A/B/n Testing with Control in the Presence of Subpopulations

Yoan Russac[1], Christina Katsimerou[2], Dennis Bohle[2], Olivier Cappé[1], Aurélien Garivier[3] & Wouter M. Koolen[4]

[1] CNRS, [2] Booking, [3] ENS Lyon, [4] CWI

October 2021

Roadmap

A/B/n Testing

The pure exploration setting in which $K \geq 2$ options are (blindly) proposed to users, in order to identify the subset of competitive options

In the *fixed confidence* setting: for a risk level δ the probability of returning a wrong answer must be $\leq \delta$

We aim to optimize adaptively both

- the allocation of options to users
- the stopping time of the A/B/n experiment

A/B/n testing with Control

The baseline is given by an additional control arm (with index $k = 0$), whose value is also unknown.

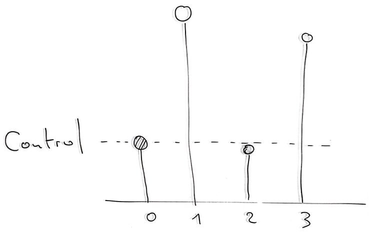

$\neq$ best arm identification
$\neq$ thresholding bandit (identify the arms above a known level)

...in the Presence of Subpopulations

The user at time t belongs to an subpopulation $I_t \in \{1, \ldots, J\}$.

$\rightarrow \alpha_i$ natural proportion of subpopulation i

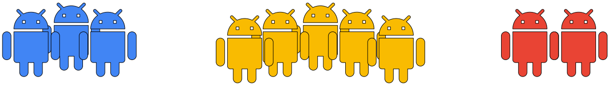

$\alpha_1 = 3/10$ $\alpha_2 = 5/10$ $\alpha_3 = 2/10$

$\rightarrow \beta = (\beta_i)_{i=1,\ldots,J}$ are known user-defined population weights defining the value of an arm

$$\mu_a = \sum_{i=1}^{J} \beta_i \mu_{a,i} \ .$$

The ABC-S Problem

Objective: Identifying all the set of Arms that are Better than the Control in the presence of Subpopulations (ABC-S)

Formally, identification of

$$\mathcal{S}_\beta(\boldsymbol{\mu}) = \left\{ a \in \{1, \ldots, K\} \text{ s.t } \sum_{i=1}^{J} \beta_i \mu_{a,i} > \sum_{i=1}^{J} \beta_i \mu_{0,i} \right\} .$$

Use Cases

Use case 1: Data with Known Seasonality Effects

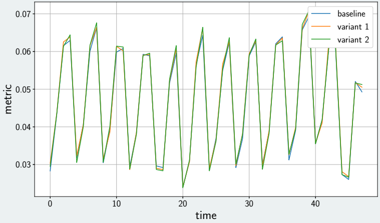

Booking webpage data: click-through rate per 6 hours for 12 days.

Value of the variant a

$$\mu_a = \sum_{i=1}^{4} \beta_i \mu_{a,i}$$

Use Cases

Use case 2: Marketing

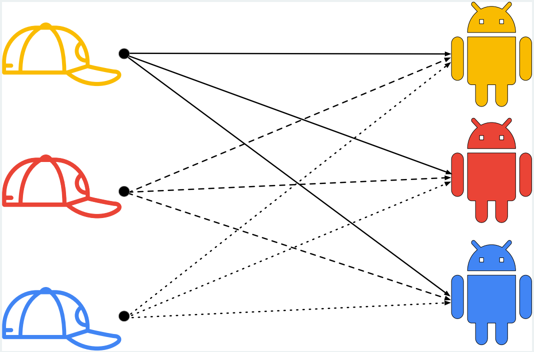

- Different subpopulations behave differently for the same product.
- The seller can value differently the subpopulations.

Value of the blue cap

$$\mu_a = \beta_1 \mu_{a,1} + \beta_2 \mu_{a,2} + \beta_3 \mu_{a,3}$$

Modes of Interaction

Different modes of interaction with the subpopulations.

(a) *Oblivious*

1. Pick A_t
2. Don't see $I_t \sim \alpha$
3. See $X_t \sim \nu_{A_t,I_t}$

(b) *Agnostic*

1. Pick A_t
2. See $I_t \sim \alpha$
3. See $X_t \sim \nu_{A_t,I_t}$

(c) *Proportional*

1. See $I_t \sim \alpha$
2. Pick A_t
3. See $X_t \sim \nu_{A_t,I_t}$

(d) *Active*

1. Pick A_t and I_t
2. See $X_t \sim \nu_{A_t,I_t}$

Roadmap

Theorem 1

For any strategy, the expected number of rounds for the ABC-S problem satisfies

$$\liminf_{\delta \to 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_\delta]}{\ln(1/\delta)} \geq T^\star(\boldsymbol{\mu}) \quad (1)$$

where

$$T^\star(\boldsymbol{\mu})^{-1} = \sup_{\mathbf{w} \in \mathcal{C}} \min_{b \neq 0} \inf_{\lambda \in \mathcal{L} : \lambda_0 < \lambda_b} \sum_{a \in \{0,b\}} \sum_{i=1}^{J} w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) \quad (2)$$

- A min-max optimization problem (as usual) but with specific constraints through $\mathcal{C}$ in each mode
- Optimal weights $\mathbf{w}^* =$ target relative frequencies of draws

Link between characteristic problems.

