rundhammer
The Department of Economics
Andrew Young School of Policy Studies
Georgia State University
PO Box 3992
Atlanta, GA 30302-3992
florian.rundhammer@gmail.com

I. Introduction

The textbook approach to dynamic decision making assumes away any direct dependence between choices made over time. Instead, past decisions interact with future choices through state variables, such as the stock of physical capital (Ramsey, 1928) or natural resources (Hotelling, 1931). The tractability of this framework has led to remarkable progress in our understanding of dynamics ranging from inflation and unemployment (Lucas, 1972) to energy policy and climate change (Nordhaus, 1994). Yet, there are many economic settings where the utility an agent experiences may not prove separable from one time period to the next. Prime examples include choices motivated by addiction, habit, or tradition.

Dating back to Pollak (1970), researchers have modeled these motivations by reimagining state variables as stocks of addiction, habit, or tradition capital. Broadly deemed habit formation, this modeling framework has proven invaluable for explaining a wide range of behaviors that are outside a parsimonious neoclassical formulation of behavior. Examples include consumption of tobacco products (Becker et al., 1994), behavior in financial markets (Constantinides, 1990 and Campbell and Cochrane, 1999), and the dynamics of economic growth and the business cycle (Carrol et al., 2000 and Boldrin et al., 2001).

But in spite of the explanatory power of habit formation models, it has proven remarkably difficult to design programs that successfully induce habitual behavior. For example, social programs that targeted habits in the 1960s and 1970s were notoriously unsuccessful (see, e.g., Rossi, 1987). More recent work has, by and large, rediscovered this fact via field experiments designed to test the theoretical predictions of habitual behavior derived in Becker and Murphy (1988). Generally, these studies go as follows. A financial incentive is introduced to motivate behavioral change. Under a habit formation model, the change in behavior spurred by the incentive will lead to an increase in habitual capital for the behavior. The incentive is then removed and the extent to which the behavioral change persists is observed.

Figure 1 presents a summary of this work by plotting the proportion of the incentive effect that persists after the incentive is removed.[1,2] Across domains ranging from charitable giving (Meier, 2007 and Landry et al., 2010), education (Levitt et al., 2016), exercise (Charness and Gneezy, 2009; Milkman et al., 2014; Royer et al., 2015; and Acland and Levy, 2015), smoking cessation (Volpp et al., 2006 and Volpp et al., 2009), and weight loss (Volpp et al., 2008), two stark patterns emerge. First, treatment effects rarely persist. Only four of the ten studies find behavior consistent with habit formation. Second, when effects do persist, they decay rapidly. Only two of the ten studies find more than 25 percent of the initial effect persisting after a month, and only one study finds any persistence after six months.

One remarkable exception to this trend is a program called the Home Energy Report (HER), a social nudge that promotes energy conservation through periodic mailers comparing a household's energy use to that of similar neighbors.[3] Using data from a series of natural field experiments, Allcott and Rogers (2014), AR henceforth, consider the lifecycle of the behavioral response to the HER. Perhaps the most provocative finding AR report is the persistence of the conservation behavior induced by the HER. As Figure 1 illustrates, the HER offers an unparalleled dividend years after the program has ended.

To understand the mechanisms underlying the long-run impacts, AR consider the extent to which the HER's persistence is driven by changes in habit capital or physical capital, such as increasing investments in energy efficient technology. To do so, they assess the effect of the HER on participation in utility-run energy efficient technology programs. They find that the HER induces no appreciable increase in participation in these programs, which suggests that the persistence of the HER is driven by changes in habits as opposed to changes in capital stock within the home. AR are careful to note,

[1] Studies were limited to field experiments that observe behavior during and after a financial incentive is used to promote behavioral change. If more than one incentive was used then the largest is reported. The time dimension reflects the end of the time period of the reported estimate. Nulls are reported if effects are not statistically significant at the 0.05 level. Studies that utilize commitment devices are excluded, as their predictions are outside of a strict habit formation model. See Online Appendix Table 1 for estimates.

[2] There is also a nascent literature on the persistence of interventions that target margins outside of the Becker and Murphy (1988) model. We point the interested reader to Rogers and Frey (2015) and the citations therein.

[3] The HER is a product offered by the company Opower. Opower is the world leader in software-based solutions for utilities. Due to implementation as natural field experiments, HERs have been studied extensively, including Allcott (2011; 2015), Ayres et al. (2013), and Costa and Kahn (2013).

however, that unobserved adoption of energy efficient technology outside of these programs may play a role in their estimates of persistence. Hence the channels through which persistence arises—changes in habits or changes in capital stock—remains an open question.

Yet, if even a small fraction of the persistence were driven by a change in habit capital, AR would represent the most compelling existent evidence of a program spurring habit formation. From a positive perspective, these results would point to a strong complementarity in social norms and the formation of habit capital. Importantly, such complementarity would validate the predictions of models that have previously considered such (see, e.g., Becker and Murphy, 2000, pp. 18-20, 152-156) and would offer a potential explanation for the past failure of programs designed to affect habits using only financial incentives. From a normative perspective, these results highlight the value of programs that change habits because they provide a long-lived stream of benefits at little to no direct or indirect costs, which significantly enhances the cost effectiveness and gains to welfare.[4]

We complement AR by taking a new approach to identifying the mechanisms driving the HER's persistence. We start by detailing a wrinkle in Opower's administration of the HER that allows us to isolate the role of technology adoption on long-run patterns of energy use: upon the sale of a home, Opower ceases all messaging for both the incoming and outgoing households but continues to observe energy consumption.

We then develop a multi-period model of household energy consumption and capital investment behavior that captures this wrinkle. Our model assumes that receipt of the HER changes the shadow price of energy consumption and that investments in physical capital are immobile. Through the lens of the model, we learn that distinguishing between the household and the home points to an empirical strategy for estimating unobserved adoption of energy efficient technology using only observations of energy

[4] For example, AR show that accounting for persistence more than doubles the cost effectiveness of the HER. Although AR discuss welfare consequences of the social costs of responding to the HER, it's unclear if there is actually a gap between choices and welfare in the case of Opower, as households are free to opt-out of the mailer (DellaVigna et al., 2012). For a broader discussion of the welfare effects of the HER see Allcott and Kessler (2015).

consumption.[5] Intuitively, our strategy identifies capital investments induced by the HER via a comparison of the HER's effect before and after the home changes hands, with any persistence attributed to capital investments inherited from the initial residents.

We then link our economic model to an empirical model of energy consumption and estimate the model using more than 9 million observations of monthly electricity consumption. These observations come from more than 250,000 homes that see a change in ownership over 38 natural field experiments implemented by Opower to test the HER. We find that homes that change ownership react to the HER in the same way as the full sample of all homes exposed to the HER: they reduce their electricity consumption by 2.4 percent, on average. Interestingly, approximately 35 to 55 percent of that effect persists in the home after the household moves and HER delivery ends. We further show that this estimated persistence in the home is robust to a battery of empirical specifications, data inclusion rules, and approaches to incorporating heterogeneity. Additionally, we show that our estimates of persistence in the home conform to higher order predictions of our theoretical model and we find no evidence to support alternative models of behavior that stress the role of sorting. Taken jointly, we interpret these findings as suggesting that the HER serves to induce changes in capital stock and that such changes are important drivers of persistence.

Relating our findings to AR, we test whether our estimates of persistence in the home due to physical capital differ from AR's estimates of persistence in the household due to habits and physical capital. We find only weak evidence that our estimates are statistically different from the estimates in AR; in fact, for the subsample most similar in spirit to AR, initial residents that were exposed to treatment for at least two years, we find persistence indistinguishable from AR. We interpret this evidence as supporting an important role for physical capital in driving persistence and casting uncertainty over the complementarity between social nudges and habit formation. Finally, from a policy perspective, we consider the importance of incorporating both direct and indirect costs into estimates of cost-effectiveness. We find that previous estimates of the HER's cost

[5] To the best of our knowledge, the only studies to previously consider an identification strategy like this are Costa and Kahn (2010) and Bernedo et al. (2014).

per kWh saved are more than doubled when the costs of physical capital induced by the HER are incorporated.

The remainder of our study proceeds as follows. In Section II we review the administration of the HER by Opower and detail the sample underlying our investigation. Section III then considers our theoretical framework, which motivates the empirical strategy for identifying unobserved investments in physical capital. Section IV presents our empirical results, considers robustness, and assesses alternative explanations. Section V discusses the implications of our findings and Section VI concludes.

II. Data and Summary Statistics

Our empirical analysis is based on a collection of 38 natural field experiments, henceforth waves, implemented by Opower across 21 utilities between 2008 and 2014. Each wave utilizes the same experimental design. Households with at least twelve billing months of energy service are randomly assigned to a Treatment or Control group.[6] Households in the Control group are left untouched, while households assigned to the Treatment group receive a periodic mailer from Opower called the Home Energy Report (HER).[7] Figure 2 presents an example of the HER, which utilizes a comparison of household and neighborhood energy usage, energy conservation tips, and information on energy-efficient technologies to motivate energy conservation.

For each wave, we observe four pieces of administrative data via a data sharing agreement with Opower. First, we observe monthly electricity consumption in kilowatt hours (kWh) for each home. Second, we observe Opower's assignment of each household to the Treatment or Control group. Third, we observe the timing of the intervention in each wave, which allows us to partition home-month observations of electricity consumption into a Pre-Treatment or Treatment time period. Fourth, we observe the timing (if any) of a household's deactivation of their energy service. Importantly, after

[6] Opower shared with us a total of 41 waves. We exclude one wave where there are no movers assigned to the Control group. We also exclude two waves that fail a Kolmogorov-Smirnov test of equality in the distribution of average pre-treatment usage at the five percent level. Online Appendix Tables 2-4 provide more details and summary statistics. Excluding waves based on a threshold of ten percent leads to three waves being excluded and does not change the results qualitatively.

[7] Across the 38 waves, the HER is received monthly, bimonthly, or quarterly. Previous work suggests that frequency does not substantially affect the effectiveness of the intervention (Allcott, 2011) and we adopt the approach of AR to pool across different frequencies.

deactivation Opower still observed monthly electricity consumption in the home but immediately ceased transmission of the HER and excluded the home and household from subsequent waves of the HER. We deem the subsample of homes with an account that is deactivated *movers* and further partition time for these homes with a Post-Move time period.[8]

Table 1 provides summary statistics for the full sample, the non-movers subsample, and the movers subsample. Comparing across the columns in Table 1, we see that slightly more than 10 percent of the homes in the full sample change hands, about two-thirds of the homes are assigned to the Treatment group regardless of the subsample, and monthly electricity usage is comparable but lower for the movers subsample relative to the non-movers subsample. Moving down the summary statistics on the movers subsample, we also see the average number of monthly electricity consumption observations per home across the Pre-Treatment, the Treatment and Pre-Move, and Post-Move time periods. In particular, there is an average of about 12 observations per home and time period, with more variation in this number of observations in the Treatment and Pre-Move and Post-Move time periods relative to the Pre-Treatment time period, which reflects the heterogeneity in time of move.

III. Conceptual Framework

To motivate our empirical strategy, we present a model of household energy consumption, investment in energy efficient technology, and moral suasion for households in our sample. Our goal in presenting this model is to link the parameters identified by the reduced form empirical model estimated in Section IV with a richer economic model of behavior. This allows us to be precise about the identifying assumptions of our empirical model and provides a foundation for assessing alternative interpretations of the data.

[8] Based on discussions with Opower's staff, we determined that some account deactivations did not constitute a move due to one of four reasons: i) accounts become inactive because of changes to the name or marital status of the account holder, e.g. marriage or divorce; ii) the deactivation date falls before the date of the first HER; iii) the deactivation date falls after the last usage observation for a given home, effectively making the household a non-mover account; and iv) the deactivation occurred because the household switched energy service providers. We exclude these households from the movers subsample. See Online Appendix Table 5 for results that include these households in the movers sample. Results are qualitatively unchanged.

IIIA: Setup

The setup of our models goes as follows. We imagine a household in time period t allocating a financial endowment, m, between a numeraire consumption good, c_t, and inputs for household production of a final consumption good, z_t.[9] We assume household production is described by an increasing and concave function, $f(\cdot)$, with energy, e_t, and a stock of technology, k_t, as the inputs. That is, $z_t = f(e_t, k_t)$. Furthermore, we assume households purchase energy at a fixed per-unit price, p_e, and vary their technology stock by investing in new technology, I_t, at a fixed per-unit price, p_I, with the technology stock evolving according to $k_t = I_t + k_{t-1}$.[10] In the spirit of the findings in Allcott (2011) and others we assume that energy consumption behavior partly reflects pro-social motivations, with households paying a moral suasion cost, s_t, for their consumption of energy. We assume this moral suasion cost varies with receipt of the HER via a_t, with the relationship between messaging and moral suasion costs described by $s_t = g(e_t, a_t)$, where a_t is a metapreference parameter in the vein of Becker and Murphy (1993), with $g(\cdot)$ increasing in its arguments.[11, 12] With this setup established, we next consider the optimization problem that households face during the Pre-Treatment, Treatment and Pre-Move, and Post-Move time periods.

