Context-dependent response-adaptive randomization for continuous endpoints and applications

Luca Rondano, Gianmarco Caruso, Pavel Mozgunov

09/06/2025

Response-adaptive randomization (RAR)

Patients enter the trial and are treated sequentially

Randomization probabilities **change during the trial** according to

- Accumulating outcomes
- Previous allocations

Compared to equal randomization, RAR aims to

- Increase patient benefit
- Better resource allocation (possibly resulting in more power)

An application to multi-arm phase II clinical trials with continuous endpoint

Given K arms, the goal is often the identification of the **best one**

The proposed RAR methodology adaptively changes the randomization probabilities to:

- Increase the allocations to the best arm
- Increase exploration of the most promising arms

based on a **context-dependent information measure** which gives a greater weight to those treatment arms which have characteristics close to a pre-specified clinical target

Bayesian inference framework

We consider indipendent normally distributed treatment responses

$X_{i,j} \sim N(\mu_{j,}\ \sigma_j^2);\ j = 1, \dots, K$ arms; $i = 1, \dots, n_j$ patients per arm

Joint prior distribution

$(\mu_{j,}\sigma_j^2) \sim NIG(\mu_0, \nu, \alpha_0, \beta_0)$

Equivalently, $\mu_j |\ \sigma_j^2 \sim N(\mu_{0,}\ \sigma_j^2/\nu)$ and $\sigma_j^2 \sim IG(\alpha_0, \beta_0)$

Joint posterior distribution after n_j responses

$$(\mu_{j,}\sigma_j^2) \sim NIG\left(\frac{n_j\overline{x_{n_j}} + \nu\mu_0}{n_j + \nu}, n_j + \nu,\ \alpha_0 + \frac{n_j}{2}, \beta_0 + \frac{n_j - 1}{2}\ \bar{s}^2 + \frac{n_j\nu}{n_j + \nu}\frac{\left(\mu_0 - \overline{x_{n_j}}\right)^2}{2}\right)$$

Context-dependent RAR

Shannon entropy

$$h(p_j) = -\int p_j(\mu_j, \sigma_j^2) \log p_j(\mu_j, \sigma_j^2)\, d\mu_j d\sigma_j^2$$

Information needed to estimate the parameters

Weighted Shannon entropy

$$h^{\phi}(p_j) = -\int \phi(\mu_j, \sigma_j^2) p_j(\mu_j, \sigma_j^2) \log p_j(\mu_j, \sigma_j^2)\, d\mu_j\, d\sigma_j^2$$

Weights our interest on specific values of the parametric space

Information gain

$$\Delta(p_j) = h(\mu_j, \sigma_j^2) - h^{\phi}(\mu_j, \sigma_j^2)$$

Additional information that is required **when considering the context-dependent** instead of the traditional estimation problem

$\boxed{p_j(\mu_j, \sigma_j^2) \text{ is the posterior density of } (\mu_j, \sigma_j^2)}$

Choice of the weight function ϕ

An investigator is typically interested in finding the arm with treatment effect closest to some desirable or optimal target values

Assume that γ and ξ are **pre-specified target values** for the **mean response** and **its variance**, determined based on their clinical relevance

The choice of $\phi\left(\mu_{j,}\sigma_j^2\right)$ should reflect the context of the trial, i.e., it should be higher in a neighbourhood of (γ, ξ)

The **density function of a NIG** with mode = (γ, ξ) naturally does that!
Moreover, we obtain a **closed-form expression** of $\Delta\left(p_j\right)$

Allocation rule

Burn-in phase to have a first estimate of the treatment effects

Then, patients are randomized **sequentially**

Posterior distributions $p_j, j = 1 \dots, K,$ are updated **after each patient response**

Assign patient i to arm j^* with probability:

$$P(a_i = j^*) = \frac{\Delta\left(p_{j^*}\right)}{\sum_{\mathrm{j}=1,\dots,m} \Delta\left(p_j\right)}$$

Exploration VS exploitation

Remember: $\phi(\mu_{j,}\sigma_j^2)$ is proportional to the density function of a NIG with mode = (γ, ξ)

NIG has 4 parameters → We introduce 2 penalization parameters κ and ω

By penalizing the arms who received more patients,
κ tailors our interest in estimating the means
ω tailors our interest in estimating the variances

