Sure It Works, But How Long Does It Last? Persistence of Savings After Short-term Participation In Behavioral Programs

Chris Russell, NMR Group, Inc., Bryan, TX
Lisa Wilson-Wright, NMR Group, Inc., Melrose, MA
Kim Oswald, Viridian LLC, Austin, TX

ABSTRACT

Evaluators have documented the success of behavioral programs administered through regular delivery of paper energy reports in inducing energy savings. However, one remaining question about the programs involves the persistence of savings after program intervention is discontinued. This paper examines whether short term exposure to program intervention also results in continued energy savings after report cessation.

In 2011, Connecticut Light & Power (CL&P) piloted a behavioral program that varied intervention duration among the participant group: One sub-group received reports for eight consecutive months (the discontinued sample), while other participant groups received reports for a full calendar year. The research compares the energy reduction rates of the discontinued sample to those receiving reports for a full year, examining energy savings for both the first eight months of the program when all groups received intervention and the remainder of the year after the discontinued sample stopped receiving reports.

The analysis relies on monthly billing data for both program participants' and program control group households. The evaluators examined electricity consumption for the baseline year of 2010 and the treatment period of 2011 and the first quarter of 2012. Regression analysis was used to isolate program savings and determine whether short term program exposure leads to long term energy savings. The results show that the discontinued group achieved statistically significant energy savings while receiving the reports and these savings persisted—albeit at a lower level—for four months after the cessation of treatment. However, by the fifth month post-treatment all significant energy savings were gone.

Introduction

In recent years, numerous program administrators across North America have implemented residential behavioral programs that encourage energy savings through social comparison and built in evaluation through a reliance on experimental design. The last 60 years of behavioral research has shown that people do not always act in a perfectly rational manner but instead display systematic patterns in their decision-making process and judgments. This paper focused on a program based on the concept of "relative context." This concept describes a process in which actors make decisions about their behavior based on their own interpretation of the behavior of their peers (Gilovich, Griffin, and Kahneman 2002). A household acting in a completely rational manner with regard to energy consumption would adopt every energy efficient behavior and measures available to them to reduce their energy use and subsequent cost.[1] However, households do not always adopt the most energy efficient option available to them for a number of reasons: mistrust of technology, dislike of product output, high costs of short term investment trumping the longer term benefits. Recent residential energy programs, however, have documented the usefulness of relative context in motivating energy-saving behavior. These programs give households a relative context to which

[1] Availability includes having the financial means to afford adopting the behavior or measures and having the technical capability of doing so.

they can compare their energy use—specifically a similar household's energy use—while also providing suggestions to change their behaviors in order to reduce their own energy consumption.

In these residential behavioral programs, randomly selected treatment households (participants) regularly receive a home energy report that compares their energy use to that of similar households and then suggests ways the treatment households can reduce their energy use. Impact evaluations of these programs demonstrate that they have been successful at inducing substantial energy savings for program administrators across the United States. The behavioral programs yield small energy savings at the household level, but the implementer sends the relatively inexpensive reports to thousands of households, meaning that the collective participant group generates thousands of kilowatt hours of savings attributable to the program. Despite this success, questions remain regarding how long savings will persist after households stop receiving reports: what, if any, is the lasting treatment effect? This paper addresses the question of the persistence of savings after report cessation. Specifically, the paper explores the persistence of savings for households that received program intervention for eight months, with the duration selected by the program administrator and implementer and not the evaluators.