IIIB: Pre-Treatment and Treatment and Pre-Move Time Periods

Consider a household in time period $\tau \in \{0, 1\}$ with additively separable preferences in utility from the numeraire, $u(c_\tau)$, and household production, $v(z_\tau)$:

[9] A specific example of z_t relevant to our analysis would be comfortable ambient temperatures. For a more general treatment of household input decisions see Becker (1965), which focuses on time as an input.
[10] I_t can be thought of as investments net of depreciation.
[11] We focus on a model with s_t increasing in a_t for all households. This precludes so-called boomerang effects. This choice is supported by the data, see Allcott (2011), Ferraro and Price (2013), and Ferraro and Miranda (2013) for a discussion, but this assumption is not restrictive in the sense that if we observe treatment prompting energy increases, our model would still describe how to relate that to technology investments.
[12] As has been studied in Herberich et al. (2011) and Allcott and Taubinsky (2015), households could also receive utility from green technology purchases. Our framework ignores this motivation, although inclusion would only strengthen the predicted dynamics.

$$
\begin{aligned}
&\max_{c_\tau, e_\tau, I_\tau} \quad u(c_\tau) + v(z_\tau) - s_\tau \\
(1) \qquad &s.t. \quad m = c_\tau + p_I I_\tau + p_e e_\tau \\
&\qquad z_\tau = f(e_\tau, k_\tau) \\
&\qquad k_\tau = I_\tau + k_{\tau-1} \\
&\qquad s_\tau = g(e_\tau, a_\tau)
\end{aligned}
$$

where k_{-1} given. With utility linear in the numeraire and increasing and concave in z_τ, the solution to (1) is given by the following first-order conditions:

$$
\begin{aligned}
(2) \qquad & v'(z_\tau) f_e(e_\tau, k_\tau) = p_e + g_e(e_\tau, a_\tau) \\
& v'(z_\tau) f_k(e_\tau, k_\tau) = p_I
\end{aligned}
$$

which tells us that households choose e_τ and I_τ to balance the marginal benefits and costs of z_τ. Importantly, the marginal cost of energy includes $g_e(e_\tau, a_\tau)$, which acts like a shadow price on energy consumption. Varying a_τ is then akin to varying the price of energy (perhaps non-linearly) and because households consume z_τ, not e_τ or k_τ, (2) tells us they will respond to a price change by varying both inputs, e_τ and I_τ.

The dynamics of energy consumption, pro-social incentives, and technology investment become clearer with a flexible parameterization of $v(\cdot), g(\cdot)$, and $f(\cdot)$. In particular, we assume isoelastic utility in z_t, $v(z_t) = (z_t^{1-\sigma} - 1)/(1-\sigma)$ with $\sigma > 1$,[13] a linear moral suasion cost function, $g(e_t, a_t) = e_t a_t$, and a Cobb-Douglas production function of z_t, $z_t = e_t^\alpha k_t^{1-\alpha}$ with $\alpha \in (0,1)$.[14] Under this parameterization (2) yields the following demand functions for energy and technology investments in period $\tau \in \{0,1\}$:

[13] The law of demand imposes $\sigma > 0$ on the typical isoelastic utility function. Following the bulk of research on in-home production that uses energy inputs we assume that implied demand for z_t is inelastic, giving us $\sigma > 1$. See, e.g., Small and Van Dender (2007), Davis (2008), Hughes et al. (2008), and Davis and Kilian (2011) for estimates.

[14] The Cobb-Douglas production function is meant to capture an unrestrictive parameterization of the household's production function in the sense that it's a first-order approximation of a general production function. See Syverson (2011) for a discussion.

$$(3) \qquad e_\tau(a_\tau) = \left(\frac{\alpha}{1-\alpha}\frac{p_I}{p_e + a_\tau}\right)^{\frac{\alpha+\sigma(1-\alpha)}{\sigma}} \left(\frac{1-\alpha}{p_I}\right)^{\frac{1}{\sigma}}$$

$$I_\tau(a_\tau) = \left(\frac{\alpha}{1-\alpha}\frac{p_I}{p_e + a_\tau}\right)^{\frac{\alpha(1-\sigma)}{\sigma}} \left(\frac{1-\alpha}{p_I}\right)^{\frac{1}{\sigma}} - k_{\tau-1}.$$

The demand functions in (3) demonstrate two important dynamics. First, there is an unambiguous energy demand response from an increase in the moral suasion cost (a_τ), with $e'_\tau(a_\tau) < 0$. Second, investment and moral suasion (a_τ) are positively related for $\sigma > 1$. That is, if preferences imply inelastic demand for z_τ then $I'_\tau(a_\tau) > 0$. While this first prediction has been extensively tested there is very little evidence on the responsiveness of technology investment to changes in moral suasion. The best estimates use data on participation in energy efficiency programs run by electric utilities and find $I'_\tau(a_\tau) \approx 0$ (see AR and citation therein). However, as electric utility programs often focus on a limited subset of all available technologies these estimates could be severely attenuated.

To summarize, Figure 3 plots a time series of energy consumption. During $t = 0$, the Pre-Treatment time period, all households consume the same amount of energy, $e_0(a_0)$. Then once treatment starts in $t = 1$, the Treatment and Pre-Move time period, households are randomly assigned to $a_1 \in \{a_1^{ctrl}, a_1^{trt}\}$, where $a_1^{trt} > a_1^{ctrl} = a_0$. Then $\delta^{trt} = e_1(a_1^{trt}) - e_1\left(a_1^{ctrl}\right) < 0$ is a reduced form measurement of the effect of treatment (i.e., a reduction) on energy use through behavioral and technological adjustments.

IIIC: Post-Move Time Period

Now consider the decision problem in the Post-Move time period, $t = 2$. We assume that new residents inherit the capital and investments of the past homeowner, i.e., $k_2 = k_1(a_1) = I_1(a_1) + k_0$, but face a new meta-preference parameter, a_2^{new}, and treat capital as fixed in the Marshallian sense of the short-run.[15,16] This leads to the following decision problem in the Post-Move time period:

$$
\begin{aligned}
(4) \qquad & \max_{c_2,e_2} \quad u(c_2) + v(z_2) - s_2 \\
& s.t. \quad m = c_2 + p_e e_2 \\
& \qquad\quad z_2 = f\big(e_2, k_1(a_1)\big) \\
& \qquad\quad s_2 = g(e_2, a_2^{new}).
\end{aligned}
$$

Then given the same parameterization as above, solving (4) yields the following energy demand function for new residents:

$$
(5) \qquad e_2(a_2^{new}, k_1(a_1)) = \left(\frac{\alpha k_1(a_1)^{(1-\alpha)(1-\sigma)}}{p_e + a_2^{new}}\right)^{\frac{1}{(1-\alpha)+\sigma\alpha}}.
$$

This demand function illustrates how technology investments induced by a_1 in period $t = 1$ can persist in the energy consumption decision of a new resident in period $t = 2$. To see this, first recall that $t = 1$ investments are increasing in a_1, i.e., $dI_1/da_1 > 0$. Furthermore (5) shows us that $de_2/dI_1 < 0$. Combining we get $de_2/da_1 < 0$.

Figure 3 illustrates these predictions in the context of a time series of energy consumption choices. Moving from $t = 1$ to $t = 2$, a home assigned to the Control group does not see their energy consumption change, i.e., $e_1\big(a_1^{ctrl}\big) = e_2(a_2^{new}, k_1\big(a_1^{ctrl}\big))$, but homes assigned to Treatment see an increase, i.e., $e_1(a_1^{trt}) < e_2\big(a_2^{new}, k_1(a_1^{trt})\big)$ because all new residents faces the same shadow price on energy consumption. Nonetheless, the Treatment home's energy consumption remains below control because of the inherited technology investments, with $\delta^{move} = e_2(a_2^{new}, k_1(a_1^{trt})) - e_2(a_2^{new}, k_1\big(a_1^{ctrl}\big)) < 0$ measuring the extent of treatment

[15] Here we assume a_2^{new} is orthogonal to a_1. This assumption seems justified, as the HER only causes about $25 of savings per year in energy costs. Nonetheless, in Section IVD we consider the consequences of relaxing this assumption.

[16] While our short-run assumption is a strong one, it only attenuates the effect of the HER in the post-move period.

persisting through those investments. Importantly δ^{move} gives us a measure of moral suasion's effect on household technology investment that does not suffer from the attenuation bias of earlier estimations because it indirectly observes all investments instead of using direct observation of a small subset of possible investments.[17]

IV. Empirical Results

We start our investigation by considering an event study of the homes in the movers sample. Figure 4 plots the difference in average energy consumption between homes assigned to Treatment and Control in the Pre-Treatment, Treatment and Pre-Move, and Post-Move time periods normalized to the Pre-Treatment time period.[18] The figure highlights two important results. First, the HER reduces energy consumption in our sample of movers. Second, nearly half of the reduction caused by the HER persists in the home after the initial household moves.

The remainder of this section considers the robustness of these findings. We start by analyzing the full movers sample in a difference-in-differences framework. This allows us to control for characteristics of the time series and cross-section that differ across waves and to assess the sensitivity of our results to different definitions of the Post-Move time period. We then follow Allcott (2015) and utilize an empirical strategy that treats each of the 38 waves as a separate experiment. This framework proves valuable both for establishing robustness and assessing predictions derived in Section III. Finally, we test alternative interpretations of the results in Figure 4.

IVA. General Results

Consider a simple empirical model of energy demand:[19]

$$e_{ijt} = \beta^T T_i + \beta^H H_t + \delta^{trt}\, T_i H_t + \beta^M M_t + \delta^{move}\, T_i M_t + \omega_j + \tau_t + U_{ijt} \quad (6)$$

[17] Utilities also often utilize rebate programs directly at a retailer. In other words, utilities subsidize appliances and other technology purchased at a retailer without being able to attribute purchases to a particular home. Furthermore, buyers from surrounding municipalities can also take advantage of these offers, further complicating attribution to the utility's customer stock.

[18] See Online Appendix Table 4 for levels of consumption over these event study time periods.

[19] Note that here we have changed the meaning of the time index t from event time periods to months because in (6) the indicator functions act to divide the sample into the three event time periods.

where e_{ijt} is energy use in month t for home i in experiment wave j, T_i is an indicator for the home's treatment status, H_t is an indicator denoting the time period in which treatment households receive HERs, M_t denotes the post-move event period for each home, ω_j is an indicator for the experimental waves, τ_t is an indicator for month of sample across all waves, and U_{ijt} is a home/wave/month varying unobservable that is orthogonal via randomization. The parameters of interest, $\theta = (\delta^{trt}, \delta^{move})$, correspond to the quantities described in Section III and summarized in Figure 3 with δ^{trt} measuring the effect of treatment on the original resident and δ^{move} capturing the extent to which the treatment persists in the home with a new resident.

To estimate the parameters of interest, θ, we run ordinary least squares on (6) and conduct inference via standard errors robust to heteroskedasticity and arbitrary autocorrelations within clustering unit i. Table 2 reports estimates with (6) presented in Column 1. We reject the null for θ, finding that receipt of HERs reduces energy use with $\hat{\delta}^{trt} < 0$ (consonant with, e.g., Allcott 2011, 2015). A point estimate of -25 kWh represents average reductions of about 2.4 percent relative to the baseline, somewhat in the upper range of ATEs in previous work.[20] Importantly, we find that a significant portion of the treatment effect persists in the home with $\hat{\delta}^{move} < 0$. Previously treated homes use approximately 11 kWh less than control homes after initial occupants have moved out. Calculated with the same counterfactual usage as above, this corresponds to about a one percent reduction.[21] To place these estimates into perspective, $\hat{\delta}^{trt}$ is equivalent to turning off two traditional incandescent lightbulbs for eight hours every day or not using a high-end AC window unit (1500W) for 16 hours. $\hat{\delta}^{move}$ is slightly less than half of these changes. In terms of technology investments, 11 kWh are the savings associated with substituting one incandescent with a CFL lightbulb for a 220 hour timespan.[22]

[20] Our sample mainly consists of the earlier Opower interventions in Allcott (2015). Consequently, we would expect future interventions to lead to smaller reductions because most promising sites were selected first. This argument follows through to persistence.
[21] The careful reader will notice a very large (absolute) β^M. We treat this potential problem in the following subsection.
[22] These numbers are extracted from www.energyusecalculator.com. We want to stress that we do not take a stance on whether such reductions are sizable in terms of the environmental impacts – we simply want to provide novel evidence of technology adoption. Unfortunately, the nature of average treatment effects and

Other columns in Table 2 augment the specification by including additional control variables that could affect energy consumption. In Column 2, we use the average pre-treatment usage, i.e. $\bar{e}_{it}$ when $H_t = 0$ and $M_t = 0$, to control for general usage patterns of a given property influenced by factors like property size. Column 3 directly controls for climate conditions by including monthly cooling and heating degree days (CDD, HDD) from the nearest weather station.[23] Columns 4 to 6 utilize measures describing the local housing market and environmental sentiment of the local population. For the former, we add the vacancy rate, i.e. percent of empty housing units, at the ZIP level. For the latter, we draw from the League of Conservation Voters' environmental scorecard for congressional representatives (Kahn and Morris, 2009). This index aggregates representative's law making decisions with respect to bills related to the environment on a scale from zero to 100.[24] Furthermore, we compile publicly available data on donations to Green Party committees from 2008 to 2015 (Kahn, 2007; Kahn and Vaughn, 2009; Wang and Xu, 2016).[25] From these data, we calculate the proportion of households in every ZIP code that donated at least once to any Green Party campaign. Lastly, Column 7 controls for the full set of additional covariates. Across all specifications, we find point estimates and inference to be extraordinarily robust.

These estimates are also economically meaningful, suggesting that households are forward looking in their response to variation in moral suasion, adjusting their marginal use of energy as well as their durable technology investments. To see this, consider a measure of the proportion of the treatment effect that persists after the original resident moves out, $\gamma = \delta^{move}/\delta^{trt}$. Estimates of γ from Table 2 range from 0.43 to 0.55 and our interpretation of these estimates is that 43 to 55 percent of the treatment effect induced by Opower's HER remain through the physical capital of the home after the original owner

---

the lack of direct household-level data do not allow us to speak to what types of adoptions were undertaken by households in the sample. Furthermore, we cannot identify whether few households adopted a wide range of technologies or many households made small changes. However, the effect sizes typically found in such RCTs can easily be explained by relatively small changes in the capital stock by the average household.

