Bayesian Designs for Phase I-II Clinical Trials

Peter F. Thall and Ying Yuan
Department of Biostatistics
M.D. Anderson Cancer Center

A Half-Day Short Course SC3

2020 ASA Biopharmaceutical Section

Regulatory-Industry Statistics Workshop

Tuesday, September 22, 2020

Textbook: Yuan, Nguyen, and Thall

Bayesian Designs for Phase I-II Clinical Trials

Chapman & Hall/CRC Biostatistics Series, 2016

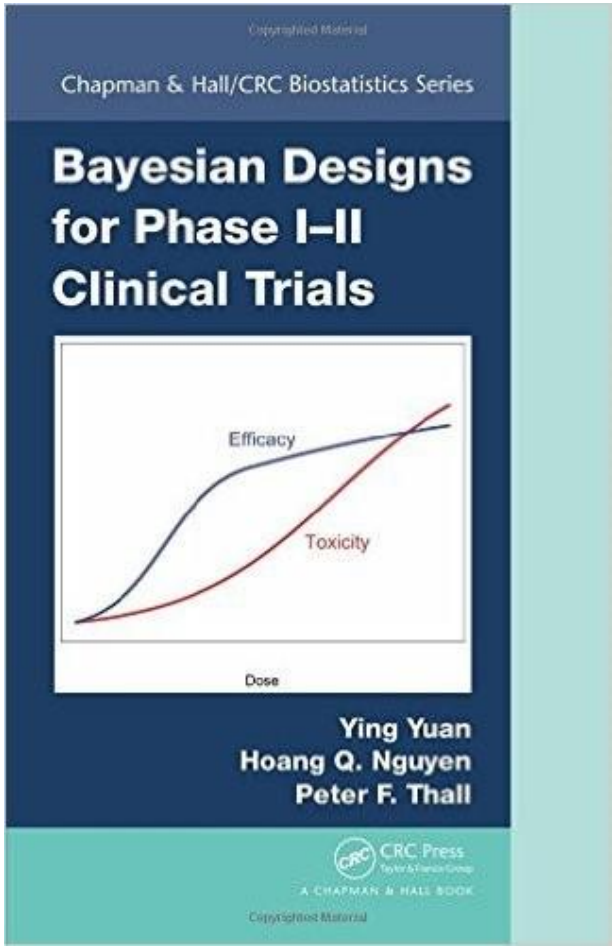

Lecture Schedule (Eastern Standard Time)

Lecture 1: 2:00 – 2:50 **Thall**

- Problems with the Phase I → Phase II paradigm
- Phase I-II designs

(5 minute break)

Lecture 2: 2:55 – 3:40 **Thall**

- The EffTox phase I-II design
- Utility based phase I-II designs

(10 minute break)

Lecture 3: 3:50 – 4:40 **Yuan**

- Model assisted phase I-II designs
- Phase I-II designs for immunotherapies

(5 minute break)

Lecture 4: 4:45 – 5:30 **Yuan**

- Tissue agnostic phase I-II designs
- Handling late onset toxicity and efficacy

The Conventional Phase I → Phase II Paradigm

Phase I

Goal: Determine a "safe" dose (or MTD = maximum tolerated dose) of an experimental agent, ***which may or may not have anti-disease efficacy,*** for use in later phase II or phase III trials

- Do this based on DLT = Dose Limiting Toxicity, usually a binary indicator of one or more specific adverse events within a given follow up period
- Choose doses for successive cohorts of 1, 2, or 3 patients.
- Usually done using a "3+3" algorithm → Escalate until "Too much Toxicity" then de-escalate, **and never re-escalate**
- Sometimes done using a variant of the Continual Reassessment Method (CRM)
- Ignore Efficacy in the dose-finding algorithm
- Often, treat a large "expansion cohort" at the selected MTD

The Conventional Phase I → Phase II Paradigm

Phase II

Goal : Determine whether the new agent, administered at the MTD chosen in phase I, is sufficiently "promising" to motivate a large randomized phase III trial

- Do this based on the probability of "Response", usually a binary indicator of an Efficacy event, compared to some fixed "standard" or "null" response probability p_0 , often using one of the Simon (1989) 2-stage designs.
- Ignore Toxicity in the design.
- Have a Data Monitoring Committee to review the data and keep an eye on adverse events (AEs), **but do not specify any formal, objective safety rules for stopping the trial early if the observed AE rate is to high.**

Common Protocol Description of "the" 3+3 Algorithm

[Number of patients with DLT] / [Number of patients evaluated] at a given dose level	Action
0 / 3	Treat 3 pats at the next higher dose level (Escalate)
1 / 3	Treat at least 3 more pats at the current dose level: • If 0/3 DLTs → Escalate • If ≥ 1/3 DLTs → De-escalate
≥ 2/3	Stop escalation. If only 3 pats were treated at the next lower dose, treat 3 more at that dose.

MTD: The highest dose at which ≤ 1/6 pats had DLTs. 6 pats must be treated at a dose before it is declared the MTD

Logical Problems with this 3+3 Algorithm

1. If a MTD does not exist, the algorithm does not say what to do. E.g. if 2/3 DLTs are seen at the lowest dose, or 0/6 DLTs are seen at the highest dose.
2. The "≥2" in the left column is ambiguous: 2/3, 2/6, and 2/9 have very different meanings.
3. Absence of a stopping rule creates ambiguity:

- If you observe 0/3 DLTs at d=1, 0/3 at d=2, then 1/3 + 1/3 = 2/6 at d=3, so de-escalate to d=2, and then observe 0/3 for a total of 0/6 at d=2, should you

  → treat 3 more patients at d=2, or

  → stop and declare d=2 the MTD?

- If you treat 3 more patients at d=2, you may end up observing 0, 1, 2, or 3 DLTs in 9 patients. The algorithm does not say what to do, or what to conclude, in these cases.

Two commonly used phase I trial 3+3 algorithms.

General Rules

1. Never re-escalate to a level after de-escalating from that level
2. If decision is to de-escalate or choose one level lower but current level is lowest, stop and choose no level
3. If decision is to escalate above highest level, stop and choose no level.
4. If decision is to stop and choose one level lower, but one level lower has 3 or fewer patients, treat 3 more at that lower level

# toxicities/ # patients	Decision
0/3	Escalate one level, if allowed by General Rule 1, otherwise treat 3 more at current level.
0/3 + [0/3 or 1/3]†	Stop, choose *current* level as MTD
0/3 + 2/3†	3+3 A: Stop, choose *one level lower* as MTD
0/3 + 2/3†	3+3 B: Stop, choose *current level* as MTD
0/3 + 3/3†	Stop, choose *one level lower* as MTD
1/3	Treat 3 more at *current* level
1/3 + 1/3	3+3 A: Stop, choose *one level lower* as MTD
1/3 + 1/3	3+3 B: Stop, choose *current* level as MTD
2/3 or 3/3	De-escalate one level
1/3 + 0/3	Escalate one level if allowed by General Rule 1, otherwise choose *current* level as MTD
1/3 + [2/3 or 3/3]	Stop, choose *one level lower* as MTD

† after de-escalating back to this level from a higher level

Typical Data from a Phase I Trial after 3+3

Dose mg/m^2	# Toxicities / # Patients	Posterior 95% Credible Interval
100	0 / 3	.00 — .43
200	1 / 6	.01 — .52
300	2 / 6	.06 — .69
400	—	—

Usual claim: "The MTD is 200 mg/m^2 "

Reality: These trial results all are very unreliable

- A 95% CI for Pr(Tox | d=MTD) runs from **.01 to .52**
- **Toxicity severity level** is ignored.
- **Efficacy** is ignored. What if

Pr(response | d=200) = **.25** and Pr(response | d=300) = **.50** ?

