Principles for Response-Adaptive Randomization

William F. Rosenberger

Professor and Chairman
Department of Statistics
George Mason University
Fairfax, VA
website: statistics.gmu.edu

July 27, 2011

Basic Setup:

- Two treatments, A and B. Randomization sequence is a random vector $\mathbf{T} = (T_1, ..., T_n)$, where $T_j = 1$ or $0(A$ or $B)$. Patient responses are given by $\mathbf{X} = (X_1, ..., X_n)$.
- Numbers allocated to each treatment are random variables N_A and N_B, total sample size fixed at n.
- Allocation rule is $\phi_j = E(T_{j+1}|\mathcal{F}_j)$.
- If $\mathcal{F}_j = \sigma(T_1, ..., T_j)$, then the randomization procedure is *restricted*.
- If $\mathcal{F}_j = \sigma(T_1, ..., T_j, X_1, ..., X_j)$, then the randomization procedure is *response-adaptive*.

Basic Setup:

- The sigma algebra $\mathcal{F}_j$ could also involve covariates. This is an important area for current research, but will not be considered in this talk.
- The theory of randomization is established for $\phi_j \in (0, 1)$, which we call *fully randomized* procedures. This means each subject should be randomized with probability less than 1. This avoids selection bias. Note that this is *not* considered a particularly important criteria in the practice of clinical trials, where blocks are often used with deterministic tails.

Response-Adaptive Randomization

Rosenberger, W. F.

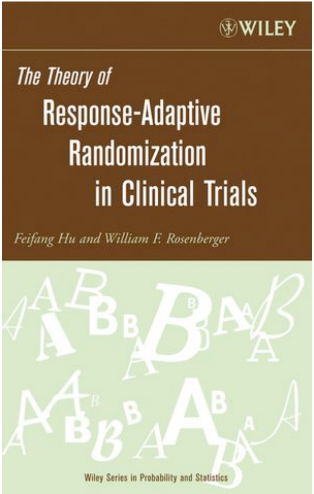

Principle 1: Optimality Criteria

What are our objectives for the experiment? Randomized experiments can have multiple, often competing objectives:

- Maximize power
- Minimize expected failures
- Minimize cost
- Achieve perfect balance

We assume that the underlying responses arise from an exponential family with parameter vector θ. Based on some optimality criteria, we wish to "target" an optimal allocation function $\rho(\theta)$. Our goal is to find a "good" randomization procedure that converges to the target, i.e.,

$$\frac{N_A}{n} \to \rho(\theta),$$

in probability or almost surely.

Example 1: Bernoulli Responses

Assume $\mathbf{X}$ arises from Bernoulli distributions with parameters $\theta = (p_A, p_B)$, where p_A (p_B) is the "success" probability on treatment A (B).

- Balance as a criterion:

$$\rho(\theta) = 1/2$$

- Maximize power as a criterion:

$$\rho(\theta) = \frac{\sqrt{p_A q_A}}{\sqrt{p_A q_A} + \sqrt{p_B q_B}}$$

- For fixed power, minimize the expected treatment failures:

$$\rho(\theta) = \frac{\sqrt{p_A}}{\sqrt{p_A} + \sqrt{p_B}}$$

(Rosenberger, et al., 2001)

Example 1: Bernoulli Responses

- Note that urn models like the randomized play-the winner rule (Wei and Durham, 1978) and the drop the loser rule (Ivanova, 2003) are ad hoc procedures that target

$$\rho(\theta) = \frac{q_B}{q_A + q_B},$$

which is a property of the procedure, not the solution to an optimality problem.

Example 2: Normal Responses

Here we consider **X** to arise from normal distributions: $\theta = (\mu_A, \sigma_A, \mu_B, \sigma_B)$, where μ_A (μ_B) and σ_A (σ_B) are the mean and standard deviation on treatment A (B).