Persistence of Behavioral Program Savings

The persistence of savings from behavioral programs has received minimal research attention to date. In the one public study of which the authors are aware, Allcott and Rodgers (2012) followed a four-year home energy report behavioral program that included a discontinued treatment group, a continued treatment group (authors refer to it as the continued group), and a control group. The discontinued treatment group in this study received monthly reports for two years while the continued treatment group received monthly reports for four years. Alcott and Rodgers showed that when both treatment groups received reports (the first two years of the study) the energy savings of the continued and discontinued treatment groups were fairly similar when compared to control group energy usage (**Figure 1**). However, after the discontinued treatment group stopped getting reports, their energy savings were similar to the continued treatment group for only about three months. After four months of not getting reports the discontinued treatment groups' energy savings started to deviate from the continued treatment group, who was still receiving reports. Specifically, the discontinued treatment group still showed statistically significant energy savings for two years after the cessation of treatment but the savings were less than those achieved by the continued group that received treatments consistently over the four year study period.

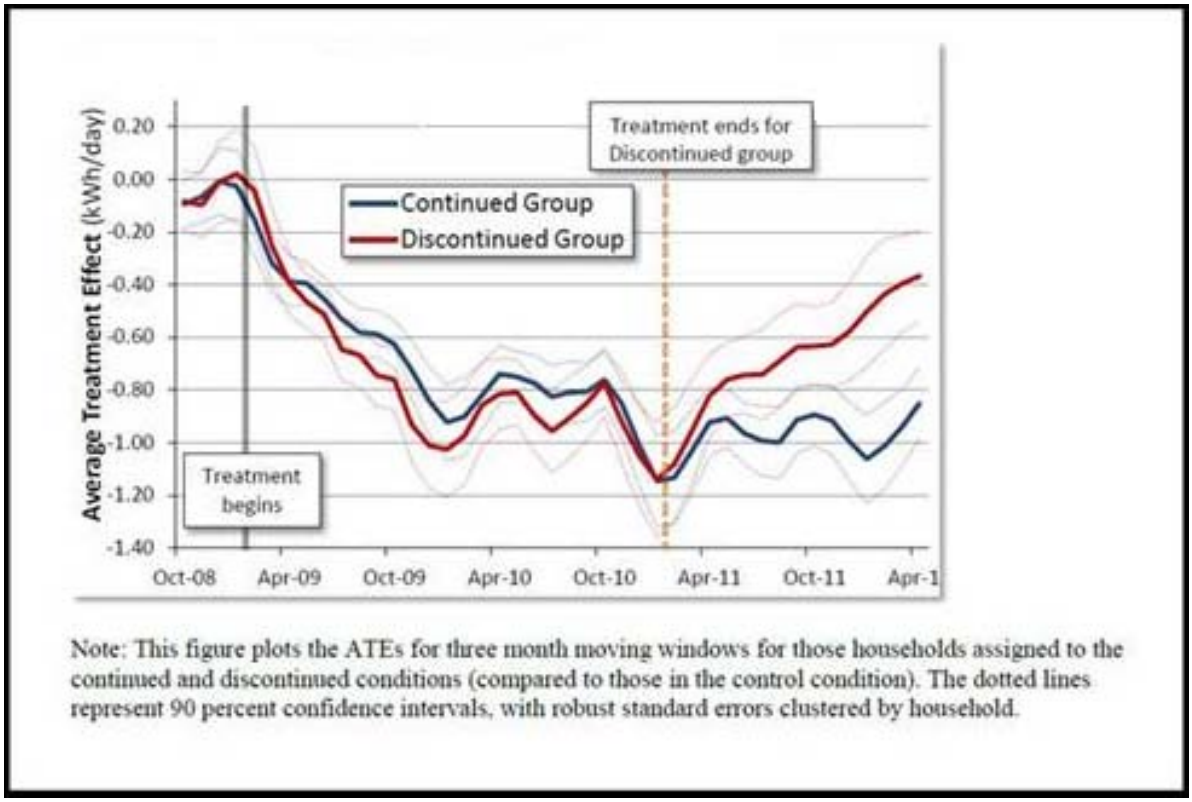

Figure 1. Persistence and Durability[2]

Allcott and Rogers demonstrate that the treatment effect (energy savings) of home behavioral reports persist after households stop receiving reports even though the magnitude of the treatment effect diminishes; yet, it is also the case that these households received report for two years prior to cessation. Do these same effects persist with a shorter treatment period? The authors of this paper address this question by looking at a shorter period of treatment and the corresponding treatment effect over time, and the results may help the energy efficiency community establish the minimum effective dose of treatment reports that translate into continued energy savings even after households stop receiving reports.