For each of the four datasets below, posterior 95% Credible Intervals (CIs) for Pr(Toxicity | MTD) all include the interval [.07 - .41]

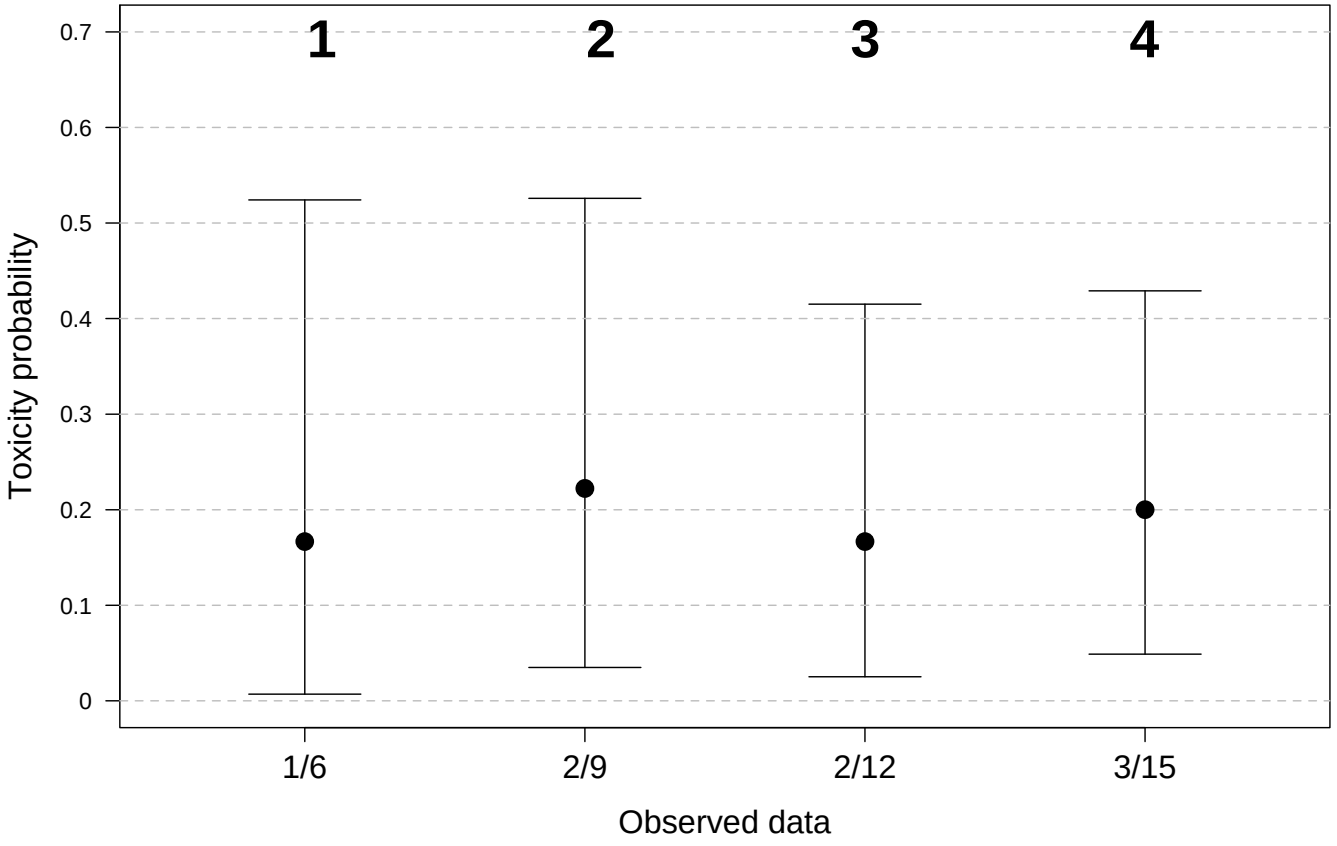

Actual Properties of 3+3 Algorithms

- Produce very small samples → Very unreliable
- Very short memory → They waste data
- Many different versions.
- Many decisions are left unspecified.
- No explicit target Pr(DLT)
- No explicit upper limit on Pr(DLT) → Likely to choose an unsafe MTD
- Ignore Efficacy → Likely to choose an ineffective MTD
- Do not allow re-escalation after de-escalation from a "toxic" dose, based on a tiny amount of data →

→ A dose above the MTD that may actually be safe and have higher **Efficacy** is likely to be missed.

Continual Reassessment Method (CRM, 1990)

1. Define a binary DLT that is scored quickly
2. Set N= maximum sample size, c=cohort size = 1, 2, or 3
3. Assume a simple model for Pr(DLT | d=dose)
4. Choose a fixed target p* = Pr(DLT)
5. For each cohort, use all (d, DLT) data to choose a dose d^{new} with E{ Pr(DLT | d^{new}) | data} closest to p*
6. When N is reached, the last choice is the "MTD"

Implicit Assumption Underlying All "Phase I Toxicity Only" Dose-Finding Designs (3+3 or CRM):

There is an Efficacy outcome for which Pr(Efficacy | d) *increases with dose. If not, then why not treat all patients at d = 0, (do not treat) to ensure Pr(DLT) = 0 ?*

Typical assumption: *Pr(PFS > t | d) ↑ in d for all t.*

Computer Simulations: 3+3 A, 3+3 B, and CRM with p* = .25 were simulated under each assumed dose-toxicity curve. 1000 trials simulated for each (curve, method) pair.

