Multi-objective dose-finding

Thomas Jaki

Medical and Pharmaceutical Statistics Research Unit,
Department of Mathematics and Statistics, Lancaster University, UK

December 7, 2018

Acknowledgement: This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 633567.

Dose escalation

- Limited prior knowledge about toxicities in humans
- Range of m regimes (doses, combinations, schedules)
- n patients

Goal:

- Find the maximum tolerated regime that corresponds to a controlled level of toxicity, usually $\gamma \in (0.2, 0.35)$ in oncology trials

TD20

$$p(d) = P(\text{toxicity}|\text{dose } d)$$

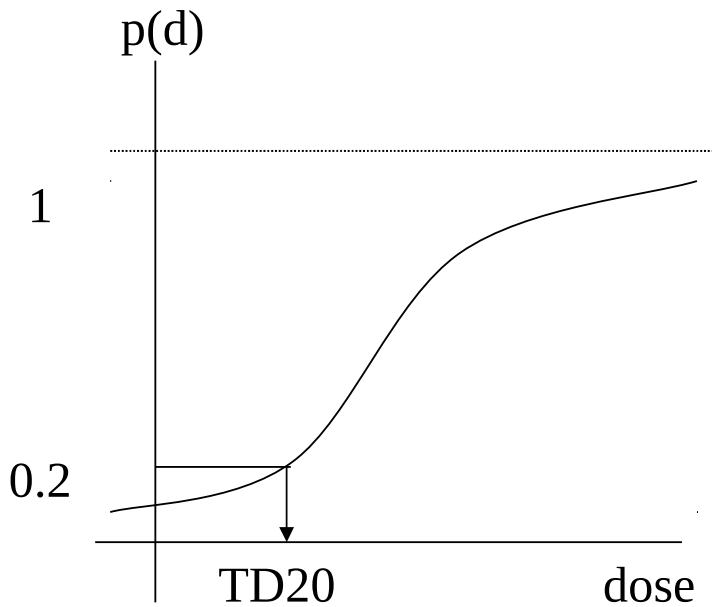

Assume that a 20% risk of toxicity is an acceptable risk to pay for a chance of benefit

General (Bayesian) approach

1. Make assumptions about the form of $p(d)$
2. Impose a prior distribution for the parameters that determine $p(d)$
3. Choose next dose to optimise some form of expected gain
4. Stop once target dose level can be estimated accurately enough

Bayesian continual reassessment method

1. $p(d_i) = d_i^{\exp(\beta)}$
2. $\beta \sim N(0, 1.34)$
3. $d^* = \min_i \mathbb{E}\left((p(d_i) - \gamma)^2\right)$
4. Stop after N patients have been recruited

Single agent dose-escalation designs

Model-based methods

- CRM
- EWOC

Algorithm based methods

- '3+3' design
- Biased Coin Design

Fundamental assumption: a **monotonic** dose-response relationship

Cannot be applied to:

- Combination trials with many treatments
- Scheduling of drugs
- Non-monotonic dose-toxicity relations

Unknown ordering problem. Example (I)

Let us consider drugs combination dose-escalation trial with

- 3 dose levels of drug A: A_1, A_2, A_3
- 3 dose levels of drug B: B_1, B_2, B_3

$(A_1; B_3)$	$(A_2; B_3)$	$(A_3; B_3)$
$(A_1; B_2)$	$(A_2; B_2)$	$(A_3; B_2)$
$(A_1; B_1)$	$(A_2; B_1)$	$(A_3; B_1)$

Even assuming monotonicity one drug being fixed, we cannot order
$(A_1; B_2)$ and $(A_2; B_1)$;
$(A_1; B_3)$ and $(A_2; B_1)$;
$(A_1; B_3)$ and $(A_3; B_1)$ and so on...

Unknown ordering problem. Example (II)

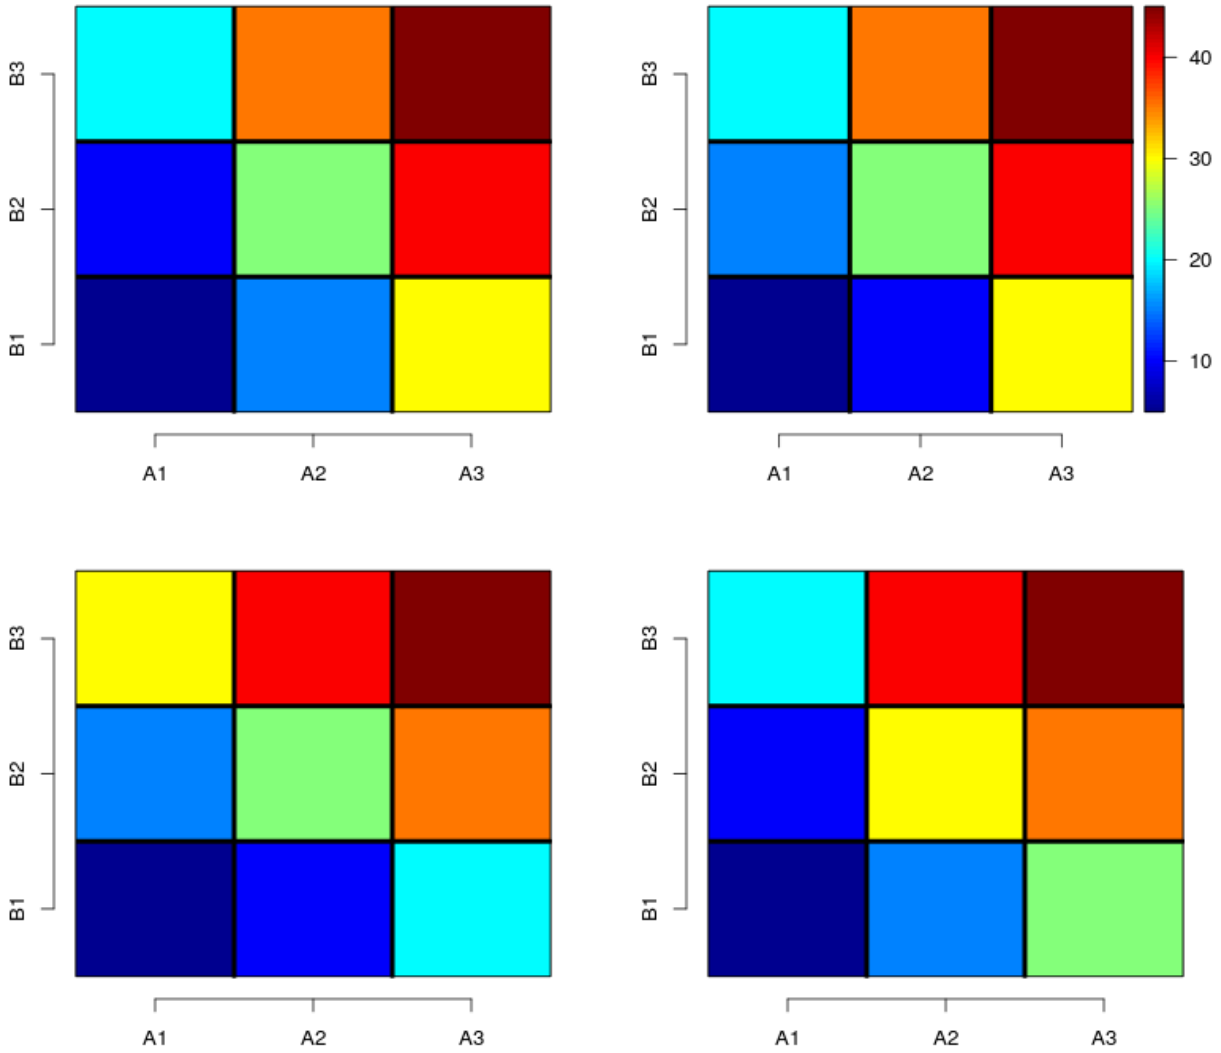

Unknown ordering problem. Example (III)

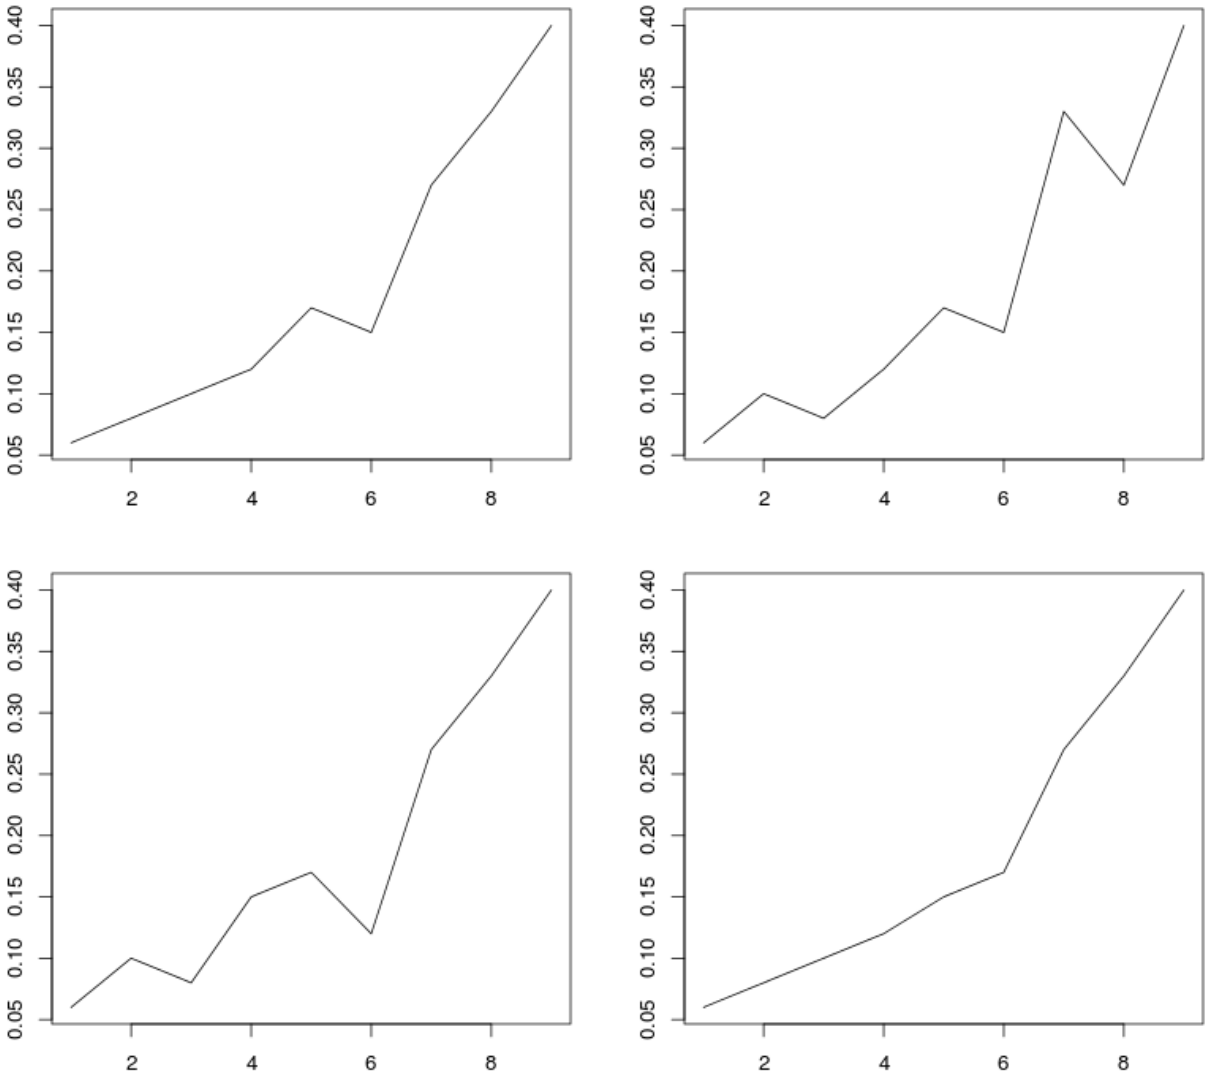

Method for drug combinations

- **Six-parameter model** *(Thall P. et al, 2003)*
- **Up-and-down design** *(Ivanova A, Kim S., 2009)*
  Using the T -statistic
- **Copula regression** *(G.Yin, Y.Yuan, 2009)*
  Parametrization of drug-drug interactive effect
- **POCRM** *(N.Wages, M. Conoway, J. O'Quigley, 2011)*
  Choose several ordering and randomize between them during the trial

General restrictions:

- Strong model assumptions are usually needed
- No diagonal switching is allowed
- Synergistic effect is usually assumed
- Only two combinations only

Goal

To propose an escalation procedure that **does not require any parametric assumptions** (including monotonicity between regimes).

Problem formulation

- Toxicity probabilities $Z_1, \ldots, Z_m$ are random variables with Beta prior $\mathrm{B}(\nu_j + 1, \beta_j - \nu_j + 1)$, $\nu_j > 0, \beta_j > 0$
- n_j patients assigned to the regime j and x_j toxicities observed
- Beta posterior f_{n_j} $\mathrm{B}(x_j + \nu_j + 1, n_j - x_j + \beta_j - \nu_j + 1)$
- Let $0 < \alpha_j < 1$ be the unknown parameter in the neighbourhood of which the probability of toxicity is concentrated
- Target toxicity γ

Information theory concepts

A statistical experiment of estimation of a toxicity probability.

The Shannon differential entropy (DE) $h(f_n)$ of the PDF f_n is defined as

$$h(f_n) = -\int_0^1 f_n(p)\log f_n(p)\mathrm{d}p$$

with the convention $0\log 0 = 0$.

Information theory concepts

A statistical experiment of estimation of a toxicity probability.

The Shannon differential entropy (DE) $h(f_n)$ of the PDF f_n is defined as

$$h(f_n) = -\int_0^1 f_n(p)\log f_n(p)\mathrm{d}p$$

with the convention $0\log 0 = 0$.

It shows the amount of information needed to answer the question

What is the toxicity probability?

