The Fields Institute

Workshop on Adaptive Designs

September 25-27, 2003

Biased Coin vs. Ehrenfest Urn: An Analysis of Randomness, Balance, and Power

by

Yung-Pin Chen
Department of Mathematical Sciences
Lewis & Clark College
Portland, Oregon

Outline

1. Desirable features of a design: randomness and balance

2. Two adaptive designs for balancing treatment assignments:

   BCDWIT: Biased Coin Design With Imbalance Tolerance

   EUD: Ehrenfest Urn Design

3. Analyzing BCDWIT and EUD: randomness and balance properties; power analysis

Statistical design:

For studies where

- non-response data are collected;
- the treatment assignment is not necessarily adapted to the responses;
- instantaneous responses are not available, or a long delay in reporting responses may occur;

a statistical design is a well-defined rule that specifies how each treatment is assigned in the future given the current allocation status.

Concerns about a statistical design:

- randomness: we don't want to know in advance which treatment will be allocated next;
- balance: we want each treatment to be allocated nearly equally;
- trade-offs between balance and randomness;
- power for detecting treatment effects.

Why randomness?

Experimenters in a less random design are more capable of predicting which treatment will be chosen, and may consciously or unconsciously choose an element favoring their expectation through predictability. Randomization can reduce bias in the choice of a sample and provide a basis for valid statistical inferences.

Why balance?

- In comparative clinical trials, imbalance affects efficiency of statistical inference procedures (Soares and Wu, 1983).

- Many policy makers are unhappy about a badly disproportionate group to treatments. (Shapiro and Louis, 1983)

We will consider allocating two treatments.

Let

$$\begin{aligned} X_{n,i} &= \text{number of times the } i\text{th treatment} \\ &\quad \text{is assigned up to epoch } n,\ i = 1, 2. \\ \Delta_n &= X_{n,1} - X_{n,2} \text{ (net difference).} \\ D_n &= |\Delta_n| \text{ (absolute difference).} \\ \pi_{n,i} &= \text{probability that the } i\text{th treatment} \\ &\quad \text{is assigned at epoch } n,\ i = 1, 2. \end{aligned}$$

Note that the sequence $\{(\pi_{n,1}, \pi_{n,2})\}$ of probability vectors determines the statistical design.

Simple random sampling (SRS):

Under repeated SRS, $(\pi_{n,1}, \pi_{n,2}) = (\frac{1}{2}, \frac{1}{2})$. Repeated SRS owns the attractive property of complete randomization (so free of any selection bias), but suffers from the disadvantage of imbalance:

- In the short run, there is high probability of large differences between treatment assignments.

- In the long run, the difference Δ_n grows at the rate of the square root of the number of assignments. Precisely,

$$\mathcal{L}(\frac{\Delta_n}{\sqrt{n}}) \longrightarrow \text{N}(0,\ 1).$$

Criterion for balance:

Define

$$\delta_n = \frac{1}{n}\sum_{k=1}^{n} \mathsf{E}(D_k)$$

to be the *average imbalance*.

Under repeated SRS, $\delta_n \approx O(\sqrt{n})$.

Criterion for randomness:

Definition (Blackwell and Hodges, 1957) The *excess selection bias* of a design is defined to be the expected number of optimally correct guesses minus $\frac{1}{2}$.

For two treatments, the quantity

$$\mathsf{E}[\sum_{k=1}^{n} \max(\pi_{k,1}, \pi_{k,2})] - \frac{n}{2}$$

is the excess selection bias of the design represented by $\{(\pi_{n,1}, \pi_{n,2})\}$.

Define

$$\beta_n = \frac{1}{n}\mathsf{E}[\sum_{k=1}^{n} \max(\pi_{k,1}, \pi_{k,2})] - \frac{1}{2}$$

to be the *average excess selection bias*.

Under repeated SRS, $\beta_n = 0$, for each n.