[23] We map the geographic center of a home's ZIP code to the closest weather station in terms of geometric distance. Some weather stations have missing observations of climate variables.

[24] These data are extracted from http://scorecard.lcv.org/. The motives for using these measures are explained in more detail in Section VD.

[25] The Federal Election Commission (FEC) publishes individual contributions of at least $200 in every year (see http://www.fec.gov/finance/disclosure/ftpdet.shtml). We aggregate these data over time and restrict them to donations to any committee affiliated with the Green Party.

moved out.[26] AR find a similar trend when looking at a different sample of households in Opower's data. In particular, they find that if Opower's messaging is discontinued, 60 to 70 percent of the treatment effect persists for households that do not move.[27] In relative terms, such an effect corresponds to about a two percent reduction in usage after treatment cessation in their sample.

Using utility energy efficiency program participation data, AR attribute this persistence to habit formation as opposed to technology investments. Yet our indirect approach to observing technology investments suggests an important role for changes in the capital stock of the home. That said, our results do not rule out a role for habit formation as we do not observe the behavior of treated households after they move. Rather, they complement existing findings and provide first evidence that at least part of the story can be explained by rational capital investment. We conduct a more nuanced comparison of persistence estimates across the two studies in Section V.

IVB. Robustness to Exclusion of Low-Usage Months

Estimates in Table 2 also suggest substantial reductions in energy consumption across all homes after move (a large absolute post-move indicators, β^M). Despite the seemingly unrealistic magnitude of the coefficient, this finding can be explained by taking into account that some homes remain vacant for extended periods of time before being sold.[28] Indeed, a look at the raw data reveals that many homes exhibit a dip in usage right after move, including multiple months with zero or very low usage. Generally, such trends do not pose problems for our identification strategy as long as they are constant (or parallel) across treatment and control homes. However, if we worry about a correlation between treatment status and the distribution of unoccupied months, our estimates can be biased.

[26] Please note that this is only one way of defining persistence. AR, for example, mainly contrast the persistent treatment effect of households no longer receiving HERs to the counterfactual of continued receipt of HERs – such a group does not naturally arise in our setting.
[27] AR also estimate a decay parameter to investigate how persistence develops over time. Online Appendix Table 6 for results from a similar approach in our setting. In particular, we interact the post-move treatment indicators with a measure of time-since-move. Unfortunately, due to the non-experimental nature of our data and large variation in time-of-move, our estimates are imprecise and do not allow strong conclusions about potential decay.
[28] This is even more important in our time frame, which includes the Great Recession and its housing market collapse for many waves.

To mitigate concerns, we test if results are sensitive to different rules of eliminating observations where the home was unoccupied. Unfortunately, we only observe the date of account deactivation but not a new resident's account opening date. Consequently, we rely on (potentially imperfect) assumptions about what constitutes unoccupied homes. Because these measures are crude, we do not directly compare the time on the housing market but rather attempt to show robustness across subsamples of the data. Online Appendix Figure 1 provides a graphical representation of observations with low usage before and after move-out. We find that both treatment and control households experience a significant spike in low-usage and zero-usage observations right around move. The proportion of these observations declines quickly but continues to be relatively pronounced for up to six months after move. Consequently, we provide a range of more or less conservative exclusion rules and re-estimate (6) on the restricted samples without low-usage observations as defined below. In doing so, we hope to provide evidence that our persistence findings are not an artifact of data limitations.

The four exclusion rules are: i) exclusion of all observations with $e_{ijt} < 150$ in the first six months after deactivation of the account and exclusion of homes for which the average energy usage in the post-move period is smaller than the pre-move average minus two standard deviations (i.e. $\bar{e}_{ijt=M} < \bar{e}_{ijt!=M} - 2\sigma_e$), ii) exclusion of the first six months after account closure regardless of usage, iii) exclusion of all observations after account closure that are smaller than the smallest pre-move energy consumption minus 20 percent (i.e. $e_{ijt=M} < min\, e_{ijt!=M} \cdot 0.8$), and finally iv) exclusion of all post-move observations with $e_{ijt} < 200$.

Table 3 presents results from these rules and various combinations. Column 1 depicts the baseline case, which applies (6) to the full sample with no exclusion rule applied. A few observations stand out. First, different sample sizes reflect the varying degrees of stringency across rules. Second, every rule reduces β^M significantly compared to the baseline case. In Column 6, for example, the point estimate of β^M is no longer distinguishable from zero. Other rules reduce β^M by between 50 and 85 percent. Third, and most importantly, the coefficients of interest are remarkably stable across the different samples. We find negative and statistically significant $\hat{\delta}^{move}$ even after artificially pushing β^M to zero. While some point estimates decrease in magnitude, we still estimate γ to be between 22 and 40 percent. We conclude that results are not driven by differential patterns of home occupancy after move.

IVC. Considering Heterogeneity in Treatment Effects

In Sections IVA and IVB, we estimated an empirical model that assumed homogeneous treatment and post-move effects. This made identifying the proportion of persistence, γ, easy because we could simply take the ratio of δ^{move} and δ^{trt}. However, if there is heterogeneity (as Allcott, 2015, clearly shows), this approach could lead to a biased estimate of γ. In this section we present a two-stage procedure to estimate γ that incorporates heterogeneity both in the time series and cross section.

In the first stage we estimate an augmented version of (6):

$$(7) \qquad e_{ijct} = \beta_{jc}^{T} T_i + \beta_{jc}^{H} H_t + \delta_{jc}^{trt}\, T_i H_t + \beta_{jc}^{M} M_t + \delta_{jc}^{move}\, T_i M_t + \tau_t + V_{ijct}$$

where we allow for heterogeneity in each wave of the experiment, j, and each cohort of movers, c, with c capturing the number of months between reception of the first HER and move-out.[29] Conceptually, (7) is just a DD model for each of the wave-cohort combinations in the data. We estimate (7) via ordinary least squares for each wave separately and conduct the following second stage empirical model with all first stage estimates:

$$(8) \qquad \delta_{jc}^{move} = \gamma \delta_{jc}^{trt} + W_{jc}$$

where γ measures the average proportion of the treatment effect that persists in the post-move period across the waves and cohorts in the data. We estimate γ via weighted least squares according to the inverse variance of $\delta_{jc}^{move}/\delta_{jc}^{trt}$ and conduct inference with standard errors clustered within each wave, j. (8) is then akin to a meta-analysis of the DD estimates found in (7).

[29] c is the number of months that a household assigned to the treatment group would have received Opower's reports. Please note that this is a measure of length of exposure rather than intensity because frequency of moral suasion messaging varies across waves and utilities (monthly, bi-monthly, quarterly).

Figure 5 plots the parameters of interest from our first-stage regression, (7), in $(\delta_{jc}^{trt}, \delta_{jc}^{move})$ space.[30] To convey the variation in the length of the different cohorts each point in Figure 3 is shaded according to the number of years a household was in the treatment time period before moving (the median time of treatment exposure is slightly over one year). Furthermore, the size of each point reflects the number of homes in a wave-cohort with many including well over hundred homes (the median wave-cohort is 330). Figure 5 also illustrates how (8) estimates γ by plotting the best-fit line. Table 4 presents that estimate in Column 1. The estimate of $\hat{\gamma}^{prst}$ suggests that about 35 percent of the treatment effect persists in the post-move period, rejecting γ of zero at conventional levels of statistical significance.[31]

Table 4 also considers the relationship between persistence and the timing of the moving decision. If a household assigned to treatment enters the treatment period with plans to move soon, they are more likely to behave like a short-run household that holds technology fixed.[32] By a simple envelope argument our model in Section IIIB predicts $\delta_{jc}^{move}/\delta_{jc}^{trt}$ to be increasing in the magnitude of c. That is, households that received HERs for a short period of time should see smaller (in magnitude) per-month energy usage reductions from treatment than households that received HERs for longer periods of time. Furthermore, homes exposed to only few HERs should see a smaller proportion of the treatment effect persisting after the original resident moves relative to households receiving treatment for an extended period.

We test this prediction in Columns 2 and 3 of Table 4 by estimating

$$(9) \qquad \delta_{jc}^{move} = \gamma_{short}^{prst}\delta_{jc}^{trt}1(short) + \gamma_{long}^{prst}\delta_{jc}^{trt}1(long) + \widetilde{W}_{jc},$$

where $1(\cdot)$ denotes the indicator function for households exposed to treatment for a short or long period of time, respectively. We define the cut-off in two ways: i) one year, the

[30] For cosmetic reasons we limit the figure to $|\delta_{jc}^{move}| < 250$ and $|\delta_{jc}^{trt}| < 150$.
[31] See Online Appendix Table 7 for estimates that consider robustness over the different exclusion rules discussed above.
[32] Some have suggested that households may make investments right before moving to increase the value of their home. When we look at the energy efficiency program participation data used in AR one can see a steep decline in participation starting 19 months before a resident moves. This evidence is available upon request from the authors.

mean length of exposure in our data, in Column 2 and ii) two years of HER exposure, mirroring the time frame in AR, in Column 3. We estimate this model without a constant and statistically test the difference between $\hat{\gamma}_{short}^{prst}$ and $\hat{\gamma}_{long}^{prst}$. Across the two specifications we find that $\delta_{jc}^{move}/\delta_{jc}^{trt}$ is indeed increasing in the length of exposure.[33] In fact, persistence in homes with initial residents who were exposed to HERs for a substantial time is approximately twice the persistence in short-term properties. These differences between the short-run and the long-run effect are significant at high levels of confidence and coefficients are precisely estimated.

IVD. Alternative Interpretations

In Section III we assumed that new residents do not sort into homes based on investment differences caused by treatment. To assess the validity of this assumption, we sketch out a partial equilibrium sorting model in the post-move time period ($t = 2$), derive predictions, and then test those predictions using proxy variables of housing market conditions. Across all proxies we reject the predictions of the sorting model.

Consider an agent looking to purchase a home in housing market h at the start of event time period $t = 2$. We assume the agent faces a market where there is a continuum of identical homes that vary only according to their technology stock, which we represent with k_1 just as we did in previous sections. The price of a home in market h varies according to the extent of its technology stock, $p_{k,h}(k_1)$, with price increasing in k_1. We also assume the agent solves a two-period decision problem. In the first time period ($t = 2$) the technology stock of the home they purchase is fixed but in the second period ($t = 3$) they can make investments to vary their new home's technology stock. The maximization problem is:

[33] We also estimate a simple model with a linear time trend in cohort, i.e. length of exposure to treatment. We find that persistence significantly increases in length of exposure (at the five percent level). Online Appendix Table 8 reports results.

$$\max_{\{c_t\},\{e_t\},k_1,I_3} u(c_2) + v(z_2) - g(e_2, a^{new}) + \beta u(c_3) + \beta v(z_3) - \beta g(e_3, a^{new})$$

$$(10) \qquad s.t. \quad m + \frac{m}{1+r} = c_2 + p_e e_2 + p_{k,h}(k_1) + \frac{p_I I_3}{1+r} + \frac{c_3}{1+r} + \frac{p_e e_3}{1+r}$$

$$z_2 = f(e_2, k_1)$$

$$z_3 = f(e_3, I_3 + k_1)$$

where β is the agent's discount rate, r is the interest rate they face, and a^{new} is the agent's exogenous pro-social meta-preference parameter that is constant over the relevant decision period. Imposing the same assumptions from Section III on (10) and assuming linearity in the price of existing technology we see the following by combining first-order conditions:

$$(11) \qquad \frac{k_1}{I_3} = \frac{p_I}{p_{k,h} - p_I}.$$

(11) tells us that the agent solving (10) will choose existing technology versus new technology according to their relative prices. Put differently, (11) shows that agents will sort into houses with more existing capital when the price of capital is low relative to the price of investments. If the price of investing is stable across housing markets, then we have a simple way of testing (10): Compare δ_{move} in housing markets with high $p_{k,h}$ to the same parameter in markets with low $p_{k,h}$, with the sorting model predicting δ_{move} is decreasing in $p_{k,h}$ and the model in Section III predicting a null effect to changes in $p_{k,h}$.

While we do not directly observe $p_{k,h}$ we assume that it varies with housing market conditions. Firstly, we consider a scenario where demand for k_1 is fixed and supply-side conditions vary according to the housing vacancy rate in each zip code (our proxy for h) with demand fixed. If high vacancy levels correspond to supply shifts of k_1 then $p_{k,h}$ will be decreasing in the extent of housing vacancies.[34] Secondly, we consider a scenario where supply of k_1 is fixed and demand-side conditions vary according to the environmental sentiments in a housing market.

[34] Alternatively, vacancies could weaken the bargaining power of the home seller causing the buyer to pass more of the cost of k_1, $p_{k,h}$, onto the seller when vacancy levels are high.

We proxy for environmental sentiments by linking a home to their Congressman's National Environmental Scorecard rating which is published annually by the League of Conservation Voter. Additionally, we utilize the Federal Election Commission's individual contributions data to compile a county and district measure of giving to any Green Party committee between 2008 and 2015. These data include all individual donations of at least $200 and we calculate the proportion of households in a given ZIP or county who gave at least once during our sample period. If high environmental sentiment markets correspond to demand shifts for k_1 then $p_{k,h}$ will be increasing in these proxies.