Weighted information

Consider a two-fold experiment:

(i) what is the toxicity probability

(ii) is the probability of toxicity close to a target, γ

Weighted information

Consider a two-fold experiment:

(i) what is the toxicity probability

(ii) is the probability of toxicity close to a target, γ

A: The *weighted Shannon information*

$$h_\phi(f) = -\int_{\mathbb{R}} \phi(z) f(z) \log f(z) \mathrm{d}z.$$

Weight Function

The Beta-form weight function

$$\phi_n(p) = \Lambda(\gamma, x, n) p^{\gamma\sqrt{n}} (1-p)^{(1-\gamma)\sqrt{n}}.$$

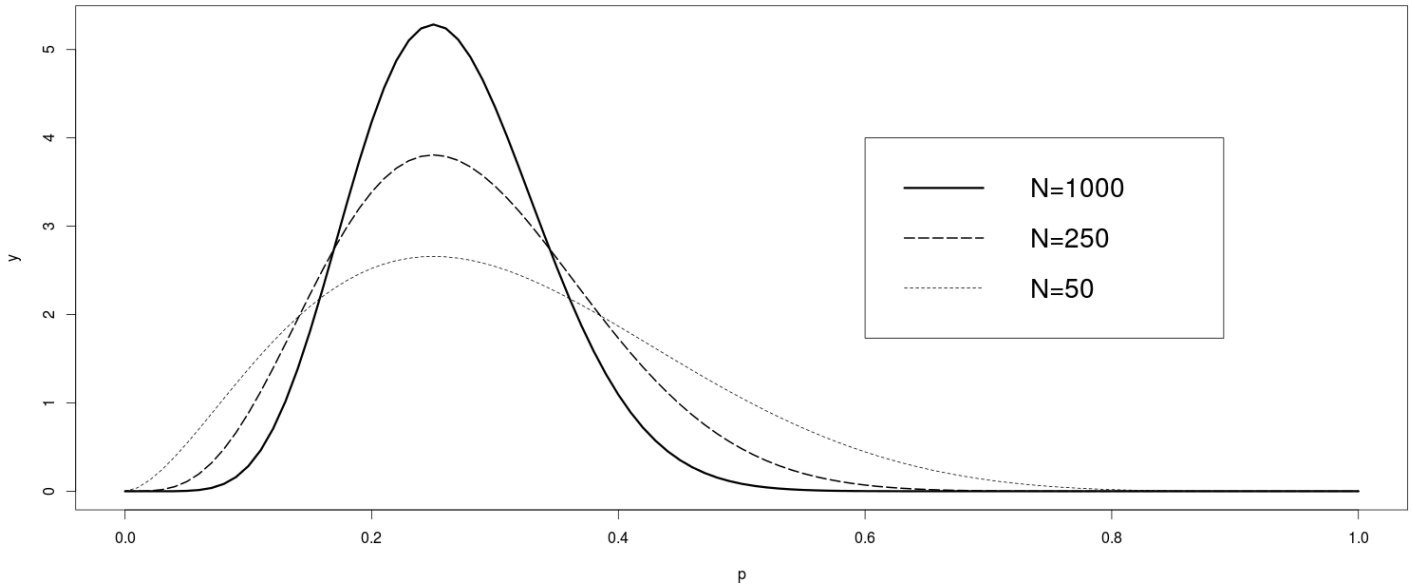

Escalation criteria

Theorem

Let $h(f_n)$ and $h^{\phi_n}(f_n)$ be the DE and WDE corresponding to PDF f_n when $x \sim \alpha n$ with the weight function ϕ_n given in (15). Then

$$\lim_{n\to\infty} \left(h^{\phi_n}(f_n) - h(f_n)\right) = \frac{(\alpha-\gamma)^2}{2\alpha(1-\alpha)}$$

Escalation criteria

Theorem

Let $h(f_n)$ and $h^{\phi_n}(f_n)$ be the DE and WDE corresponding to PDF f_n when $x \sim \alpha n$ with the weight function ϕ_n given in (15). Then

$$\lim_{n\to\infty} \left(h^{\phi_n}(f_n) - h(f_n)\right) = \frac{(\alpha-\gamma)^2}{2\alpha(1-\alpha)}$$

Therefore, for a regimen d_j, $j = 1, \ldots, m$, we obtained that

$$\Delta_j \equiv \frac{(\alpha_j-\gamma)^2}{\alpha_j(1-\alpha_j)}.$$

Criteria:

$$\Delta_j = \inf_{i=1,\ldots,m} \Delta_i.$$

Estimation

Consider the mode of the posterior distribution f_{n_j}

$$\hat{p}_j^{(n)} = \frac{x_j + \nu_j}{n_j + \beta_j}.$$

Then the following "plug-in" estimator $\hat{\Delta}_j^{(n)}$ may be used

$$\hat{\Delta}_j^{(n)} = \frac{(\hat{p}_j^{(n)} - \gamma)^2}{\hat{p}_j^{(n)}(1 - \hat{p}_j^{(n)})}.$$

Escalation design

Let $d_j(i)$ be a regime d_j recommended for cohort i.

- The procedure starts from $\hat{\Delta}_j^{(0)}$
- l cohorts were already assigned

The $(l+1)^{th}$ cohort of patients will be assigned to regime k such that

$$d_j(l+1): \ \hat{\Delta}_k^{(l)} = \inf_{i=1,\ldots,m} \hat{\Delta}_i^{(l)}, \ l = 0, 1, 2, \ldots, C.$$

We adopt regime $d_j(C+1)$ as the final recommended regime.

Alternative angle

One can consider

$$\hat{\Delta}_j^{(n)} = \frac{(\hat{p}_j^{(n)} - \gamma)^2}{\hat{p}_j^{(n)}(1 - \hat{p}_j^{(n)})}$$

as a **loss function** for a parameter defined on $(0, 1)$.

- Loss function penalize $\hat{p}_j^{(n)}$ close to 0 to 1 and **'pushes' the allocation away from bounds** to the neighbourhood of γ
- Does not include any definition of safety $\rightarrow$ safety constraint is needed

Safety constraint

Considers regime d_j as safe if at the moment n its PDF satisfies

$$P(\text{regime is overly toxic}) = \int_{\gamma^*}^{1} f_{n_j}(p)\mathrm{d}p \leq \theta_n$$

where

- γ^* is some threshold after which all regimes above are declared to have excessive risk, $\gamma^* = \gamma + 0.2$
- θ_n is the level of probability that controls the overdosing
  - Note that this depends on n

Why is a time-varying SC is needed?

If $\beta = 1$ and $\theta_n = \theta = 0.50$ then regimes with prior mode ≥ 0.40 will never be considered since

$$\int_{0.45}^{1} f_0(p|x = 0)\mathrm{d}p = 0.5107 > 0.50$$

Requirements to the function θ_n

- θ_n is a decreasing function of n
- $\theta_0 = 1$
- $\theta_N \leq 0.3$
- $\rightarrow \theta_n = 1 - rn$

Choice of SC parameters

	r							
	0.010	0.015	0.020	0.025	0.030	0.035	0.040	0.045
$\gamma^* = 0.55$	0.00	0.32	4.32	18.47	36.15	49.06	61.49	75.70
	26.47	26.65	26.40	26.05	26.85	25.03	24.10	20.23
$\gamma^* = 0.50$	0.15	2.50	17.76	38.75	52.74	63.06	74.94	87.22
	26.27	26.22	26.53	27.24	25.46	23.30	19.35	17.10
$\gamma^* = 0.45$	1.13	12.72	35.72	56.49	67.16	**77.55**	86.53	93.49
	26.15	26.02	26.81	25.18	22.26	**21.75**	15.16	11.05
$\gamma^* = 0.40$	7.47	37.95	59.49	70.52	80.53	88.32	94.18	97.63
	26.04	25.91	24.90	21.98	17.66	14.47	8.05	3.51
$\gamma^* = 0.35$	33.98	58.22	74.42	84.14	90.52	94.86	97.90	99.20
	25.65	24.54	20.45	15.55	13.77	7.21	3.25	0.70
$\gamma^* = 0.30$	55.51	77.02	87.21	92.99	96.50	98.55	99.37	99.83
	24.21	18.09	14.40	11.42	7.13	0.95	0.08	0.04

Table: Top row: Proportion of no recommendations for toxic scenario. Bottom row: Proportion of correct recommendations. 10^6 simulations.

Simulations

For simulations below the following parameters were chosen:

- The cohort size $c = 1$
- Total sample size $N = 20$
- Number of regimes $m = 7$
- The target probability $\gamma = 0.25$
- Safety constraint

$$\theta_n = \begin{cases} 1 - 0.035n, \ \textit{if} \ 0.035 \times n \leq 0.7; \\ 0.3, \ \textit{otherwise}. \end{cases}$$

Investigated scenarios

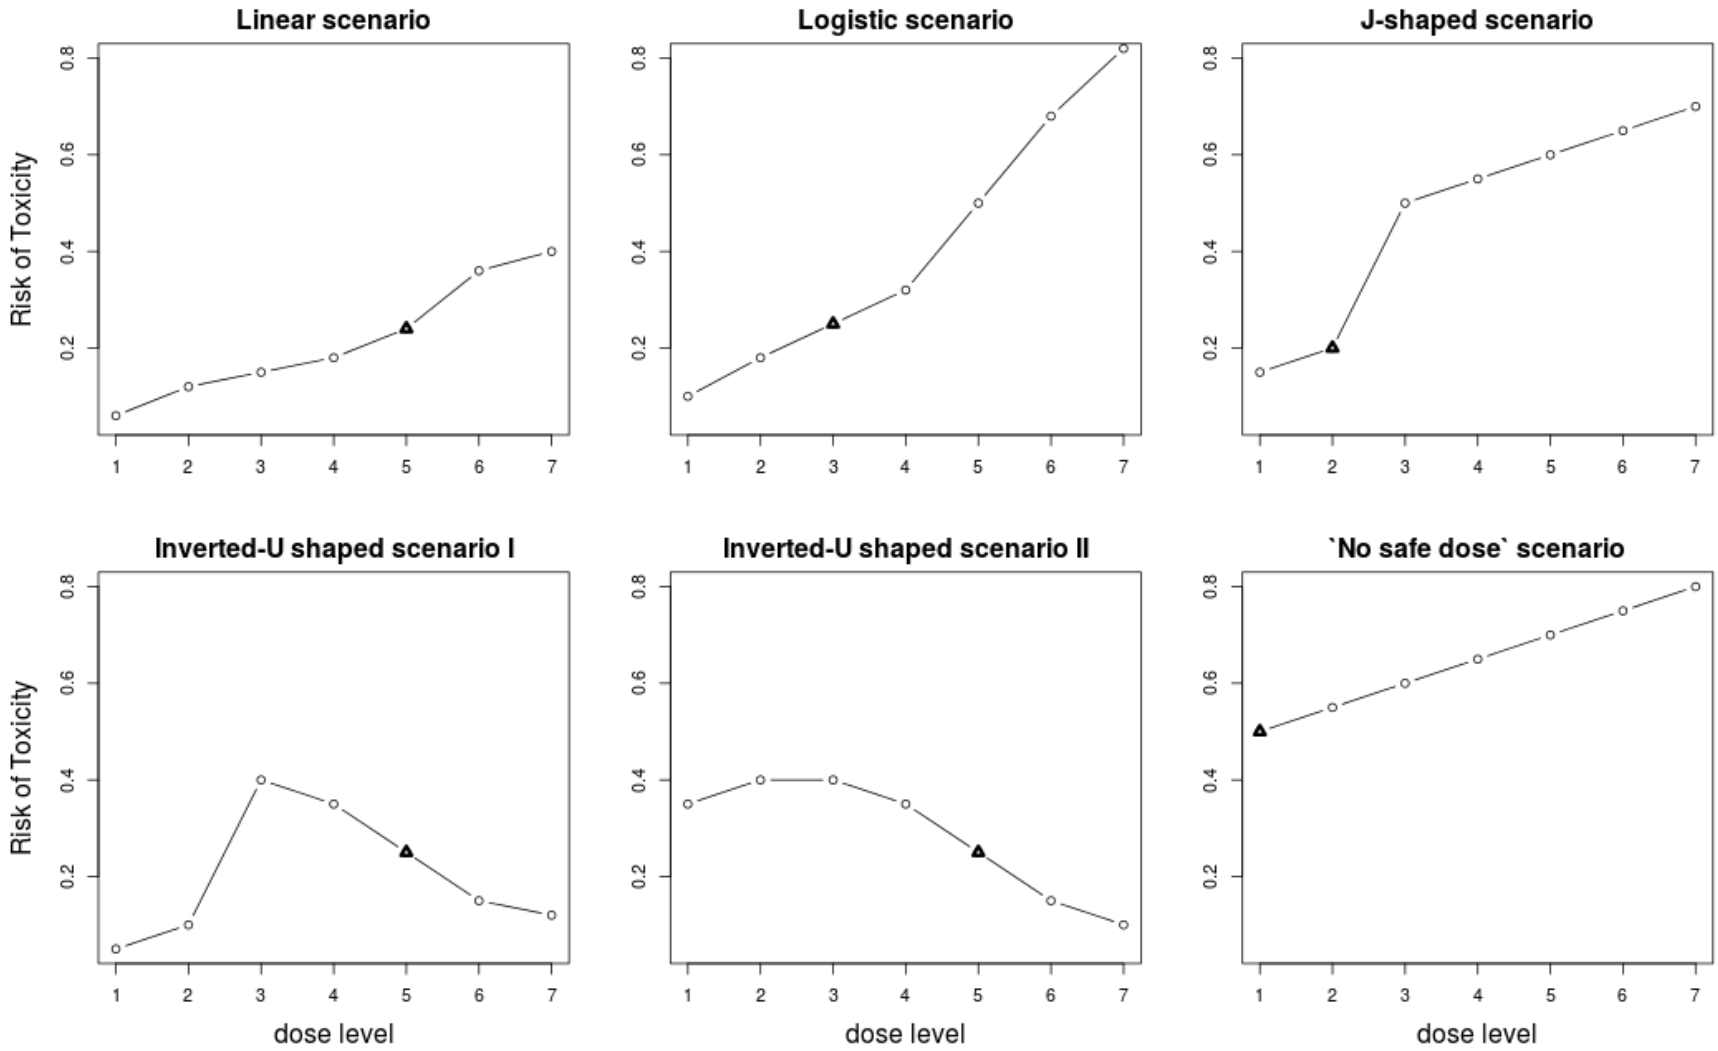

Figure: Considering response shapes. The TD is marked as triangle.

Specifying the prior

Assumptions:

- Vague beliefs about toxicity risk
- Prior belief: regimes have been correctly ordered monotonically
- A escalation to be started from d_1

The prior for regime d_j $(1 \leq j \leq 7)$ is specified thought the mode $\hat{p}_j^{(0)} = \frac{\nu_j}{\beta_j}$.
Starting from the bottom: $\hat{p}_1^{(0)} = \gamma$.
The vector of modes $\hat{\mathbf{p}}$ for all regimes is defined

$$\hat{\mathbf{p}} = [0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55]^{\mathrm{T}}.$$

Vague prior $\rightarrow \beta_j = \beta = 1$ for $j = 1, \ldots, m$.

Is there a unique set of prior parameters that lead to the equivalent performance?