Biased Coin Design or BCD (Efron, 1971):

Whenever $D_n = 0$, toss a fair coin.
Whenever $D_n \neq 0$, toss a biased p-coin, $p < \frac{1}{2}$. This design is denoted BCD(p).

Big Stick Design or BSD (Soares and Wu, 1983):

Choose a prescribed imbalance tolerance $b > 0$.
Whenever $D_n < b$, toss a fair coin.
Whenever $D_n = b$, make a deterministic allocation.

Design 1. BCDWIT(p, b): (Chen, Y.-P., 1999)
Biased Coin Design With Imbalance Tolerance, where $0 < p \le \frac{1}{2}$ and $1 \le b \le \infty$.

This design is a synthesis and generalization of BCD(p) and BSD(b). It is represented by

$$(\pi_{n+1,1}, \pi_{n+1,2}) = \begin{cases} (1,0) & \text{if } \Delta_n = -b. \\ (1-p,p) & \text{if } -b < \Delta_n < 0. \\ (1/2,1/2) & \text{if } \Delta_n = 0. \\ (p,1-p) & \text{if } 0 < \Delta_n < b. \\ (0,1) & \text{if } \Delta_n = b. \end{cases}$$

BCD(p): $0 < p < \frac{1}{2}$ and $b = \infty$.

BSD(b): $p = \frac{1}{2}$ and $1 < b < \infty$.

Analyzing BCDWIT

Under the BCDWIT(p, b), the absolute difference process D_n forms a Markov chain on the state space $\{0, 1, \cdots, b\}$ with $(b+1) \times (b+1)$ probability transition matrix

$$P = \begin{bmatrix} 0 & 1 & 0 & \cdot & \cdot & \cdot & 0 \\ q & 0 & p & \cdot & \cdot & \cdot & 0 \\ 0 & q & 0 & \cdot & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & & & & \cdot \\ \cdot & \cdot & \cdot & & & & \cdot \\ \cdot & \cdot & \cdot & & & & \cdot \\ 0 & 0 & 0 & \cdot & \cdot & 1 & 0 \end{bmatrix}.$$

Analyzing BCDWIT (continued)

Theorem 1.

Under **BCDWIT**(p,b), the difference process D_n forms a Markov chain on the state space $\{0,1,\cdots,b\}$ with the stationary distribution

$$\begin{aligned}\eta_0 &= \frac{(q-p)q^{b-1}}{2(q^b-p^b)},\\ \eta_m &= \frac{p^{m-1}}{q^m}\eta_0,\ m=1,\cdots,b-1,\\ \eta_b &= \frac{(q-p)p^{b-1}}{2(q^b-p^b)},\end{aligned}$$

and has high order transition probabilities

Analyzing BCDWIT (continued)

$$
\begin{aligned}
P_{l,m}^{(j)} &= \eta_m + (-1)^{j+l+m}\,\eta_m + \\
&\quad \frac{\eta_m}{\eta_0} \times \sum_{k=1}^{b-1} \frac{(2\sqrt{pq}\cos\frac{k\pi}{b})^j\,\phi_{kl}\,\phi_{km}}{\frac{b(q-p)^2}{2p^2q} + \frac{2b}{p}\sin^2\frac{k\pi}{b}},
\end{aligned}
$$

where

$$
\begin{aligned}
\phi_{km} &= (\frac{q}{p})^{(m+1)/2}\sin\frac{(m-1)k\pi}{b} \\
&\quad -(\frac{q}{p})^{(m-1)/2}\sin\frac{(m+1)k\pi}{b},
\end{aligned}
$$

$$1 \le k \le (b-1)\ 0 \le m \le b.$$

Design 2. EUD(w):
Ehrenfest Urn Design, where $1 \leq w$.

- An urn initially contains w type-1 balls and w type-2 balls.
- A ball is randomly chosen. The selection of a type-i ball corresponds to the selection of the ith treatment, $i = 1, 2$.
- Instead of replacing the chosen ball, add one ball of the opposite type into the urn.
- Repeat these operations at each trial.

