Glossary of Econometrics and Statistics Terms

Table of Contents

Nansana Field Experiment (Buntaine et al., 2024)....1
Observation (units) & Sample....1
Matched-Pairs Design (Clustered Random Assignment)....1
Underlying mechanisms....2
Preregistration....2
Randomized field experiment....2
Hiding....3
Spillover (no interference)....3
Robust estimators (robustness checks)....4
Stable Unit Treatment Value Assumption (SUTVA)....4
Quasi-Experimental Design (Larsen et al., 2019)....5
Random variables (statistics)....5
Exogeneity Assumption....5
Endogeneity Bias....6
Difference-in-Difference (DiD)....6
Parallel Time Trends....6
Within Estimator (Fixed Effects Model)....7

Nansana Field Experiment (Buntaine et al., 2024)

Observation (units) & Sample

The observation or unit of analysis for this experiment is the “neighborhood (zone)”. The sample consists of 44 neighborhoods, specifically those with “verified access to formal waste collection services”. Therefore, the population we are making inferences about is specifically neighborhoods with access to formal waste collection services. The population of inference does not include neighborhoods without waste collection services.

Matched-Pairs Design (Clustered Random Assignment)

This word “paired” refers to a treatment assignment strategy where neighborhood units were paired based on contiguity and proximity. Half of the pairs were randomly assigned to treatment and half to the control group. For example, in the Nansana experiment there are 44 neighborhoods and 22 paired strata– 11 of the pairs (22 neighborhoods) were randomly assigned to treatment and 11 pairs (22 neighborhoods) were assigned to the control group. The purpose of the pair-matching design, in the context of this experiment, was to foster competition. The rationale being that paired contiguous neighborhoods will have an increased likelihood of being competitive.

This study design introduces a dependency structure within the data, paired neighborhoods will likely be more alike to each other and less alike than other paired neighborhoods with respect to the treatment effect. This introduces dependencies in the error term that needs to be accounted for in the model (see; i.i.d assumption).

Underlying mechanisms

The intervention (competition) involves multiple components which may drive the overall causal effect. Underlying mechanisms refer to the identification of the specific components of the intervention which are responsible for causing the effect. These mechanisms are sometimes referred to as *intermediate outcomes (or mediators)*, implying a chain of causality. For example, potential mechanisms identified by the authors as explaining in competition interventions effectiveness include “social comparison”, “leader and resident efforts”, “neighborhood pride”, and “sense of communal purpose” (p.4; Buntaine et al., 2024).

Preregistration

Preregistration is a procedure in which the hypothesis, methods, and research design are publicly documented using a pre-registration service before the study is conducted. This is done to ensure that researcher expectations (bias) does not influence the studies results– enhancing the credibility, transparency, and reproducibility of research findings. This has been identified as a critical component of rigorous experimental inference after the replication crisis was found to impact nearly every discipline across the sciences, including environmental science (Baker, 2016; Kelly, 2009).

Randomized field experiment

Randomized field experiments are conducted in real-world settings (i.e., ‘in the field’) where the researcher determines both the intervention (manipulation) and random assignment to treatment and control groups. In contrast, a *randomized controlled trial*

(RCT) is conducted in a controlled or semi-controlled setting (e.g., laboratory or hospital).

A third type of causal inference study design we will encounter in this course is called a *natural experiment* or *quasi-experimental design*. A natural experiment is a causal inference strategy where the researcher identifies an external factor or event which creates *as-if random variation* in the assignment of a 'treatment'. In a natural experiment, the *treatment* may be a naturally occurring event or policy, and the assignment to treatment and control groups is determined by external circumstances rather than by the researcher.

In econometrics, the term *treatment* is used broadly to refer to a variable of focal interest hypothesized to exert a causal effect on the outcome (i.e., dependent variable).

Hiding

If residents hide waste burning or don't follow the treatment program as expected, it can violate a key causal inference assumption *excludability*. The excludability assumption requires that the intervention influences the outcome (reduction in waste burning) only through the proposed mechanisms, such as leader efforts and community coordination, and not through other unintended pathways.