Choice of prior

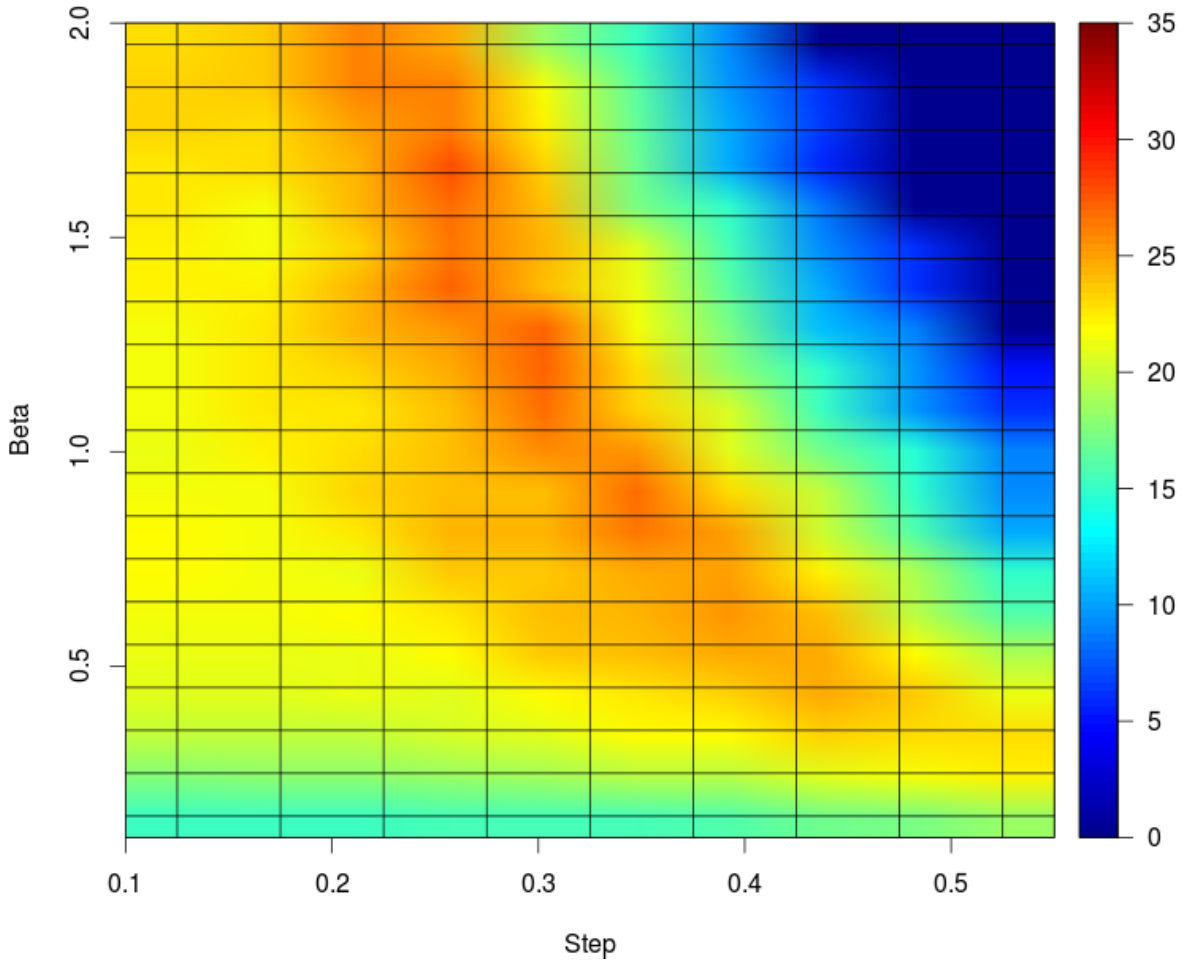

Figure: Proportion of correct recommendations: $\beta =$ number of patients and difference between the risk of toxicity on lowest and highest dose across six scenarios.

Alternative methods

We have also investigated

- Continual reassessment method (CRM)
- Partial ordering continual reassessment method (POCRM)
  All correct orderings used in simulation are incorporated in the model.
- Escalation with overdose control (EWOC)
  A target 25^{th} *percentile is used.*
- Non-parametric optimal benchmark

Simulation results. Ordering is correctly specified

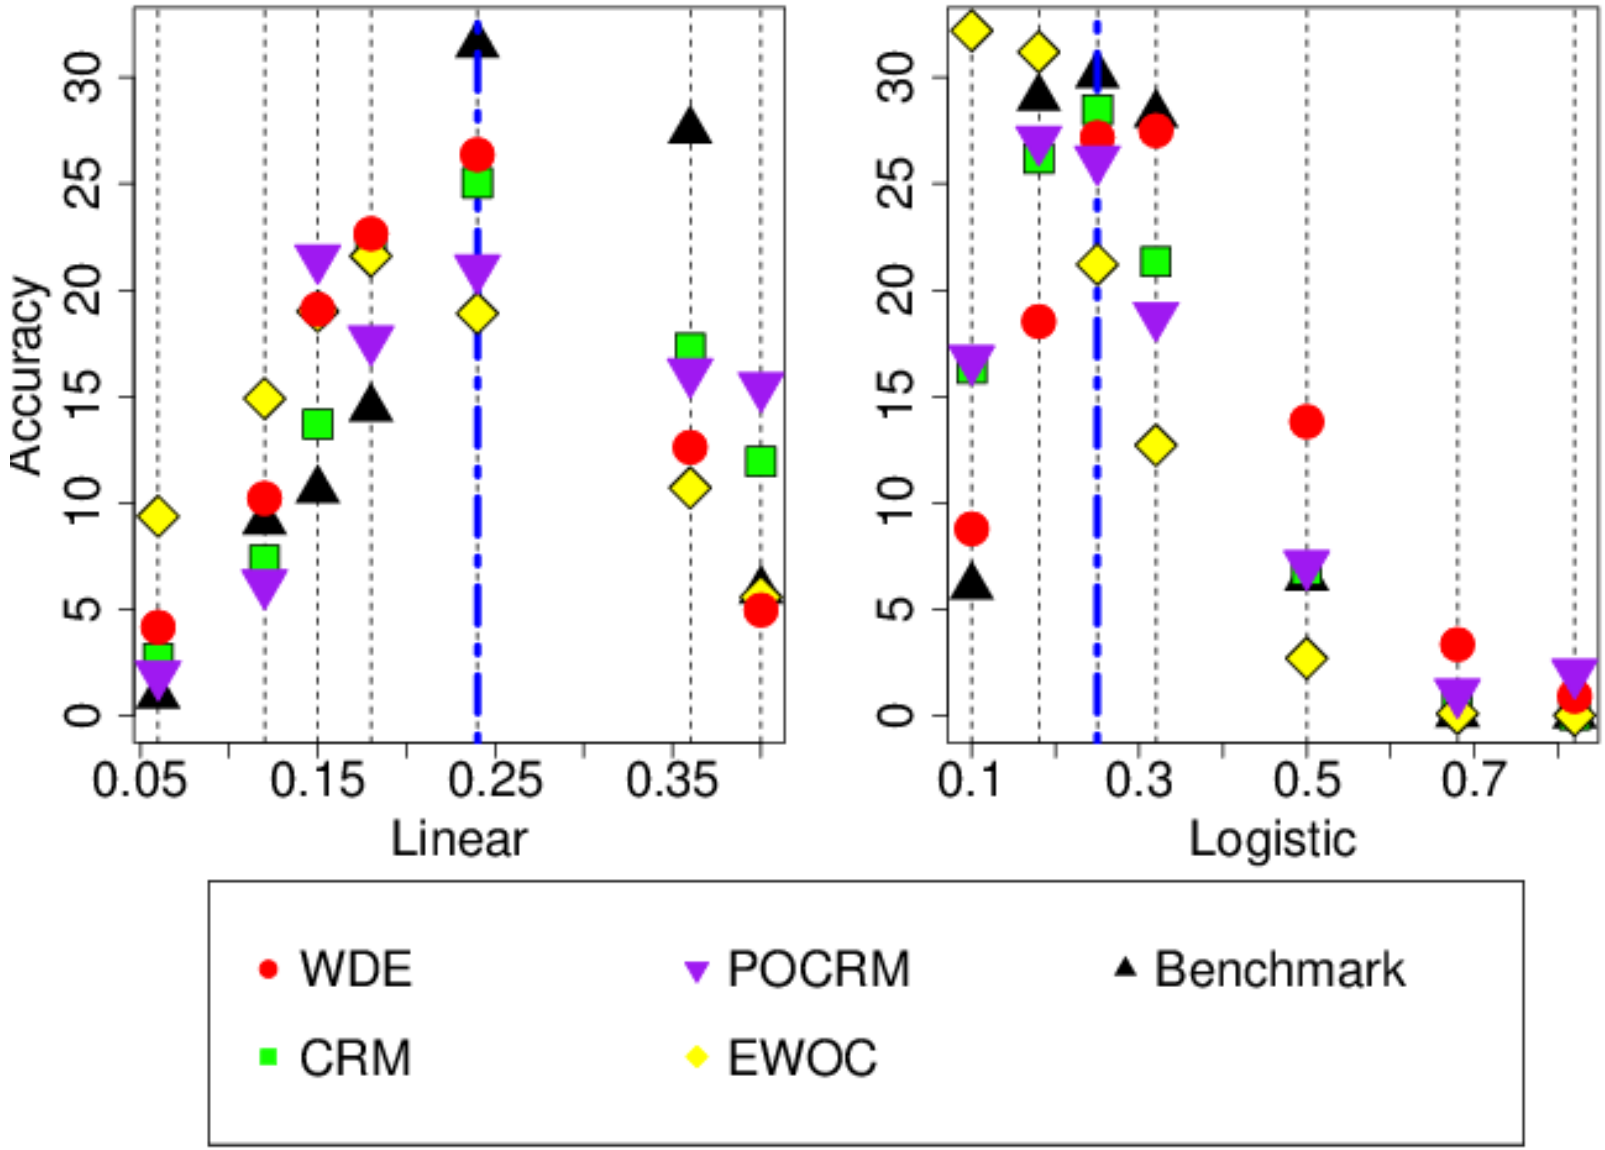

Simulation results. Ordering is wrongly specified.

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	No	TR	$\bar{N}$
True	0.05	0.10	0.40	0.35	**0.25**	0.15	0.12			
WDE_{SC}	14.11	19.13	11.77	18.27	**27.90**	8.50	0.23	0.15	**4.26**	19.99
CRM_{SC}	4.26	**19.90**	17.70	6.31	**2.84**	3.00	**46.10**	0.31	**3.26**	19.92
POCRM_{SC}	2.87	11.39	11.75	9.32	**19.11**	33.94	11.62	0.24	**4.29**	19.99
EWOC_{SC}	7.18	**24.90**	18.60	3.79	**2.52**	3.79	**30.60**	6.62	**2.73**	18.89

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	No	TR	$\bar{N}$
True	0.35	0.40	0.40	0.35	**0.25**	0.15	0.10			
WDE_{SC}	15.57	12.65	13.31	18.27	**27.92**	8.90	0.58	9.96	**5.81**	19.73
CRM_{SC}	**47.41**	2.51	0.97	0.48	**0.72**	0.40	**30.10**	27.30	**4.27**	15.96
POCRM_{SC}	16.81	5.98	5.66	12.42	**20.10**	**23.13**	10.23	9.67	**5.14**	19.46
EWOC_{SC}	**30.75**	1.26	0.78	0.47	**0.47**	0.31	9.78	**56.15**	**3.30**	11.02

Simulation results. Highly toxic scenarios.

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	No	TR	$\bar{N}$
True	0.15	**0.20**	0.50	0.55	0.60	0.65	0.70			
WDE_{SC}	38.07	**44.65**	6.59	3.44	1.48	0.28	0.02	5.47	**5.94**	19.77
CRM_{SC}	37.47	**37.85**	17.41	2.92	0.36	0.07	0.00	3.92	**5.10**	19.41
POCRM_{SC}	33.57	**37.76**	13.27	2.55	0.54	1.33	6.04	4.95	**6.06**	19.82
EWOC_{SC}	**51.00**	**26.11**	11.01	0.88	0.13	0.00	0.00	10.87	**3.60**	16.82
True	**0.50**	0.55	0.60	0.65	0.70	0.75	0.80	**No**		
WDE_{SC}	13.63	5.53	2.45	0.88	0.27	0.06	0.00	**77.17**	**8.02**	14.28
CRM_{SC}	**32.24**	0.32	0.08	0.00	0.00	0.00	0.00	**67.36**	**5.33**	10.30
POCRM_{SC}	15.18	0.57	0.12	0.04	0.01	3.06	0.08	**80.94**	**7.12**	12.59
EWOC_{SC}	16.17	0.00	0.12	0.00	0.00	0.00	0.00	**83.71**	**3.07**	6.05

Conclusions - toxicity only

The WDE-based method

- **performs comparably** to the model-based methods **when the ordering is specified correctly** scenarios
- **outperform** them in **wrongly specified** setting

However, WDE-based method

- **experience problems** in scenarios with **no safe doses** or with sharp jump in toxicity probability at the bottom.
- **The time-varying safety constrain** in the proposed form *can overcome overdosing problems* and increase the accuracy of the original method

Motivating trial - Dual endpoint

Immunotherapy (Molecularly Targeted Agent, MTA) + Chemotherapy:

- 2/3 days immunotherapy AFTER chemotherapy (S_1/S_2),
- 4 days immunotherapy OVERLAP chemotherapy for 1/2 days (S_3/S_4);
- binary toxicity and efficacy endpoints.

Regimen	R_1	R_2	$\mathbf{R_3}$	$\mathbf{R_4}$	$\mathbf{R_5}$	R_6
Cycle 1		S_1	S_2	S_3	S_3	S_4
Cycle 2	S_1	S_2	S_2	S_3	S_4	S_4

Motivating trial - Dual endpoint

Immunotherapy (Molecularly Targeted Agent, MTA) + Chemotherapy:

- 2/3 days immunotherapy AFTER chemotherapy (S_1/S_2),
- 4 days immunotherapy OVERLAP chemotherapy for 1/2 days (S_3/S_4);
- binary toxicity and efficacy endpoints.

Regimen	R_1	R_2	$\mathbf{R_3}$	$\mathbf{R_4}$	$\mathbf{R_5}$	R_6
Cycle 1		S_1	S_2	S_3	S_3	S_4
Cycle 2	S_1	S_2	S_2	S_3	S_4	S_4

The aim: to find the **optimal** regimen (maximum efficacy, least toxicity)
correct regimen (maximum efficacy, acceptable toxicity)

Current approaches

Two perspectives for model-based designs:

- to include parameters for each term (agent, cycle, interaction)
  see e.g. Riviere et al (2016) for a Phase I/II single-agent design.

  Challenge: many parameters to be estimated.

- to include all possible orderings of regimens according to toxicity/efficacy
  see e.g. Wages and Tait (2015) for a Phase I/II single-agent design.

  Challenge: many orderings to be considered.

Alternative: a design **relaxing parametric/monotonicity assumptions**

Derivation of selection criterion (I)

Finding a measure of uncertainty in a Phase I/II trial with 3 outcomes.

