Adaptive Experimental Designs

Julian Gerez

Columbia University

POLS-GU4726: Topics in Political Methodology Spring 2021

Agenda

1 Review of Randomized Experiments















Overview of Experiments

- A randomized experiment is the gold standard for making causal inferences
- Randomization of the treatment will make the treatment and control groups similar on average with respect to *observed* and *unobserved* covariates
- Advantage 1: Identification is justified by design of experiments

 \star We control the treatment assignment mechanism

- * We do not need to make "ignorability"-type assumptions
- Advantage 2: Estimation is simple
 - * Difference-in-means (DiM) or some weighted averages of DiM

• Advantage 3: Inference is simple

 \star We can again use the known treatment assignment mechanism as a "reason basis for inference"

Basic Setup for Randomized Experiment

- Units: $i \in \{1, \dots, N\}$
- Treatment: $T_i \in \{0, 1\}$, randomly assigned
- Potential outcomes: $Y_i(0)$ and $Y_i(1)$
- Observed outcome: $Y_i = T_i Y_i(1) + (1 T_i) Y_i(0)$ (consistency)
- Treatment Assignment Mechanism

(1) Complete randomization: Exactly N_1 units are treated

(2) Bernoulli (simple) randomization: Each unit is independently assigned to treatment with probability \boldsymbol{p}

• Randomization (complete or simple) implies

 $\{Y_i(1), Y_i(0)\} \perp T_i$

Identification of Average Treatment Effect

• Causal Estimand: Average Treatment Effect

 $\tau_{ATE} \equiv \mathbb{E}\{Y_i(1) - Y_i(0)\}$

- Not directly estimable as we don't observe $Y_i(1) Y_i(0)$ for each unit
- Identification Question:
 Can we write down τ_{ATE} only with observed data (Y_i, T_i)?

$$\mathbb{E}\{Y_i(1) - Y_i(0)\}$$

$$= \mathbb{E}\{Y_i(1)\} - \mathbb{E}\{Y_i(0)\} \quad (\because \text{ Linearity of } \mathbb{E})$$

$$= \mathbb{E}\{Y_i(1) \mid T_i = 1\} - \mathbb{E}\{Y_i(0) \mid T_i = 0\} \quad (\because \text{ Randomization of } T)$$

$$= \mathbb{E}(Y_i \mid T_i = 1) - \mathbb{E}(Y_i \mid T_i = 0) \quad (\because \text{ Consistency of PO})$$

• Estimation Question: Can we estimate $\mathbb{E}(Y_i \mid T_i = 1) - \mathbb{E}(Y_i \mid T_i = 0)$?

$$\frac{1}{N_1} \sum_{i=1}^{N} T_i Y_i - \frac{1}{N_0} \sum_{i=1}^{N} (1 - T_i) Y_i$$





3 Adaptive Experiments



Considering Many Different Treatments

- What if we want to learn about the effects of many separate treatments?
- Multi-arm trials: instead of testing 1 treatment, test K treatments
- Units: $i \in \{1, \ldots, N\}$
- Treatments: $T_i \in \{1, 2, \dots, K\}$, randomly assigned
- Potential outcomes: $Y_i(1), Y_i(2), \ldots, Y_i(K)$
- Observed outcome: $Y_i = \sum_{k=1}^{K} \mathbb{1}[T_i = k] Y_i(k)$

Identification of Average Treatment Effects

• Here, we can make pairwise comparisons between arms:

$$\tau_{ATE} \equiv \mathbb{E}\{Y_i(k) - Y_i(k')\}$$

- Where k and k' are two separate treatment groups
- Not directly estimable as we don't observe $Y_i(k) Y_i(k')$ for each unit
- Identification: can we write down τ_{ATE} only with observed data (Y_i, T_i) ?

$$\mathbb{E}\{Y_i(k) - Y_i(k')\}$$

$$= \mathbb{E}\{Y_i(k)\} - \mathbb{E}\{Y_i(k')\} \quad (\because \text{ Linearity of } \mathbb{E})$$

$$= \mathbb{E}\{Y_i(k) \mid T_i = k\} - \mathbb{E}\{Y_i(k) \mid T_i = k'\} \quad (\because \text{ Randomization of } T)$$

$$= \mathbb{E}(Y_i \mid T_i = k) - \mathbb{E}(Y_i \mid T_i = k') \quad (\because \text{ Consistency of PO})$$

• Estimation: can we estimate $\mathbb{E}(Y_i \mid T_i = k) - \mathbb{E}(Y_i \mid T_i = k')$?

$$\frac{1}{N_k} \sum_{i=1}^{N_k} \mathbb{1}[T_i = k] Y_i(k) - \frac{1}{N_{k'}} \sum_{i=1}^{N_{k'}} \mathbb{1}[T_i = k'] Y_i(k')$$

Notice that this is an analogous generalization of the simple experiment!