Table 5 presents estimates of our coefficients of interest from (6) and their interaction with proxies of supply and demand conditions. The results in Row 1 indicate that δ_{move} is increasing in vacancy rates, one of our proxies for $p_{k,h}$. An increase in the vacancy rate of one percentage point is associated with a reduction in persistence of about one kWh. This rejects the sorting model, which predicts the opposite relationship. Rows 2 to 5 of Table 5 present specifications where the coefficients of interest are interacted with our four definitions of environmental sentiments. Across all specifications, we cannot reject the null hypothesis of no additional impact of green sentiment. Again, this does not match the predictions of the sorting model.[35]

The evidence presented in this section indicates that sorting consonant with a simple price-based model is not supported by the data. While this does not formally rule out sorting, three additional factors speak against an important impact of sorting. Firstly, homes with extensive capital investments plausibly attract residents with high baseline use because large savings can be realized. If such people were sorting into treatment homes, $\hat{\delta}_{move}$ is likely to be a lower bound on persistence. Secondly, more efficient technology can lead to a rebound effect because household production of energy-related goods becomes cheaper. Again, such an effect would strengthen our results. Thirdly, for sorting to be an issue, there has to be technology adoption in the first place. Otherwise, no signal of efficient capital could be observed by home buyers. In other words, if sorting of a form not captured by our proxies is a common occurrence, it can only be due to prior capital investment by initial residents. The only difference then lies in interpretation: the

[35] Table 5 is based on DD specifications with interaction terms. The full regression tables are provided in Online Appendix Table 9.

persistent effect estimated by us would capture a combination of technology and crowding in of green residents. Together, these observations suggest that our estimates pick up direct effects of technology investment.

V. Implications

In this section, we relate our estimates of persistence due to capital investment to the estimates of persistence in AR. While our study and AR both assess the long-run effects of the HER, the samples have no overlap, making a direct comparison of behavior impossible. Instead, we conduct the following exercise. First, we obtain the preferred estimates of persistence in AR across each of their three sites.[36] Second, we test the hypothesis that our estimates of persistence, which we attribute to capital investments in the home, are equal to the estimates of persistence in AR, which are attributed to a combination of investments in the home and habits. In doing so, we hope to formally assess the extent to which accumulation of habit capital is responsible for the estimates of persistence in AR.

Panel A of Table 6 presents results. In Column 1, we see that the persistence from technology observed over our entire sample is significantly lower than the persistence estimated across the three sites in AR. However, when we focus on just the subsample of homes in our sample that are most comparable to AR—households in AR receive the HER for two years—in Columns 2 and 3, we see that estimates across the two studies become quite similar. While future work is clearly needed to decompose the mechanisms that lead to long-run effects in response to the HER, our estimates suggest that a significant portion (or nearly all) of the persistence observed in AR is due to unobserved investments in energy efficient technology rather than habit formation.

This finding also has important implications for the normative assessment of the HER. Economists have long emphasized the importance of incorporating both direct and indirect costs when assessing program performance (see, e.g., Heckman and Smith,

[36] In particular, we use point estimates from Table 4 of Allcott and Rogers (2014, pp. 3024) that compare the proportion of the HER effect that persists for households that have the HER discontinued versus households that do not have the HER discontinued. We choose these estimates because they correspond most closely to the discussion of persistence in AR, "The point estimates…suggest that continuing the intervention increases the treatment effects in the post-drop period by a remarkable 50 to 60 percent" (Allcott and Rogers, 2014, pp. 3024). Please note that the data used by AR are not part of our sample.

1998). If the HER induces adoption of costly capital, earlier estimates of its cost-effectiveness may be unduly optimistic due to the omission of these adoption costs. To estimate the extent to which cost-effectiveness is influenced by the costly capital investment, we revisit earlier estimates of the HER's cost-effectiveness in light of our findings.

Panel B of Table 6 presents the results of this exercise for the treatment and pre-move period. Consistent with AR, we start by assuming the direct cost of administering the HER to one household is \$1 per report. Following the standard approach in the literature, we simply compare this direct cost to savings achieved by the HER as estimated in Column 1 of Table 2 (-24.98 kWh) (*No Technology*). To also account for the indirect costs of the HER, we use a lower-bound estimate of the cost of capital per kWh of electricity saved in Allcott and Greenstone (2012, pp. 17).[37] With this cost estimate, we convert savings achieved by capital investment—as a proportion of total savings using our estimates of persistence in the first row—into dollars.[38] Applying this conversion, we see that the HER induces an indirect cost that ranges from \$0.74 to \$1.15 per report and household. The resulting cost-effectiveness is reported in the last row for all three subsamples (*Technology*).

Comparing the two approaches to estimating cost-effectiveness in Panel B of Table 6, we see that incorporating the indirect cost of investments in capital more than doubles the cost per kWh in two out of three cases. From this perspective, after accounting for direct and indirect costs of technology adoption, alternative programs to

[37] Allcott and Greenstone (2012) estimate a cost of energy efficient technology per kWh saved of 8.5¢. They reach this conclusion by reviewing the literature on demand-side management programs, which uses subsidies and other economic incentives to encourage uptake of energy-efficient technology. Assuming a discount rate of 5 percent and installation and purchase costs of 70 percent, they conclude that the cost of energy efficient technology per kWh is about 8.5¢. An alternative approach would use the marginal cost of a kWh of electricity as an upper bound on the cost of technology. The national average for residential users since 2010 is about 12¢. Consequently, using this strategy would only strengthen our point.

[38] Allcott and Greenstone (2012) provide a detailed discussion of other indirect costs, e.g. households could experience utility losses due to a less comfortable ambient temperature, warm glow from contributing to a public good, etc. Consonant with previous work on cost-effectiveness of Opower's programs, we ignore these costs and benefits. Furthermore, we do not include estimates of discounted long-term benefits exceeding the sample period or capital depreciation over time. Allcott and Kessler (2015) find supporting evidence for effects driven by moral utility.

the HER discussed in Allcott and Mullainathan (2010) and Allcott and Greenstone (2012) appear much more attractive.[39]

VI. Conclusion

Policies motivated by behavioral insights have increased in popularity across many different environments, based in part on experimental evidence of their effectiveness. Yet, relatively little is understood about the underlying mechanisms – the *how* and *why* of the average treatment effects observed. In particular, we know little about whether such policies change habits in ways that persist when the behavioral intervention ends. Our lack of understanding is caused by a lack of rigorous theories of nudges (an area that is improving quickly) and by missing household-level data that would permit the analyst to draw conclusions about actual steps taken by the people who are nudged. Importantly, depending on the underlying mechanisms, conclusions about short- and long-term costs and benefits can change significantly.

This paper develops a simple short- and long-run theoretical model of household production in the Beckerian tradition. In our model, households produce consumption good in their home using energy and capital that varies in its energy efficiency. Moral suasion from an intervention like Opower's standard Home Energy Reports is represented by a shadow price on energy use. Moral suasion thus induces relative price changes that lead rational agents to adjust their behavior – a phenomenon widely documented in the literature – and to invest in energy efficient capital. We then utilize data from 38 natural field experiments and employ a novel identification strategy to show that a significant proportion of the initial treatment effect persists after treated households move out of their homes. This effect is robust across different specifications and alternative models of behavior that emphasize sorting are rejected by the data.

Although we do not invalidate or rule out behavioral adjustments, our findings suggest a previously understated role of capital investments in response to social nudges. This channel is widely missing from existing calculations of cost-effectiveness and cost-

[39] Interestingly, Allcott and Kessler (2015) find that household willingness to pay for continued receipt of the HER is less than estimates of the HER's cost-effectiveness that assumes no investments in physical technology. The authors attribute this gap to moral utility. Clearly more work is needed to parse the costs of motivating conservation via moral utility from the costs of energy efficient technology upgrades.

benefit analyses. As such, our findings have clear policy relevance and highlight the importance of theory and creative empirical strategies to identify parameters of interest. Unfortunately, current technology and data availability do not allow us to speak to the open question of what types of technologies are adopted. With the ascent of smart technology solutions and better ways to monitor energy usage, we are optimistic that future work will be able to provide answers.

Earlier research on habit formation presents a pessimistic perspective on the ability of policies and programs to induce persistent changes in habits. Our study does little to overturn this view. Despite our findings, we believe that there is still much that we do not understand about habit formation and ways to induce changes in such. Recent research suggests that one particularly promising dimension is to focus investments on programs that target habitual behavior at early ages (see, e.g., Almlund et al., 2011). We would encourage future work to explore such programs and whether habits are easier to form at a young age.

Finally, in the context of energy policy goals, our research does point to a promising alternative to habit formation, such as the dissemination of energy efficient technologies (see, e.g., Acemoglu et al., 2012). Our research suggests that while social nudges may have little impact on the formation of new habits, they do appear to provide an effective way to induce the adoption of technologies that obviate the need for changes in habits. We imagine a useful exercise for future work is to leverage the identification strategy developed in this study to parse the role of habits induced by different policy changes from changes in other state variables, such as physical capital.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change," *American Economic Review*, 102(1): 131-166.

Acland, Dan, and Matthew R. Levy. 2015. "Naiveté, Projection Bias, and Habit Formation in Gym Attendance," *Management Science*, 61(1): 146-160.

Allcott, Hunt, and Sendhil Mullainathan. 2010. "Behavior and Energy Policy," *Science*, 327(5970): 1204-1205.

Allcott, Hunt. 2011. "Social Norms and Energy Conservation," *Journal of Public Economics*, 95(9): 1082-1095.

Allcott, Hunt. 2015. "Site Selection Bias in Program Evaluation," *Quarterly Journal of Economics*, 130(3): 1117-1165.

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

Allcott, Hunt, and Todd Rogers. 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation," *American Economic Review*, 104(10): 3003-037.

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market," *American Economic Review*, 105(8): 2501-2538.

Allcott, Hunt, and Judd B. Kessler. 2015. "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," NBER Working Paper No. 21671.

Almlund, Mathilde, Angela Lee Duckworth, James J. Heckman, and Tim D. Kautz. 2011. "Personality Psychology and Economics," *Handbook of the Economics of Education*, 4: 1–181.

Ayres, Ian, Sophie Raseman, and Alice Shih. 2013. "Evidence from Two Large Field Experiments that Peer Comparison Feedback Can Reduce Residential Energy Usage." *Journal of Law, Economics, & Organization*, 29(5): 992-1022.

Becker, Gary S. 1965. "A Theory of the Allocation of Time," *Economic Journal*, 75(299): 493-517.

Becker, Gary S., Michael Grossman, and Kevin M. Murphy. 1994. "An Empirical Analysis of Cigarette Addiction," *American Economic Review*, 84(3): 396-418.

Becker, Gary S., and Kevin M. Murphy. 1988. "A Theory of Rational Addiction," *Journal of Political Economy*, 96(4): 675-700.

Becker, Gary S., and Kevin M. Murphy. 1993. "A Simple Theory of Advertising as a Good or Bad," *Quarterly Journal of Economics*, 108(4): 941-964.

Becker, Gary S., and Kevin M. Murphy. 2000. *Social Economics*, Harvard University Press, Cambridge, MA.

Bernedo, Maria, Paul J. Ferraro, and Michael K. Price. 2014. "The Persistent Impacts of Norm-Based Messaging and Their Implications for Water Conservation," *Journal of Consumer Policy*, 37(3): 437-452.

Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher. 2001. "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, 91(1): 149-166.

Charness, Gary, and Uri Gneezy. 2009. "Incentives to Exercise," *Econometrica*, 77(3): 909-931.

Campbell, John Y., and John H. Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2): 202-251.

Carroll, Christopher D., Jody Overland, and David N. Weil. 2000. "Savings and Growth with Habit Formation," *American Economic Review*, 90(3): 341-355.

Constantinides, George M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98(3): 519-543.

Costa, Dora L., and Matthew E. Kahn. 2010. "Why Has California's Residential Electricity Consumption Been So Flat since the 1980s? A Microeconometric Approach," NBER Working Paper No. 15978.

Costa, Dora L., and Matthew E. Kahn. 2013. "Energy Conservation "Nudges" and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment," *Journal of the European Economic Association*, 11(3): 680-702.

Davis, Lucas W. 2008. "Durable goods and residential demand for energy and water: evidence from a field trial," *RAND Journal of Economics*, 39(2): 530-546.

Davis, Lucas W., and Lutz Kilian. 2011. "The Allocative Cost of Price Ceilings in the U.S. Residential Market for Natural Gas," *Journal of Political Economy*, 119(2): 212-241.

DellaVigna, Stefano, John A. List, and Ulrike Malmendier. 2012. "Testing for Altruism and Social Pressure in Charitable Giving," *Quarterly Journal of Economics*, 127(1): 1-56.

Ferraro, Paul J., and Juan Jose Miranda. 2013. "Heterogeneous treatment effects and mechanisms in information-based environmental policies: Evidence from a large-scale field experiment," *Resource and Energy Economics*, 35(3): 356-379.

Ferraro, Paul J., and Michael K. Price. 2013. "Using Nonpecuniary Strategies to Influence Behavior: Evidence from a Large-Scale Field Experiment," *Review of Economics and Statistics*, 95(1): 64-73.

Ferraro, Paul J., Juan Jose Miranda, and Michael K. Price. 2011. "The Persistence of Treatment Effects with Norm-Based Policy Instruments: Evidence from a Randomized Environmental Policy Experiment," *American Economic Review*, 101(3): 318-322.

Heckman, James J., and Jeffrey A. Smith. 1998. "Evaluating the Welfare State," *Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial*, Strom, Steiner, ed., Cambridge: Cambridge University Press, 1998, pp. 241-318.