Outcome	Probability	Optimal characteristics
Efficacy + No Toxicity	θ_1	γ_1
No Efficacy + No Toxicity	θ_2	γ_2
Toxicity	$\theta_3 = 1 - \theta_1 - \theta_2$	$\gamma_3 = 1 - \gamma_1 - \gamma_2$

Derivation of selection criterion

Using the same arguments as before we base our criterion on

$$\delta(\cdot) = \lim_{n\to\infty} h^{\phi_n}(f_n) - h(f_n)$$

which, for a Dirichlet distribution f_n, and a Dirichlet form weight ϕ_n yields

$$\delta\left(\boldsymbol{\theta}, \boldsymbol{\gamma}\right) := \frac{\gamma_1^2}{\theta_1} + \frac{\gamma_2^2}{\theta_2} + \frac{(1-\gamma_1-\gamma_2)^2}{1-\theta_1-\theta_2} - 1.$$

Trade-off function

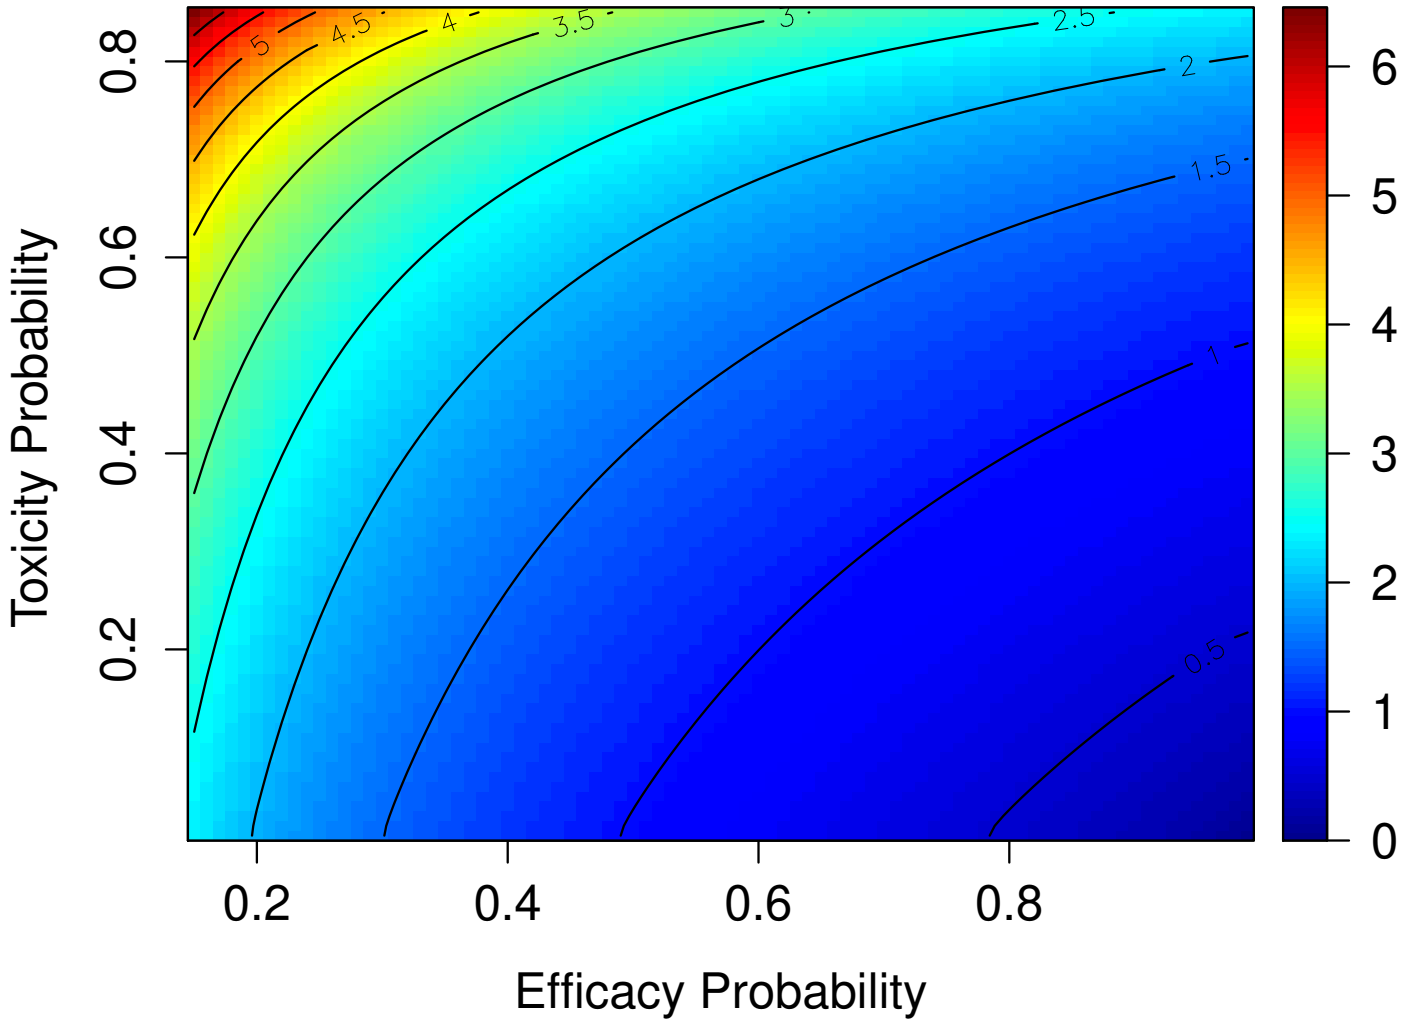

Figure: $\gamma_t = 0.01$, $\gamma_e = 0.99$

Regimen-finding design

As before but with **randomization** between best two regimens with probabilities proportional to

$$1/\hat{\delta}_l^{(k)} \quad l = i, j$$

Application to the motivating trial

$M = 6$ regimens and $N = 36$ patients

We study

1. the proportion of **optimal** selections (maximum efficacy, least toxicity)
2. the proportion of **correct** selections (maximum efficacy, acceptable T)

Application to the motivating trial

$M = 6$ regimens and $N = 36$ patients

We study

1. the proportion of **optimal** selections (maximum efficacy, least toxicity)
2. the proportion of **correct** selections (maximum efficacy, acceptable T)

Scenarios:

8 scenarios for single-MTA studies $\rightarrow$ **six permutations** wrt toxicity orderings.

	1	2	3	4	5	6
1.1	(.005;.01)	(.01;.10)	(.02;.30)	(.05;.50)	**(.10;.80)**	(.15;.80)

Application to the motivating trial

$M = 6$ regimens and $N = 36$ patients

We study

1. the proportion of **optimal** selections (maximum efficacy, least toxicity)
2. the proportion of **correct** selections (maximum efficacy, acceptable T)

Scenarios:

8 scenarios for single-MTA studies $\rightarrow$ **six permutations** wrt toxicity orderings.

	1	2	3	4	5	6
1.1	(.005;.01)	(.01;.10)	(.02;.30)	(.05;.50)	**(.10;.80)**	(.15;.80)
1.2	(.005;.01)	(.01;.10)	(.02;.30)	**(.10;.80)**	(.05;.50)	(.15;.80)

Application to the motivating trial

$M = 6$ regimens and $N = 36$ patients

We study

1. the proportion of **optimal** selections (maximum efficacy, least toxicity)
2. the proportion of **correct** selections (maximum efficacy, acceptable T)

Scenarios:

8 scenarios for single-MTA studies $\rightarrow$ **six permutations** wrt toxicity orderings.

	1	2	3	4	5	6
1.1	(.005;.01)	(.01;.10)	(.02;.30)	(.05;.50)	**(.10;.80)**	(.15;.80)
1.2	(.005;.01)	(.01;.10)	(.02;.30)	**(.10;.80)**	(.05;.50)	(.15;.80)
1.3	(.005;.01)	(.01;.10)	(.05;.50)	(.02;.30)	**(.10;.80)**	(.15;.80)
1.4	(.005;.01)	(.01;.10)	**(.10;.80)**	(.02;.30)	(.05;.50)	(.15;.80)
1.5	(.005;.01)	(.01;.10)	(.05;.50)	**(.10;.80)**	(.02;.30)	(.15;.80)
1.6	(.005;.01)	(.01;.10)	**(.10;.80)**	(.05;.50)	(.02;.30)	(.15;.80)

Practical considerations

- Delayed efficacy response
  e.g. toxicity is evaluated in 1st cycle and efficacy in a 2nd
- Missing efficacy response
  no efficacy data for patients with toxic response
- Coherence principles
  Escalation/De-escalation restrictions

Comparator:

Partial Ordering CRM with 6 toxicity and 48 efficacy orderings.

Results

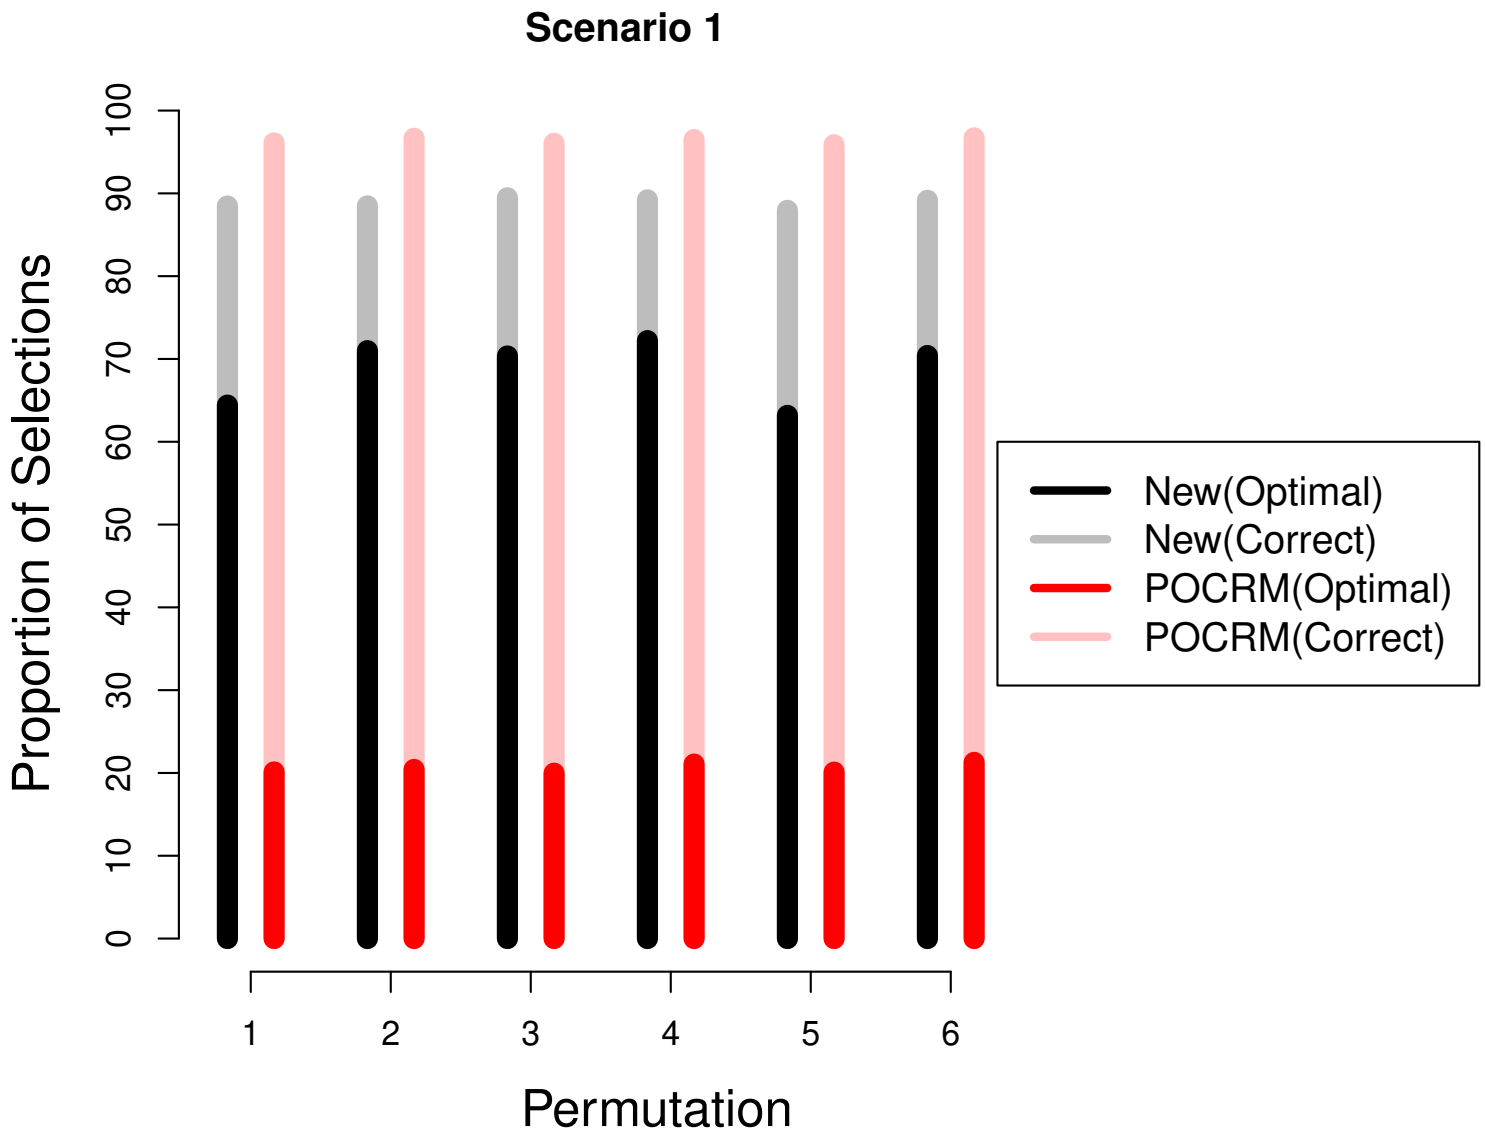
Scenario 1
Proportion of Selections
0
10
20
30
40
50
60
70
80
90
100
1
2
3
4
5
6
Permutation
New(Optimal)
New(Correct)
POCRM(Optimal)
POCRM(Correct)

