Incentivizing Participation in Randomized Clinical Trials (RCTs)

Yingkai Li (Yale)   Aleksandrs Slivkins (Microsoft Research NYC)

Incentivize participation in RCTs

Use **information asymmetry** rather than money.

RCT as a bandit algorithm with T rounds and n arms:

- in each round t, a new patient arrives
- algorithm chooses distribution p_t over arms, draws an arm from p_t, observes outcome $\omega_t \in \Omega$
- patients do not observe previous arms & outcomes, RCT commits to the algorithm (both are **realistic!**)

Goal 1: incentivize each patient to participate (**new!**)
Goal 2: counterfactual estimates (**standard**)

Model: patients and their incentives

Heterogenous agents with subjective utilities:

- public type $\theta_{\texttt{pub}}$: medical history, affects outcomes;
- private type $\theta_{\texttt{pri}}$: subjective utilities over the outcomes, strategically reported by the patient

Patients' beliefs (*not reality!*):

- type $(\theta_{\texttt{pub}}, \theta_{\texttt{pri}})$ is drawn from fixed distribution $F_{\texttt{type}}$;
- outcome is drawn from fixed distribution $\psi^*(\text{arm}, \theta_{\texttt{pub}})$;
- ψ^* is **state**, initially drawn from known prior $\mathcal{P}$.

The algorithm must incentivize each patient to:

- report $\theta_{\texttt{pri}}$ (Bayesian Incentive-Compatibility, BIC)
- participate (Bayesian Individual Rationality, BIR)

Model: statistics (standard for RCTs)

- initially, an **adversary** chooses, for each patient t: public & private types, outcome $\omega_{a,t} \in \Omega$ for each arm a
- each outcome $\omega \in \Omega$ assigned a known $\texttt{score}(\omega) \in [0,1]$
- **average score** over a subset S of rounds:
$$f_{\texttt{adv}}(\text{arm } a) := \tfrac{1}{|S|} \textstyle\sum_{t \in S} \texttt{score}(\omega_{a,t}).$$
- initially, algorithm chooses $S \subset T$; after round T, it computes estimators $\widehat{f}(\cdot)$ for $f_{\texttt{adv}}(\cdot)$, for all arms.
- Objective: maximal mean-squared error
$$\texttt{MaxMSE} = \max_{\text{all arms } a, \text{ all adversaries}} \mathbb{E}[\ (\ \widehat{f}(a) - f_{\texttt{adv}}(a)\ )^2\ ],$$
where $\mathbb{E}[\,\cdot\,]$ is over the randomness in the algorithm.
- Standard IPS estimators suffice for our algorithm. Then:
$$\texttt{MaxMSE} \le |S|^{-2} \cdot \max_{\text{arms } a} \textstyle\sum_{t \in S} 1/p_t(a).$$
- Incentives aside, **uniform randomization** is best.

Benchmark: what if state ψ^* is known?

- focus on a **single round** of the algorithm
- policy class Π: input patient's type $(\theta_{\texttt{pub}}, \theta_{\texttt{pri}})$ and state ψ, output distribution over arms, satisfy BIC and BIR.
- **Informally:** policy $\pi \in \Pi$ and state ψ determine $\texttt{MaxMSE}$;
$$\texttt{bench}(\mathcal{P}, F_{\texttt{type}}) := \inf_{\pi \in \Pi} \sup_{\psi \in \texttt{support}(\mathcal{P})} \texttt{MaxMSE}(\pi, \psi).$$
- homogeneous patients $\Rightarrow$ dependence on $F_{\texttt{type}}$ is not needed.

Our focus: "two-stage algorithms"

- initially, algorithm chooses $T_0 < T/2$.
- warm-up stage: first T_0 rounds, collects "warm-up data",
- main stage: the remaining rounds $S = [T] \setminus [T_0]$.
- main-stage data ignored in the main stage; "warm-up data" ignored by the estimators $\widehat{f}(\cdot)$.
- **Realistic!** all main-stage patients can be treated simultaneously; adapting to data is challenging in practice

Our results (homogeneous patients)

Our algorithm: if "warm-up data" contains $\ge N_{\mathcal{P}}$ samples from each arm, with $N_{\mathcal{P}}$ determined by prior $\mathcal{P}$ (but constant in T)
$$\texttt{MaxMSE} \le 2 \cdot \texttt{bench}(\mathcal{P})/(T - T_0)$$

- "main stage" approximates $\texttt{bench}(\mathcal{P})$ given "warm-up data"
- "warm-up data" collected via techniques from "incentivized exploration", with T_0 determined by $\mathcal{P}$ (and constant in T)

Impossibility: for any BIR mechanism there is an adversary s.t.
$$\texttt{MaxMSE} \ge \text{constant} \cdot \texttt{bench}(\mathcal{P})/T.$$

Heterogeneous patients: similar results, more complex to state.

Working paper at arXiv:2202.06191

Comparison to the work on "incentivized exploration": ask us!