- Maximize power (D-optimality or "Neyman allocation"):

$$\rho(\theta) = \sigma_A/(\sigma_A + \sigma_B)$$

- Minimize the maximum eigenvector of the inverse of Fisher's information (E-optimality):

$$\rho(\theta) = \sigma_A^2/(\sigma_A^2 + \sigma_B^2)$$

  (Baldi Antognini and Giovagnoli, 2005)
- For fixed power, minimize the expected mean response:

$$\rho(\theta) = \max\left\{\frac{1}{2}, \frac{\sigma_A\sqrt{\mu_B}}{(\sigma_A\sqrt{\mu_B} + \sigma_B\sqrt{\mu_A})}\right\},$$

  if $\mu_A \leq \mu_B$ (Zhang and Rosenberger, 2006)

Principle 2: Randomizing with Low Variability

With the exception of $\rho(\theta) = 1/2$, the target allocation depends on unknown parameters. A response-adaptive randomization procedure establishes an allocation function that depends on the current estimates of the parameters, $\phi_j = p(N_A(j), \rho(\hat{\theta}_j))$, where $\hat{\theta}_j$ are estimated after j patient responses. One such obvious function is $\phi_j = \rho(\hat{\theta}_j)$. Melfi and Page (2000) were the first to show that $N_A/n \to \rho(\theta)$ almost surely.

Unfortunately, this function induces a lot of variability. Extra variance in N_A and N_B can negatively impact the optimal properties of $\rho(\theta)$.

Principle 2: Randomizing with Low Variability

In order to reduce variability, a more general family of functions can be found by finding an allocation that minimizes a metric between $N_A(j)$ and $\rho(\hat{\theta}_j)$ in some sense:

$$
\begin{aligned}
p(x,y) &= \frac{y(y/x)^{\gamma}}{y(y/x)^{\gamma}+(1-y)((1-y)/(1-x))^{\gamma}}; \\
p(0,y) &= 1; \\
p(1,y) &= 0,
\end{aligned}
$$

where x$=N_A(j)/j$ and $y=\rho(\hat{\theta}_j)$, and $\gamma \geq 0$ is a tuning parameter. This is Hu and Zhang's (2004) version of the doubly-adaptive biased coin design (DBCD).

Principle 2: Randomizing with Low Variability

- When $\gamma = 0$, this reduces to $\phi_j = \rho(\hat{\theta}_j)$, which has the highest variability
- The least variability is $\gamma = \infty$, but this design is not fully randomized
- Somewhere between 0 and ∞ is a suitable trade-off value
- When $\rho(\theta) = 1/2$, this design is no longer response-adaptive, and reduces to

$$\phi_j = \frac{N_B^\gamma}{N_B^\gamma + N_A^\gamma},$$

which is Smith's generalized biased coin design (GBCD), used for restricted randomization. For Smith's design, $\gamma = 0$ reduces to complete randomization, $\gamma = 1$ is Wei's urn design, $\gamma = 2$ is Atkinson's D_A-optimal design; Smith suggested $\gamma = 5$.

Principle 3: Inference

While the likelihoods look the same for response-adaptive randomization as for non-response-adaptive randomization, they are quite different. In restricted randomization, N_A may be random, but it is ancillary since it is independent of θ. In response-adaptive randomization, it is no longer ancillary. Consider the binomial case with number of successes S_A and S_B on treatments A and B, respectively. Then (N_A, S_A, S_B) are jointly sufficient, and valid test procedures must incorporate their joint distribution, which may be quite complicated.

Hu and Rosenberger (2006) state as a "guiding principle" that, for response-adaptive randomization to be practical, it must be assured that standard inferential tests can be used at the conclusion of the trial.

Bias of the MLE

The MLE is biased following response-adaptive randomization. For example, Coad and Ivanova (2002) show that for a estimating the binomial probability p_A, the bias is given by

$$E(\hat{p}_A - p_A) = p_A(1 - p_A)\frac{\partial}{\partial p_A} E\left(\frac{1}{N_A(n)}\right).$$

Asymptotic Distribution of the MLE

Assume the following regularity conditions for the general K-treatment case:

- The parameter space $\mathbf{\Theta}_j$ is an open subset of $\mathcal{R}^d, d \geq 1$, for $j = 1, ..., K$.
- The distributions $f_1(\cdot, \boldsymbol{\theta}_1), ..., f_K(\cdot, \boldsymbol{\theta}_K)$ follow an exponential family.
- For limiting allocation $\boldsymbol{\rho}(\boldsymbol{\theta}) = (\rho_1(\boldsymbol{\theta}), ..., \rho_K(\boldsymbol{\theta})) \in (0, 1)^K$,

$$\frac{N_j(n)}{n} \rightarrow \rho_j(\boldsymbol{\theta})$$

almost surely for $j = 1, ..., K$.