Results

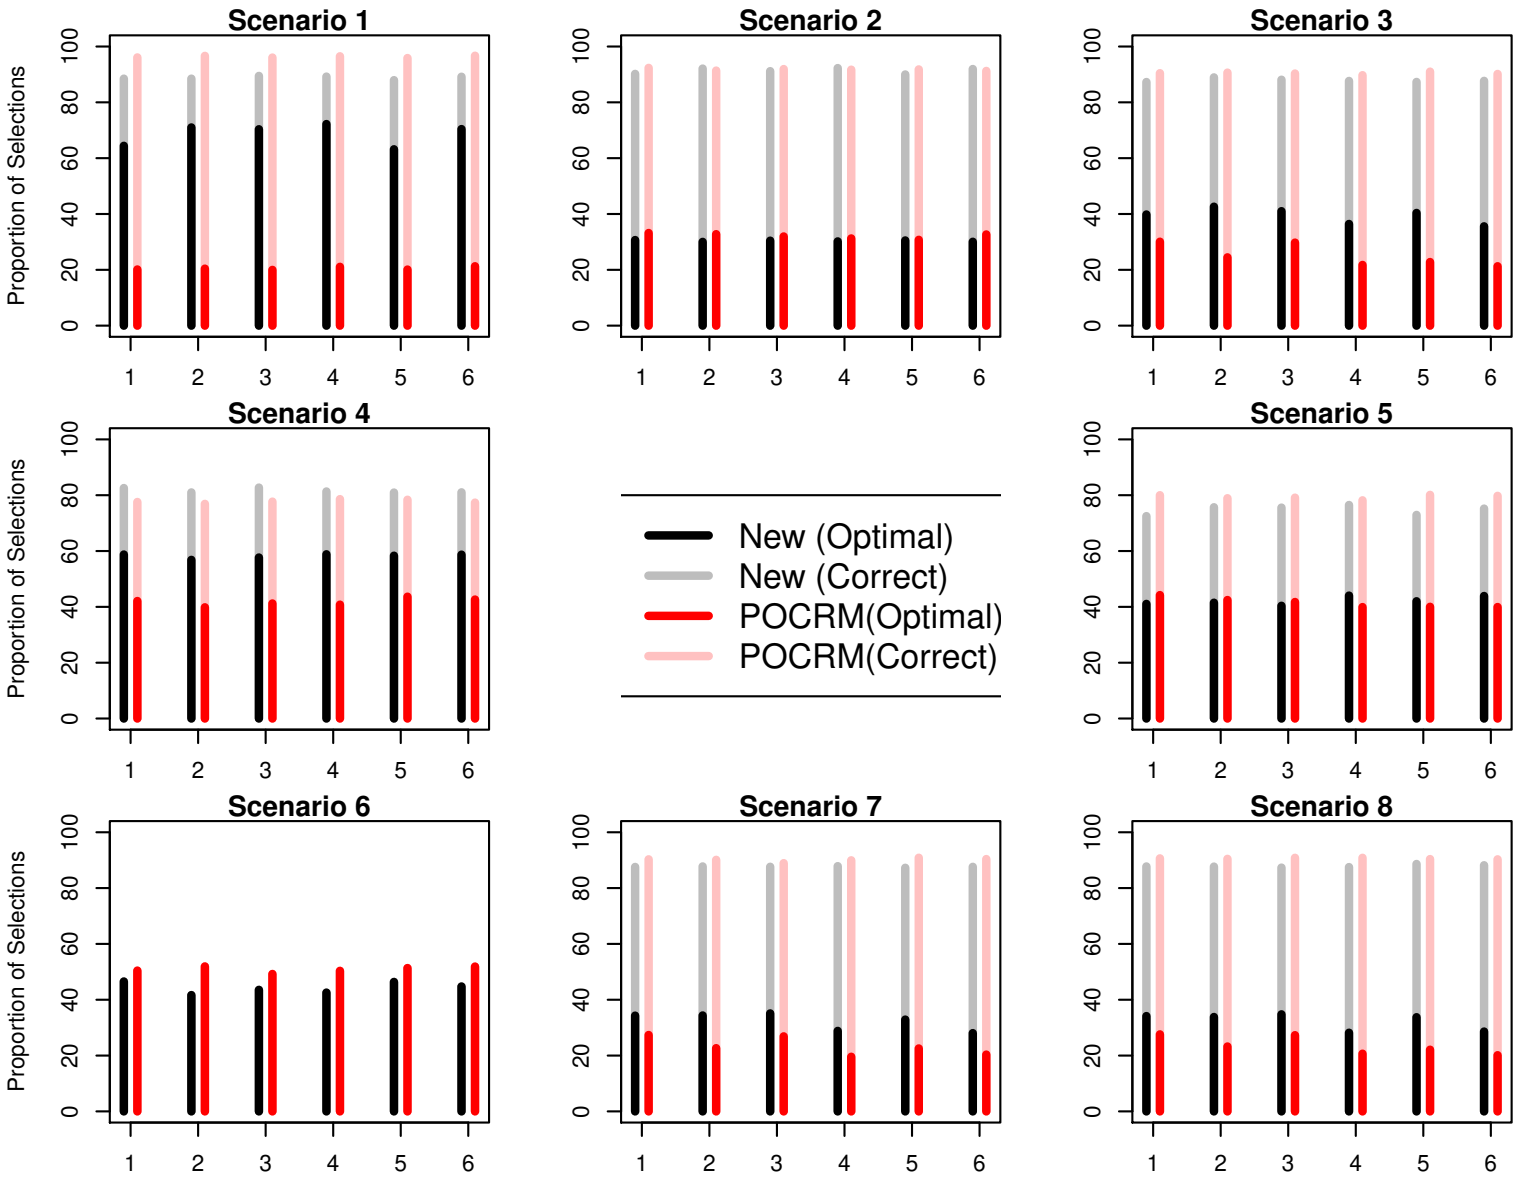
Scenario 1
Scenario 2
Scenario 3
Scenario 4
Scenario 5
Scenario 6
Scenario 7
Scenario 8
Proportion of Selections
0
20
40
60
80
100
1
2
3
4
5
6
New (Optimal)
New (Correct)
POCRM(Optimal)
POCRM(Correct)

Conclusions - dual endpoint

- The intuitively **clear** and **simple** trade-off function
- Performs **comparably or better** than model-based alternatives in majority of scenarios
- **Robust** to true ordering
- Results in **fewer** toxicities and **comparable** number of efficacies

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2

Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.

10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2
Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.
10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

$$(\hat{p}_i - \gamma)^2$$

The next patient can be assigned to either of doses, but one can argue that doses are not 'equal' for two reasons.

Motivation

Consider a dose-finding trial with binary responses and two doses: d_1, d_2
Goal is to find the maximum tolerated dose (MTD): $\gamma = 0.30$.
10 patients were assigned to each dose, 2 and 4 toxicities observed

Q: Which dose should be administered to the next patient?

$$(\hat{p}_i - \gamma)^2$$

The next patient can be assigned to either of doses, but one can argue that doses are not 'equal' for two reasons.

1. The squared distance ignores the randomness of the estimates.

$$\mathbb{P}\left(p_2 \in (0.25, 0.35)\right) > \mathbb{P}\left(p_1 \in (0.25, 0.35)\right).$$

2. $\hat{p}_2 = 0.4$ is an unacceptably high toxicity.

Motivation

It is usually of interest to balance both aims in a Phase I clinical trial

Current solutions

Safety:

Escalation with Overdose Control (EWOC) design (Babb et al., 1998):

$$\mathbb{E}\left(\alpha(\gamma - P_i)^+ + (1-\alpha)(P_i - \gamma)^+\right)$$

+ Low average number of DLTs
− Underestimation of the MTD
- Modifications: α_n by Tighiouart et al. (2010) and Wheeler et al. (2017)

Safety & Uncertainty

Bayesian Logistic Regression Model (Neuenschwander et al., 2008).
uses the distribution of DLT probabilies. For example, for $\gamma = 0.33$

$$L = \begin{cases} 1 \text{ if } p \in (0.00, 0.26); & 0 \text{ if } p \in (0.26, 0.41); \\ 1 \text{ if } p \in (0.41, 0.66); & 2 \text{ if } p \in (0.66, 1.00) \end{cases}$$

Goal

We propose a new criterion for selecting doses in dose-escalation trials that accounts for

1. Uncertainty in the estimates
2. Ethical constraints

and requires only **one additional parameter** to be specified.

Goal

We propose a new criterion for selecting doses in dose-escalation trials that accounts for

1. Uncertainty in the estimates
2. Ethical constraints

and requires only **one additional parameter** to be specified.

We incorporate the proposed criterion to the one-parameter Bayesian continual reassessment method

Novel Criterion

The main object of estimation is the probability of DLT $p_i \in (0, 1)$

We propose a distance satisfying the desirable properties

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p(1 - p)}.$$

- $\delta(\cdot) = 0$ at $p = \gamma$
- $\delta(\cdot) \to \infty$ as $p \to 0$ or $p \to 1$
- The variance in denominator (Criterion is a score statistic)

Novel Criterion

The main object of estimation is the probability of DLT $p_i \in (0, 1)$

We propose a distance satisfying the desirable properties

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p(1 - p)}.$$

- $\delta(\cdot) = 0$ at $p = \gamma$
- $\delta(\cdot) \to \infty$ as $p \to 0$ or $p \to 1$
- The variance in denominator (Criterion is a score statistic)

In the illustration example above

$$\delta(\hat{p}_1 = 0.2, \gamma = 0.3) = 1/16 \quad \text{and} \quad \delta(\hat{p}_2 = 0.4, \gamma = 0.3) = 1/24$$

(!) Single point estimate summarizes the information about uncertainty.

Introducing safety compound

The target toxicity γ is always less than 0.5.

Then for estimates $\hat{p}_1 = \gamma - \theta$ and $\hat{p}_2 = \gamma + \theta$, symmetric criterion favours $\hat{p}_2$.

Introducing safety compound

The target toxicity γ is always less than 0.5.
Then for estimates $\hat{p}_1 = \gamma - \theta$ and $\hat{p}_2 = \gamma + \theta$, symmetric criterion favours $\hat{p}_2$.

We introduce an asymmetry parameter a:

$$\delta(p, \gamma) = \frac{(p - \gamma)^2}{p^a(1 - p)^{2-a}}.$$

$0 < a < 1$ implies more severe penalty for more toxic doses.

(!) Selection of under toxic doses remain to be undesirable as well.

In the illustration example above, for $a = 0.5$

$$\delta(\hat{p}_1 = 0.2, \gamma = 0.3, a = 0.5) < \delta(\hat{p}_2 = 0.4, \gamma = 0.3, a = 0.5).$$

Bayesian continual reassessment method

- DLT probability modelled as $p(d_i) = d_i^{\exp(\beta)}$
- $\beta \sim N(0, 1.34)$
- Then, the dose d_k minimising

$$\mathbb{E}\left(\frac{(p(d_i) - \gamma)^2}{p(d_i)^a(1 - p(d_i))^{2-a}}\right)$$

among all $d_1, \ldots, d_m$ is recommended for the next group of patients

Bayesian continual reassessment method

- DLT probability modelled as $p(d_i) = d_i^{\exp(\beta)}$
- $\beta \sim N(0, 1.34)$
- Then, the dose d_k minimising

$$\mathbb{E}\left(\frac{(p(d_i) - \gamma)^2}{p(d_i)^a(1 - p(d_i))^{2-a}}\right)$$

among all $d_1, \ldots, d_m$ is recommended for the next group of patients

Convex Infinite Bounds Penalization with parameter a as CIBP(a).

Numerical Study

Setting by [Wheeler et al.(2017)Wheeler, Sweeting and Mander].

- $n = 40$ patients; $m = 6$ doses; $c = 1$ cohort size; target $\gamma = 0.33$
- $\beta \sim \mathcal{N}(0, 1.34)$
- $a = \{0.5, 0.25, 0.10\}$.

We study the performance of designs in terms of
(i) *Accuracy*

$$\mathcal{A} = 1 - m\frac{\sum_{i=1}^{m}(p_i - \gamma)^2 \pi_i}{\sum_{i=1}^{m}(p_i - \gamma)^2}$$

(*ii*) mean number of toxic responses (DLTs) and focus on the mean performance.

Scenarios

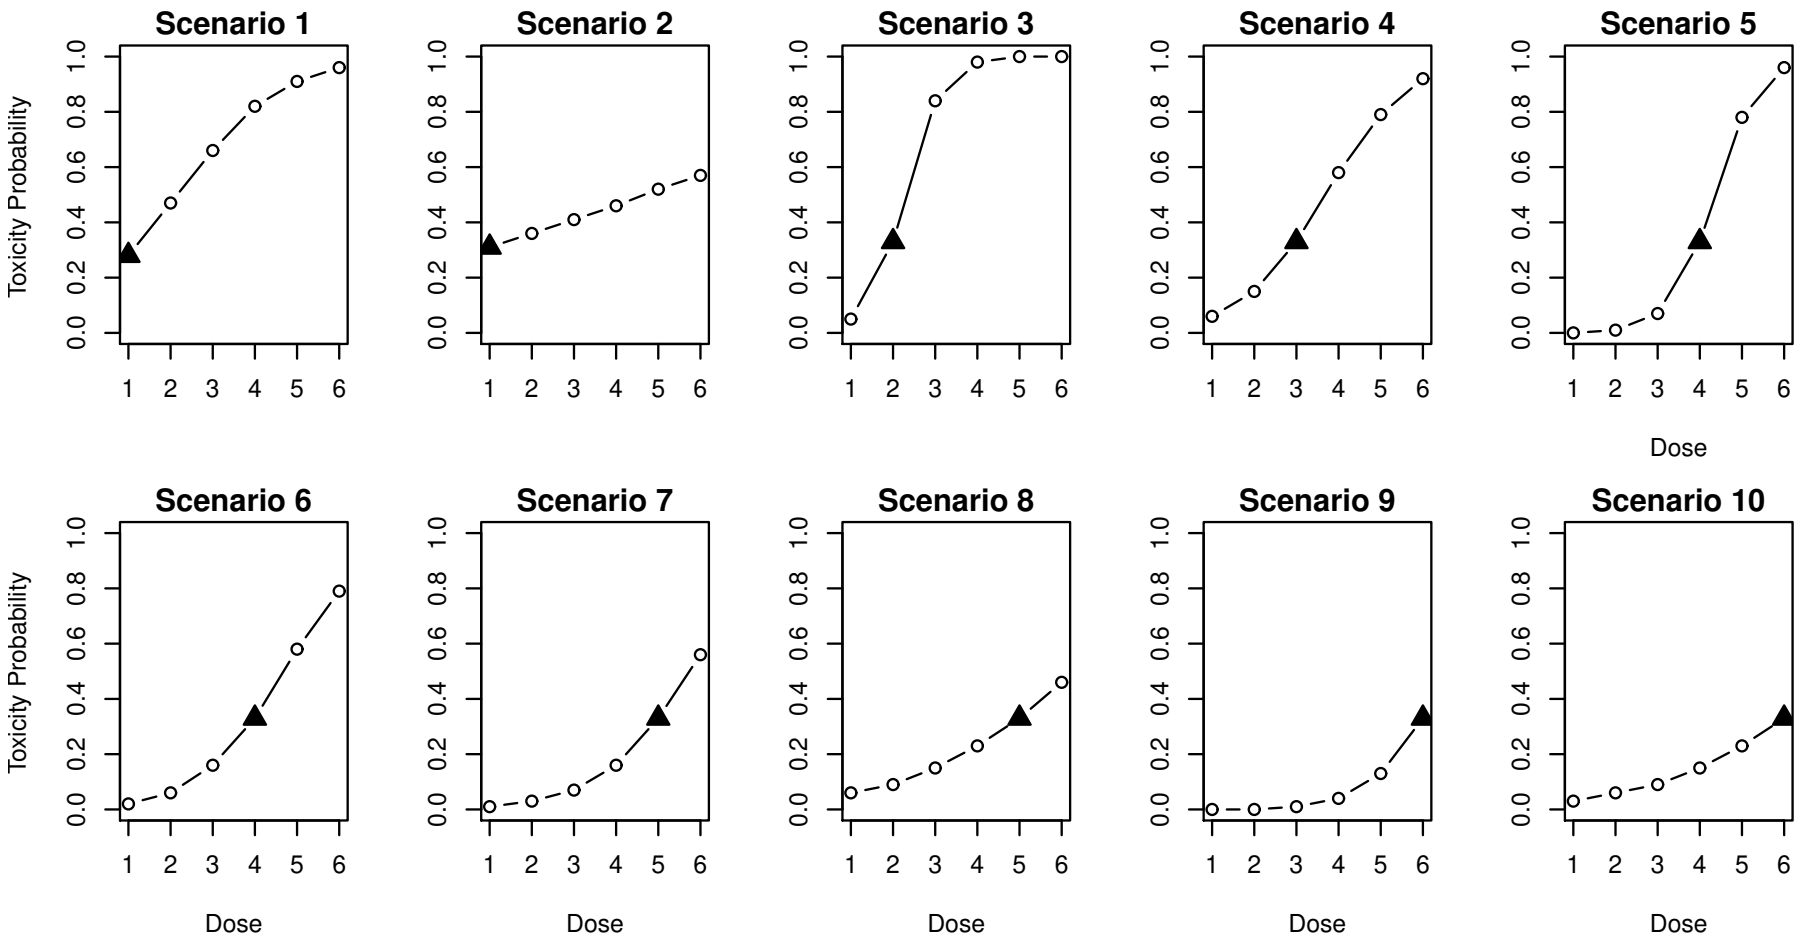
Scenario 1
Scenario 2
Scenario 3
Scenario 4
Scenario 5
Scenario 6
Scenario 7
Scenario 8
Scenario 9
Scenario 10
Toxicity Probability
Dose
0.0
0.2
0.4
0.6
0.8
1.0
1
2
3
4
5
6

Comparators

We compare the performance of the proposed approach to

- **EWOC**
- **TR** design by Tighiouart et al. (2010)
- Toxicity-dependent feasibility bound (**TDFB**) by Wheeler et al. (2017)
- **BLRM** by Neuenschwander et al. (2008)

