

Not Your Mom's Matched Control Group

Jennifer Gai, Nexant

Current behavioral energy conservation programs are targeted and implemented at the individual household level. With the ability to implement interventions on large subsets of a target population, these programs are typically designed as randomized control trials (RCTs). While RCTs have become the standard for impact evaluations, they are difficult to implement when the intervention of interest is designed to treat social groups using actions that are expected to affect the entire group being treated (e.g., cities, schools, etc.), rather than a random subset of individuals.

Such is the case with the Energy Improvement Corporation's (EIC) Energize NY program—a community-based energy efficiency program. In partnership with the New York State Research Development Authority (NYSERDA), EIC has focused the Energize NY program on three communities in the Hudson Valley with the goal of increasing household participation in NYSEDA's Home Performance retrofit program. Because the intervention is implemented at the community-level, it is practically impossible to randomly assign some subset of the community to treatment, while withholding it from others. In cases where a RCT design is inappropriate, another common evaluation method is matched control groups, which finds an individual untreated unit that looks very similar to each treated unit prior to the onset of treatment. This method can provide reliable estimates of the program's impact when the treated units are very similar to their corresponding matched control unit, but it requires a large enough number of treatment and control units where the expected differences between the two groups prior to treatment is zero.

The Energize NY program will only be implemented in three of 15 communities in the Hudson Valley. It is unlikely to find a matched control community for each of the three randomly selected treatment communities from such a small pool of potential control communities. An extension of the matching approach that overcomes this limitation is called the synthetic control method (SCM). The SCM was developed for use in cases where the units of analysis are small numbers of large entities (e.g., countries, states, communities, etc.). In these cases, there are often a small number of treated units, but a somewhat larger number of potential control units, as is the case in the Energize NY program. The SCM estimates treatment impacts by constructing a single prediction of the outcome variable of interest for each treated community based on statistical weighting. Each synthetic control is a single weighted average of communities in the control pool selected for its ability to match as closely as possible to the variables of interest in the treated communities during the period prior to treatment. This poster discusses the challenges of designing a research and evaluation plan for community-level behavioral interventions, as well as the application of the SCM to evaluating the impact of the Energize NY intervention on participation in the Home Performance retrofit program.

In a series of additional graphics, we present the basic mechanics of constructing a synthetic matched control group, using the Energize NY program as an example. We also present the types of evaluations that are most conducive to the SCM. While the data are currently unavailable for us to show the estimated impacts from the intervention using the SCM, our poster provides a detailed background of this unique methodology that can provide reliable impact estimates for future behavioral programs conducted at the community/regional level.