A potential threat to the validity of the study's main finding, that the treatment reduced the number of waste piles burned, is the possibility that residents concealed burn piles from the auditors measuring waste pile counts. One way to reduce the likelihood of this form of treatment non-compliance, as implemented in this study's experimental design, is to keep residents un-informed (blinded) to the timing and location of audits.

To further address this concern, the authors conducted a series of empirical tests to detect whether there was evidence that residents (or leaders) were *hiding* waste pile burns. One such test involved estimating waste pile counts as a function of the interaction between proximity to roads and treatment. This interaction term, if significant, would suggest that residents actively displaced waste piles to conceal burning from the research team. For example if the relationship between proximity and waste pile counts was stronger for the treated than for the control group (slope of higher magnitude) this might imply that treated neighborhoods are burning farther from the roads to conceal burns.

While this type of active non-compliance appears unlikely, given the tradeoff between effort and incentive, it remains critical to rule out alternative explanations to strengthen the causal claim and make it more difficult to refute. The authors make a compelling case that there is no evidence of hiding based on targeted models designed explicitly to detect such violations. Across estimators they found the treatment effect remained consistent (i.e., robust).

Spillover (no interference)

Spillover is a special case of the violation of the causal inference assumption Stable Unit Treatment Value Assumption (SUTVA). Spillover specifically refers to the possibility that one unit's treatment affects a control unit's outcome. For example, in the Nansana experiment it is not hard to imagine that a neighborhood (and associated community leader) assigned to the control condition would observe the benefits conferred by the treated neighborhoods and devise their own treatment to realize a similar outcome. Note, that when the treatment effect is a positive impact (e.g., reducing waste pollution) such spillover effects are desirable in consideration of overall impact. However, such spillover impedes the ability for researchers to measure a causal effect between conditions (i.e., it reduces the difference between the control and treatment group means). In the Nansana experiment spillover was observed, the municipal government independently organized their own competition in the control neighborhoods and reduced waste burning (Hurrah!- this is a fantastic result from an impact perspective). Luckily, for the research team, the spillover effect occurred after they had a chance to measure the treatment effect during the treatment competition period. This independent organization, however, precluded the research team from measuring long-term impacts of the waste burning experiment. As seen in figure 2 (panel A) the gap between treatment and control groups was reduced by this unexpected social action affecting control neighborhoods (8-months post-award).

Robust estimators (robustness checks)

In econometrics and statistics a finding is said to be 'robust' when reasonable alternative model specifications are tested and the treatment estimate remains stable. This may include testing alternative outcomes, sub-samples, or statistical assumptions– with findings said to be robust if the treatment effect estimate remains consistent/stable across conditions. If an estimator is found not to be robust, or the size/direction of the effect varies due to changes in model specification the result is sometimes referred to as fragile. In the Nansana experiment a series of alternative specifications were tested, these results are documented in the supplemental materials. If findings are not robust to specification changes, such as the inclusion/exclusion of pre-treatment controls or interaction coefficients, it makes it difficult to argue that selection bias/OVB can be ruled out.

Stable Unit Treatment Value Assumption (SUTVA)

This causal inference assumption implies that the treatment affects each treated observation unit directly and consistently. Said differently, SUTVA means the outcome of any unit is unaffected by the treatment status of other units. SUTVA includes two conditions:

1. *No interference:* One unit's treatment does not affect another unit's outcome.
2. *No hidden variation*: The treatment is implemented consistently for all units (i.e., all units receive the same treatment)

For example, SUTVA would be violated if observations in an experiment that didn't receive the treatment (control group) were indirectly affected by those who did (treated group). This kind of "spillover effect" obscures or contaminates the identification of the true impact of the treatment.

Quasi-Experimental Design (Larsen et al., 2019)

Random variables (statistics)

In this article, the term "random variable" is used to describe the *treatment variable* in the context of natural experiments, which rely on observational data. This might seem counterintuitive because in natural experiments, treatment is not randomly assigned or controlled by the researcher (as it is in a randomized experiment).

