Causal Inference from Competing Treatments

Ana-Andreea Stoica[1], *Vivian Y. Nastl*[1,2], *Moritz Hardt*[1]

[1]Max Planck Institute for Intelligent Systems, Tübingen and Tübingen AI Center

[2]Max Planck ETH Center for Learning Systems

ICML International Conference On Machine Learning

arXiv:2406.03422

Problem Formulation

Problem: Estimate a treatment effect with multiple experimenters competing for the same population of subjects. When multiple experimenters compete, the subjects see multiple treatments in some order, at different ranks (e.g. a user seeing multiple ads on a online platform).

What would the effect have been had the treatment been *first*?

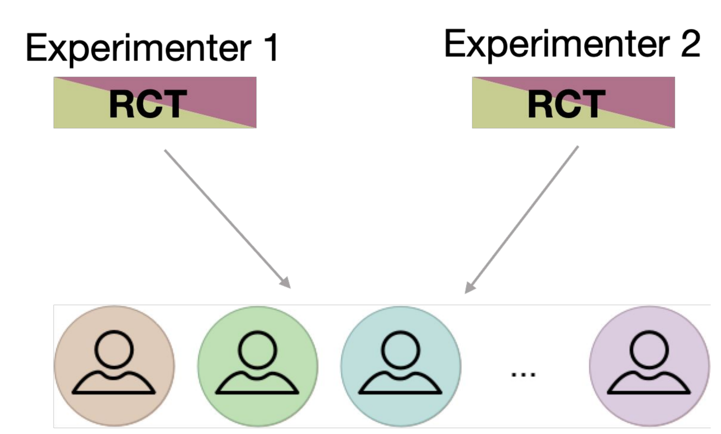

subject population

Estimand: $\tau = \mathbb{E}\left[Y_i(1) - Y_i(0)\,|\,\text{do}(\text{rank}_i) = 1\right]$

outcome of user i

- We assume lower ranks have lower effects $\tau_r = \alpha_r \cdot \tau$, for known discount factors $\alpha_r \in (0,1]$
- *Intuition*: samples at rank 1 are most valuable, but samples at lower ranks may still be useful
- We model competition over subjects through an auction, where experimenters can bid on subjects