Scope

During the 2011 calendar year a CL&P administered a behavior program. The program was set up as an experimental design and had both randomly selected treatment groups (three in total) and a control group. One treatment group consisted of 10,000 participants—here called the continued treatment group (equivalent to Allcott's continued group) —and received reports every month for the duration of the study (**Table 1)**. A second treatment group, the discontinued group, included 4,000 participants who received monthly treatments for an abbreviated portion of the program. A third treatment group of 10,000 participants that received quarterly reports also existed in the original study, but the nature of their savings is beyond the scope of the present paper and will not be discussed further. The treatment groups were compared to a randomly selected, control group consisting of 24,000 households that were not aware that their energy use was being observed as part of the program.

[2] The Persistence and Durability figure was originally found in H. Allcott and T. Rogers 2012 paper "How long do treatment effects last? Persistence and durability of a descriptive norms intervention's effect on energy conservation (Faculty Research Working Paper)." http://web.hks.harvard.edu/publications.

Table 1. Sample Sizes

Sub-treatment group	Treatment Group	Control Group
Continued	10,000	n/a
Discontinued	4,000	n/a
Quarterly*	10,000	n/a
Total	24,000	24,000

*A third treatment group that received reports quarterly for a year is not examined in this paper.

Treatment households in all groups received their first report in the first quarter of 2011.[3] The discontinued treatment group received reports through August or September 2011, while the monthly treatment group received reports for a full twelve months (i.e., through early 2012) (Figure 2). This paper summarizes two separate but related analyses. First, it establishes the energy savings of the discontinued treatment group and for the continued treatment group, when compared to the control group, during the period of program intervention. The authors then examine how long the discontinued treatment group continued to save energy compared to the continued treatment group after the discontinued group stopped receiving energy reports.