Motivating Adaptive Designs

If we're interested in many treatments, we can create several treatment groups:

- Identification is justified by design (randomization)
- Estimation and inference are simple for pair-wise comparisons across arms

So why not just divide our N subjects into K treatment groups?

- Trade-off between exploration and exploitation
- Illustration: the multi-armed bandit



Motivating Adaptive Designs: Examples

- Get-out-the-vote: normatively, may want to increase turnout
 - Test many interventions, but maximize subjects that get best arm
 - Ethics: maybe this will help you convince an implementing partner to do an RCT (participant welfare maximized)
- Business: company wants to experiment with ads, but maximize click-rate
 - Useful to discard bad ads over time
- Medical: test different drugs, ensure most participants healed as possible
 - Once you know a drug doesn't work, stop giving it to participants
- Other social science applications? Open questions!

Can narrow number of treatments with theory, but what if we miss something?

- Let the data decide which arms are most promising
- Key idea: adapt our treatment assignments based on the outcomes over time
- Increase the speed and efficiency of finding best arm









Adaptive Designs: General Workflow



Setup: Adaptive Designs

- Units: $i \in \{1, \ldots, N\}$
- Treatments: $T_i \in \{1, 2, \dots, K\}$, randomly assigned
- Potential outcomes: $Y_i(1), Y_i(2), \ldots, Y_i(K)$
- Observed outcome: $Y_i = \sum_{k=1}^{K} \mathbb{1}[T_i = k] Y_i(k)$
- New! Time periods (waves): w
- For ease of exposition, consider binary outcomes $Y \in \{0,1\}$

K arms have success rates $\theta_1, \theta_2, \ldots, \theta_k$, where $\theta_k = \Pr(Y = 1 | T = k)$

- Likelihood: $Y_1, \ldots, Y_{n,k} \mid \theta_k \stackrel{\text{i.i.d.}}{\sim} \text{Bern}_{n,k}(\theta_k)$, where n, k is the number of units in arm k
- Can set priors to the success rate of each arm $\theta_k \sim \text{Beta}(\alpha_0^k, \beta_0^k)$
- Let's set agnostic priors: for all k, $\theta_k \sim \text{Beta}(1,1)$

Recalling the Beta Distribution



Assigning Treatment Probabilities

In each period w, treatment assigned by the probability of each arm being best

- In the first period (wave 0), for agnostic priors $\theta_k \sim \text{Beta}(1, 1)$, this means equal probability of assignment across all arms k
- For later periods, we can calculate based on observed outcomes
 - For example, which arm had the maximum outcome for each draw of the posterior distribution divided by total number of draws
 - Depends on what you want to maximize, so define carefully
 - For more, see Kasy, Maximilian, and Anja Sautmann. "Adaptive treatment assignment in experiments for policy choice." *Econometrica* 89.1 (2021): 113-132.
- At the end of each period, the posterior is updated
- ullet \implies Probability that each arm is best is recalculated

Updating on Best Treatment

Each period w, treatment is assigned, observations recorded for each arm k

- Cumulative assignment to arm k: $n_{k,t}$
- Responses under k in wave w

Distribution of each θ_k given data $Y_k^{\{n_{k,t}\}}$ is then:



- Beta distribution is conjugate for binomial likelihoods
- So posteriors follow beta distribution—we can just count!
- Posterior distribution: $f(\theta_k | Y_k^{\{n_{k,t}\}}) \sim \text{Beta}(\alpha_w^k, \beta_w^k)$
- α_w^k is α_0^k + total successes (1) observed from that arm
- β_w^k is β_0^k + total failures (0) observed from that arm

Practical Considerations

- Weight by inverse treatment assignment probabilities for estimation
- Large burn-in period could be useful to ensure results are not driven by noise
- Stop after a fixed number of periods or number of subjects have participated
- If interested in control comparison, useful to do a control-augmented design
 - e.g., Offer-Westort, Molly, Alexander Coppock, and Donald P. Green. "Adaptive experimental design: Prospects and applications in political science." *American Journal of Political Science* 65.4 (2021): 826-844.
 - At each step, parity between best performing condition and control
- $\bullet\,$ In practice, fit a model using MCMC and generate the posterior in that way
- Or use quasi-Bayesian approach: generate "posteriors" based on estimates





3 Adaptive Experiments