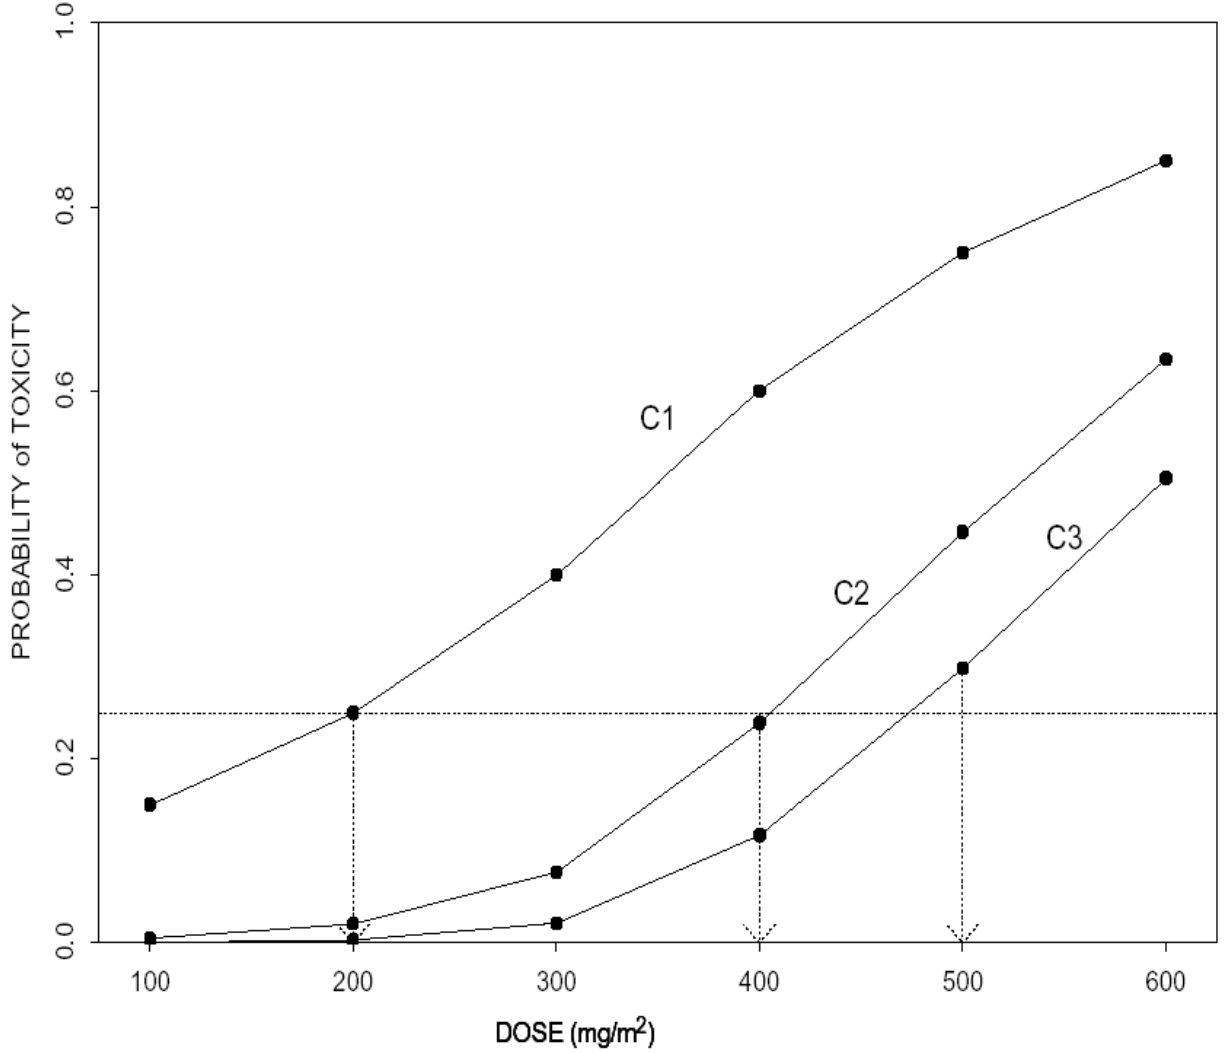

Selection Percentages Under C1

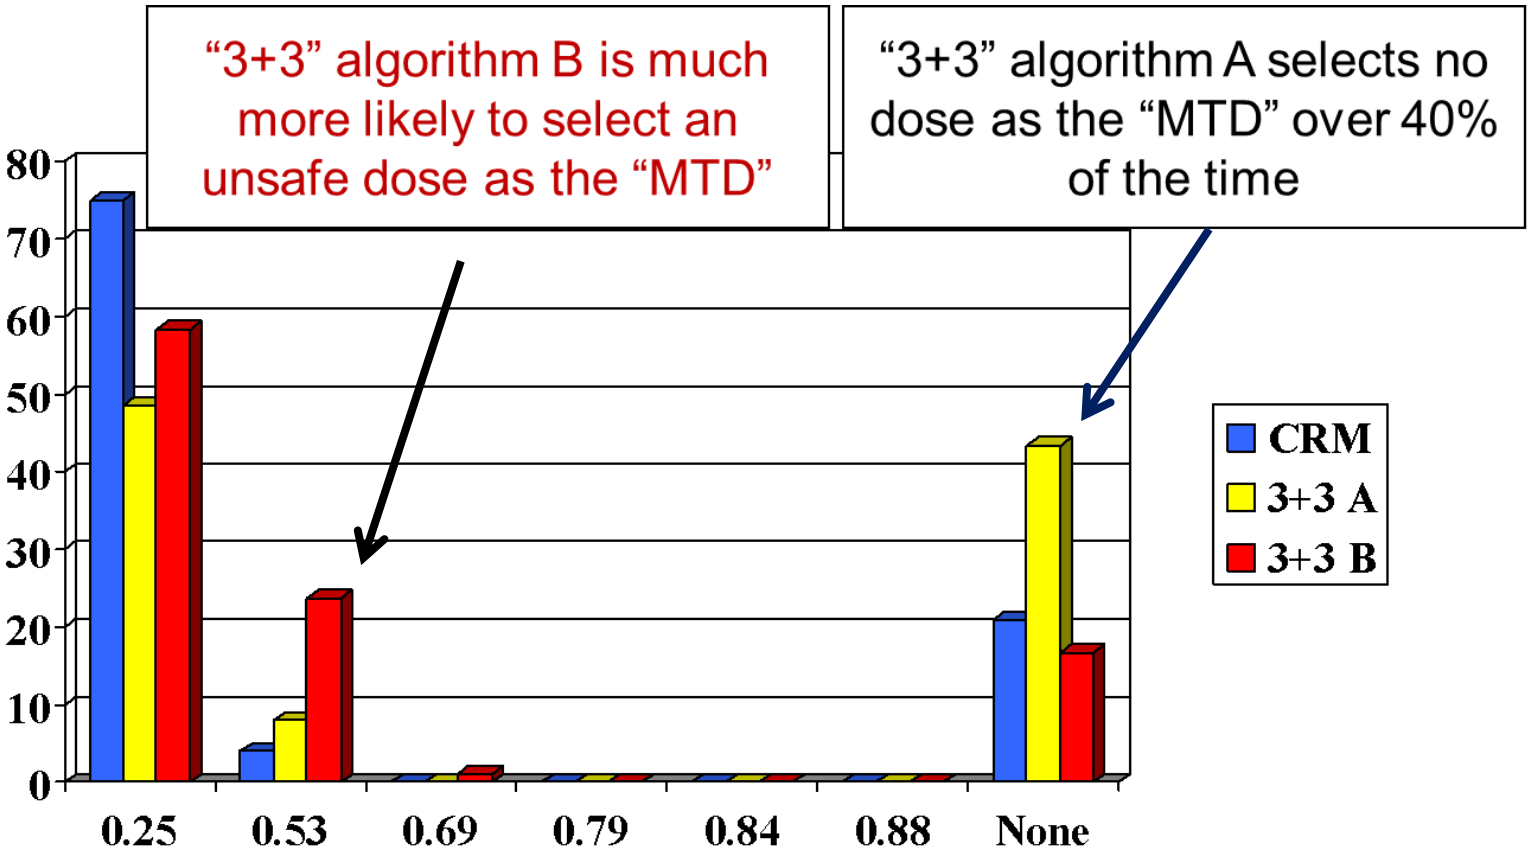

Dose	100	200	300	400	500	600

Selection Percentages Under C2

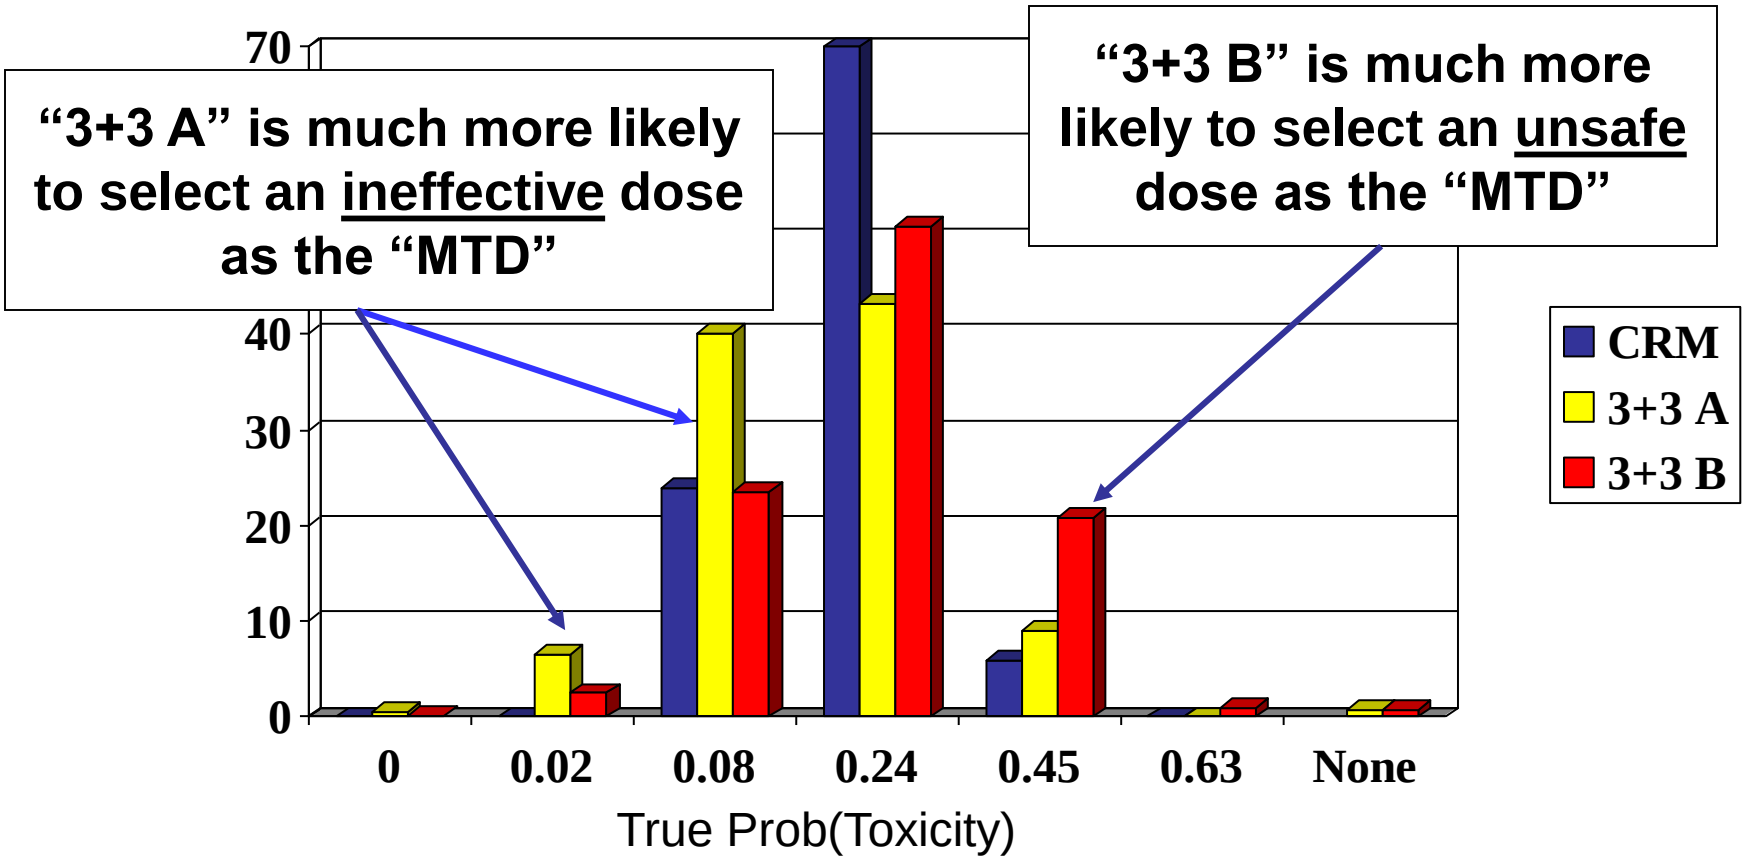

Dose	100	200	300	400	500	600

Selection Percentages Under C3

"3+3 A" is more likely to select a dose with ptox = .12 as the "MTD"

"3+3 B" is more likely to select a dose with ptox = .30 as the "MTD"