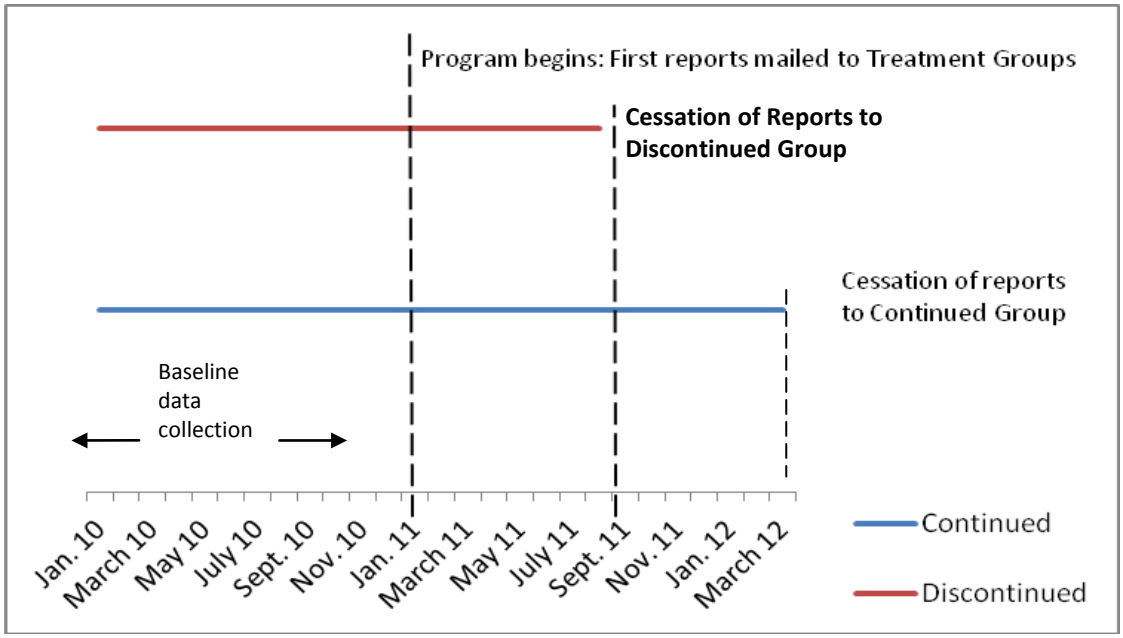

Figure 2. CL&P Behavioral Program Timeline

Methodology

The team performed a billing analysis of customer electricity bills to determine the discontinued of savings. The program administrator provided 27 months of billing data for the participant and control group households covering a year-long pre-treatment period used to establish baseline energy use, the eight month program intervention (treatment) period for the discontinued group, and a seven month post-

[3] The implementer staggered deliver of reports based on billing cycles, with the majority of first reports being sent in January and February but a few were sent as late as April. Likewise, continued report households stopped receiving reports on a rolling basis twelve months after treatment began, with the last batch being delivered in March 2012.

treatment period. The study covers the months of January 2010 through March 2012. The evaluators prepared a dataset containing billing, program, rate code, and weather data and then analyzed the data in STATA. The billing analysis relied on a statistical technique known as ordinary least squares (OLS) robust regression, which is resistant to any imbalances in pre-program use between treatment and control groups and also to data point outliers; thus, OLS ensures that the method does not over-estimate or underestimate treatment effects.

The analyses relied on the estimating equation below:

$$\text{Estimated Average Energy Savings} = \beta_0(\text{Avg. Post-Treatment Energy Use}) + \beta_1(\text{Dichotomous Treatment}) + \beta_2(\text{Avg. Pre-Treatment Energy Use}) + \beta_3(\text{Dichotomous Electric Heat}) + \beta_4(\text{Dichotomous Single Family Home}) + \beta_5(\text{Heating Degree Days}) + \beta_6(\text{Cooling Degree Days})$$

All reported results have also been multiplied by negative one (-1.0) for ease of interpretation; this step converts a measure of decreased use—a negative number—to a measure of savings—a positive number.

To assess the persistence of savings, the team compared the energy savings of the discontinued treatment group to those of the continued treatment group before and after cessation of the reports. To assess treatment effect while both treatment groups received monthly reports, the team aggregated savings for January through August. Then, to examine the persistence of savings explicitly, the team turned to two analyses. First, the team measured the treatment effect after report cessation for the discontinued treatment group by looking at the combined months of September 2011 through March 2012 to verify whether or not there were any statistically significant energy savings for the discontinued and continued treatment groups during the months after treatment, all the while controlling for baseline usage. In the second approach the team examined savings for each individual month between September 2011 and March 2012, while controlling for baseline energy usage, providing a way to identify not only if savings persist but also how long they persist.

Results

Table 2 presents the energy savings for both continued and discontinued treatment groups during the period when the discontinued treatment group was receiving monthly reports (from January 2011 through August 2011) and then again for the months immediately following the period when the discontinued group stopped receiving reports. The first column shows that during the first eight months of the program the average continued treatment household saved one kWh daily while the average discontinued household saved a similar 0.8 kWh during the first eight months of the program. The second column demonstrates that a gap in energy savings appears between the discontinued group and continued report groups during the six months after the discontinued group stopped receiving reports; that is, while the continued group savings increased to 1.1 kWh per day, the discontinued group savings decreased to 0.5 kWh per day. Thus, the discontinued treatment group achieved statistically significant but comparatively smaller savings during the time period.