Results. Accuracy

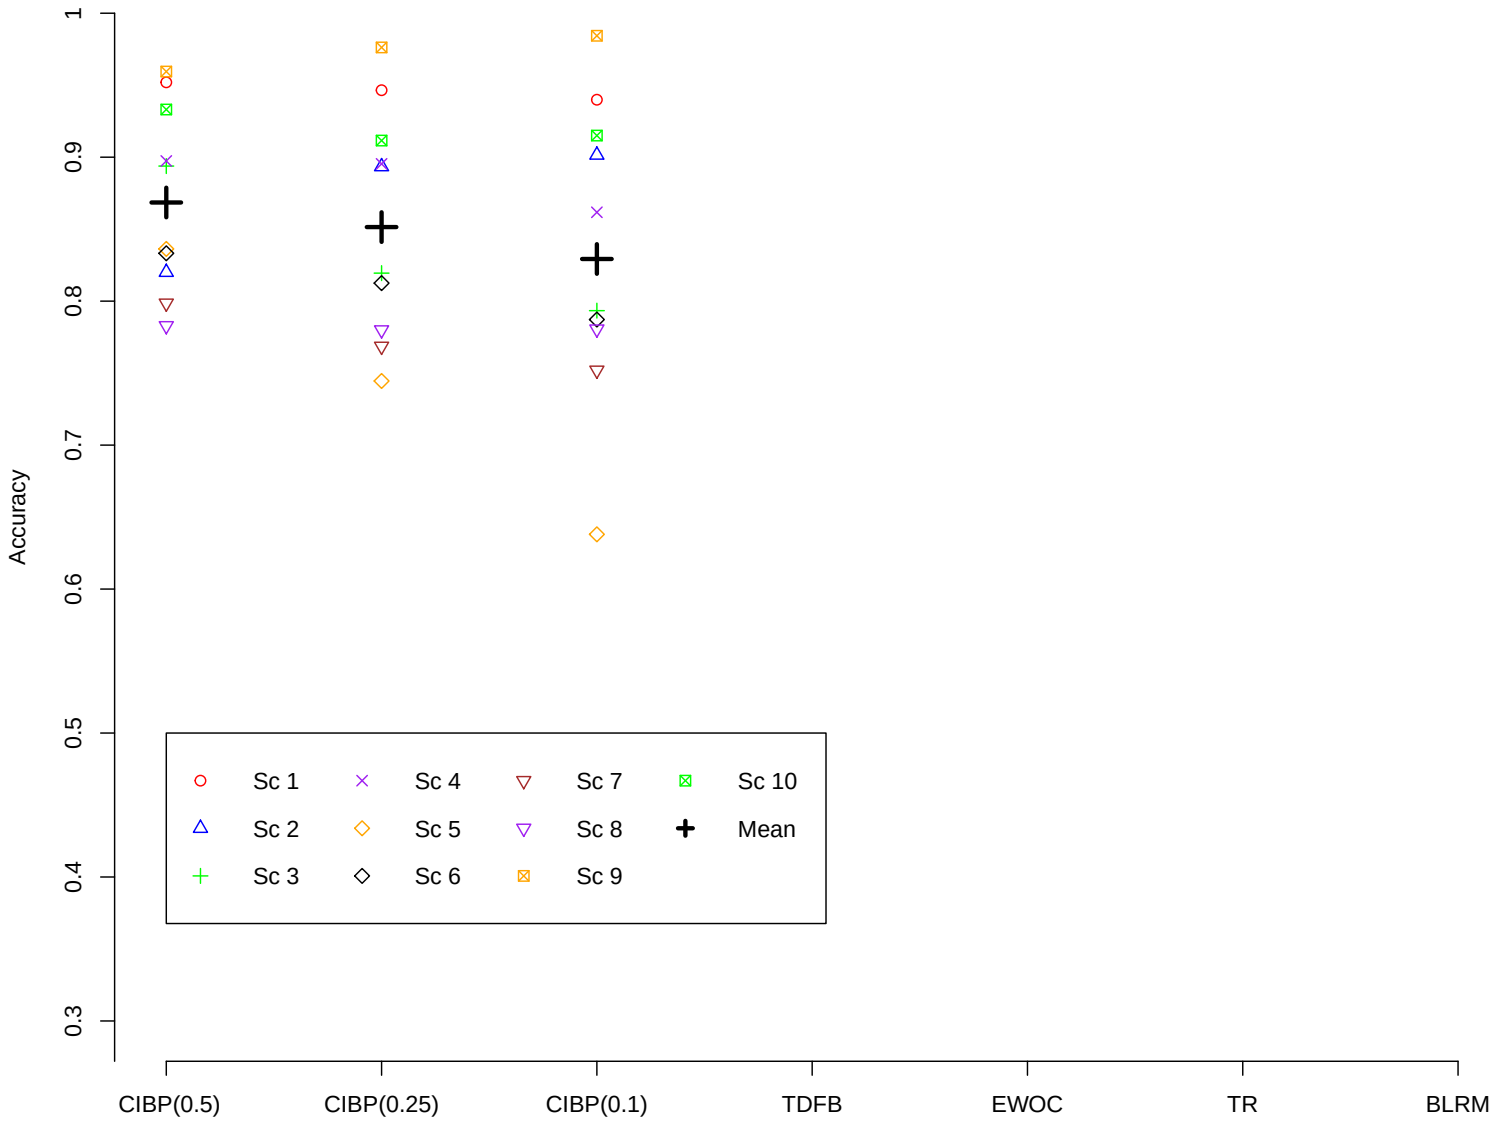

Results. Accuracy

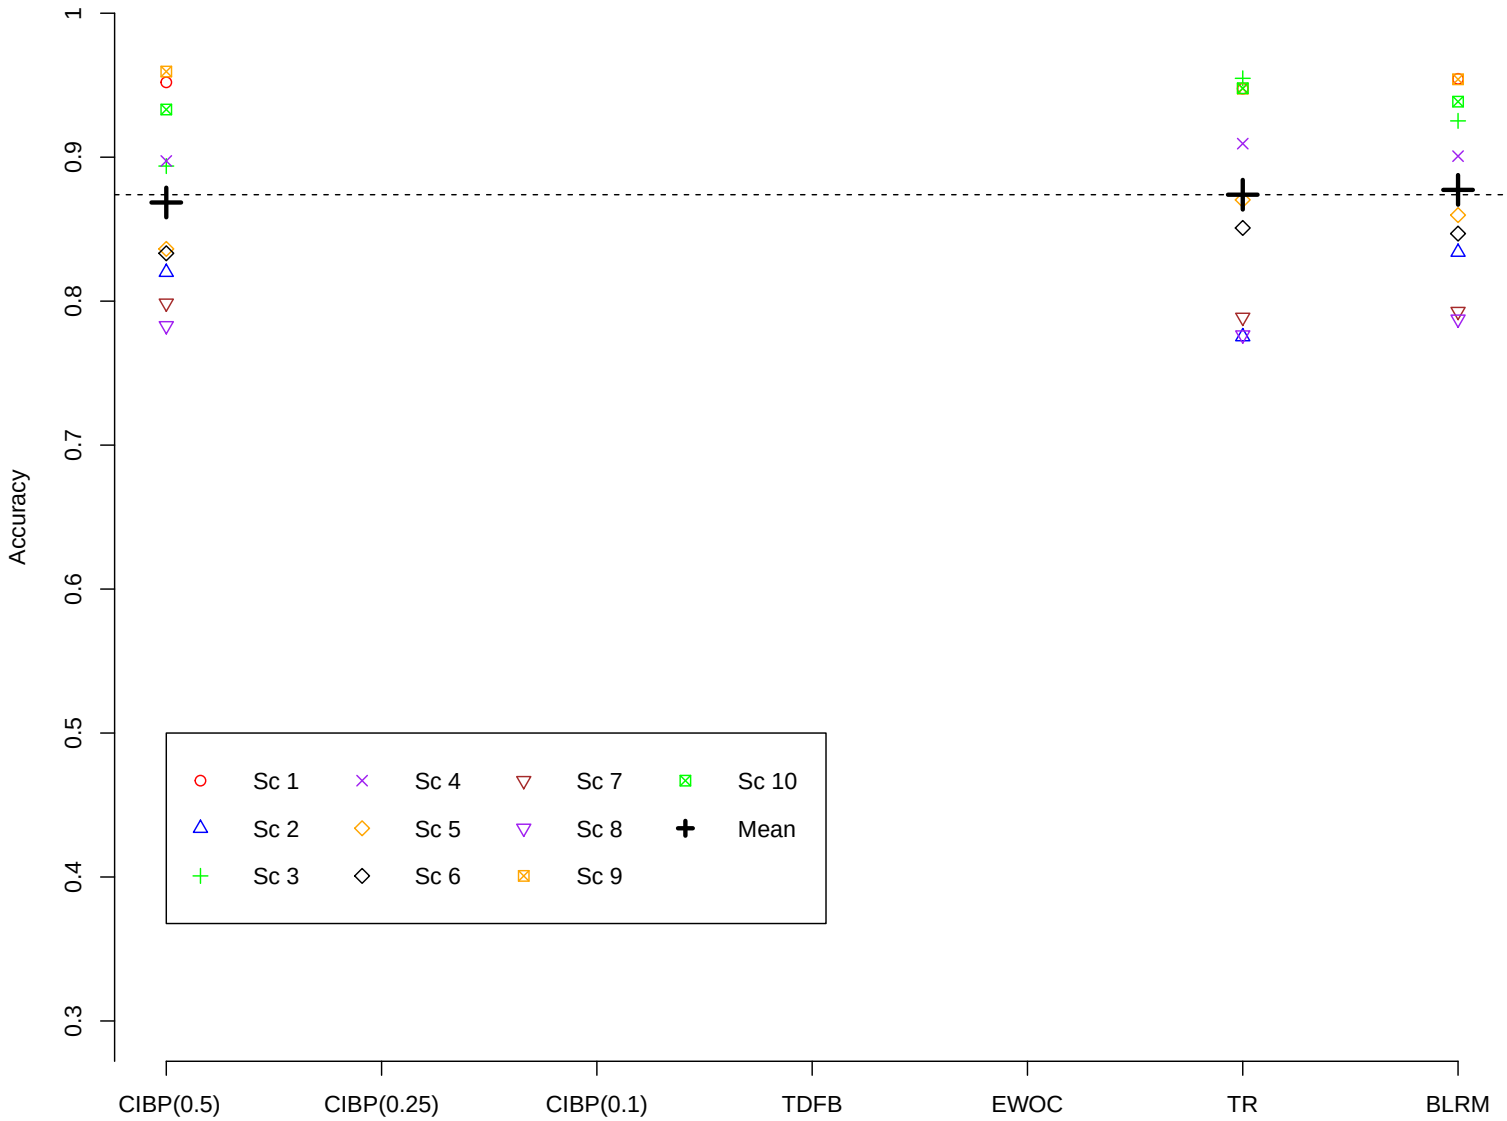

Results. Accuracy

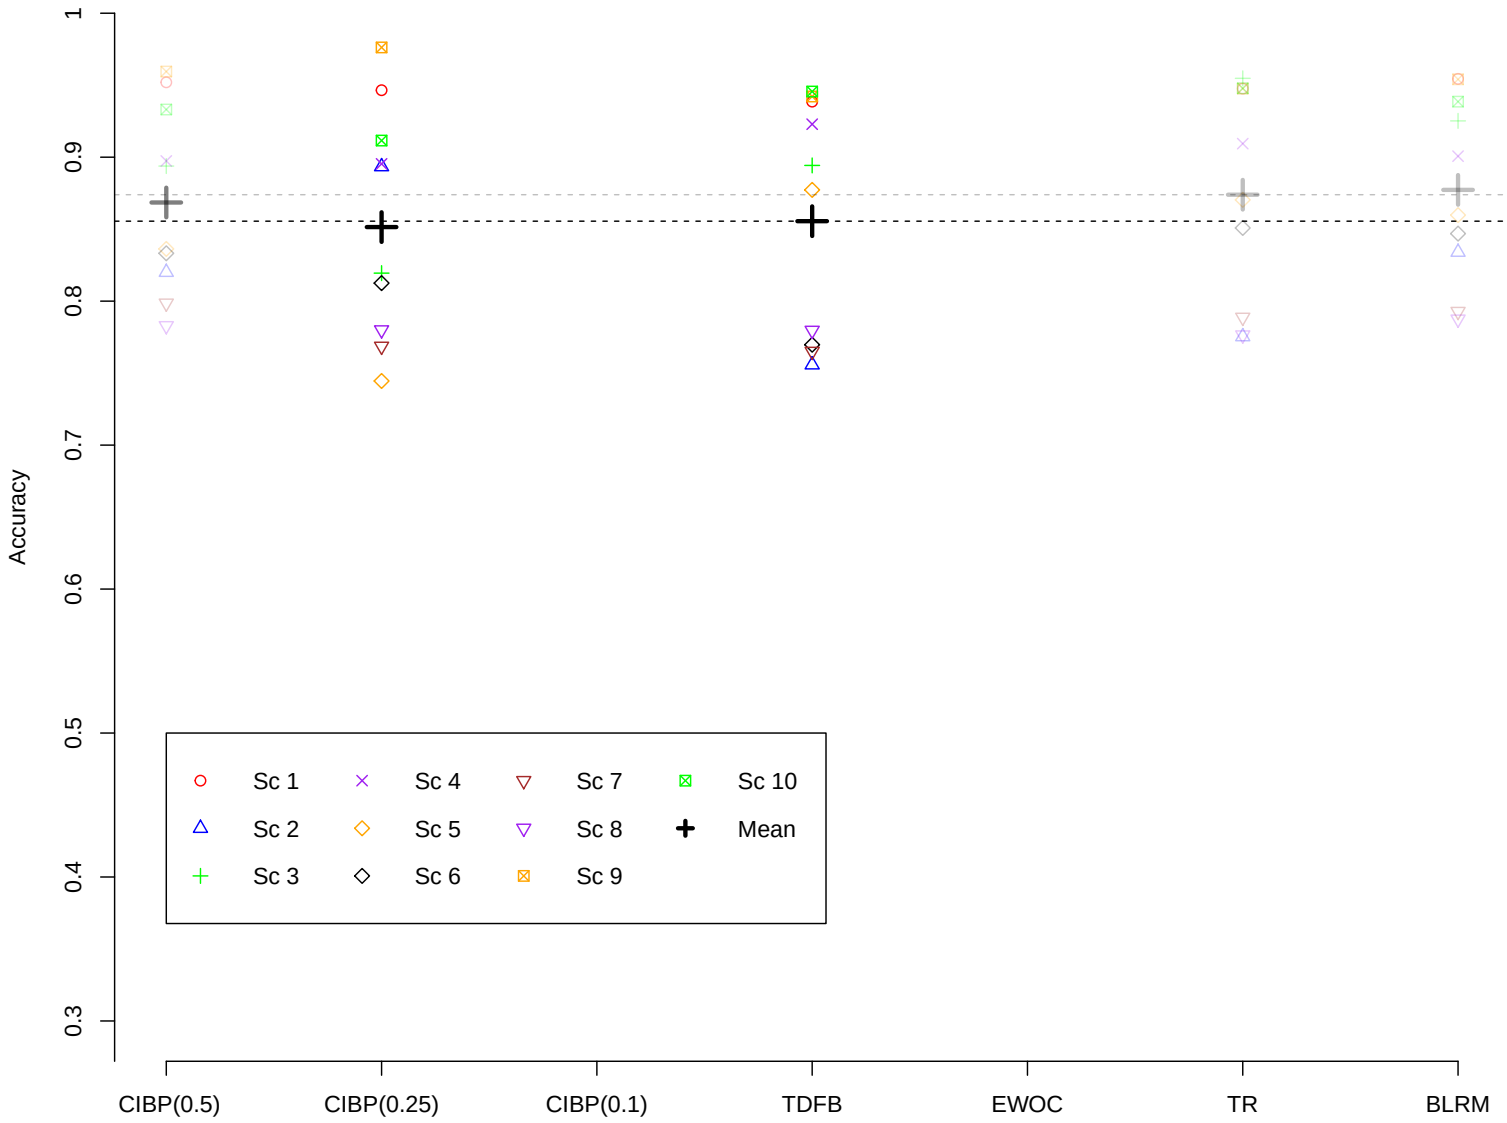

Results. Accuracy

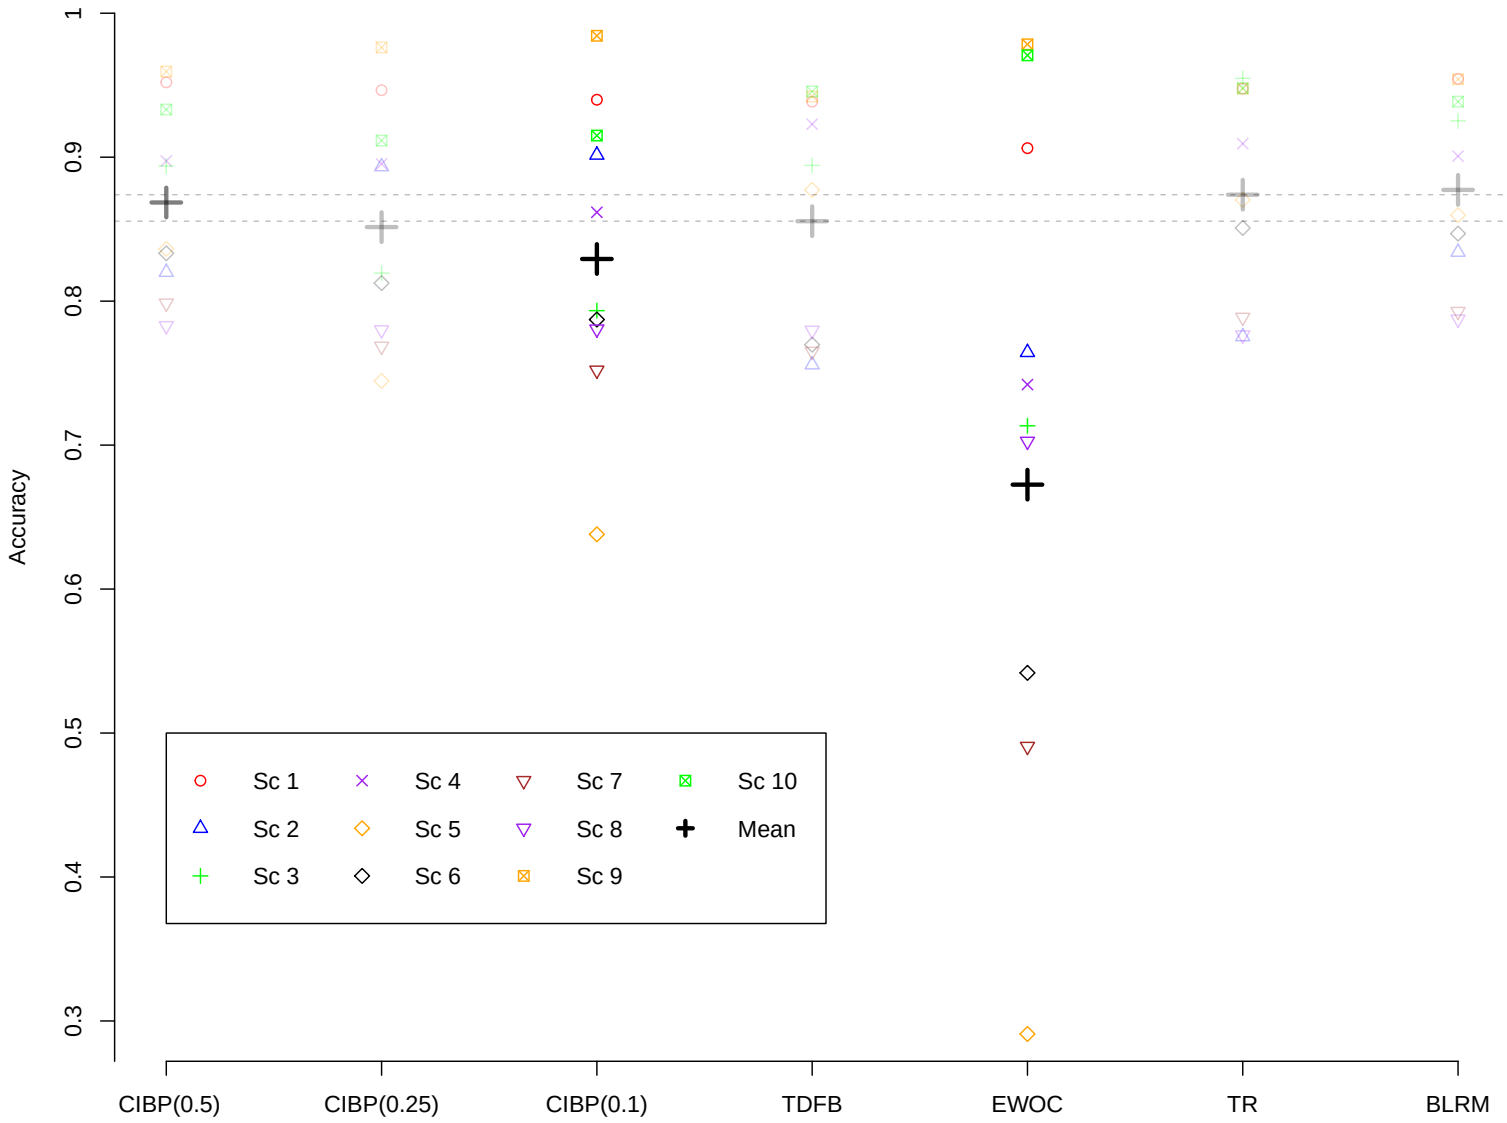

Results.DLTs

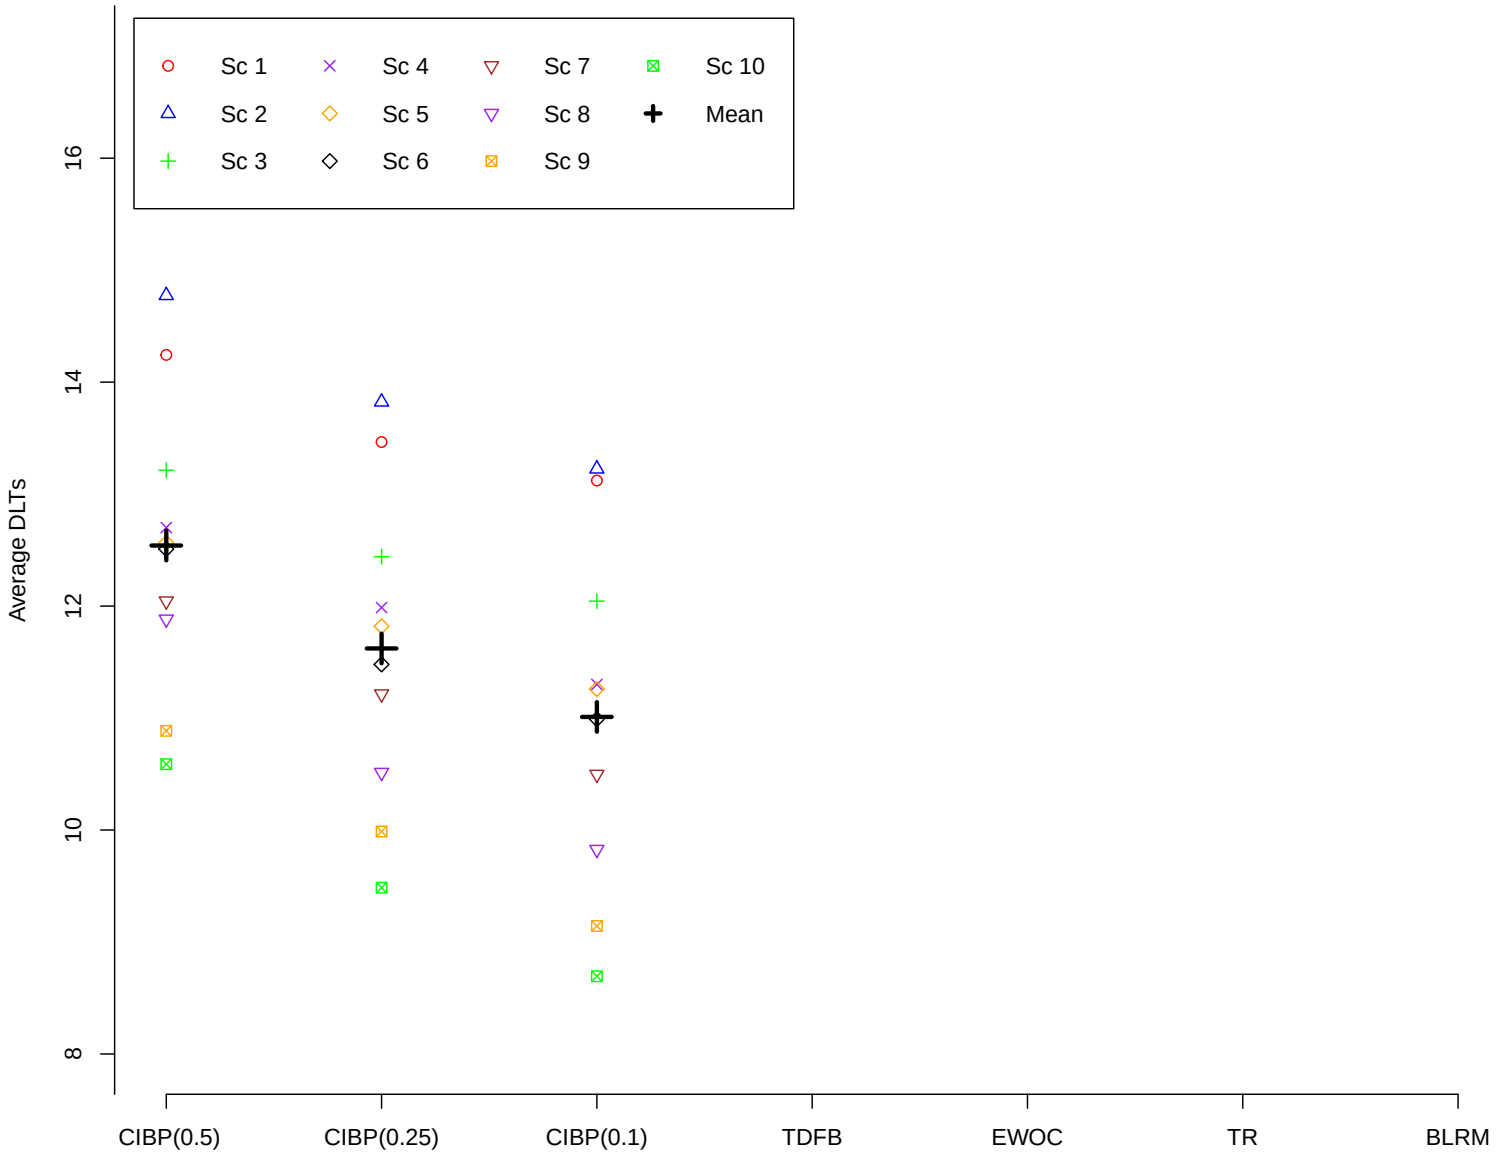

Results.DLTs

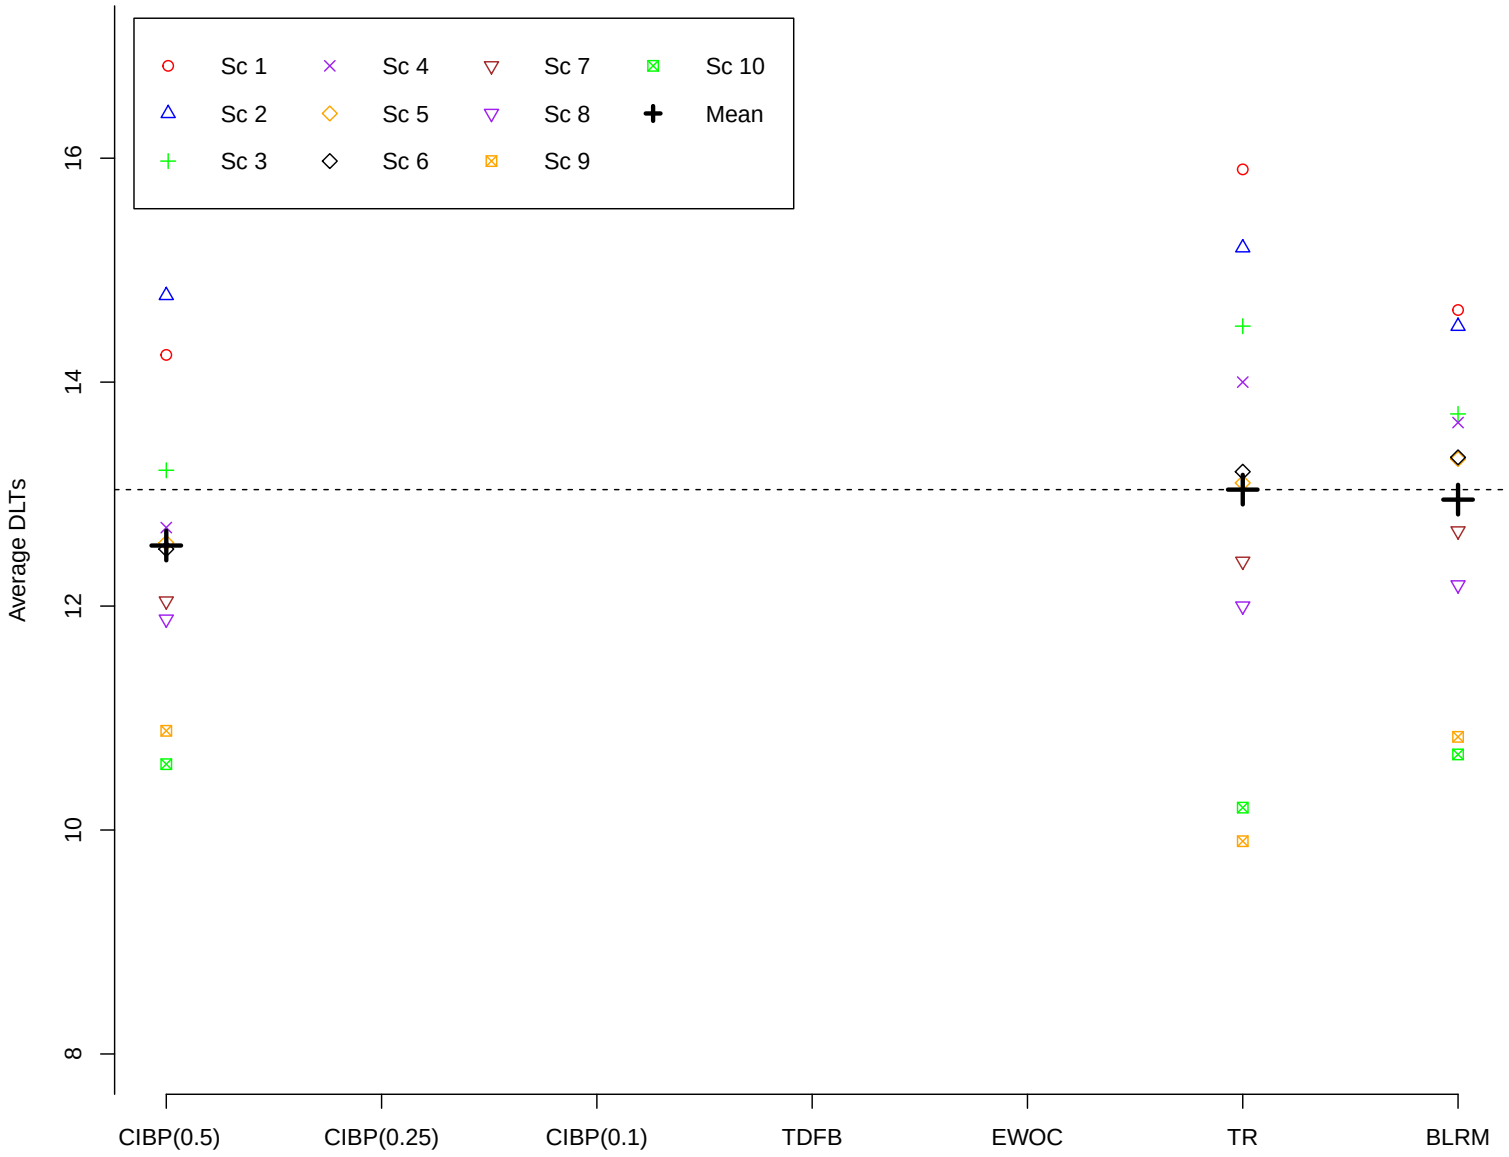

Results.DLTs

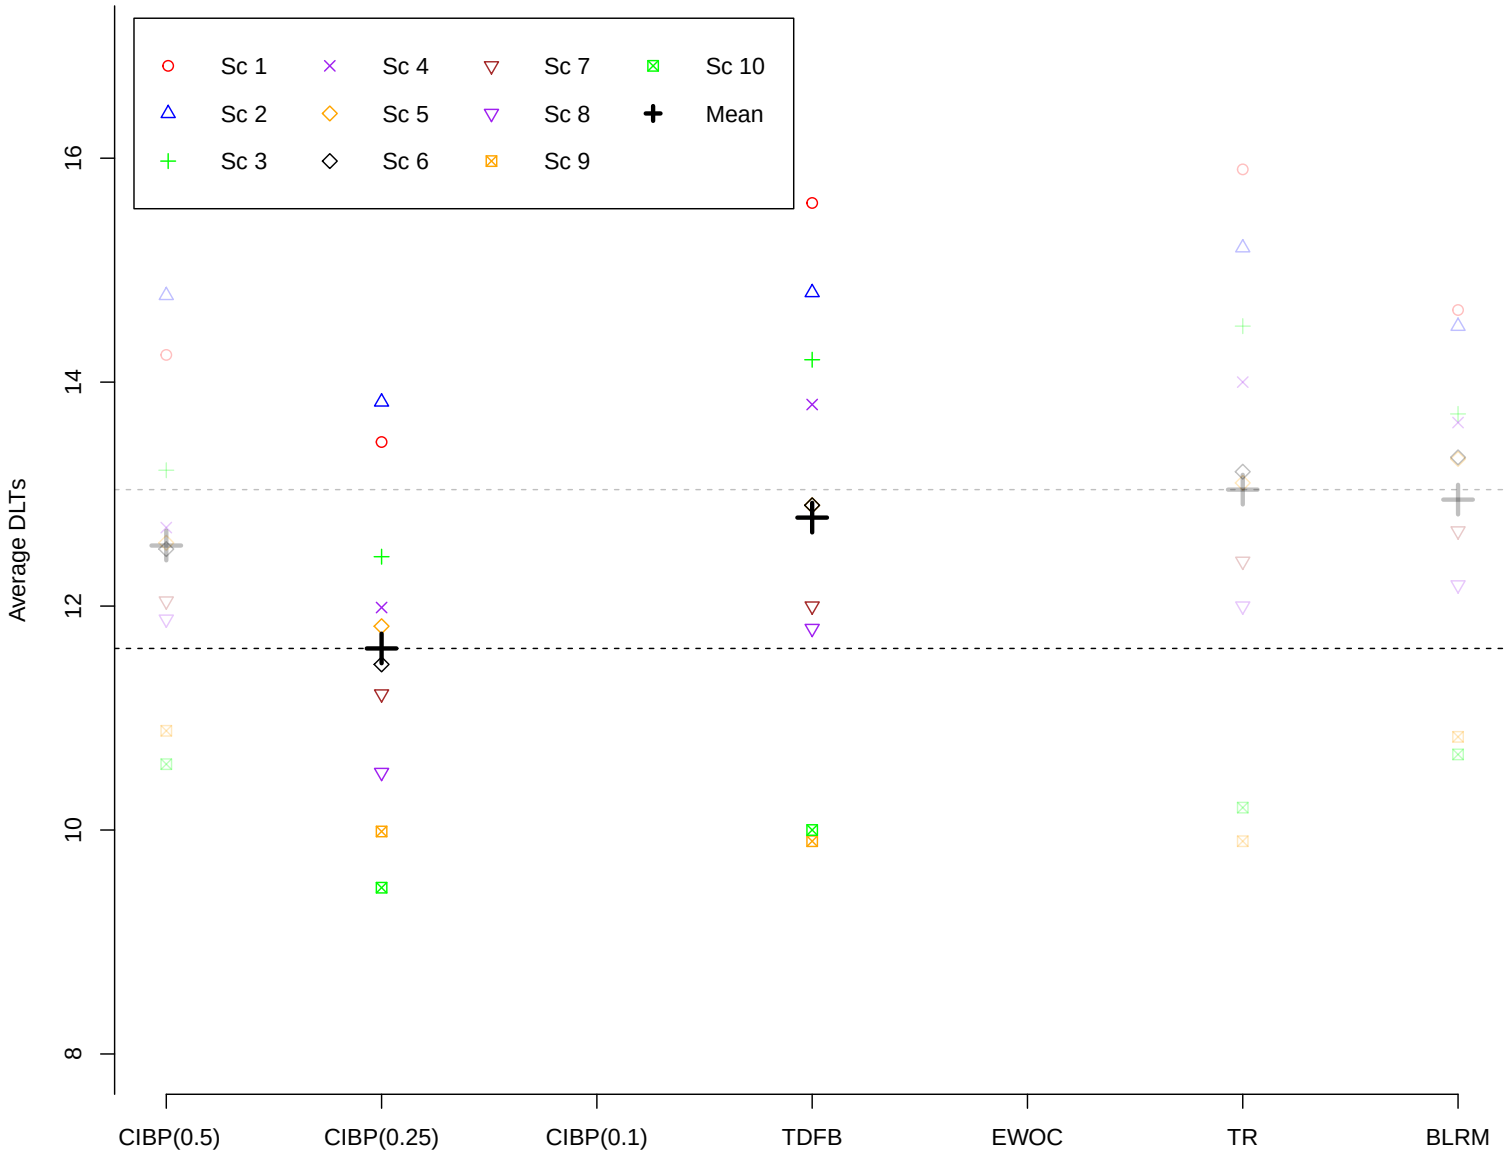

Results.DLTs

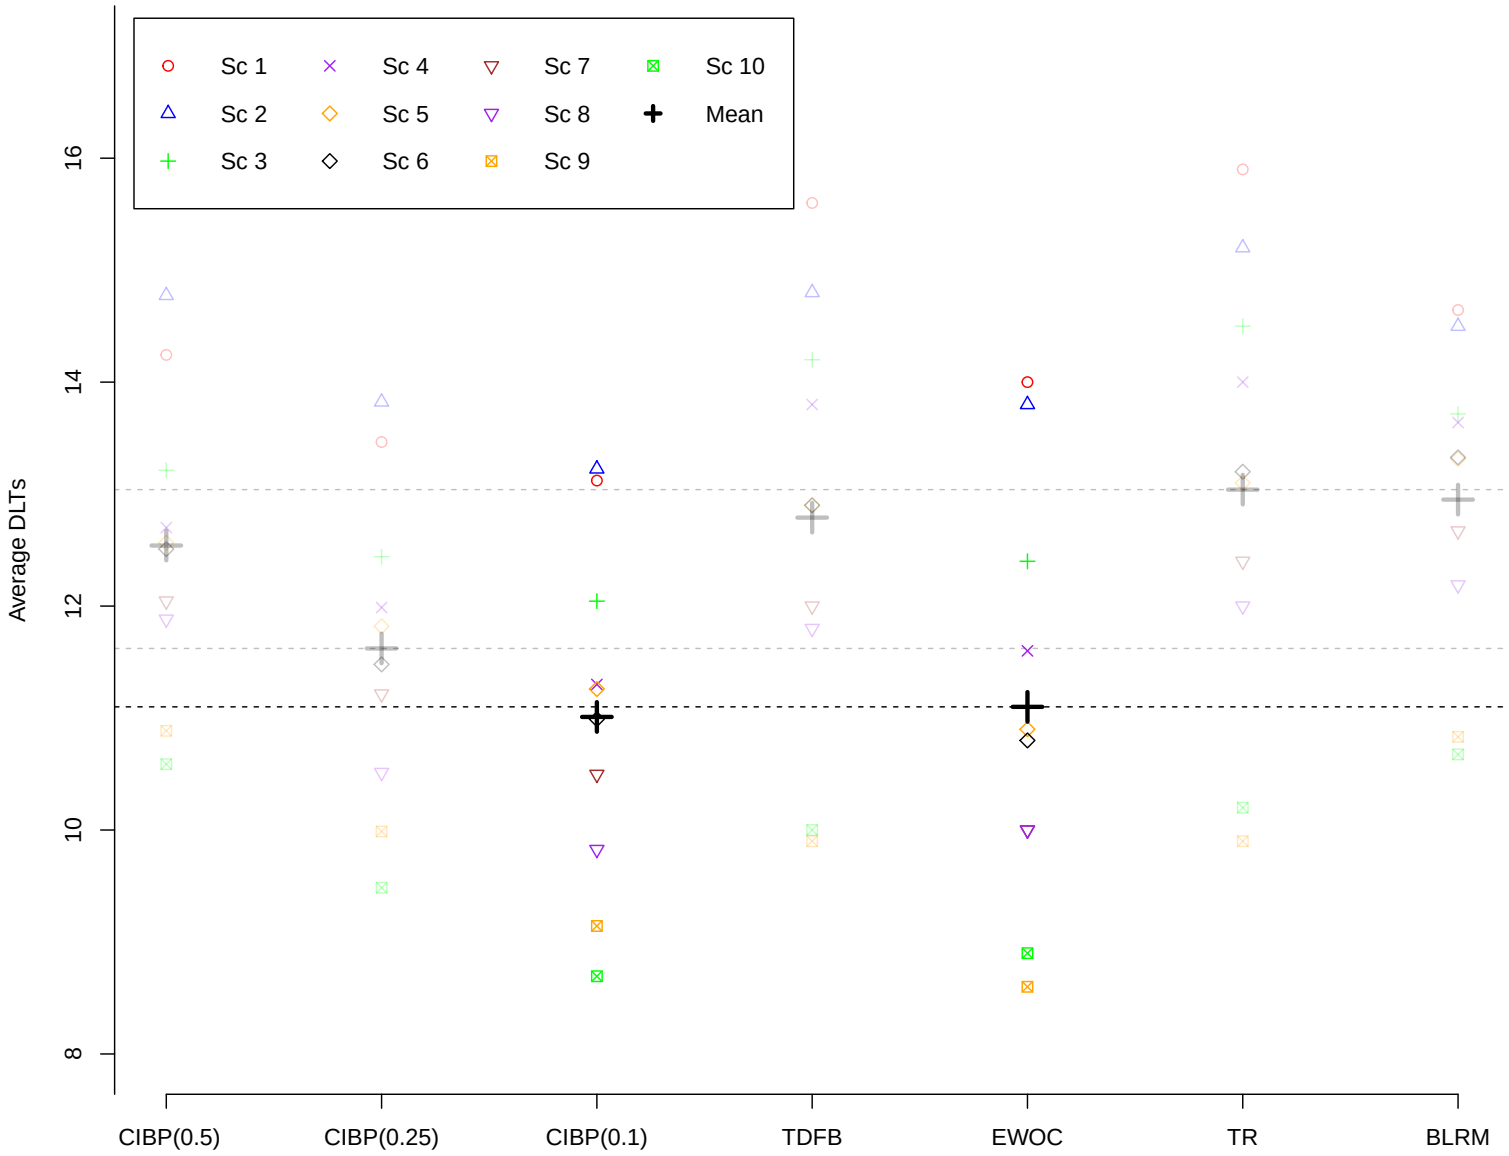

Conclusions - Safety

- The novel criterion requires **one additional parameter only**.
- The criterion incorporated into the one-parameter CRM method is found to result in
  1. **Similar** accuracy, but **fewer** mean number of DLTS.
  2. **Greater** accuracy, but **similar** mean number of DLTs.
- The new criterion allows to make model-based design **more ethical** as it does not lead to any decrease in accuracy.

Discussion

- Information theory can be useful in dose-finding