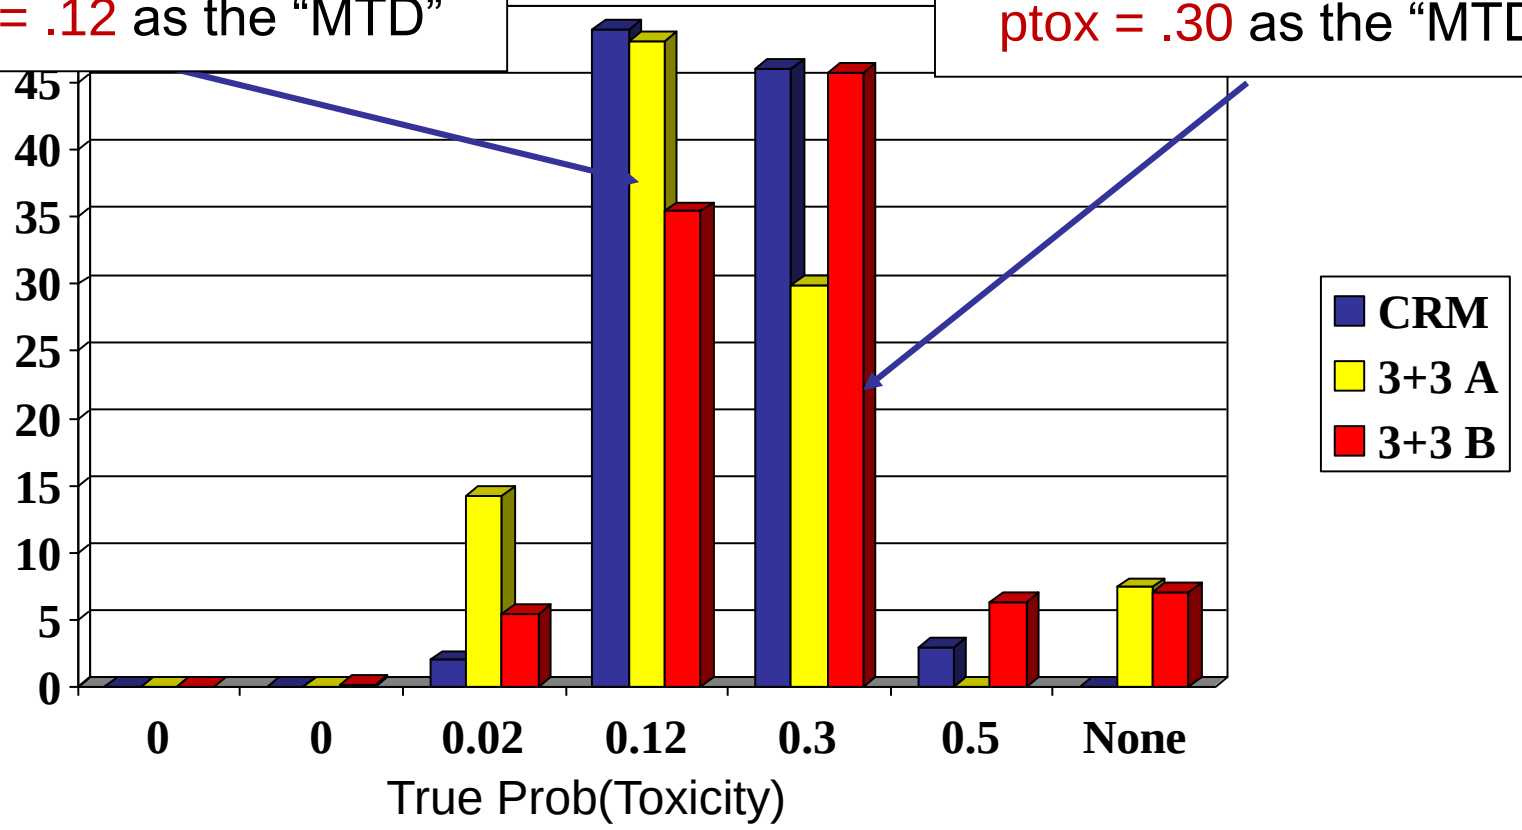

Dose	100	200	300	400	500	600

An Example of the Inherent Nuttiness of the CRM

The CRM with target $p_T^* = .25$ considers a dose d_1 with $p_T(d_1) = .30$ superior to a dose d_2 with $p_T(d_2) = .05$, because

$|.30 - .25| = .05 < .20 = |.05 - .25|$ → **Using the CRM with target $p_T^* = .25$ implies that you believe it is better to have a dose with 30% toxicity than a dose with 5% toxicity.**

'*The Tragedy of Agent X*' : *A True Story*

- At a large, well known cancer center, a phase I trial was conducted to optimize dose of ***Agent X***, a histone deacetylase inhibitor that enhances cytotoxicity when combined with nucleoside analogs: added to Fludarabine + Clofarabine + Busulfan as a preparative regimen in allogeneic stem cell transplantation for acute leukemia.
- 6 doses of ***X*** were studied using the TiTE-CRM.
- Toxicity was defined, over 30 days of follow up, as time to any of the very severe, very unlikely events
  - Graft Failure
  - Grade 4 or 5 (fatal) non-hematologic, non-infectious toxicity, Mucositis, or Diarrhea

  Grade 3 was not counted as a "DLT"

- Very few of the very unlikely DLTs were seen, so the TITE-CRM quickly escalated to d=6, producing final sample sizes (3, 3, 3, 4, 4, 51) at doses (1, 2, 3, 4, 5, 6).
- To do survival analyses, doses {1,2,3,4,5} were combined as a “Low” dose group (n=17), with dose 6 “High” (n=51).
- **Painful Surprise**: The 51 patients who got the “optimal” dose 6 had worse survival than the 17 patients who got doses 1 – 5.

Survival with the “optimal” dose X(6)

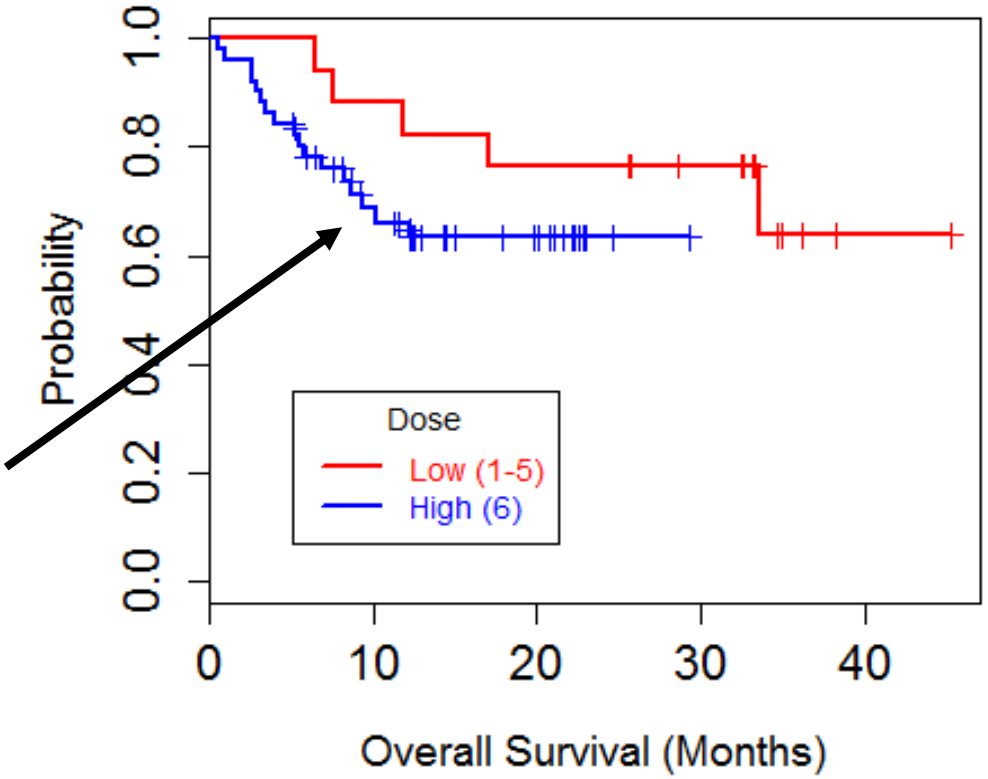

But what about covariate effects? Fitted piecewise exponential Bayesian survival regression model

	Posterior Estimates		
	Hazard Ratio	95% Credible Interval	Probability of a Harmful Effect
High vs Low Dose	2.74	0.82 - 7.66	0.92
Age	1.03	0.98 - 1.09	0.89
MRD	3.63	1.17 - 11.63	0.98
Not in CR1	0.77	0.21 - 3.46	0.36
High Risk	1.40	0.39 - 5.52	074
AML	0.88	0.23 - 2.99	0.42
Comorbidity Score	1.05	0.78 - 1.40	0.67
Maintenance Therapy	0.16	0.03 - 0.85	0.02
MUD	2.12	0.62 - 7.21	0.92
Graft Source = Bone Marrow	0.38	0.06 - 1.77	0.14

Now What?

Based on the survival time data, it would be unethical to run a phase II or phase III trial with ***Agent X*** at dose 6.

There is far too little information to determine which of doses 1 – 5 might be “optimal” in some sense.

Running another dose-finding trial in this setting is far from feasible, given the cost and time to run a second trial, and numerous competing agents.

Unanswerable Question:

Was dose 6 of ***Agent X*** too much of a good thing?

Some Examples of Early Treatment Efficacy Events

1. ≥50% **shrinkage of a solid tumor** within 4-8 weeks

2. **Complete remission of leukemia** within 42 days

3. **Dissolve the blood clot** that caused a stroke within 24 hours

4. **Engraftment of a stem cell transplant** within 4 weeks

Three Examples of Nutty Flaws with the Phase I (Toxicity Only) → Phase II (Efficacy Only) Paradigm

		d=1	d=2	d=3	d=4	d=5
Case	Pr(Toxicity)	**.05**	**.10**	**.25**	**.35**	**.50**
1	Pr(Efficacy)	.20	.50	.50	.50	.50
2	Pr(Efficacy)	.20	.25	.30	.60	.65
3	Pr(Efficacy)	.00	.01	.01	.02	.02

3+3, or CRM with target $p_T^* = .25$, are most likely to choose d=3 in all 3 cases <u>because they ignore efficacy</u>

<u>Case 1</u>: $p_E(2) = p_E(3) = .50$, but $p_T(2) = .10 < .25 = p_T(3)$

<u>Case 2</u>: [$p_E(3) = .30$, $p_T(3) = .25$] vs [$p_E(4) = .60$, $p_T(4) = .35$]

<u>Case 3</u>: *All doses are inefficacious, with* $p_E(d) \leq .02$

Three Examples of Nutty Flaws with the Phase I (Toxicity Only) → Phase II (Efficacy Only) Paradigm

In Words: A phase I design that uses Toxicity but ignores Efficacy when choosing an "optimal" dose d, like the 3+3 or CRM, is very likely to . . .