Table 2. Estimated Average Daily Energy Savings for Discontinued Treatment Group

	From January 2011 Through August 2011	From September 2011 Through March 2012
Continued Treatment Effect	0.98 kWh (1.93%)	1.11 kWh (2.30%)
Discontinued Treatment Effect	0.80 kWh (1.58%)	0.51 kWh (1.06%)
Treatment Sample Size	23,592	22,815
Control Sample Size	23,702	22,923
Explained Variance	78%	72%

Breaking the savings down by discontinued treatment group post-report month reveals even more information about the persistence of savings. Table 3 shows that for the first two months after the discontinued treatment group stopped receiving reports (columns A and B), its estimated average energy savings were similar to those of the continued group and even exceeded the continued group in September of 2011—this is most likely due to simple variation in energy use and is not systematically related to the fact that households in this group stopped receiving reports. However, starting in the third month after report cessation, the discontinued treatment group savings decreased substantially (column C to Column G). In November through December (Columns C through D), the discontinued treatment group still saved energy, but at less than one-half the amount they saved when receiving monthly reports. By January (Column E), the discontinued households were no longer achieving statistically significant savings.

Table 3: Estimated Average Daily Energy Savings for Discontinued Treatment Group by Month

	A	B	C	D	E	F	G
	Sept. 2011	Oct. 2011	Nov. 2011	Dec. 2011	Jan. 2012	Feb. 2012	March 2012
Continued Treatment Effect	1.13 kWh (2.38%)	0.90 kWh (2.16%)	0.83 kWh (2.14%)	1.21 kWh (2.49%)	1.46 kWh (2.53%)	1.29 kWh (2.33%)	1.25 kWh (2.56%)
Discontinued Treatment Effect	1.34 kWh (2.81%)	0.84 kWh (2.03%)	0.34 kWh (0.86%)	0.39 kWh (0.80%)	0.23 kWh* (0.40%)	0.16 kWh* (0.29%)	0.26 kWh* (0.54%)
Treatment Sample Size**	22,259	21,804	22,045	22,356	22,348	22,029	22,240
Control Sample Size	22,985	22,170	22,152	22,478	22,488	22,194	22,369
Explained Variance	54%	57%	56%	58%	61%	61%	59%

* Not statistically different from the comparison group at the 90 percent level indicating no measureable savings.
** Includes the quarterly treatment group, whose savings are not reported here.

Figure 3 is a graphic representation of savings for the continued and discontinued treatment groups. Both the discontinued and continued treatment groups show statistically significant energy savings from June 2011 (six months after treatment had begun) until after the discontinued treatment group was no longer receiving reports (December 2011). Though the November and December energy savings are statistically significant for both treatment groups, these months are characterized by a large gap between the amount of energy saved between the two groups—namely the discontinued treatment group shows a steep decrease in energy savings two months after they stopped receiving reports. The evidence shows that in the absence of reports the discontinued treatment group energy savings persisted for two months at their pre-cessation level, then declined sharply and, by the fifth month of not receiving a report, the discontinued treatment group no longer displayed any measurable energy savings over the control group.