- Coherent framework with little tuning necessary
- Useful in itself or in combination with traditional model based ideas

References

J. Babb, A. Rogatko, S. Zacks. *Cancer phase I clinical trials: efficient dose escalation with overdose control*. (1998). Statistics in Medicine, 17(10), 1103–20.

Mozgunov and Jaki (2018) An information-theoretic phase i/ii design for molecularly targeted agents that does not require an assumption of monotonicity. *JRSS C (Applied Statistics)*, **68**, 1–24, Epub.

Neuenschwander, B., Branson, M. and Gsponer, T. (2008) Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in Medicine*, **27**, 2420–2439.

J. O'Quigley, M. Pepe, L. Fisher, *Continual reassessment method: A practical design for phase I clinical trials in cancer*, 1990, Biometrics 46 33–48.

O'Quigley J, Paoletti X, MacCario J., *Non-parametric optimal design in dose finding studies*, (2002) Biostatistics; 3: 51–56.

Riviere, M.-K., Yuan, Y., Jourdan, J.-H., Dubois, F. and Zohar, S. (2016) Phase I/II dose-finding design for MTA: Plateau determination. *Stat. Meth. in Medical Research*, **27**, 466–479.

Tighiouart, M., Rogatko, A. et al. (2010) Dose finding with escalation with overdose control (ewoc) in cancer clinical trials. *Statistical Science*, **25**, 217–226.

Wages N., Conaway M., O'Quigley J. (2011a). *Continual reassessment method for partial ordering. Biometrics* **67(4)**, 1555-1563.

Wages, N. A. and Tait, C. (2015) Seamless phase i/ii adaptive design for oncology trials of molecularly targeted agents. *Journal of Biopharmaceutical Statistics*, **25**, 903–920.

Wheeler, G. M., Sweeting, M. J. and Mander, A. P. (2017) Toxicity-dependent feasibility bounds for the escalation with overdose control approach in phase I cancer trials. *Statistics in Medicine*.

Mozgunov, P and Jaki T (2017) An information-theoretic approach for selecting arms in clinical trials.. *arXiv preprint*. arXiv:1708.02426 (2017).