Herberich, David H., John A. List, and Michael K. Price. 2011. "How Many Economists does it take to Change a Light Bulb? A Natural Field Experiment on Technology Adoption," University of Chicago Working Paper.

Hotelling, Harold. 1931. "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39(2): 137-175.

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," *Energy Journal*, 29(1): 113-134.

Kahn, Matthew E. 2007. "Do greens drive Hummers or hybrids? Environmental ideology as a determinant of consumer choice," *Journal of Environmental Economics and Management*, 54(2): 129-145.

Kahn, Matthew E., and Eric A. Morris. 2009. "Walking the Walk: The Association Between Community Environmentalism and Green Travel Behavior," *Journal of the American Planning Association*, 75(4): 389-405.

Kahn, Matthew E., and Ryan K. Vaughn. 2009."Green Market Geography: The Spatial Clustering of Hybrid Vehicles and LEED Registered Buildings," *B.E. Journal of Economic Analysis & Policy*, 9(2): 1-24.

Landry, Craig E., Andreas Lange, John A. List, Michael K. Price, Nicholas G. Rupp. 2010. "Is a Donor in Hand Better Than Two in the Bush? Evidence from a Natural Field Experiment," *American Economic Review*, 100(3): 958-983.

Levitt, Steven D., John A. List, and Sally Sadoff. 2016. "The Effect of Performance-Based Incentives on Educational Achievement: Evidence from a Randomized Experiment," NBER Working Paper No. 22107.

Lucas, Robert E. 1972. "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4(2): 103-124.

Meier, Stephan. 2007. "Do Subsidies Increase Charitable Giving in the Long Run? Matching Donations in a Field Experiment," *Journal of the European Economic Association*, 5(6): 1203-1222.

Milkman, Katherine L., Julia A. Minson, and Kevin G. M. Volpp. 2014. "Holding the Hunger Games Hostage at the Gym: An Evaluation of Temptation Bundling," *Management Science*, 60(2): 283-299.

Nordhaus, William D. 1994. *Managing the Global Commons*, MIT Press, Cambridge, MA.

Pollak, Robert A. 1970. "Habit Formation and Dynamic Demand Functions," *Journal of Political Economy*, 78(4): 745-763.

Small, Kenneth A., and Kurt Van Dender. 2007. "Long Run Trends in Transport Demand, Fuel Price Elasticities and Implications of the Oil Outlook for Transport Policy." OECD/ITF Joint Transport Research Centre Discussion Papers, No. 2007/16.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature*, 49(2): 326-365.

Ramsey, Frank P. 1928. "A Mathematical Theory of Saving," *Economic Journal*, 38(152): 543-559.

Rogers, Todd, and Erin Frey. 2015. “Changing Behavior Beyond the Here and Now,” *The Wiley Blackwell Handbook of Judgment and Decision Making*, John Wiley & Sons, Ltd, Chichester, UK.

Rossi, Peter H. 1987. "The Iron Law of Evaluation and Other Metallic Rules," *Research in Social Problems and Public Policy*, 4: 3-20.

Royer, Heather, Mark Stehr, and Justin Sydnor. 2015. "Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company," *American Economic Journal: Applied Economics*, 7(3): 51-84.

Wang, Zhongmin, and Cheng Xu. 2016. "Using Donations to the Green Party to Measure Community Environmentalism," RFF Discussion Paper 16-25.

Volpp, Kevin G., Andrea Gurmankin Levy, David A. Asch, Jesse A. Berlin, John J. Murphy, Angela Gomez, Harold Sox, Jingsan Zhu and Caryn Lerman. 2006. "A Randomized Controlled Trial of Financial Incentives for Smoking Cessation," *Cancer Epidemiology, Biomarkers, & Prevention*, 15(1): 12-18

Volpp, Kevin G., Andrea B. Troxel, Mark V. Pauly, Henry A. Glick, Andrea Puig, David A. Asch, Robert Galvin, Jingsan Zhu, Fei Wan, Jill DeGuzman, Elizabeth Corbet, Janet Weiner, and Janet Audrain-McGovern. 2009. “A Randomized, Controlled Trial of Financial Incentives for Smoking Cessation,” *New England Journal of Medicine*, 360 (7): 699–709.

Volpp, Kevin G., Leslie K. John, Andrea B. Troxel, Laurie Norton, Jennifer Fassbender, and George Loewenstein. 2008. "Financial Incentive-Based Approaches for Weight Loss: A Randomized Trial," *Journal of the American Medical Association*, 300(22): 2631-2637

Figure 1: Persistence in Habit Formation Literature

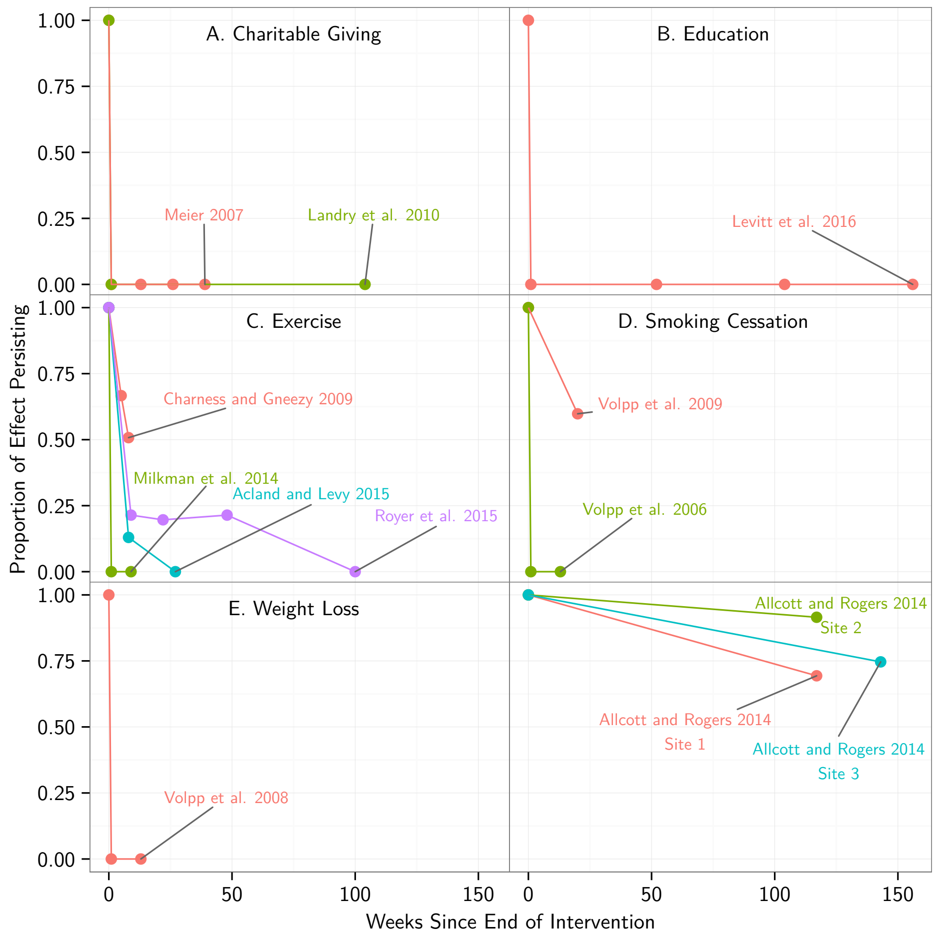

Notes: Each point represents the proportion of the initial treatment effect that persists for a given amount of time since the end of a given intervention. All observations are based on point estimates presented in the corresponding studies with insignificance at the five percent level constituting persistence of zero.

Figure 2: Opower's Home Energy Report

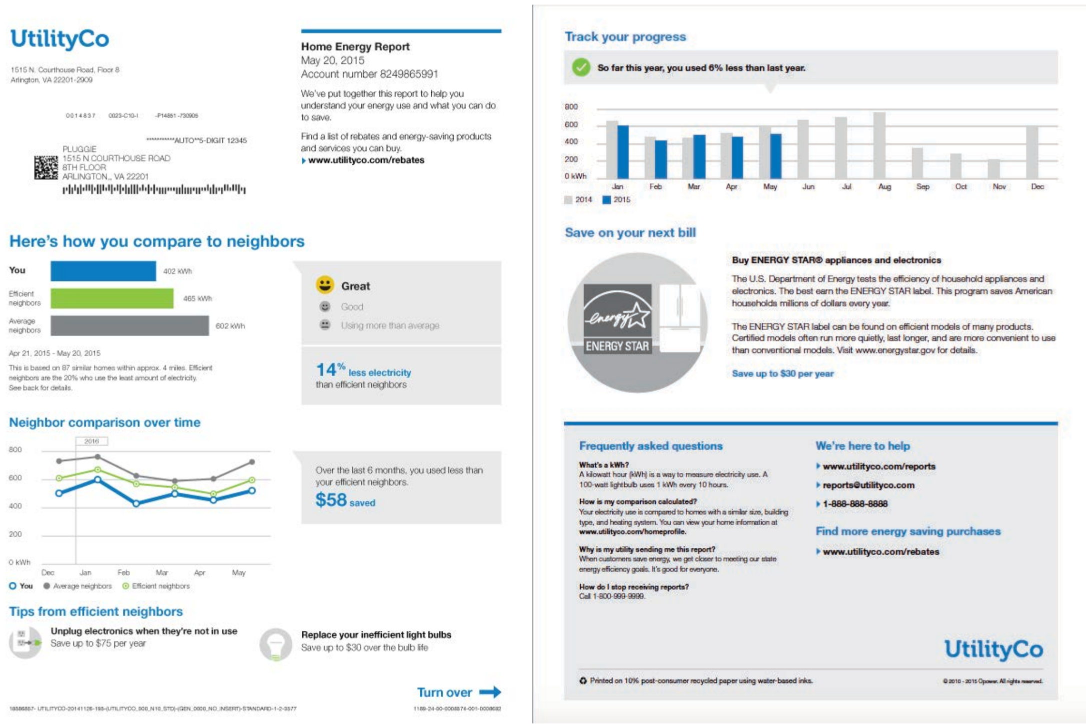

(a) Front (b) Back

Notes: The figure presents the front and back of a typical Home Energy Report (HER). Treatment households receive reports regularly (monthly, bi-monthly, or quarterly) before move-out. The example depicts a household below the 20th percentile of usage in a given month.

Figure 3: Summary of Theoretical and Empirical Model

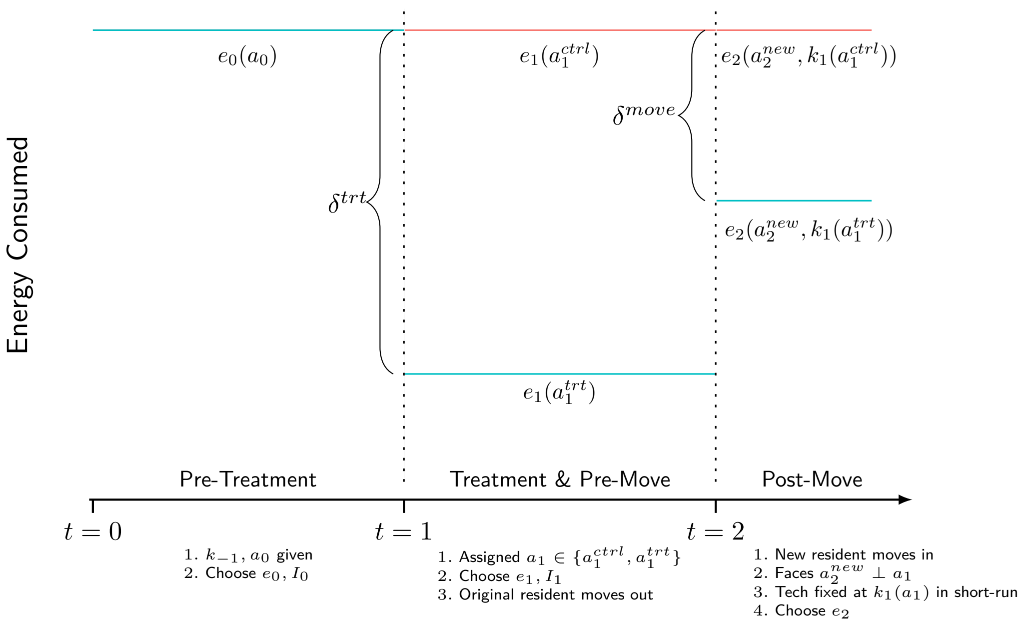

Notes: Demand for energy is plotted during the three event time periods according to the solutions derived in Section III for $(\sigma, \alpha, k_{-1}, a_0, a_1^{trt}, a_2^{new}, p_I, p_e) = (5, 0.5, .1, 0, 1, 0, 1, 1)$. Implied regression coefficients of interest for corresponding empirical model are plotted next to curly brackets. The x-axis plots the order of events during each time period.