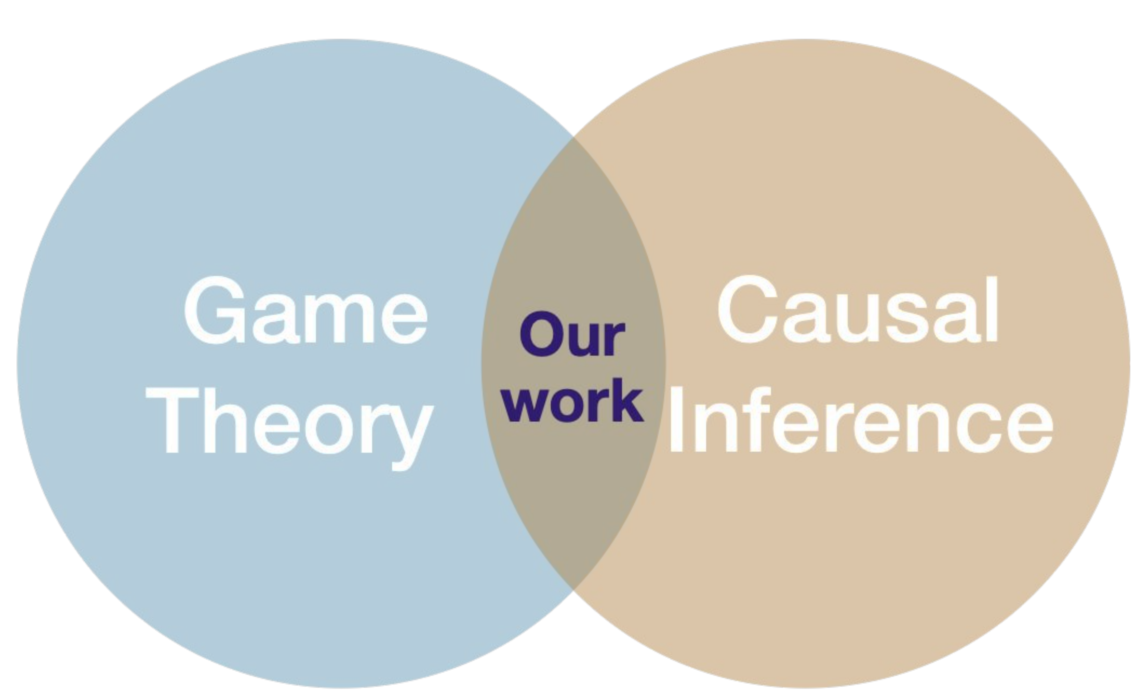

Objectives and Results

Ideal goal: minimize the MSE

$$\inf_{\hat{\tau}, bids} \mathbb{E}_{Z \sim X(bids)}\left[(\hat{\tau}(Z) - \tau)^2\right]$$

Hard!

Theorem: $\forall \hat{\tau}\ \exists M$ instance of the potential outcome model instances s.t.

$$\mathbb{E}_M\left[(\hat{\tau} - \tau)^2\right] \geq \min\left(\frac{\sigma^2}{\text{ct} \cdot S}, 1\right)$$

Realistic goal: estimation error objective

Define the sample value by discounting the samples at each rank appropriately: $S = \sum_r n_r \cdot \alpha_r^2$

$$\mathbb{E}_{\overline{r} \sim \mathcal{A}(bids)} \min\left(\mathbb{E}_{X(\overline{r})}\left[1/S\right], 1\right)$$

Still hard!

- Auction design + estimation error obj $\Rightarrow \mathcal{G}_{\text{MSE}}$ game

Theorem [main result]: Any NE for the sample value game $\mathcal{G}_{\text{SV}}$ (k experimenters, n subjects) is an (ϵ, η)-approximate NE for the error estimation game $\mathcal{G}_{\text{MSE}}$, with

$$\epsilon = O\left(k/\sqrt{B}\right) \text{ and } \eta = O\left(exp\left(-B/k^2\right)\right)$$

where $B = \min_a B^{(a)}$ and the max bid per experimenter is bounded.

**in particular, for $B^{(a)} = \omega(k^2)$ this is a (o(1),o(1))-approximation*

Sample value objective provides a good approx:

$$\mathbb{E}_{\overline{r} \sim \mathcal{A}(bids)} \mathbb{E}_{X(\overline{r})}\left[S\right]$$

Tractable!

- Auction design + sample value obj $\Rightarrow \mathcal{G}_{\text{SV}}$ game

Theorem: The $\mathcal{G}_{\text{SV}}$ game with k admins and n subjects has a pure NE, characterized by:

- Experimenters bid all their budget $B^{(a)}$
- Experimenters bid (almost) equally on all subjects
- Experimenters avoid competing on the same subjects

Modeling Choices

Part 1: Experimenters bid on subject slots, with a budget, and get assigned ranks based on an allocation function (auction design).

A rank allocation implies a sample profile for an experiment (# of datapoints at each rank).

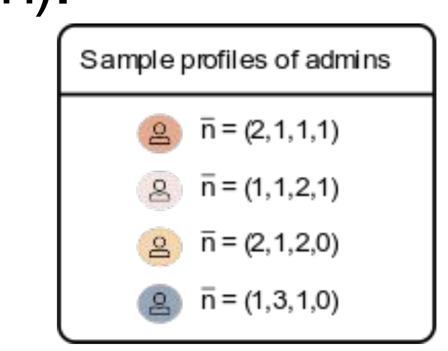

Part 2: Experimenters use the sample profiles to run an RCT on each subject and estimate the treatment effect (defined as a counterfactual estimand, had the treatment been shown at rank 1 for all subjects).

