Constructing Credible Counterfactuals in Programme Evaluation

As conservation interventions such as Payments for Ecosystem Services (PES) programmes, Agri-Environment schemes, Protected Areas (PAs), and Eco-certification schemes continue to grow in popularity and scale, it is becoming increasingly important to look at what we actually know about the effectiveness and impacts of these programmes in order to be able to identify any gaps or weaknesses in the evidence base. This understanding requires a critical assessment of the evaluation methods currently being used to evaluate conservation interventions.

Various methodological challenges exist in the design of studies to assess the impacts of conservation interventions. These challenges are not unique to the conservation policy field, but exist for research on most types of policy interventions. In order to be able to unambiguously attribute impacts to a conservation intervention, evaluation methods must construct a counterfactual outcome, which is an estimate of what would have happened in the absence of the conservation intervention. Using eco-certification as an example, the counterfactual is an estimate of what the environmental or socio-economics outcomes would have been for a certified farm or firm had it never been certified. Failure to construct a credible counterfactual can lead to biased or weak evidence regarding the true impacts of conservation interventions and is thus an important consideration when it comes to designing evaluation studies.

The following examples of different approaches to evaluating conservation interventions are based on Blackman and Rivera (2010).

Before-and-After Comparison Approach

One common approach to constructing a counterfactual is to simply compare outcomes for the same entity (village, individual, forest, farm, firm, etc.) before and after it has received the conservation intervention. In this case, the pre-intervention outcome is used as the counterfactual outcome. The implicit assumption here is that if entities that received the intervention had not received the intervention, that their outcomes on average, would have stayed the same. This assumption is violated when outcomes change during the study period because of confounders, i.e., factors that are unrelated to the intervention but affect the outcome. For example, imagine that an impact study of certified coffee production uses certified farmers' pre-certification productivity as the counterfactual outcome, thus measuring impact as the difference between average pre-certification and post-certification productivity levels. Furthermore, imagine that this difference is large and positive, leading the evaluator to conclude that certification raised average productivity levels. This estimate of certification impact would be biased upward (and the evaluator's finding of a causal effect would be misleading) if farmers' productivity levels rose after certification for reasons that had nothing to do with certification. A reason might include advantageous weather conditions over the time period studied, for example.

With-and-Without Comparison Approach

A second common approach to constructing a counterfactual is to simply compare entities that received the intervention with entities that have not received the intervention. In this case, the outcome from the non-intervention entity serves as the counterfactual outcome, that is, the entity that has not received the intervention serves as the control group. The implicit assumption here is that if entities that received the intervention had not received the intervention, their outcomes would be the same, on average, as those of the non-participating entities. This assumption is violated when entities with characteristics that affect outcome select themselves (or are purposively selected) into the programme, a problem known as selection bias. For instance, imagine that an impact study of certified coffee production on soil erosion uses a measure of soil erosion on non-certified farms as the counterfactual outcome and, therefore, calculates impact as the difference between average soil erosion measures for certified and non-certified farms. Furthermore, imagine that this difference is negative, significant, and large, leading the evaluator to conclude that certification drove reductions in soil erosion. This estimate of certification's impact would be biased upward (and the evaluator's finding of a causal effect would be misleading) if farmers with lower soil erosion rates self-selected into the certification scheme. This might happen if a disproportionate number of farms who had already adopted soil conservation measures sought certification, recognizing that they would not have to invest in additional conservation measures to meet certification standards.

Randomised or 'experimental' Approach

In order to be able to guarantee reliable (unbiased) inferences about causal effects, the evaluation literature emphasises two alternative approaches to impact evaluation that can be used to construct credible counterfactuals: experiments and quasi-experiments.

Experimental methods identify the effect of a conservation intervention by randomly distributing alternative causes over experimental conditions. These experiments, also known as 'randomised evaluations' or 'randomised field experiments,' require gathering a set of entities equally eligible to receive the conservation intervention and randomly dividing them into two groups, those that will receive the intervention and those that will not receive the intervention. The former is known as the treatment group and the latter is known as the control group. By randomly assigning the intervention among eligible beneficiaries, the random assignment itself creates comparable participating and non-participating groups that are statistically equivalent to one another (given appropriate sample sizes). As a result, the control groups generated through randomised selection serve as a perfect counterfactual outcome, free from selection bias. For certification projects, this amounts to randomly selecting entities to receive certification from among a group of qualified and interested candidates. The outcome for the randomly constituted (non-certified) control group is then used as the counterfactual outcome for certified entities. This approach requires building evaluation into conservation project design.

Matching or 'quasi-experimental' Approach

When it is not possible to use an experimental design, 'quasi-experimental' methods, grounded in theory and statistics, can be used to isolate and identify the causal effect of a conservation intervention. One common quasi-experimental approach, is to construct the control group using propensity matching or 'pair-wise' comparisons. The idea is to match participating entities with non-participating entities that have very similar, if not identical, observable characteristics that plausibly affect outcomes, and to use outcomes for this matched control sample as the counterfactual outcome. For example, in a study of the water quality impacts of banana certification, certified farmers would be matched with non-certified farmers of similar size, education, and previous history of adopting conservation practices. Measures of water quality for this matched control group would be used as the counterfactual. This approach depends on the dual assumptions that no unobservable characteristics of the entities in question (e.g., management skill) affect both selection into the certification program and outcomes, and that all non-certified entities in the matched control sample have characteristics that make them suitable for certification. The introduced uncertainty when using matching methods requires larger sample sizes compared to experimental designs.

References:

Blackman, A. & Rivera, J. (2010). The evidence base for environmental and socioeconomic impacts of "sustainable" certification. Washington DC: Resources For the Future.