Case 1: ***Choose a dose that is too high*** if the $p_E(d)$ curve has a plateau

Case 2: ***Choose a dose that is too low*** if the $p_E(d)$ curve increases sharply for doses near the (Toxicity based) MTD

Case 3: ***Fail to stop the trial early*** if all doses are ineffective with very small $p_E(d)$

A Likely Catastrophe in Case 2

		d=1	d=2	d=3	d=4	d=5
Case 2	Pr(Toxicity)	.05	.10	**.25**	**.35**	.55
	Pr(Efficacy)	.20	.25	**.30**	**.60**	.65

1. A 3+3 algorithm, or CRM with p_T^* = .30, are most likely to choose d=3 or d=2 **because both methods ignore Efficacy** →
2. **d=4 is discarded, despite the fact that it DOUBLES the response rate from .30 to .60** →
3. Phase II then shows that the agent **at d=3** is "promising" →
4. **A large, expensive phase III trial concludes that the new agent at d=3 does not improve survival.** This is a disastrous false negative.

Expansion Cohorts
or
"There is No Design Like No Design"

1. Use a toxicity-based phase I design (3+3, CRM, etc.) to determine an MTD.

2. Behave as if the MTD is known with certainty to be the "right" dose (according to either explicit or unstated criteria)

3. Treat a fixed number of patients (10, 50, 100, or whatever) at the MTD.

Expansion Cohorts
or
“There is No Design Like No Design”

4. Do not bother with any experimental design, or specific monitoring/stopping rules for either poor safety or low efficacy

5. Once all patients have been treated, analyze the data any way you like, if possible, cherry picking a patient subset with a high response rate.

6. Use the data, and your analyses, to submit a New Drug Application to the FDA.

Why Use Expansion Cohorts?

Usual **Stated** Motivation

- Once the MTD has been "determined" treating more patients at the MTD will give a more reliable estimator of Pr(DLT | MTD), Pr(Response | MTD), and PK data
- Since the MTD is "safe" treating more patients at the MTD is perfectly ethical

Actual Motivation

- It avoids designing a phase II trial, especially futility rules that might say a new drug isn't any good
- **It pretends that the MTD is known with certainty to be the "best" dose**
- **It avoids the painful process of thinking**
- **A statistician is not needed**

Some Problems with Expansion Cohorts

1. With a typical phase I design, Pr(Toxicity | MTD) is estimated very unreliably →

There is a non-trivial probability that the MTD is too toxic, since the sample size at the MTD is very small

2. No 3+3 algorithm has any criterion for "right dose." It is just a very vaguely described algorithm.

3. The CRM has the "optimal dose" criterion that $E\{\pi_T(\text{dose})|\text{data}\}$ is close to π^* , but it ignores $\pi_E(\text{dose})$.

4. Example:

What if the trial ends with 1/6 toxicities at the MTD and then 7 DLTs occur in the first 10 expansion cohort patients, for a total of 8 / 16 (50%) toxicities.

Does any sensible physician want to treat 90 more patients at a dose where 8 / 16 DLTs were observed ?

Expansion Cohorts:
A *reductio ad absurdum*

What if you observe 0 responses in phase I ?

Dose mg/m^2	# Toxicities / # Patients	# Responses / # Patients
100	0 / 3	0 / 3
200	1 / 6	0 / 6
300	2 / 6	0 / 6
400	—	—

At the MTD of 200 mg/m^2 , where 0/6 responses were observed, does any sensible physician really want to treat 100 more patients at that dose?

Expansion Cohorts: Yet Another Problem!!

True Toxicity and Efficacy probabilities

	d=1	d=2	d=3	d=4	d=5
π_T^{true}	.05	.10	.20	.30	.40
π_E^{true}	.05	.10	.25	.35	.70

- The 3+3, or CRM with target .20 to .30, are most likely to choose d=3 or d=4, and much less likely to choose d=5.

- If the MTD chosen in phase I is d=4, then an expansion cohort of 100 patients are treated at a suboptimal dose with $\pi_E^{true}(4) = .35$, which is half $\pi_E^{true}(5) = .70$.

General Phase I-II Paradigm

1. Evaluate the effects of treatment regime ρ = dose, dose pair, or (dose, schedule) on a 2- or 3-dimensional outcome **Y** including both Efficacy and Toxicity variables.

2. Choose optimal ρ for each new patient cohort adaptively based on all data observed thus far.

3. Base the adaptive decisions on an explicit criterion function $\phi(\rho$, data), such as an Efficacy-Toxicity trade-off or posterior mean utility.

General Phase I-II Paradigm

4. Impose regime/dose admissibility rules based on marginal rates of Toxicity and Efficacy.

5. **Tailor the design** (treatment regimes, outcomes, sample size, cohort size, decision rules) **to the actual trial at hand.**

6. Use **computer simulation** to calibrate the design and establish frequentist operating characteristics :

"It is better to kill computer generated patients rather than real ones when calibrating design parameters."

Peter F. Thall

Establishing a Prior

Y = outcome vector, including binary, categorical, ordinal, or event time outcomes

τ = "regime" = treatment, dose, schedule, (dose, schedule), a multi-cycle regime

θ = model parameter vector

ξ = fixed hyper-parameters: Usually $\theta_j \sim N(\mu, \sigma^2)$, *Gamma*$(\alpha, \beta)$, or *Beta*$(\alpha, \beta)$ for each j.

Bayes Theorem : For likelihood $f(Y \mid \tau, \theta)$ and prior $p(\theta \mid \xi)$, the posterior is

$$f(\theta \mid Y_1, \ldots, Y_n, \tau, \xi) = c\ f(Y_1 \mid \tau_1, \theta) \ldots f(Y_n \mid \tau_n, \theta)\ p(\theta \mid \xi)$$

Establishing a Prior

General Strategy

1. Elicit prior means of various probabilities.

2. Use the elicited means to solve for the means $\mu_1, \ldots \mu_p$ in ξ

3. Use prior effective sample size (ESS) and preliminary trial simulations to calibrate the hyper-variances $\sigma_1^2, \ldots, \sigma_p^2$ in ξ

Accounting for both **Efficacy** and **Toxicity**

Dose 1	No Efficacy	**Efficacy**	
No TOX	.45	**.25**	
TOX	.25	.05	**.30**
		.30	

Dose 2	No Efficacy	**Efficacy**	
No TOX	.45	**.25**	
TOX	.05 →	.25	**.30**
		.50	

<u>Implication</u>: Looking at Pr(**TOX** | d) is not enough.

Accounting for both **Efficacy** and **Toxicity**

1. Toxicity Only: $Pr(TOX|d_1) = Pr(TOX|d_2) = .30$
   → $d_1 \sim d_2$ (A usual phase I design's conclusion)

2. Optimist: Define "Response" = [Efficacy, No TOX] →
   Since $Pr(Response|d_1) = Pr(Response|d_2) = .25$
   → $d_1 \sim d_2$ (An optimist's conclusion)

3. Reality: $Pr(TOX|d_1) = Pr(TOX|d_2) = .30$ and
   $Pr(Efficacy|d_2) = .50 > .30 = Pr(Efficacy|d_1)$
   → **d_2 obviously is MUCH MORE DESIRABLE than d_1**

But . . . how should one quantify dose desirability?

