How much of that savings is really attributable to the program?

Author: Stef Wayland

As evaluators, we are often asked, “Is the savings statistically significant?” a great follow-up question that we don’t often hear is, “Is that savings really attributable to the program?” The first question is about the precision of the model, assuming that it is the correct model. The second question requires answering whether the model may be biased or inaccurate because it is not the right model, or the model is not really answering the right research question.

For energy impact evaluation, our research question is, “How much did the program cause participants to save energy and reduce demand?” This is the fundamental question for attribution of savings to a program, and getting to the answer takes more than just a fixed-effects model. This question of whether a program caused a change in energy usage moves us away from standard statistical analysis into causal analysis. In energy evaluation terms, we want to estimate net impacts.

In this paper, we compare traditional methods for impact evaluation methods with several emerging methods such as latent class discrete choice (LCDC), multilevel, and Bayesian models to demonstrate where our current methods fail and how we can overcome those shortfalls to best attribute savings to programs. There are many major issues that arise in these impact evaluations, including using the baseline required by regulators, adjusting for free-riders and free-drivers, and selecting an appropriate comparison group. All of these can have a substantial effect on the way that we calculate attributable savings.

Researchers specializing in causal modeling have made substantial gains in methods for assessing causality in just the last few years. We are working to bring these methods into energy evaluation to take advantage of the latest advances and to fully use AMI data when it is available.

Different programs can require substantially different approaches to attribute savings. We have used multilevel modeling in a home energy report program to assess savings for individual participants, and LCDC for a finance program, to assess the additional impact of financing on existing programs. This paper also examines Bayesian additive regression trees for a home energy report program and a weatherization program. We discuss standard methods for attribution in these programs and compare results to show where the newer methods outperform and where we might still need better methods.