Information gain plots

sample mean

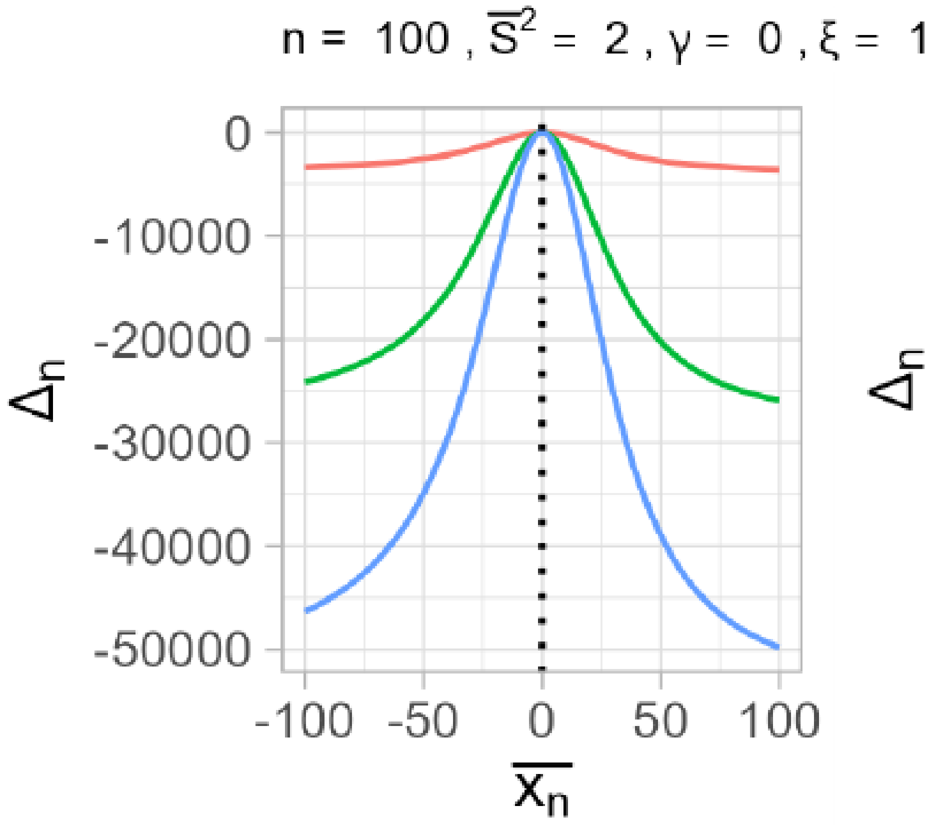

sample variance

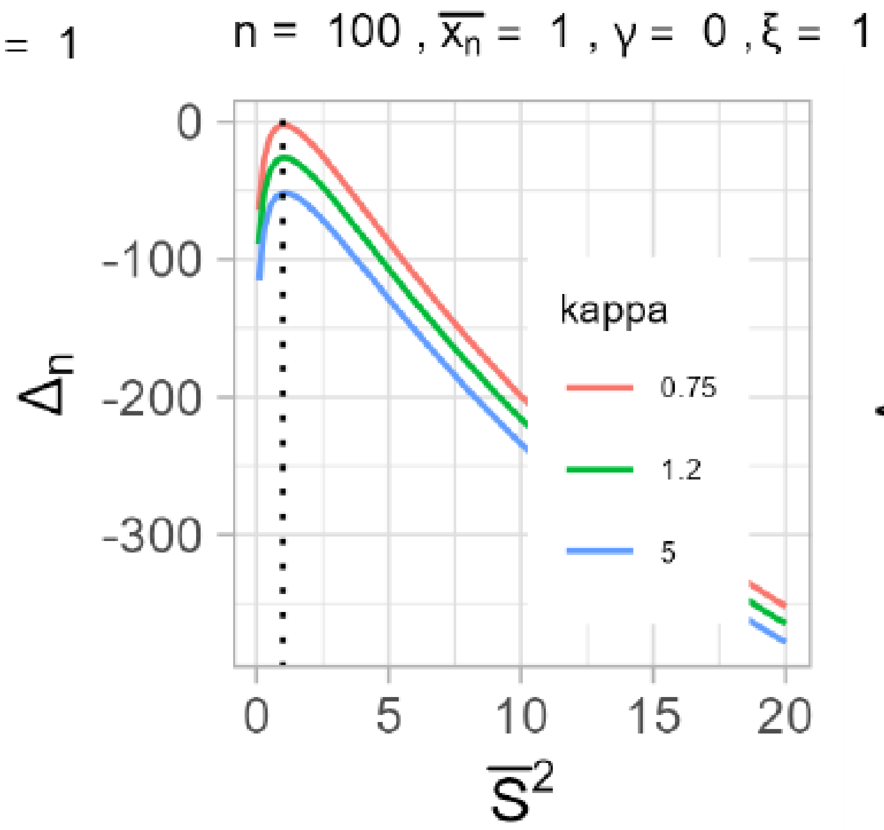

sample size

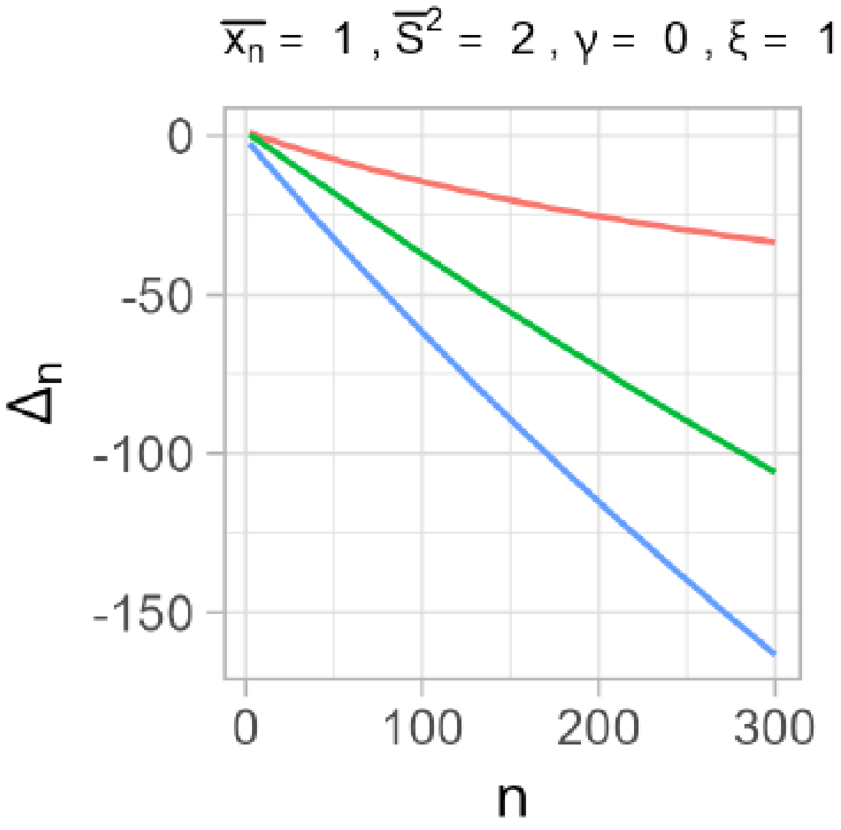

Robust strategy for selecting κ and ω

Consider a set of S plausible alternative scenarios and a grid of values for κ and ω.

1. Select an **operating characteristic of interest** (e.g., expected proportion of patients assigned to the best arm). Its value under a scenario s for fixed κ and ω is $u^s(\kappa, \omega)$

2. Define an objective function $g(u^s(\kappa, \omega))$

3. Compute $u^s(\kappa, \omega)$, for each scenario s and pair of (κ, ω)

4. Find the optimal (κ, ω) that minimizes $\sum_{s=1}^{S} g(u^s(\kappa, \omega))$

Selecting κ and ω based on patient benefit

$$g(u^s(\kappa,\omega)) = \left(u^s(\kappa,\omega) - \max_{\kappa\prime,\,\omega\prime} u^s(\kappa',\omega')\right)^2$$

$u^s(\kappa,\omega)$ is the expected proportion of patients assigned to the best arm