Design 2. EUD(w): (continued)

Let $W_{n,i}$ denote the number of type-i balls in the urn at epoch n, then

$$\Delta_n = w - W_{n,1}$$

and EUD(w) is characterized by

$$\begin{aligned}(\pi_{n+1,1}, \pi_{n+1,2}) &= (\frac{W_{n,1}}{2w}, \frac{W_{n,2}}{2w}) \\ &= (\frac{1}{2} - \frac{\Delta_n}{2w}, \frac{1}{2} + \frac{\Delta_n}{2w}).\end{aligned}$$

Analyzing EUD

Under the EUD(w), the absolute difference process D_n forms a Markov chain on the state space $\{0, 1, \cdots, w\}$ with reflecting barriers 0 and w, and it has the $(w+1) \times (w+1)$ probability transition matrix

Analyzing EUD (continued)

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 & 0 \\ \frac{w+1}{2w} & 0 & \frac{w-1}{2w} & 0 & \cdots & 0 & 0 \\ 0 & \frac{w+2}{2w} & 0 & \frac{w-2}{2w} & \cdots & 0 & 0 \\ 0 & 0 & \frac{w+3}{2w} & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix}.$$

Analyzing EUD (continued)

Theorem 2.

Under **EUD**(w), $D_n = |w - W_{n,1}|$ forms a Markov chain on $\{0, 1, \cdots, w\}$ with the stationary distribution

$$\begin{aligned}\eta_0 &= \frac{(2w)!}{w!w!}(\frac{1}{2})^{2w}, \\ \eta_m &= 2\frac{(2w)!}{(2w-m)!m!}(\frac{1}{2})^{2w},\ m = 1, \cdots, w.\end{aligned}$$

and has high order transition probabilities

Analyzing EUD (continued)

$$P^{(j)}_{l,m} \;=\; (\frac{1}{2})^{2w} \sum_{k=0}^{2w} (\frac{w-k}{w})^{j}\, a_{l,k}\, a_{k,m},$$

where the numbers $a_{k,l}$ are the coefficients of the $2w$-degree Krawtchouk polynomial

$$\sum_{l=0}^{2w} a_{k,l}\, z^{l} = (1-z)^{k}\,(1+z)^{2w-k},$$

i.e.,

$$a_{k,l} = \sum_{\nu} (-1)^{\nu} \left(\begin{array}{c} k \\ \nu \end{array} \right) \left(\begin{array}{c} 2w-k \\ l-\nu \end{array} \right),$$

where the summation is taken over

$$l+k-2w \le \nu \le \min(l,,k).$$

Randomness properties

Asymptotic Average Excess Selection Bias β_n under two designs:

<u>**Theorem 3.**</u>
BCDWIT(p, b):

$$\beta_n \longrightarrow \frac{1-(p/q)}{4[1-(p/q)^b]}.$$

EUD(w):

$$\beta_n \longrightarrow \frac{(2w-1)!}{w!(w-1)!}(\frac{1}{2})^{2w}.$$

Proof: It follows from convergence to the stationary distribution.

Balance Properties

Asymptotic Average Imbalance δ_n under two designs:

<u>**Theorem 4.**</u>

BCDWIT(p, b):

$$\delta_n \longrightarrow \frac{1}{2(q-p)} - \frac{bp^b}{q^b - p^b}.$$

EUD(w):

$$\delta_n \longrightarrow \frac{(2w)!}{w!(w-1)!}(\frac{1}{2})^{2w}.$$

Proof: It follows from the mean square ergodic theorem.

Which is better? (Chen, Y.-P., 2000)

Numerical comparison:

Setting $\beta_{\text{EUD}}(w = b) = \beta_{\text{BCDWIT}}(p, b = w)$, and solving for p, say p^*.

Setting $\delta_{\text{EUD}}(w = b) = \delta_{\text{BCDWIT}}(p, b = w)$, and solving for p, say p^{**}.