Desirability: Efficacy-Toxicity Probability Trade-Offs (Thall and Cook, 2004; Thall et al., 2014)

Patient Outcome: $Y_E = I(\text{Efficacy})$ and $Y_T = I(\text{Toxicity})$

$\pi_E(d,\theta) = \Pr(\text{Efficacy} \mid d,\theta)$

$\pi_T(d,\theta) = \Pr(\text{Toxicity} \mid d,\theta)$

Bivariate model for $\Pr(Y_T=a, Y_E=b \mid d, \theta)$ for a, b = 0, 1

Non-informative prior on θ with specified prior effective sample size close to 1

Dose-Finding Based On Efficacy-Toxicity Trade-Offs (Thall and Cook, 2004; Thall et al., 2014)

The physician must specify N_{max}, cohort size,

- a fixed lower limit π_R^* on $\pi_R(d,\theta)$
- a fixed upper limit π_T^* on $\pi_T(d,\theta)$
- several equally desirable fixed (π_R, π_T) pairs

A dose d is *Unacceptable* if

1) it is likely that d is unsafe :

$$\Pr\{ \pi_T(d,\theta) > \pi_T^* \mid \text{data} \} > .90, \text{ or}$$

2) it is likely that d is inefficacious :

$$\Pr\{ \pi_E(d,\theta) < \pi_E^* \mid \text{data} \} > .90$$

EffTox: Dose-Finding Based On Efficacy-Toxicity Trade-Offs (Thall and Cook, 2004; Thall, et al., 2014)

Goal: Choose the "best" acceptable dose

How "best" is defined constructively :

- Three equally desirable fixed (π_E, π_T) pairs are used to define Efficacy-Toxicity trade-off ("desirability") contours.
- The current most desirable acceptable dose is chosen for each new cohort.
- The final most desirable acceptable dose is selected at the end of the trial.

Establishing a Target Trade-Off Contour

1. Specify three equally desirable probability pairs $\boldsymbol{\pi}_1^* = (\pi_{1,E}, 0)$, $\boldsymbol{\pi}_2^* = (1, \pi_{2,T})$, $\boldsymbol{\pi}_3^* = (\pi_{3,E}, \pi_{3,T})$, with

$$\phi(\pi_E, \pi_T) = 1 - \|(\pi_E, \pi_T) - (1,0)\|_p =$$

$$1 - \left\{ \left(\frac{\pi_E - 1}{\pi^*_{E,1} - 1} \right)^p + \left(\frac{\pi_T - 0}{\pi^*_{T,2} - 0} \right)^p \right\}^{1/p}$$

2. Use bisection method to solve for p with $\phi(\pi_{3,E}, \pi_{3,T}) = 0$

3. The target contour is $\boldsymbol{C}_0$ where $\phi(\boldsymbol{\pi}) = 0$. For real number z, $\boldsymbol{C}_z$ = the contour of $\boldsymbol{\pi}$ values in $[0, 1]^2$ with $\phi(\boldsymbol{\pi}) = z$.