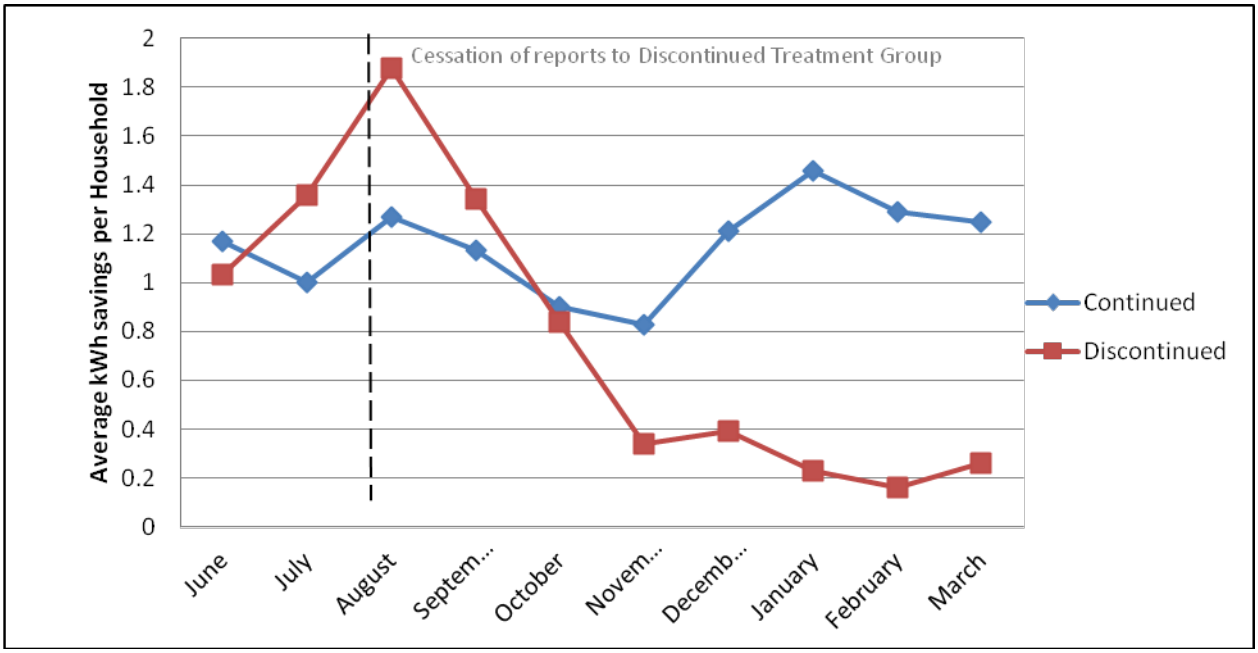

Figure 3. Estimated Average Energy Savings After the Cessation of Treatment for the Discontinued Treatment Group*

* Prior to the months represented in this graph, the discontinued group had been receiving monthly reports and their energy use was statistically similar to that of other continued report recipients.

Conclusions and Next Steps

The aggregated study results suggest that the discontinued treatment group continued to achieve measurable energy savings after it stopped receiving the home energy reports. This is an encouraging finding and would suggest that an eight month exposure to treatment reports translates into energy savings for at least six more months. However, because the evaluators took the analysis one step further and estimated monthly post-treatment energy savings, they were also able to determine that the discontinued treatment group savings for only the first four months after it stopped receiving reports. In other words, the savings persisted, but they were only significantly higher for a short period of time after cessation of the treatment.

The results of this study complement those of Allcott and Rogers (2012). After report cessation the energy savings of the current study's discontinued group were very similar to that of the continued treatment group's energy savings for two months. Allcott and Rogers found a similar length (three months) of comparable energy savings between discontinued and continued treatment groups even though the discontinued group had been receiving reports for two years prior to cessation. Just because eight months is a long enough exposure period to generate persistent results does not mean that one can conclude that this treatment period is the minimum effective dose to generate persistent savings after treatment cessation—the loss of significant savings among the discontinued treatment group four months after treatment seems to argue against the conclusion that eight months of program exposure is long enough to produce a lasting behavioral change that results in continued energy savings for the long-term. Future studies should continue to explore the optimal length of treatment to produce significant and lasting energy savings among treatment households. These studies will be strengthened by also taking the cost-effectiveness of treatment duration and persistence of savings into account in order to arrive at the optimal design of behavioral-report programs.

References

Allcott, H. and T. Rogers (2012) "How long do treatment effects last? Persistence and durability of a descriptive norms intervention's effect on energy conservation (Faculty Research Working Paper)." http://web.hks.harvard.edu/publications.

Gilovich, T., D. Griffin, D. Kahneman, eds. 2002. Heuristics and biases: The psychology of intuitive judgment. Cambridge University Press, Cambridge, U.K.