Numerical solutions for p^* and p^{**}

$w = b$	$\beta_{\text{EUD}}(w) = \beta_{\text{BCDWIT}}(p^*, b)$ p^*
2	0.250000
10	0.394870
20	0.428589
50	0.456772
100	0.470144
500	0.487061
5000	0.495979

$w = b$	$\delta_{\text{EUD}}(w) = \delta_{\text{BCDWIT}}(p^{**}, b)$ p^{**}
2	0.250000
10	0.360707
20	0.400545
50	0.437180
100	0.455633
500	0.480178
5000	0.493733

Which is better? (continued)

Now $p^{**} < p^{*}$, it implies

$$\begin{aligned}\delta_{\text{EUD}}(w=b) &= \delta_{\text{BCDWIT}}(p^{**}, b=w) \\ &< \delta_{\text{BCDWIT}}(p^{*}, b=w)\end{aligned}$$

since $\delta_{\text{BCDWIT}}(p, b=w)$ is increasing in p as b is fixed.

The numerical work suggests that:

With $b = w$ (same imbalance tolerance level), if p is chosen such that both designs have the same level of asymptotic average excess selection bias, then EUD(w) has a smaller asymptotic average imbalance than BCDWIT(p, b).

Power analysis

Consider a trial where subjects are sequentially assigned between a control and a treatment. Assume:

(i) There are $n > 0$ subjects of which n_c are assigned to the control and n_t are assigned to the treatment.

(ii) The control responses $X_1, \cdots, X_{n_c}$ are independent and normally distributed with unknown mean μ_c and variance $\sigma_c^2 = 1$, and the treatment responses $Y_1, \cdots, Y_{n_t}$ are independent and normally distributed with unknown mean μ_t and variance $\sigma_t^2 = 1$. Let $\overline{X}_{n_c}$ and $\overline{Y}_{n_t}$ denote the correponding mean responses.

(iii) We want to test H_0: $\mu_c = \mu_t$ (no treatment effect) versus H_0: $\mu_c < \mu_t$ (there is a treatment effect).

Power analysis (continued)

An α-level test is:

reject H_0, if $\overline{Y}_{n_t} - \overline{X}_{n_c} > z_\alpha \sqrt{1/n_c + 1/n_t}$;
do not reject H_0, otherwise.

This α-level test has the power function

$$\beta_\alpha(\mu_t - \mu_c | n_c, n_t) = \Phi(\frac{\mu_t - \mu_c}{\sqrt{1/n_c + 1/n_t}} - z_\alpha),$$

where $\mu = \mu_t - \mu_c \geq 0$.

Suppose $n = 20$, and $\overline{Y}_{n_t} - \overline{X}_{n_c} = 0.74$.
Under SRS,

$\|n_c - n_t\|$	0	2	4	6	≥ 8
Prob	.176	.320	.240	.148	$\geq .11$
p-value	.048	.049	.053	.057	$\geq .79$

The 5%-level test can detect the treatment effect only when $|n_c - n_t| \leq 2$, with prob

$$\sum_{k=9}^{11} \binom{20}{k} (1/2)^{20} = .4966.$$

Power analysis (continued)

In this case, it is slightly more likely to falsely miss than to correctly detect the treatment effect under repeated SRS.

The inferences may be influenced by what was not observed as well as by what was observed (R. Simon, 1977), and so is the power of the inference procedure.

Most clinicians enthusiastically want to learn more about the new treatment, and they may fail to prove a truly more effective treatment simply because of *unlucky* loss of power incurred by a less balanced randomization for allocating treatments.

How can we make a randomization achieve more balance in allocating treatments?

Power analysis (continued)

Now we use BCD with a common choice $p = 1/3$. Again, let us take $n = 20$. Under the BCD($p = 1/3$),

$\|n_c - n_t\|$	0	2	4	6	≥ 8
Prob	.508	.378	.090	.020	$< .005$

The 5%-level test can detect the treatment effect with prob

$$\Pr\{|D_n| \leq 2\} = .508 + .378 = .886.$$

Different randomization may generate the same treatment assignments with different likelihoods.