Asymptotic Distribution of the MLE

Then $\hat{\boldsymbol{\theta}}$ is strongly consistent for $\boldsymbol{\theta}$ and

$$\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \rightarrow N(\mathbf{0}, \mathbf{I}^{-1}(\boldsymbol{\theta})),$$

in distribution, where
$\mathbf{I}(\boldsymbol{\theta}) = \text{diag}\{\rho_1(\boldsymbol{\theta})\mathbf{I}_1(\boldsymbol{\theta}_1), ..., \rho_K(\boldsymbol{\theta})\mathbf{I}_K(\boldsymbol{\theta}_K)\}$ and

$$\mathbf{I}_j(\boldsymbol{\theta}_j) = -E\left(\frac{\partial^2 \log f_j(X_{1j}, \boldsymbol{\theta}_j)}{\partial \boldsymbol{\theta}_j^2}\right)$$

is the Fisher's information for a single observation on treatment $j = 1, ..., K$.

Example

For the doubly-adaptive biased coin design targeting Neyman allocation, we have

$$\sqrt{n}\left(\left[\begin{array}{c}\hat{\mu}_A\\ \hat{\mu}_B\end{array}\right]-\left[\begin{array}{c}\mu_A\\ \mu_B\end{array}\right]\right)\rightarrow N\left(\mathbf{0},\left[\begin{array}{cc}\sigma_A(\sigma_A+\sigma_B) & 0\\ 0 & \sigma_B(\sigma_A+\sigma_B)\end{array}\right]\right)$$

in distribution.

Principle 4: Asymptotically Best Procedures

Most designs have the asymptotic property that

$$\sqrt{n}\left(\frac{N_A}{n} - \rho(\theta)\right) \rightarrow N\left(0, v(\theta)\right).$$

For all the designs that we have discussed, this property holds. (One exception is Efron's biased coin design, in which the asymptotic distribution is discrete.)

Principle 4: Asymptotically Best Procedures

Then we have the following lower bound on the variance $v(\theta)$ (Hu, Rosenberger, and Zhang, 2007):

$$v(\theta) \geq \frac{\partial \rho(\theta)}{\partial \theta} I^{-1}(\theta) \frac{\partial \rho(\theta)'}{\partial \theta},$$

Example: Maximize power as a criterion:

$$\rho(\theta) = \frac{\sqrt{p_A q_A}}{\sqrt{p_A q_A} + \sqrt{p_B q_B}},$$

the asymptotically best procedure will have asymptotic variance

$$\frac{1}{4(\sqrt{p_A q_A} + \sqrt{p_B q_B})^3} \left(\frac{p_B q_B (q_A - p_A)^2}{\sqrt{p_A q_A}} + \frac{p_A q_A (q_B - p_B)^2}{\sqrt{p_B q_B}} \right).$$

Principle 4: Asymptotically Best Procedures

The doubly-adaptive biased coin design does not attain the lower bound and thus is not "asymptotically best".

Hu, Zhang, and He (2009) found an asymptotically best procedure, which they called the *efficient randomized-adaptive design (ERADE)*. Under certain regularity conditions, which are satisfied by any practical designs, the ERADE can target any optimal allocation $\rho(\theta)$ and attain the minimum variance.

The ERADE

For $0 < \alpha < 1$,

$$\phi_j = \begin{cases} \alpha\rho(\hat{\theta}_j), & \text{when} \quad N_A(j) > \rho(\hat{\theta}_j), \\ \rho(\hat{\theta}_j), & \text{when} \quad N_A(j) = \rho(\hat{\theta}_j), \\ 1 - \alpha(1 - \rho(\hat{\theta}_j)), & \text{when} \quad N_A(j) < \rho(\hat{\theta}_j). \end{cases}$$

- This creates a discontinuous allocation function, similar to Efron's biased coin design, that must be analyzed differently from continuous allocation functions, such as Hu and Zhang's allocation function and Smith's allocation function.

The ERADE

- Note that Efron's biased coin design is a special case of the ERADE when $\rho(\theta) = 1/2$ and $\alpha = p/2$, where p is Efron's biased coin parameter.
- The choice of α does not impact the asymptotic properties of the design, but it does affect the small sample variability of the design. Hu, Zhang, and He suggest selecting a value between 0.4 and 0.7.