Figure 4: Event Study

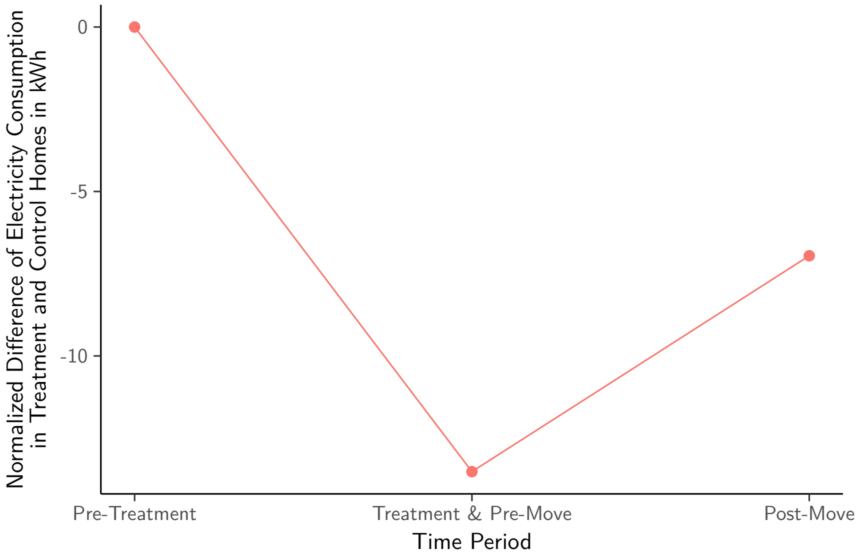

Notes: Figure plots the difference between average electricity consumption homes assigned to Treatment and Control in the Pre-Treatment, Treatment and Pre-Move, and Post-Move time periods normalized to differences in the Pre-Treatment time period. See Online Appendix Table 4 for levels.

Figure 5: Scatterplot of Treatment Effect by Wave-Cohort

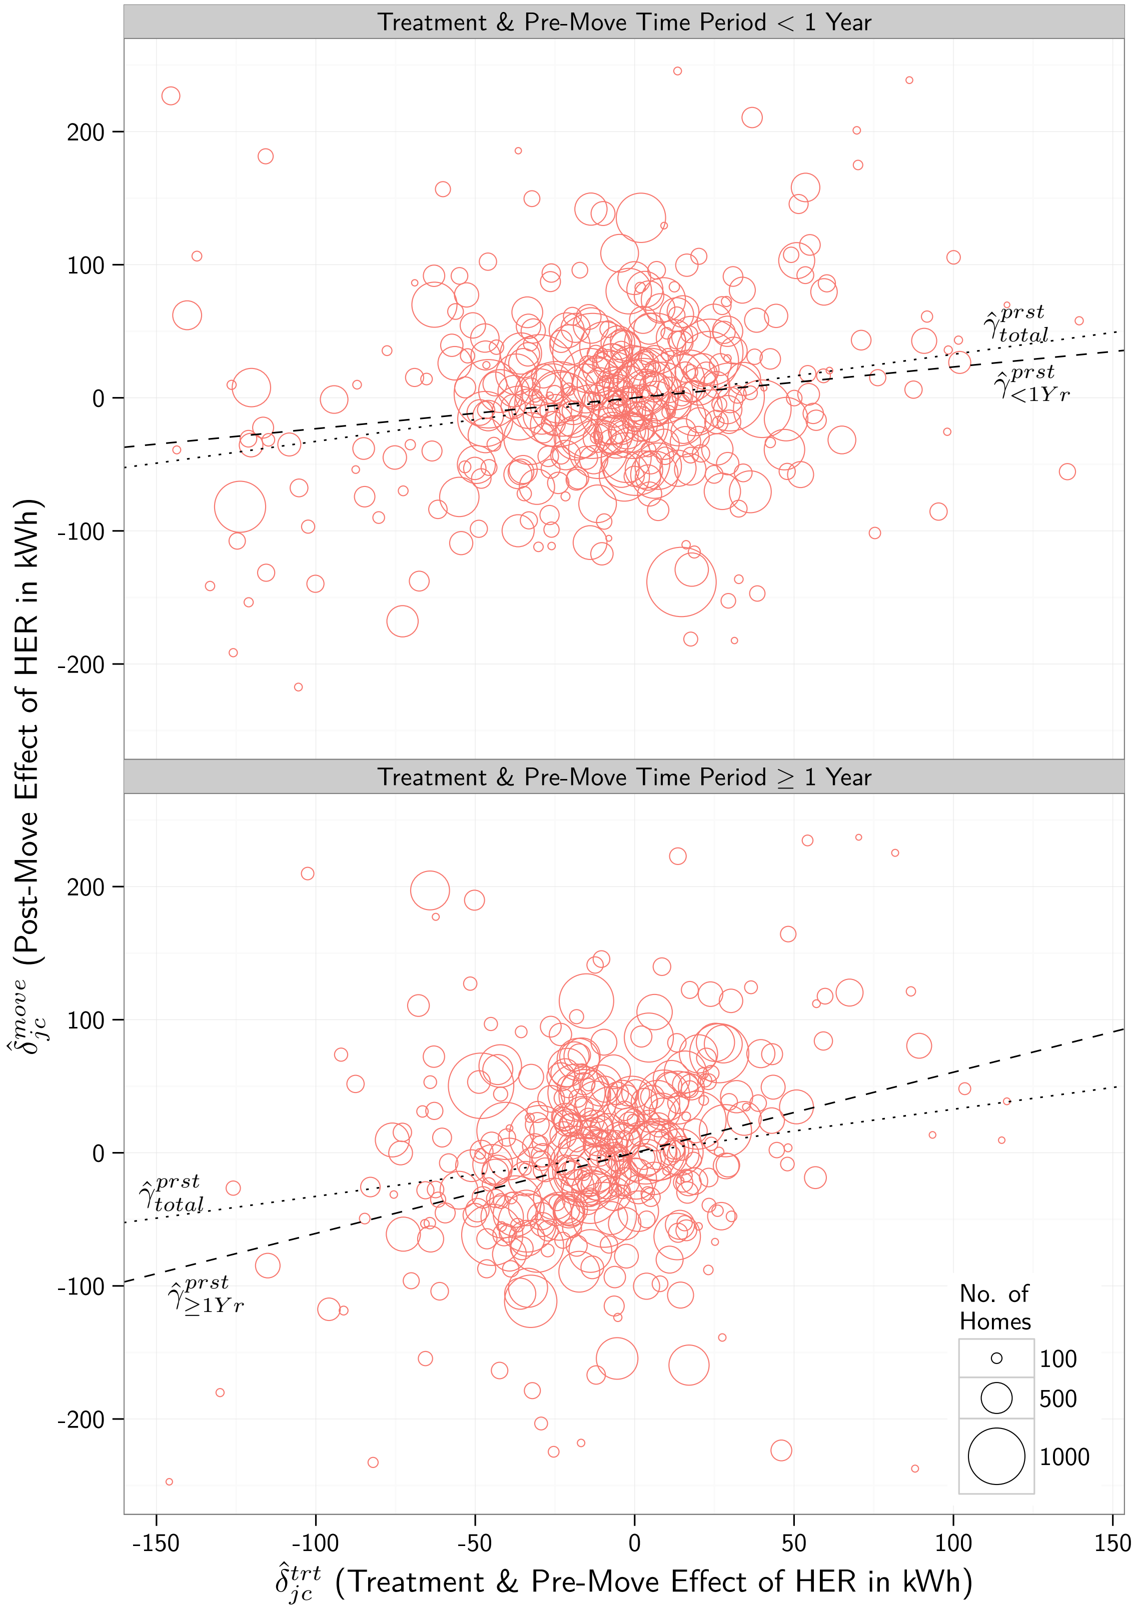

Notes: The figure plots the effect of the HER for each wave-cohort in the Post-Move time period as a function of the effect of the HER in the Treatment and Pre-Move time period; i.e., the parameters of interest from our first stage regression in (7), $\hat{\delta}^{trt}_{jc}$ and $\hat{\delta}^{move}_{jc}$. The top panel plots wave-cohorts that receive the HER in the Treatment and Pre-Move time period for less than 1 year and the bottom panel plots wave-cohorts that receive the HER in the Treatment and Pre-Move time period for more than 1 year. The size of each point reflects the number of unique homes in that wave-cohort. The best-fit line illustrates how (8) estimates γ^{prst}, γ^{prst}_{short}, and γ^{prst}_{long}. The figure excludes wave-cohorts with $|\delta^{move}_{jc}| > 250$ and $|\delta^{trt}_{jc}| > 150$ in the figure but the best-fit lines are from regression estimates on the full-sample.

Table 1: Overview of Sample

	Sample		
	Full	Non-Movers	Movers
Utilities	21	21	21
Waves	38	38	38
Households	2,516,089	2,258,185	253,383
Treatment Indicator	0.67	0.67	0.65
	(0.47)	(0.47)	(0.48)
Pre-Treatment Usage (kWh)	1,198.82	1,212.08	1,084.51
	(652)	(655.44)	(609.48)
Pre-Treatment Observations (mos.)			13.47
			(1.26)
Treatment and Pre-Move Observation (mos.)			11.67
			(9.62)
Post-Move Observation (mos.)			12.17
			(9.44)

Notes: Summary statistics for the full sample, the subsample of households that remain in the same home throughout (Non-Movers), and houses that change hands (Movers). Treatment Indicator is a binary measure of assignment to reception of Home Energy Reports (HERs). Pre-Treatment Usage describes the average monthly usage in kWh in months prior to the first HER. Pre-Treatment Observations show the number of usage reads before the treatment event period begins (i.e. the first full month after the first HER), Treatment and Pre-Move Observations represents the length of exposure to HERs, and Post-Move Observations occur after initial residents deactivate their account. The small disparity between the difference of the full sample and the non-movers sample and the number of unique households in the final movers sample is due to additional data cleaning as described in-text.

Table 2: Estimates of Treatment Effect: Pooled Difference in Difference Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\hat{\beta}^T$	6.50	7.01	8.34	6.35	6.70	6.50	7.01
	(2.05)	(0.48)	(2.37)	(2.05)	(2.09)	(2.05)	(0.62)
$\hat{\beta}^H$	-53.71	-58.45	-54.36	-53.83	-54.16	-53.72	-57.93
	(1.72)	(1.45)	(2.04)	(1.73)	(1.77)	(1.72)	(1.73)
$\hat{\delta}^{trt}$	-24.98	-22.69	-25.63	-24.93	-24.48	-24.97	-22.64
	(2.05)	(1.50)	(2.45)	(2.06)	(2.10)	(2.05)	(1.84)
$\hat{\beta}^M$	-148.30	-131.10	-154.40	-149.14	-149.13	-148.30	-136.24
	(2.73)	(2.51)	(3.24)	(2.74)	(2.79)	(2.73)	(3.03)
$\hat{\delta}^{move}$	-11.35	-12.73	-11.18	-11.25	-11.41	-11.35	-13.14
	(2.64)	(2.57)	(3.13)	(2.65)	(2.69)	(2.64)	(3.09)
Pre-Exp. Usage		0.80					0.80
		(0.00)					(0.00)
CDD			0.88				0.89
			(0.01)				(0.01)
HDD			0.14				0.15
			(0.00)				(0.00)
Vacancy Rate				-1.07			-1.08
				(0.16)			(0.11)
Env. Index					-0.28		0.05
					(0.03)		(0.02)
Green Party Donations						-0.05	0.21
						(0.03)	(0.09)
R^2	0.216	0.444	0.232	0.216	0.217	0.216	0.453
N	9,350,745	9,350,725	6,127,816	9,247,833	8,921,649	9,350,642	5,821,922

Notes: Dependent variable is monthly energy usage (kWh). The unit of observation is household-month. All models include wave-of-treatment (RCT) and month-of-sample fixed effects. Coefficients superscripted by T, M, and H denote the Pre-Treatment, Treatment and Pre-Move, and Post-Move time period, respectively. Additional controls include (2) pre-experiment average monthly usage, (3) cooling and heating degree days (CDD and HDD), (4) vacancy rate in percent at the ZIP level, (5) environmental concern index (lifetime) of congressional representatives at the ZIP level (0-100), (6) the percentage of households giving at least $200 to any Green Party committee from 2008 to 2015 on the county level, and (7) a combination of (2)-(6). Sample sizes differ because we do not observe all covariates for every ZIP code (county). Robust standard errors are clustered at the property level for all specifications.

Table 3: Estimates of Treatment Effect: Robustness of Pooled Difference in Difference Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\hat{\beta}^{T}$	6.50	6.45	5.31	6.01	5.52	5.83	5.48	6.25
	(2.05)	(2.13)	(2.05)	(2.13)	(2.05)	(2.13)	(2.05)	(2.13)
$\hat{\beta}^{H}$	-53.71	-45.70	-36.96	-44.45	-38.98	-41.32	-39.26	-45.70
	(1.72)	(1.79)	(1.74)	(1.83)	(1.71)	(1.78)	(1.71)	(1.79)
$\hat{\delta}^{trt}$	-24.98	-25.71	-25.69	-25.49	-27.71	-27.36	-26.57	-26.37
	(2.05)	(2.18)	(2.05)	(2.18)	(2.05)	(2.18)	(2.05)	(2.18)
$\hat{\beta}^{M}$	-148.30	-56.93	-81.38	-48.06	-38.59	-0.67	-44.72	-22.02
	(2.73)	(2.82)	(3.46)	(3.59)	(2.78)	(2.86)	(2.75)	(2.84)
$\hat{\delta}^{move}$	-11.35	-10.40	-8.20	-10.23	-6.20	-7.01	-7.64	-9.03
	(2.64)	(2.67)	(3.32)	(3.38)	(2.64)	(2.67)	(2.65)	(2.68)
Rule(s)		1	2	1 & 2	3	1 & 3	4	1 & 4
R^2	0.216	0.218	0.227	0.228	0.225	0.225	0.217	0.219
N	9,350,745	8,455,824	8,020,384	7,339,044	8,963,045	8,257,039	9,019,435	8,317,631

Notes: Dependent variable is monthly energy usage (kWh). The unit of observation is household-month. All models include wave-of-treatment (RCT) and month-of-sample fixed effects. Coefficients superscripted by T, M, and H denote the Pre-Treatment, Treatment and Pre-move, and Post-Move time period, respectively. Column (1) is the baseline model utilizing the full sample. Exclusion Rules: Rule 1 excludes observations with $e_{ijt} < 150$ in any or all of the first six months after move and homes for which $\bar{e}_{ijt=move} < (\bar{e}_{ijt=pre-move} - 2 \cdot SD(e_{ijt=pre-move}))$. Rule 2 excludes the first six months after move regardless of use. Rule 3 disregards post-move observations that are below the smallest pre-move observation minus 20 percent (i.e. $\min e_{ijt=pre-move} \cdot 0.8$). Rule 4 excludes all post-move observations with $e_{ijt} < 200$, regardless of when they occur. No additional controls are included. Robust standard errors are clustered at the property level for all specifications.