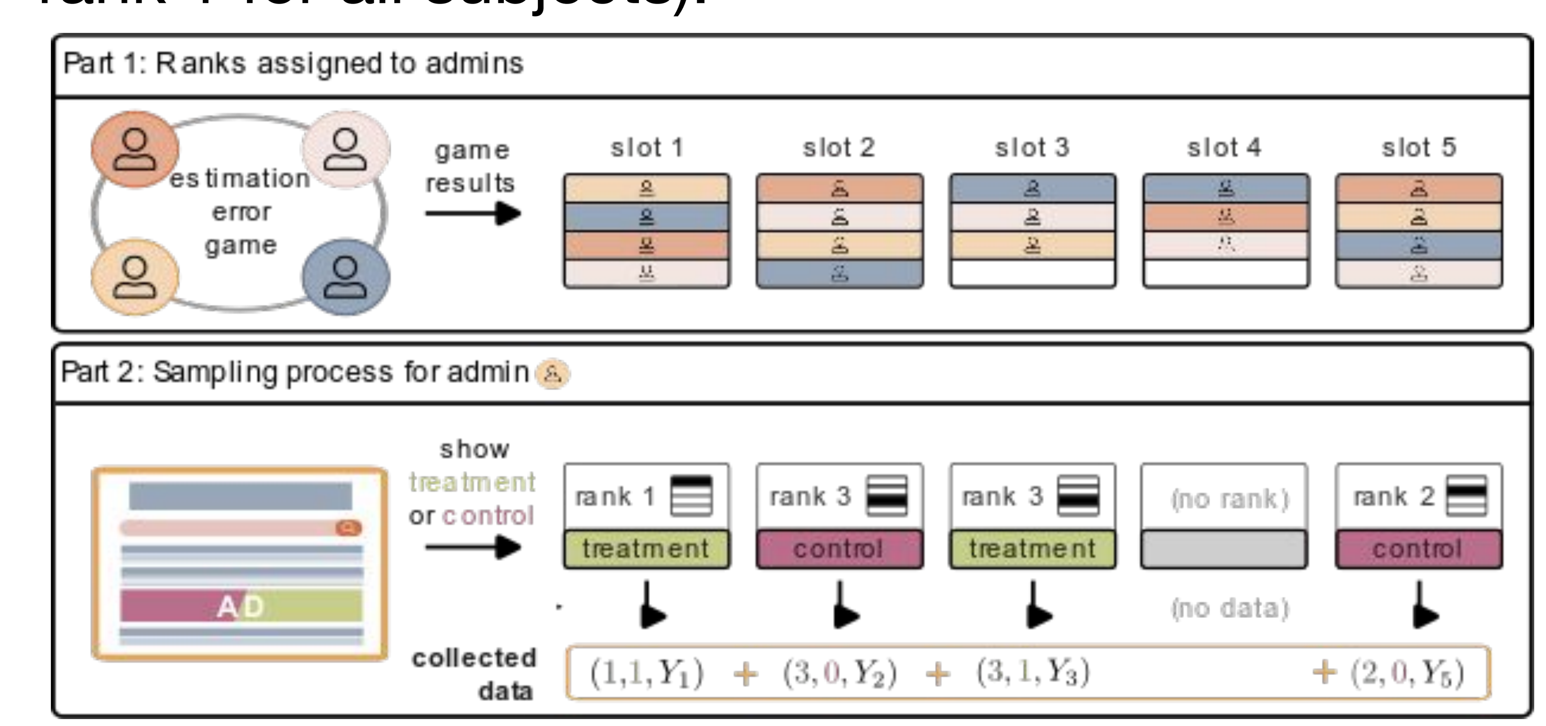

Summary:

- We used a probabilistic allocation function derived from auction theory; future work: first-/second-price auctions.
- We assumed known discount factors; if they have to be estimated from the same data, we can show it's not worth doing it (e.g. we get lower MSE with just using the data conditional at rank 1).
- Future work can generalize the problem from treatment estimation to statistical inference questions.