$\kappa = 0.5$ and $\omega = 0.3$ are the optimal values

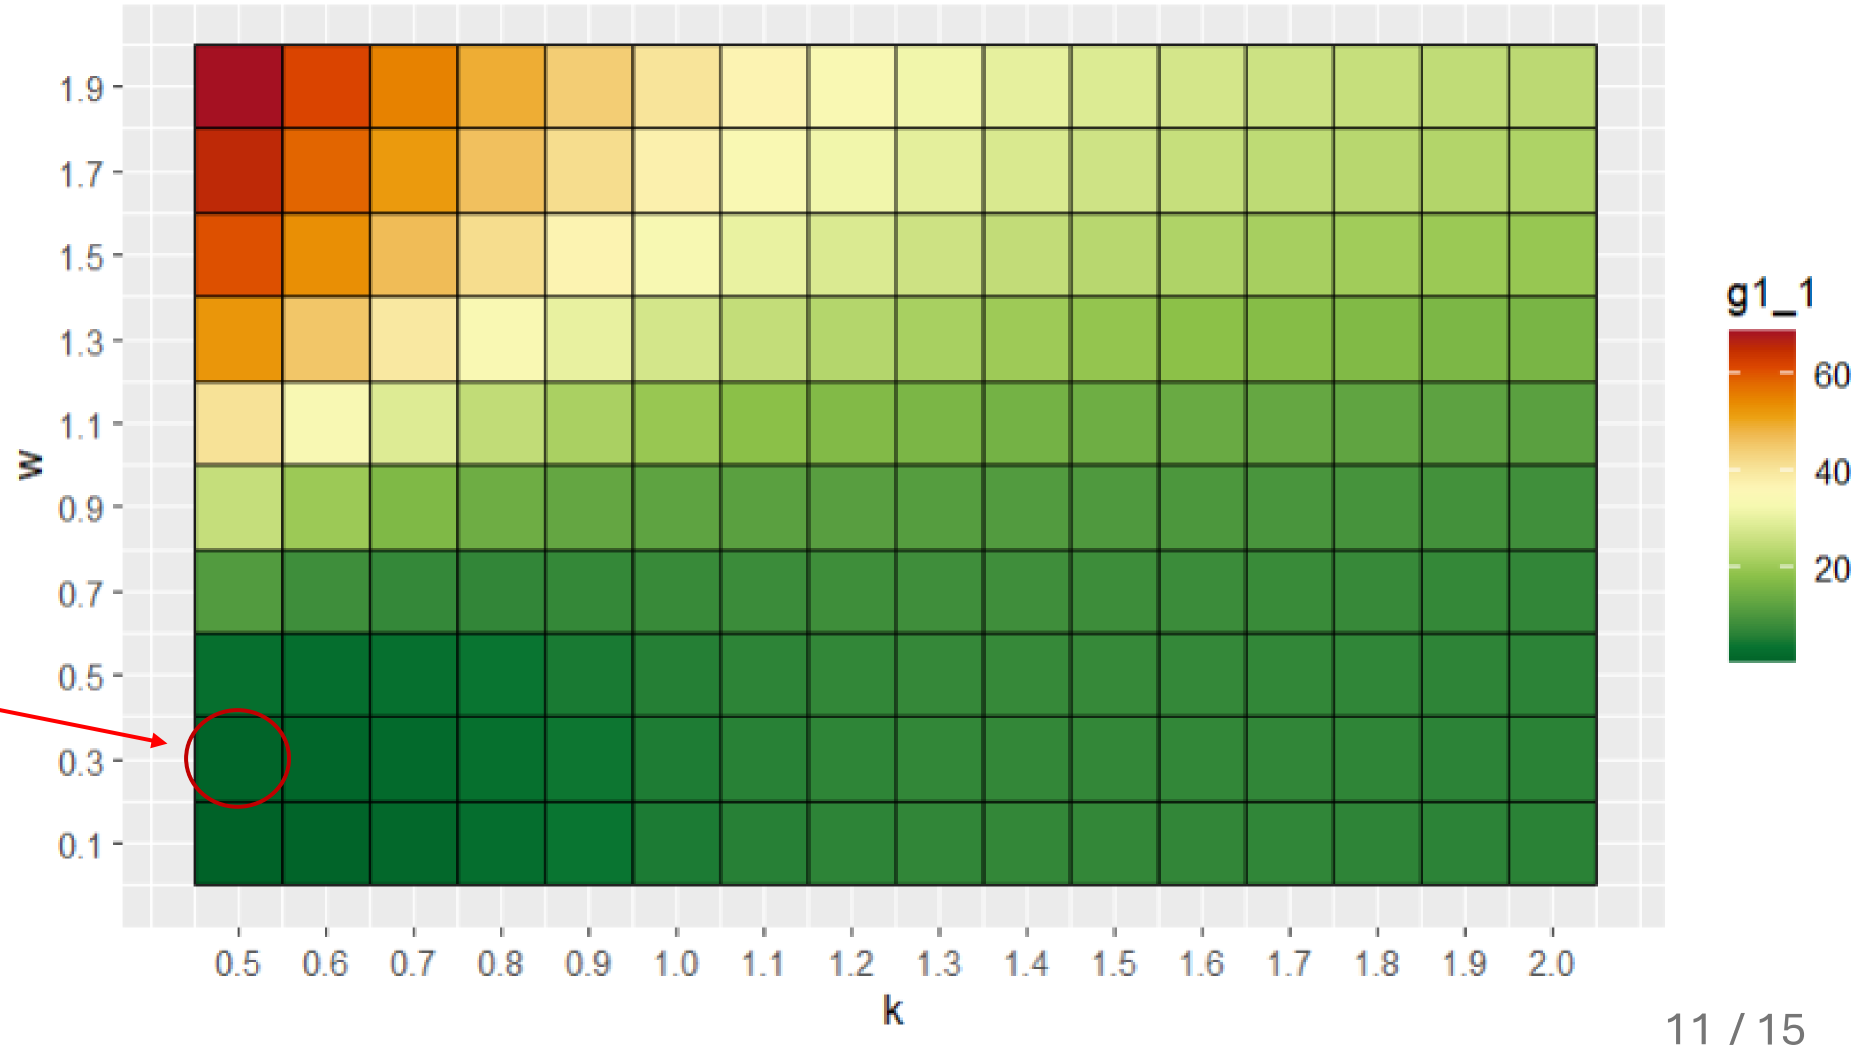

What is the best arm?

$$j^* = \underset{\mathrm{j=1,\ldots,K}}{\operatorname{argmin}} |\overline{x_{n_j}} - \gamma|$$

$$j^* = \underset{\mathrm{j=1,\ldots,K}}{\operatorname{argmax}}\, P\left(X_j \epsilon \left[\gamma - c\xi, \gamma + c\xi\right] \right)$$

All valid definitions!

$$j^* = KL\left(p_j\left(x \mid x_{j,1}, \ldots, x_{j,n_j}\right) \parallel d(x;\gamma,\xi)\right)$$

Randomization does not depend on the definition of best arm, but the evaluation of the operating characteristics does!

Simulation study – performance evaluation

PLACEHOLDER FOR TABLE

Take home messages

- Context-dependent response-adaptive randomization for continuous endpoints
- Designed to accrue more data on the treatments that have a mean effect and variance closest to pre-specified target values
- Gain more information on the best performing arms, while at the same time allocating more patients to the best treatments
- Robust selection of the penalization parameters to achieve optimal values of the desired operating characteristics

Thank you for your attention

References

David S. Robertson, Kim May Lee, Boryana C. López-Kolkovska and Sofía S. Villar (2023). Response-adaptive randomization in clinical trials: from myths to practical considerations. *Statistical Science*, 38(2):185-208.

Sofía S. Villar, Jack Bowden and James Wason (2015). Multi-armed Bandit Models for the Optimal Design of Clinical Trials: Benefits and Challenges. *Statistical Science*, Vol. 30, No. 2, 199–215.

Pavel Mozgunov and Thomas Jaki (2020). An information theoretic approach for selecting arms in clinical trials. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82, Part 5, pp. 1223–1247.

Gianmarco Caruso and Pavel Mozgunov (2024). A response-adaptive multi-arm design for continuous endpoints based on a weighted information measure. *???*.