Table 4: Estimates of Persistence: Meta-Analysis of Wave-Cohort Difference in Difference Model

	(1)	(2)	(3)
$\hat{\gamma}^{prst}_{total}$	0.3468		
	(0.0495)		
$\hat{\gamma}^{prst}_{<1Yr}$		0.2632	
		(0.0640)	
$\hat{\gamma}^{prst}_{\geq 1Yr}$		0.5295	
		(0.0614)	
$\hat{\gamma}^{prst}_{<2Yr}$			0.3300
			(0.0513)
$\hat{\gamma}^{prst}_{\geq 2Yr}$			0.5449
			(0.0693)
Null Hypothesis (H_0)			
$\hat{\gamma}^{prst}_{<t} = \hat{\gamma}^{prst}_{\geq t}$, p-value		<0.01	0.02
R^2	0.000	0.162	0.147
N	654	654	654

Notes: Coefficients in the table represent the average proportion of initial treatment effects (i.e., during the Treatment and Pre-Move time period) that persist in the Post-Move period, γ. For example, a coefficient of 0.3 means that 30% of the initial treatment persist after move. The first column presents an estimate of γ^{prst} based on (8) for the full sample. Robust standard errors are clustered at the wave level for all specifications. We excludes wave-cohorts with less than 10 unique households. We estimate γ via weighted least squares according to the inverse variance of $\hat{\delta}^{move}_{jc}/\hat{\delta}^{trt}_{jc}$. Columns 2 and 3 present estimates of γ^{prst} for two strata: movers who are exposed to one year (two years) of treatment or less and those who are exposed to more than one year (two years) for Columns 2 and 3, respectively. Robust standard errors are clustered at the wave level for all specifications.

Table 5: Testing Predictions of a Sorting Model

Sort Variable (0-100)	Sorting Prediction	Estimates $\hat{\delta}^{trt}$	$\hat{\delta}^{trt} \cdot$ Sort	$\hat{\delta}^{move}$	$\hat{\delta}^{move} \cdot$ Sort
Vacancy Rate	[−]	-27.86	0.39	-21.50	1.16
		(3.46)	(0.33)	(4.78)	(0.47)
Environmental Index (Annual)	[+]	-34.51	0.20	-11.68	0.03
		(3.97)	(0.07)	(4.69)	(0.08)
Environmental Index (Lifetime)	[+]	-35.63	0.23	-13.21	0.07
		(3.94)	(0.07)	(4.83)	(0.08)
Green Party Donations (District)	[+]	-33.36	3.86	-16.51	2.32
		(4.89)	(2.05)	(6.21)	(2.59)
Green Party Donations (County)	[+]	-25.68	0.16	-11.88	0.13
		(2.08)	(0.08)	(2.68)	(0.10)

Notes: Brackets indicate the sign on the interaction coefficient ($\hat{\delta}^{move} \cdot$ Sort) predicted by the sorting model. Dependent variable is monthly energy usage (kWh). The unit of observation is household-month. We utilize the following measures as proxies of home vacancies and environmental sentiment: i) vacancy rate in percent at the ZIP level, ii) annual environmental concern index of congressional representatives for each district, iii) lifetime index of representatives, iv) percentage of households giving at least $200 to any Green Party committee from 2008 to 2015 on the congressional district level, and v) giving to Green Party on the ZIP code level. See Online Appendix Table A4 for all model estimates. Robust standard errors are clustered at the property level for all specifications and we report coefficients using wave and month-of-sample fixed effects.

Table 6: Implications of Findings

	$\hat{\gamma}^{prst}_{tot}$	$\hat{\gamma}^{prst}_{\geq 1Yr}$	$\hat{\gamma}^{prst}_{\geq 2Yr}$
Panel A: Comparison of Persistence Estimates to AR			
Estimate	0.347	0.529	0.545
Standard Error	0.049	0.061	0.069
Null Hypothesis (H_0)			
$\hat{\gamma}^{prst}_i = 0.635$ [Site 1], p-value	0.000	0.093	0.200
$\hat{\gamma}^{prst}_i = 0.671$ [Site 2], p-value	0.000	0.027	0.077
$\hat{\gamma}^{prst}_i = 0.623$ [Site 3], p-value	0.000	0.138	0.270
Panel B: Cost-Effectiveness in Treatment and Pre-Move Period			
Direct Program Costs ($/household-month)	1	1	1
Indirect Program Costs ($/household-month)	0.74	1.12	1.16
Savings (kWh/household-month)	24.98	24.98	24.98
Cost-Effectiveness (¢/kWh)			
No Technology	4.00	4.00	4.00
Technology	6.95	8.50	8.64

Notes: Panel A provides a direct comparison of our estimates of persistence due to capital investment and persistence in AR due to technology and habit formation. We relate our estimates from Table 4 to all three sites in AR (Table 7, p. 3031). In particular, persistence in AR is defined as the proportion of the treatment effect of the dropped group in the post-drop period relative to the treatment effect of the continued group in the same time period. The comparison is based on a simple one-sample T-Test and table entries represent the corresponding p-values. Panel B reproduces a typical cost-effectiveness calculation for two cases in the treatment and pre-move period: (No Technology) assumes that reductions are solely caused by behavioral adjustments; (Technology) instead considers that a proportion of the overall treatment effect may be due to capital investment, based on our estimates of persistence. The difference between the two scenarios is the inclusion of indirect costs associated with capital investment. We follow Allcott and Greenstone (2012) and use their preferred estimate of 8.5¢ per kWh saved due to capital investment (p. 17). For both cases, we assume a cost of $1 per report and use the preferred estimate from Table 4 to determine monthly savings (-24.98 kWh).

Online Appendix Figure 1: Time Series of Low-Usage Observations

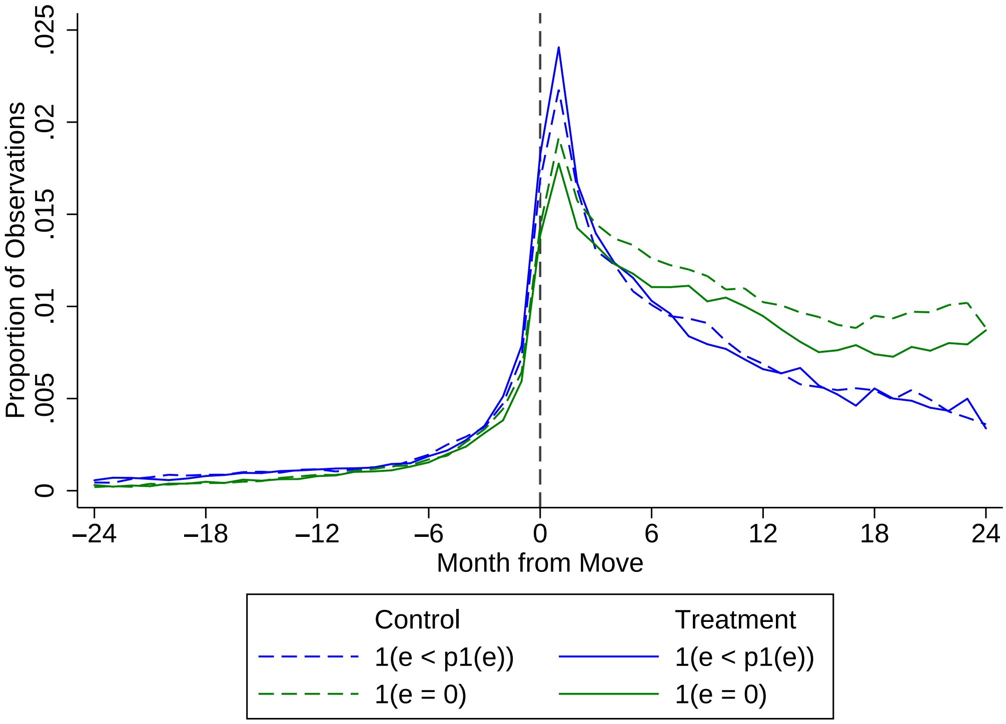

Notes: We plot the proportion of low-usage observations—defined as usage within the first percentile—and zero-usage observations for treatment and control households. Time $t = -1$ depicts the last full month before move, $t = 0$ the month in which the move occurs, $t = 1$ the first full month after move, and so on.

Online Appendix Table 1: Persistence in Studies of Habit Formation

Study	Topic	Weeks Since End of Intervention	Reported Effect	Significance Level of Reported Effect	Citation for Reported Effect	Persistence
Charness & Gneezy 2009	Exercise	0	1.32	0.01	Online Appendix & Authors' Calc.	1
Charness & Gneezy 2009	Exercise	5	0.88	0.01	Online Appendix & Authors' Calc.	0.67
Charness & Gneezy 2009	Exercise	8	0.67	0.05	Online Appendix & Authors' Calc.	0.51
Milkman, Minson, & Volpp 2014	Exercise	0	0.48	0.01	Table 3, Col. 2	1
Milkman, Minson, & Volpp 2014	Exercise	9	0.03	>0.05	Table 3, Col. 2 & Authors' Calc.	0
Acland & Levy 2015	Exercise	0	1.449	0.01	Table 1, Col. 1	1
Acland & Levy 2015	Exercise	8	0.188	0.05	Table 1, Col. 1	0.13
Acland & Levy 2015	Exercise	27	0.096	>0.05	Table 1, Col. 1	0
Royer, Stehr, & Syndor 2015	Exercise	0	0.56	0.01	Table 2, Col. 2	1
Royer, Stehr, & Syndor 2015	Exercise	9	0.12	0.05	Table 2, Col. 2	0.21
Royer, Stehr, & Syndor 2015	Exercise	22	0.11	0.05	Table 2, Col. 2	0.20
Royer, Stehr, & Syndor 2015	Exercise	48	0.12	0.05	Table 3, Col. 2	0.21
Royer, Stehr, & Syndor 2015	Exercise	100	0.12	>0.05	Table 3, Col. 2	0
Volpp et al. 2008	Weight Loss	0	9.2	0.05	Table 2 & Authors' Calc.	1
Volpp et al. 2008	Weight Loss	13	7.4	>0.05	Figure 2 & Corresponding Text	0
Levitt, List, & Sadoff 2010	Education	0	0.083	0.05	Table 7, Col. 1	1
Levitt, List, & Sadoff 2010	Education	52	0.063	>0.05	Table 7, Col. 1	0
Levitt, List, & Sadoff 2010	Education	104	-0.003	>0.05	Table 7, Col. 1	0
Levitt, List, & Sadoff 2010	Education	156	0.017	>0.05	Table 7, Col. 1	0
Volpp et al. 2009	Smoking	0	9.7	0.01	Table 2	1
Volpp et al. 2009	Smoking	26	5.8	0.01	Table 2	0.60
Volpp et al. 2006	Smoking	0	11.7	0.01	Abstract	1
Volpp et al. 2006	Smoking	13	1.9	>0.05	Abstract	0

Online Appendix Table 2: Overview of Utilities in the Sample

		All Households		Movers		
Utility	Location	Unique Homes	Observations	Unique Homes	Observations	Percentage
1	Midwest	346,480	14,450,660	72,384	3,013,774	20.89
2	Midwest	320,414	11,024,640	40,247	1,394,702	12.56
3	West	34,693	1,226,796	8,239	286,772	23.75
4	Midwest	482,902	16,291,960	44,694	1,625,511	9.26
5	Midwest	182,083	4,890,361	19,698	529,661	10.82
6	Midwest	76,721	2,175,558	2,837	83,622	3.70
7	West	119,025	4,479,500	13,046	484,373	10.96
8	Northeast	182,875	3,975,601	5,536	113,633	3.03
9	Northeast	50,374	1,092,431	3,948	81,246	7.84
10	Northeast	128,243	2,791,966	1,406	29,112	1.10
11	South	62,566	1,898,726	66	1,495	0.11
12	West	96,980	2,651,808	9,329	267,076	9.62
13	West	24,708	908,969	5,049	184,099	20.43
14	Midwest	73,918	2,122,780	6,309	177,671	8.54
15	Northeast	109,207	2,618,094	6,427	149,159	5.89
16	Midwest	40,680	1,252,132	6,123	193,201	15.05
17	South	96,434	2,731,790	6,754	208,618	7.00
18	West	26,773	556,650	1,091	22,362	4.08
19	West	83,896	5,527,008	18,319	1,206,545	21.84
20	West	39,918	680,969	12	141	0.03
21	West	122,580	2,546,808	3,783	77,302	3.09
22	Midwest	106,308	4,566,773	14,398	632,713	13.54

Notes: The table presents an overview of all utilities in the Opower sample with simple summary statistics. A nondisclosure agreement prohibits us from naming utilities but we present the geographic location. Please note that some utilities implement multiple treatment assignment waves (between one and six). For our research question, we focus on households from which initial occupants move out after receiving treatment for at least one month in both treatment and control ("Movers"). The last column provides the percentage of households that are movers in each utility, e.g. for utility one about 21 percent of households move during the sample period. All results presented in this paper are based on the movers subsample only.