Principle 5: Sample Size Computation

For a fixed design, where n_A and n_B are fixed in advance, sample size computation is a straightforward rote exercise.
However restricted randomization and response-adaptive randomization lead to random treatment numbers N_A and N_B, which themselves have distributions, in which case power is a random variable as well.
Hu and Rosenberger (2006) state a "guiding principle" that response-adaptive randomization should only be used when power is preserved.

Principle 5: Sample Size Computation

Hu and Rosenberger (2006) distinguish among three types of sample size computation.

- Type I: the naive method where N_A and N_B are assumed to be fixed in advance
- Type II: the average sample size computed over the distribution of N_A
- Type III: a quantile of the sample size distribution computed over the distribution of N_A

Type III is preferred, because in that case, 95 percent of clinical trials will be adequately powered, whereas Type II only ensures adequate power half the time.

Principle 5: Sample Size Computation

In order to compute the required Type II and Type III sample sizes, we need the randomization procedure to satisfy the following conditions:

- $\frac{N_A(n)}{n} \to \rho \in (0, 1)$ almost surely;
- $\sqrt{n}\left(\frac{N_A(n)}{n} - \rho(\theta)\right) \to N(0, v(\theta))$ in distribution.

Example 1: Restricted Randomization

From Hu and Rosenberger's book, Chapter 6.

Table: *Sample sizes for complete randomization (CR) and Smith's generalized biased coin design (GBC) ($\alpha = 0.05$, $\beta = 0.8$; comparison of normal means with effect size 1).*

(σ_1, σ_2)	(1, 1)	(1, 2)	(1, 4)	(1, 8)
Type I	25	62	211	804
Type II (CR)	26	63	212	805
Type III (CR)	28	72	234	853
Type II (GBC $\gamma = 1$)	26	63	212	805
Type III (GBC $\gamma = 1$)	26	68	223	831
Type II (GBC $\gamma = 5$)	25	62	211	804
Type III (GBC $\gamma = 5$)	25	65	217	818

Example 2: Response-Adaptive Randomization

For the doubly adaptive biased coin design targeting Neyman allocation, we have the following results:

- $\frac{N_A(n)}{n} \to \sigma_1/(\sigma_1+\sigma_2) \in (0,1)$ almost surely;
- $$\sqrt{n}\left(\frac{N_{n1}}{n} - \frac{\sigma_1}{\sigma_1+\sigma_2}\right) \to N\left(0, \frac{2+\gamma}{(1+2\gamma)}\frac{\sigma_1\sigma_2}{(\sigma_1+\sigma_2)^2}\right)$$

  in distribution

So we can use these techniques to find the requisite sample size.

Example 2: Response-Adaptive Randomization

Table: $\alpha = 0.05$, $\beta = 0.8$, *and* $\mu_1 - \mu_2 = 1$.

(σ_1, σ_2)	(1, 1)	(1, 2)	(1, 4)	(1, 8)
I (CR)	25	62	211	804
II (CR)	26	63	212	805
III (CR)	28	72	234	853
I (DBCD)	25	56	155	501
II ($\gamma = 0$)	28	58	157	504
III ($\gamma = 0$)	31	63	163	510
II ($\gamma = 1$)	26	57	156	502
III ($\gamma = 1$)	28	59	159	506
II ($\gamma = 4$)	26	57	156	502
III ($\gamma = 4$)	27	58	158	505

Conclusions

- Randomization is taken as a standard exercise, and few statisticians even give a thought to its implications.
- We provide 5 guiding principles that can be used to find an appropriate procedure and conduct inference and sample size computations.

Conclusions

Hu and Rosenberger (2006) discuss 3 properties that should always be considered in designing clinical trials using response-adaptive randomization:

- Designs should be fully randomized to protect from biases.
- Designs should allow for standard inferential tests to be used.
- Designs should not impact the operating characteristics of the trial: size, power, ethical considerations.

Conclusions

- Many of these designs involve "tuning parameters" that affect the degree of randomization. Examples are p in Efron's biased coin design, γ in the DBCD and GBCD, and α in the ERADE. Optimal values of these parameters can be found using multi-objective optimization criteria (e.g., Cook and Wong). Competing objectives may be simultaneously minimizing selection bias and variability. As the variability increases, the chance of selection bias increases. A weighted compound criterion can be established with the weights determined by the relative importance of the criteria. This is the topic of Wang's thesis.
- Covariate-adaptive randomization is unique and has its own theory; that is another talk at another time...