An interesting link

When $\alpha = \beta$, for all $\boldsymbol{\mu}$

$$T^\star_{\text{active}}(\boldsymbol{\mu}) \leq T^\star_{\text{proportional}}(\boldsymbol{\mu}) \leq T^\star_{\text{agnostic}}(\boldsymbol{\mu}) \leq T^\star_{\text{oblivious}}(\boldsymbol{\mu})$$

Roadmap

Track-and-Stop (T&S) Algorithm

For $t \geq 1$:

- Sampling rule: given the current estimates
  1. estimate the target weights $\mathbf{w}_t$
  2. pick arm $\begin{cases} \text{active:} & (A_t, I_t) \in \text{argmax}_{a,i} N_{a,i}(t-1) - t\mathbf{w}_t(a,i) \\ \text{proportional:} & A_t \in \text{argmax}_a N_{a,I_t}(t-1) - t\alpha_{I_t}\mathbf{w}_t(a|I_t) \\ \text{agnostic:} & A_t \in \text{argmax}_a N_a(t-1) - t\mathbf{w}_t(a) \end{cases}$
- Recommendation: $\mathcal{S}(\hat{\boldsymbol{\mu}}_t) = \{a \in \{1, \ldots, K\} : \hat{\mu}_a(t) > \hat{\mu}_0(t)\}$ at confidence level $\hat{\delta}_t = \min\{\delta \in (0,1) | \Lambda(t) \geq \beta(t, \delta)\}$

Theorem 2

For $\beta(t,\delta) = 6J \ln\ln t + \ln\frac{1}{\delta} + K + 2J \cdot O(\ln\ln\frac{1}{\delta})$, *T&S is safely calibrated:*

$$\forall \boldsymbol{\mu} \in \mathcal{L},\ \forall \delta \in (0,1), \quad \mathbb{P}_{\boldsymbol{\mu}}\left(\exists t \geq 1 : \hat{\mathcal{S}}_t \neq \mathcal{S}(\boldsymbol{\mu}) \ \cap\ \hat{\delta}_t \leq \delta\right) \leq \delta\ . \tag{3}$$

In practice: $\ln((1+\ln t)/\delta)$ works well.

Roadmap

Numerical Results

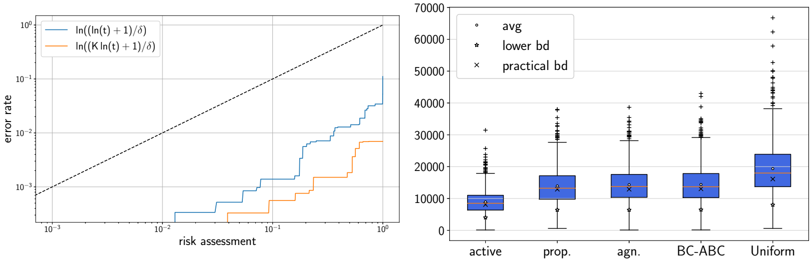

Figure: (Left) Risk assessment calibration on a log-log scale. (Right) Stopping time boxplot for $\boldsymbol{\mu} = [0.1\ 0.4\ 0.3; 0.2\ 0.5\ 0.2; 0.5\ 0.1\ 0.1]$ when $\beta = [1/3, 1/3, 1/3], \alpha = [0.4, 0.5, 0.1]$ with Bernoulli distributions.

Real Data Experiment

Booking webpage data

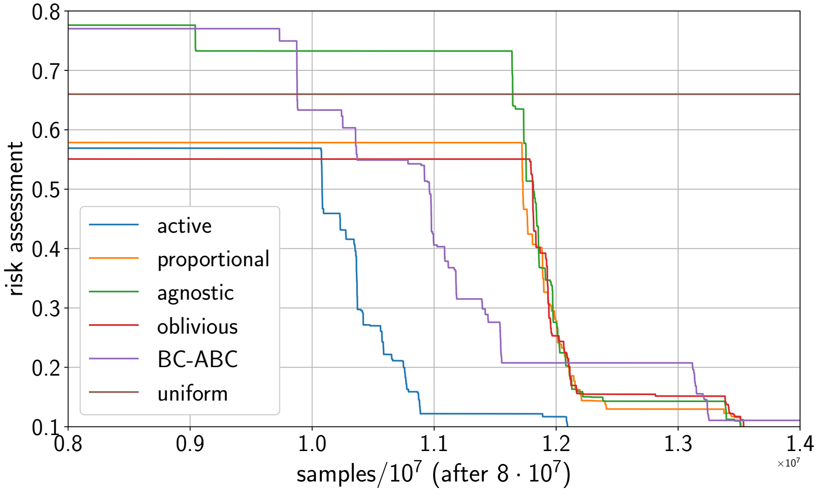

Thank you !