Online Appendix Table 3: Movers Summary Statistics by Wave-of-Treatment

	Movers		Pre-Treatment Use				Movers		Pre-Treatment Use		
Wave	Number	Observations	T	C	First Letter	Wave	Number	Observations	T	C	First Letter
1*	30,395	1,291,026	864.363	833.725	Aug 2010	22	3,948	81,246	1,066.304	1,076.468	Jul 2012
2	23,070	978,824	1,688.004	1,694.489	Aug 2010	23	1,406	29,112	1,014.776	1,011.297	Jul 2012
3	15,822	656,702	1,182.726	1,163.252	Aug 2010	24	66	1,495	1,790.643	1,724.128	May 2011
4	3,097	87,222	2,141.371	2,123.092	Nov 2011	25	4,244	139,754	744.238	736.733	Apr 2011
5	14,463	616,760	1,084.527	1,073.237	Jul 2010	26	5,085	127,322	520.897	538.172	Feb 2012
6	12,423	411,274	1,162.318	1,179.342	May 2011	27	5,049	184,099	638.862	661.12	Mar 2011
7	13,361	366,668	807.008	809.596	Nov 2011	28	6,309	177,671	978.627	977.99	Oct 2011
8	8,239	286,772	390.083	388.592	Apr 2011	29	4,489	107,770	1,068.279	1,091.592	Feb 2012
9	12,933	699,228	1,236.463	1,271.17	Jul 2009	30	1,938	41,389	847.694	844.301	May 2012
10	800	32,865	1,126.435	1,150.096	Sep 2010	31	5,812	188,273	819.573	828.201	Jun 2011
11	19,697	657,358	1,634.975	1,635.109	May 2011	32$^\diamond$	311	4,928	621.94	.	Nov 2012
12	2,314	59,800	1,080.175	1,086.286	Jan 2012	33	6,702	207,878	1,700.615	1,687.892	Jun 2011
13	6,266	124,023	525.36	528.389	Jul 2012	34	52	740	2,730.257	2,653.666	Jan 2013
14	2,684	52,237	1,576.21	1,546.465	Jul 2012	35	1,091	22,362	1,894.862	1,881.965	Jul 2012
15	8,724	287,265	852.642	859.739	May 2011	36	18,319	1,206,545	923.236	914.195	Oct 2008
16	2,431	55,822	787.757	752.441	Feb 2012	37	12	141	798.864	395.571	Jun 2011
17	3,079	67,365	787.383	800.248	Apr 2012	38	3,783	77,302	1,411.025	1,427.986	Jul 2012
18	5,464	119,209	579.811	588.76	Apr 2012	39	10,821	541,080	837.02	847.447	Dec 2009
19	2,837	83,622	793.282	780.959	Jul 2011	40	1,129	31,823	1,163.58	1,138.052	Nov 2011
20	13,046	484,373	787.648	775.622	Jan 2011	41	2,448	59,810	611.772	605.47	Mar 2012
21*	5,536	113,633	1,131.824	1,158.822	Jul 2012						

Notes: The table presents descriptive statistics for movers in every wave-of-treatment (RCT) of our sample. Stars (*) denote waves that fail a Kolmogorov-Smirnov test for equality of the distribution of pre-experiment average usage across treatment and control at the 5 percent level. Exclusion of waves at the 10 percent level would also include waves 22, 23, and 40; results do not change qualitatively if this smaller sample is used instead. Additionally, wave 32 (denoted by $^\diamond$) does not feature any control movers. We consequently exclude waves 1, 21, and 32 from all main specifications. The "First Letter" column indicates the date on which the first treatment letter was generated; households enter the sample approximately one year prior to that. T denotes pre-experiment average monthly usage of households assigned to the treatment group (Home Energy Report) while C stands for control homes (no correspondence). Multiple waves can belong to the same utility (between one and six).

Online Appendix Table 4: Summary Statistics of Pooled Movers Sample

	Treatment	Control	Total
Pre-Treatment Use	1,090.703	1,073.156	1,084.505
	(620.013)	(589.549)	(609.483)
Post-Treatment & Pre-Move Use (H)	945.846	941.816	944.418
	(636.551)	(612.991)	(628.307)
Post-Move Use (M)	884.492	873.899	880.75
	(701.722)	(674.904)	(692.385)
Months in Sample	35.509	39.456	36.904
	(11.506)	(13.193)	(12.275)
Premises	163,877	89,505	253,383
Observations	5,819,193	3,531,552	9,350,745

Notes: The table presents descriptive statistics for the pooled sample of movers in all 38 waves. Average monthly usage is reported for the Pre-Treatment, Treatment and Pre-Move, and Post-Move time periods. T denotes households assigned to the Treatment group (Home Energy Report) while C stands for Control. The last column presents the overall average for the sample.

Online Appendix Table 5: Persistence of Treatment Effects after Move with Different Samples

	(1)	(2)	(3)	(4)
$\hat{\beta}^T$	6.03***	5.95***	6.52***	6.50***
	(1.82)	(1.85)	(2.03)	(2.05)
$\hat{\beta}^H$	-41.58***	-43.77***	-52.20***	-53.71***
	(1.57)	(1.59)	(1.73)	(1.72)
$\hat{\delta}^{trt}$	-23.69***	-23.45***	-25.02***	-24.98***
	(1.86)	(1.89)	(2.06)	(2.05)
$\hat{\beta}^M$	-131.33***	-134.63***	-145.69***	-148.30***
	(2.60)	(2.64)	(2.68)	(2.73)
$\hat{\delta}^{move}$	-10.62***	-10.48***	-11.03***	-11.35***
	(2.59)	(2.62)	(2.60)	(2.64)
R^2	0.223	0.223	0.216	0.216
N	11,149,461	10,906,948	9,491,121	9,350,745

Notes: Dependent variable is monthly energy usage (kWh). The unit of observation is household-month. All models include wave-of-treatment (RCT) and month-of-sample fixed effects. Coefficients superscripted by T, M, and H denote the Pre-Treatment, Treatment and Pre-Move, and Post-Move time period, respectively. We estimate the model for four samples: (1) all observations of households initially flagged as movers for balanced waves, (2) exclusion of households that were wrongfully flagged as movers due to changes to the account holder (e.g. marriage or divorce), (3) additional exclusion of households whose move-out date was updated with a later data pull and who moved after our last observation for the given customer ID, and (4) additional exclusion of households with updated move-out date that moved before the first treatment letter was received. Robust standard errors are clustered at the property level for all specifications. *** denotes significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Online Appendix Table 6: Development of Persistence Post-Move

	(1) 12-36 Months	(2) 15-33 Months
$\hat{\beta}^T$	6.33*	7.70*
	(3.45)	(4.06)
$\hat{\beta}^H$	10.08***	31.57***
	(2.67)	(3.11)
$\hat{\delta}^{trt}$	-27.09***	-28.39***
	(2.21)	(2.51)
$\hat{\beta}^M$	-127.14***	-95.92***
	(5.11)	(6.06)
$\hat{\beta}^M \cdot$ Time	10.65***	13.68***
	(0.37)	(0.46)
$\hat{\delta}^{move}$	-18.48***	-20.44***
	(4.90)	(5.97)
$\hat{\delta}^{move} \cdot$ Time	0.89*	0.46
	(0.45)	(0.57)
R^2	0.203	0.202
N	4,069,425	2,909,426

Notes: Dependent variable is monthly energy usage (kWh). The unit of observation is household-month. All models include wave-of-treatment (RCT) and month-of-sample fixed effects. Coefficients superscripted by T, M, and H denote the Pre-Treatment, Treatment and Pre-Move, and Post-Move time period, respectively. Time denotes the time in months since move-out of the original household. Interactions with this time indicator provide a measure of decay. (1) utilizes all households that were exposed to treatment between 12 and 36 months before move, (2) those who were exposed between 15 and 33 months. Robust standard errors are clustered at the property level for all specifications. *** denotes significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Online Appendix Table 7: Heterogeneity in Persistence with a Meta-Analysis Approach and various Exclusion Rules

	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Underlying Model: Month-of-Year FEs						
$\hat{\gamma}$	0.3383***	0.3517***	0.3547***	0.3425***	0.3562***	0.2918***
	(0.0425)	(0.0417)	(0.0438)	(0.0418)	(0.0434)	(0.0525)
N	693	691	677	700	679	414
Underlying Model: Month-of-Sample FEs						
$\hat{\gamma}$	0.3409***	0.3544***	0.3576***	0.3292***	0.3591***	0.2940***
	(0.0426)	(0.0418)	(0.0439)	(0.0428)	(0.0435)	(0.0527)
N	693	691	677	700	679	416

Notes: The table presents estimates of γ based on (8). Robust standard errors are clustered at the wave level for all specifications. Model 1 utilizes a first stage regression with month-of-year fixed effects. Model 2 uses month-of-sample fixed effects. We estimate γ via weighted least squares according to the inverse variance of $\hat{\delta}_{jc}^{move}/\hat{\delta}_{jc}^{trt}$. Exclusion rules are: Rule 1: Exclude wave-cohorts with fewer than ten movers; Rule 2: Exclude wave-cohorts with 16 or fewer movers; Rule 3: Exclude first to fifth percentile of wave-cohorts in terms of number of movers; Rule 4: Exclude wave-cohorts with less than 100 observations; Rule 5: Exclude first to fifth percentile of wave-cohorts in terms of number of observations; Rule 6: Rule 1 and exclude all wave-cohorts with positive $\hat{\delta}^{trt}$, i.e. no reduction in response to the social comparison messaging. *** denotes significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Online Appendix Table 8: Linear Trend of Persistence using a Meta-Analysis Approach

	(1)
$\hat{\gamma}^{prst}_{total}$	0.2442***
	(0.0720)
Cohort	0.0108**
	(0.0046)
R^2	0.014
N	654

Notes: Coefficients in the table represent the average proportion of initial treatment effects (i.e., during the treatment and pre-move time period) that persist in the Post-Move time period, γ. For example, a coefficient of 0.3 means that 30% of the initial treatment persist after move. Cohort denotes the length of exposure to treatment pre-move. Inclusion of this variable is aking to a simple time trend in length of exposure. We excludes wave-cohorts with less than 10 unique households. We estimate γ via weighted least squares according to the inverse variance of $\hat{\delta}^{move}_{jc}/\hat{\delta}^{trt}_{jc}$. Robust standard errors are clustered at the wave level for all specifications.

Online Appendix Table 9: Sorting of Households based on Home Vacancies and Environmental Sentiment

	(1) Vacancy Rate	(2) Env. Index (Annual)	(3) Env. Index (Lifetime)	(4) Green Party (District)	(5) Green Party (County)
$\hat{\beta}^T$	7.83**	6.33*	6.87**	16.42***	6.74***
	(3.47)	(3.35)	(3.37)	(4.77)	(2.07)
Sorting Variable	-0.45*	-0.47***	-0.44***	9.96***	0.04
	(0.27)	(0.05)	(0.05)	(1.59)	(0.06)
$\hat{\beta}^T$ · Sorting	-0.18	0.00	-0.01	-4.53**	-0.06
	(0.31)	(0.05)	(0.05)	(1.99)	(0.07)
$\hat{\beta}^H$	-40.29***	-44.85***	-44.81***	-44.49***	-52.48***
	(2.89)	(3.28)	(3.25)	(3.74)	(1.75)
$\hat{\beta}^H$ · Sorting	-1.51***	-0.18***	-0.17***	-4.27***	-0.28***
	(0.27)	(0.06)	(0.06)	(1.54)	(0.07)
$\hat{\delta}^{trt}$	-27.86***	-34.51***	-35.63***	-33.36***	-25.68***
	(3.46)	(3.97)	(3.94)	(4.89)	(2.08)
$\hat{\delta}^{trt}$ · Sorting	0.39	0.20***	0.23***	3.86*	0.16**
	(0.33)	(0.07)	(0.07)	(2.05)	(0.08)
$\hat{\beta}^M$	-138.29***	-167.22***	-169.14***	-145.22***	-147.90***
	(4.38)	(4.20)	(4.27)	(5.31)	(2.76)
$\hat{\beta}^M$ · Sorting	-1.19***	0.43***	0.47***	-1.39	-0.10
	(0.40)	(0.06)	(0.06)	(2.13)	(0.09)
$\hat{\delta}^{move}$	-21.50***	-11.68**	-13.21***	-16.51***	-11.88***
	(4.78)	(4.69)	(4.83)	(6.21)	(2.68)
δ^{move}· Sorting	1.16**	0.03	0.07	2.32	0.13
	(0.47)	(0.08)	(0.08)	(2.59)	(0.10)
R^2	0.216	0.217	0.217	0.216	0.216
N	9,247,833	8,916,489	8,921,649	9,350,745	9,350,642

Notes: Dependent variable is monthly energy usage (kWh). The unit of observation is household-month. All models include wave-of-treatment (RCT) and month-of-sample fixed effects. Coefficients superscripted by T, M, and H denote the Pre-Treatment, Treatment and Pre-move, and Post-Move time period, respectively. We use five different sorting variables: (1) Vacancy Rate, i.e. the percent of empty homes in a given month and ZIP code from 0 to 100, (2) Average Annual Environmental Index for Congressional Representatives for each ZIP-year in the sample from 0 to 100, (3) Lifetime Environmental Index for every representative in all years of the sample, (4) Proportion of households donating to any Green Party committee between 2008 and 2015 in each Congressional District, and (5) Proportion of donors in each county. Sample sizes differ because data for some ZIP codes, counties or districts are missing. Robust standard errors are clustered at the property level for all specifications. *** denotes significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.