

Randomized Controlled Trials

Erick F Molina

Summer, 2024

ITAM

Randomized Controlled Trials

- Randomized controlled trials (RCTs) solve the problem of **self-selection bias** by assigning treatment at random.
- If T_i is randomly assigned, it is easier to justify the independence assumption:

$$\{Y_i^T, Y_i^C\} \perp T_i$$

- We would expect treatment effects to be distributed similarly in the treatment and control groups:

$$E(Y_i^T - Y_i^C \mid T_i = 1) = E(Y_i^T - Y_i^C \mid T_i = 0)$$

- The ATT and the ATE should be equal in expectation, and we could estimate them using β_1 from a simple regression.

The usual timeline is:

- Collect baseline data: the initial status of individuals. This helps characterize the future sample.
- Implementation: the treatment is delivered. Individual characteristics may change during this period due to real-world factors unrelated to the treatment.
- Collect endline data: the final status of individuals in the sample.

Baseline

- Allows us to produce balance tables and check the independence of treatment assignment.
- Provides key characteristics of the sample that can be used as controls to reduce omitted variable bias (OVB).
- Stratification: enables random assignment in a way that makes treatment and control groups resemble the target population.
- Allows us to measure *attrition*, i.e., the loss of observations that were present at the beginning of the RCT.

In practice, it is generally not true that **all** individuals assigned to a group actually take (or do not take) the treatment.

- This is known as a *compliance* problem.
- Let Z_i denote assignment and T_i denote take-up (treatment received).

	$T_i(0) = 0$	$T_i(0) = 1$
$T_i(1) = 0$	Never Takers	Defiers
$T_i(1) = 1$	Compliers	Always Takers

- When not everyone is a *complier*, then $\beta_1 \neq \text{ATE}$.
- However, we can compute the Intention-to-Treat (ITT), identified by γ_1 :

$$Y_i = \gamma_0 + \gamma_1 Z_i + U_i \quad (1)$$

Attrition

- Loss of observations that were originally part of the experimental design and treatment assignment.
- It is important to document the extent of attrition and whether it differs between the treatment and control groups.
- Attrition leads to:
 - Loss of efficiency in the estimators.
 - Threats to external and internal validity.
- **Random attrition:** Loss of observations is independent of treatment and individual characteristics.
- **Non-random attrition:** Loss of observations is related to treatment and/or individual characteristics.

Random attrition

- Balance on observable variables is preserved between treatment and control.
- It should not pose an external validity problem.
- It does reduce efficiency.

Define s_i as a dummy indicating whether the observation is available at the end of the RCT. Then:

$$\mathbb{E}[s_i T_i U_i | X_i] = \mathbb{E}[s_i | X_i] \mathbb{E}[T_i U_i | X_i] = 0 \quad (2)$$

We can estimate the probability that an observation is available at the end of the RCT using the same approach as with a dummy dependent variable. Let S_i^* be a latent variable; we can estimate:

$$S_i^* = \delta_0 + \delta_1 T_i + X_i' \delta_2 + Z_i' \delta_3 + V_i \quad (3)$$

- We test the joint hypothesis:

$$H_0 : \delta_1 = 0, \delta_2 = 0, \delta_3 = 0 \quad (4)$$

$$H_1 : \text{at least one of the } \delta\text{'s is different from 0} \quad (5)$$

Non-random attrition

- It can create threats to external and internal validity.
- Estimation is restricted to the observed data:

$$E(Y_i | T_i, X_i, Z_i, S_i = 1) = \beta_0 + \beta_1 T_i + X_i' \beta_2 + E(U_i | T_i, X_i, Z_i, S_i = 1) \quad (6)$$

- **Inverse probability weights:** a nonparametric approach to address non-random attrition.
- **Heckman model:** uses variables Z_i to model attrition and estimate the ATE without bias (under assumptions).

- Reweight the available observations to obtain a sample similar to the original one.
- Use a probit or logit specification to compute the weights:

$$w(X_i, T_i = 1) = \left(\frac{\Phi(X_i' \delta_1 + \delta_2)}{N_{TS}/N_T} \right)^{-1} \quad (7)$$

$$\tau_{IPW} = Y_s(1) - Y_s(0) \quad (8)$$

Sample size and statistical power

Sample size and statistical power

- Choosing an appropriate sample size helps ensure results are statistically meaningful and reliable.
- Statistical power measures the probability of detecting an effect if it truly exists.

Type I error and significance level

- Type I error (α) is the probability of rejecting the null hypothesis when it is true.
- Common significance levels are 5% (0.05), though 1% and 10% are also used.

$$\alpha = \Pr \left(\frac{\hat{\tau}}{\text{se}(\hat{\tau})} > t_{\alpha} \mid \tau = 0 \right) \quad (9)$$

Statistical power

- Statistical power (κ) is the probability of obtaining a statistically significant test statistic given that the true parameter value is τ_0 .
- Formally:

$$\kappa = \Pr \left(\frac{\hat{\tau}}{\text{se}(\hat{\tau})} > t_{\alpha} \mid \tau = \tau_0 \right) \quad (10)$$

$$\kappa = 1 - \Phi \left(t_{\alpha} - \frac{\tau_0}{\text{se}(\hat{\tau})} \right) \quad (11)$$

- To achieve statistical power at least κ :

$$1 - \kappa \leq \Phi \left(t_{\alpha} - \frac{\tau_0}{\text{se}(\hat{\tau})} \right) \quad (12)$$

$$t_{\alpha} - \frac{\tau_0}{\text{se}(\hat{\tau})} \leq -t_{1-\kappa} \quad (13)$$

Statistical power

- Assuming homoskedastic errors and a sample size N :

$$\text{se}(\hat{\tau}) = \sqrt{\frac{1}{P \cdot (1 - P)} \frac{\sigma^2}{N}} \quad (14)$$

- where P is the share of individuals assigned to treatment.

Minimum Detectable Effect (MDE)

- For given values of α , κ , P , and N , the estimator will be statistically significant if:

$$\tau_0 \geq (t_\alpha + t_{1-\kappa}) \cdot \text{se}(\hat{\tau}) \quad (15)$$

$$\text{MDE} = (t_\alpha + t_{1-\kappa}) \cdot \text{se}(\hat{\tau}) \quad (16)$$

- A larger sample size (N) reduces the standard error of the estimator. This, in turn, reduces the MDE and increases statistical power.