In statistics, a "random variable" is a technical term referring to a variable whose value can vary due to random factors or measurement error. Unlike in randomized experiments– where treatment is carefully assigned and assumed to be measured precisely (without error)– natural experiments often use data where these ideal conditions don't hold.

In observational data, it's generally unrealistic to assume that treatment is measured without error or that it is completely unrelated to other unobserved factors affecting the outcome (known as exogeneity). This introduces additional complexity when analyzing natural experiments compared to randomized ones.

...

Exogeneity Assumption

The exogeneity assumption means that the variable of interest (e.g., treatment) is unrelated to unobserved factors that could influence the outcome being studied. Exogeneity implies that the treatment is not correlated with the error term in a regression model. The error term represents all the factors we haven't measured or included in the model but that might still affect the outcome. If treatment is exogenous, we can be confident that changes in the treatment cause changes in the outcome and are not due to some hidden factor.

For example, in a randomized experiment, exogeneity is usually ensured because participants are randomly assigned to treatment or control groups. This randomization breaks any potential link between the treatment and unobserved factors, allowing researchers to isolate the true effect of the treatment.

…

Endogeneity Bias

When the exogeneity assumption is violated, it leads to what is known as endogeneity bias. This means that the treatment (or independent variable) is correlated with the error term in a regression model. In simpler terms, there is a connection between the treatment and unobserved factors that also influence the outcome, making it hard to determine the true effect of the treatment.

Endogeneity bias is closely related to selection bias and is a major concern in models that rely on observational data. For example, if motivated individuals are more likely to participate in the intervention and motivation also directly affects outcome, then motivation is an omitted variable which will introduce endogeneity bias (i.e., the treatment estimate will be biased).

…

Difference-in-Difference (DiD)

The Difference-in-Differences (DiD) estimator utilizes panel data (repeated measures) from before and after an intervention for both a treatment group and a control group. This gives us four conditions which can be used to estimate a treatment effect using the following equation,

$$\mathrm{DiD} = \left(\bar{Y}_{\text{treatment, post}} - \bar{Y}_{\text{treatment, pre}}\right) - \left(\bar{Y}_{\text{control, post}} - \bar{Y}_{\text{control, pre}}\right)$$

The key idea of DiD is to estimate what would have happened to the treatment group if it had not received the treatment. This is done by assuming that, in the absence of treatment, the treatment group and the control group would have followed the same trend over time (parallel trends assumption). The difference in outcomes between the two groups after accounting for these trends gives us the treatment effect.

…

Parallel Time Trends

The key to identification of an unbiased estimate of the treatment effect using the DiD estimator is whether the *parallel time trends assumption* holds up. To assess whether the parallel trends assumption is reasonable, measuring the trends in the period before treatment provides a counterfactual comparison. If the trends for the treatment and control groups were parallel before the intervention, it strengthens the argument that they would have remained parallel in the absence of treatment. The trends during the pre-treatment and post-treatment periods can be evaluated visually by comparing plotted trend lines for both the treatment and control groups. By focusing on differences in trends, the DiD approach effectively controls for factors that are constant over time within each group, isolating the effect of the treatment.

…

Within Estimator (Fixed Effects Model)

The “within estimator,” also known as the panel fixed effects model, uses differences within each group (e.g., year, site, or other units) to account for and remove variance that doesn’t change over time (i.e., time-invariant). By doing this, it removes selection bias concerns from unobserved factors that are constant over time, site, or unit. However, the *selection on unobservables* problem must still be considered for all time-varying endogenous variables (i.e., time-varying covariates correlated with the outcome and treatment).

This approach simplifies the problem of selection bias by focusing only on sources of OVB caused by factors that change over time. The model is estimated by either:

1. Including a separate coefficient (dummy variable) for each group (e.g., year, site, unit), or
2. Using a technique called demeaning, which calculates differences from the group’s average.

For the within estimator to give valid results, two main assumptions must hold:

1. No omitted time-varying factors: All variables that change over time and could affect both the treatment and outcome must be included in the model.
2. No reverse causality: The outcome should not influence the treatment (e.g., the dependent variable should not affect the independent variable).