We can treat the power function $\beta_\alpha(\mu_t - \mu_c | n_c, n_t)$ as the conditional probability of rejecting H_0 on the realized treatment assignments n_c and n_t.

Power analysis (continued)

Under repeated SRS, the (overall) power function of the α-level test is

$$\begin{aligned}
&\beta_{\alpha,\text{SRS}}(\mu_t - \mu_c) \\
= & \sum_{n_c+n_t=n} \beta_\alpha(\mu_t - \mu_c | n_c, n_t)\text{Pr}_{\text{SRS}}\{(n_c, n_t)\} \\
= & \sum_{n_c+n_t=n} \Phi\left(\frac{\mu_t - \mu_c}{\sqrt{\frac{1}{n_c} + \frac{1}{n_t}}} - z_\alpha\right) \binom{n}{n_c} 2^{-n}.
\end{aligned}$$

Under the BCD(p), the (overall) power function of the α-level test is

$$\begin{aligned}
&\beta_{\alpha,\text{BCD}(p)}(\mu_t - \mu_c) \\
= & \Phi\left(\frac{\mu_t - \mu_c}{\sqrt{\frac{1}{n_c} + \frac{1}{n_t}}} - z_\alpha\right) P_{0,0}^{(n)} I_{\{n_c = n_t\}} + \\
& \sum_{n_c+n_t=n} \Phi\left(\frac{\mu_t - \mu_c}{\sqrt{\frac{1}{n_c} + \frac{1}{n_t}}} - z_\alpha\right) \times \\
& \frac{1}{2} P_{0,|n_c - n_t|}^{(n)} I_{\{n_c \neq n_t\}}.
\end{aligned}$$

Power analysis (continued)

Plot of power functions: ($n = 20$, $p = 1/3$) BCD (dark) vs. SRS (gray)

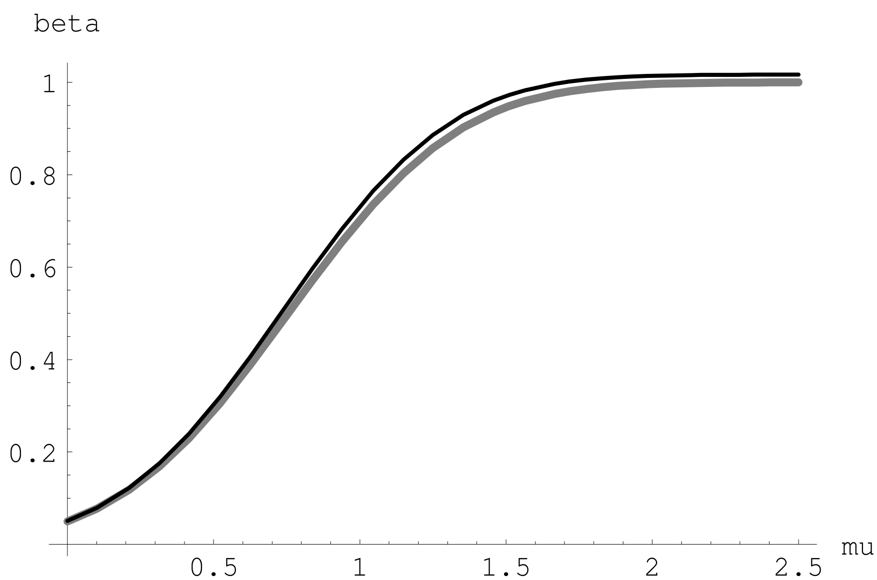

The biased coin design makes equal allocations more likely.

Further investigation and other concerns:

- Power analysis with a large sample size?
- Power analysis under various conditions?
- There is a trade-off between balance and selection bias.
- Can we find designs that are balanced enough to maintain the power of the inference procedures, and at the same time, they are random enough to minimize the selection bias?