Target Efficacy -Toxicity Trade-Off Contour

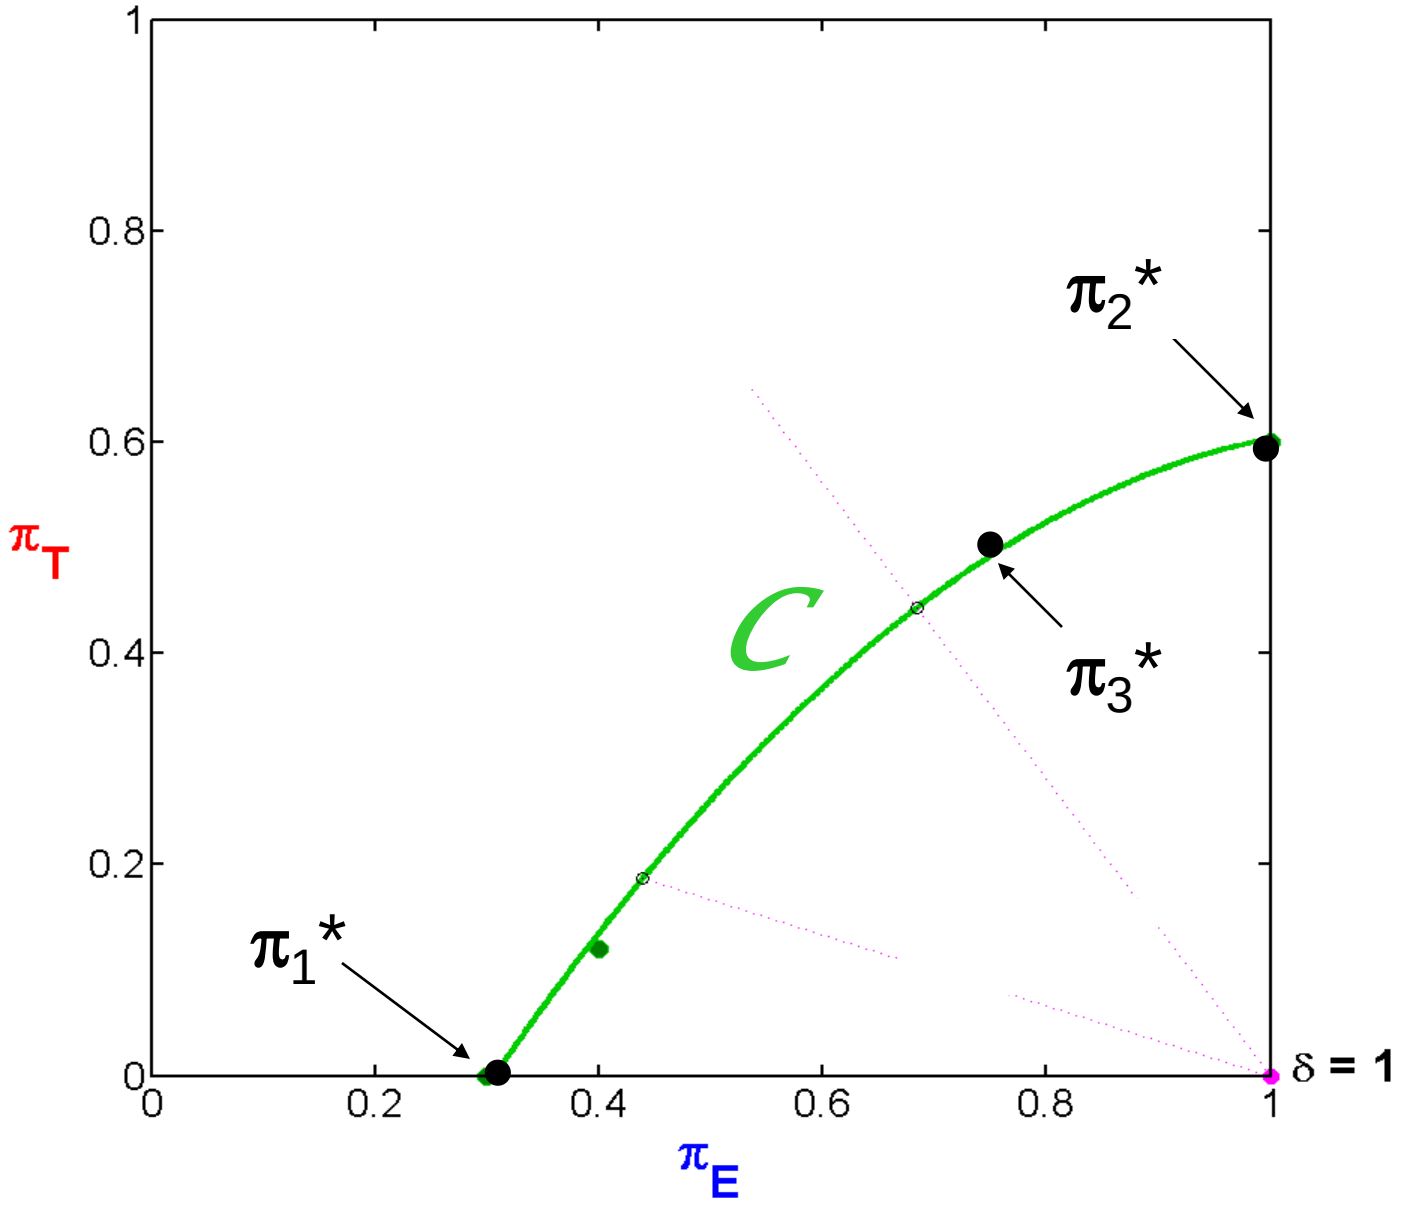

Use C to Generate a *Family of Trade-Off Contours*

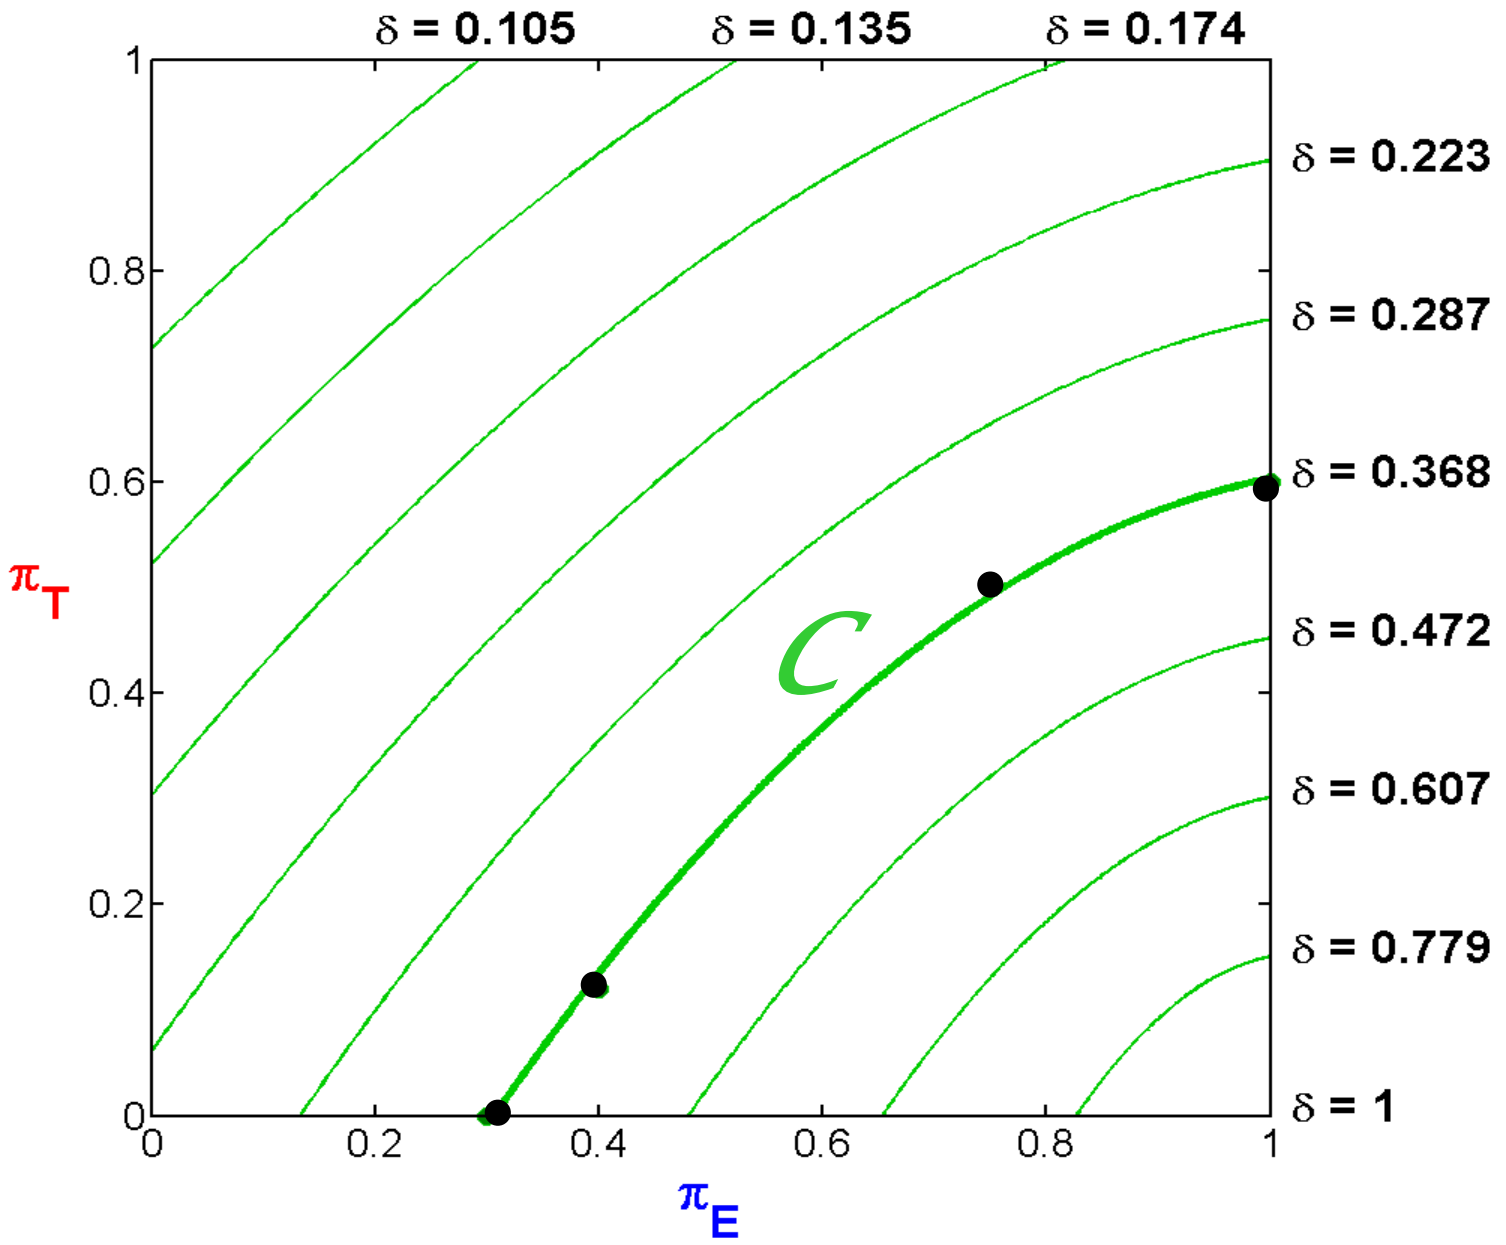

Which of these two doses is more desirable?

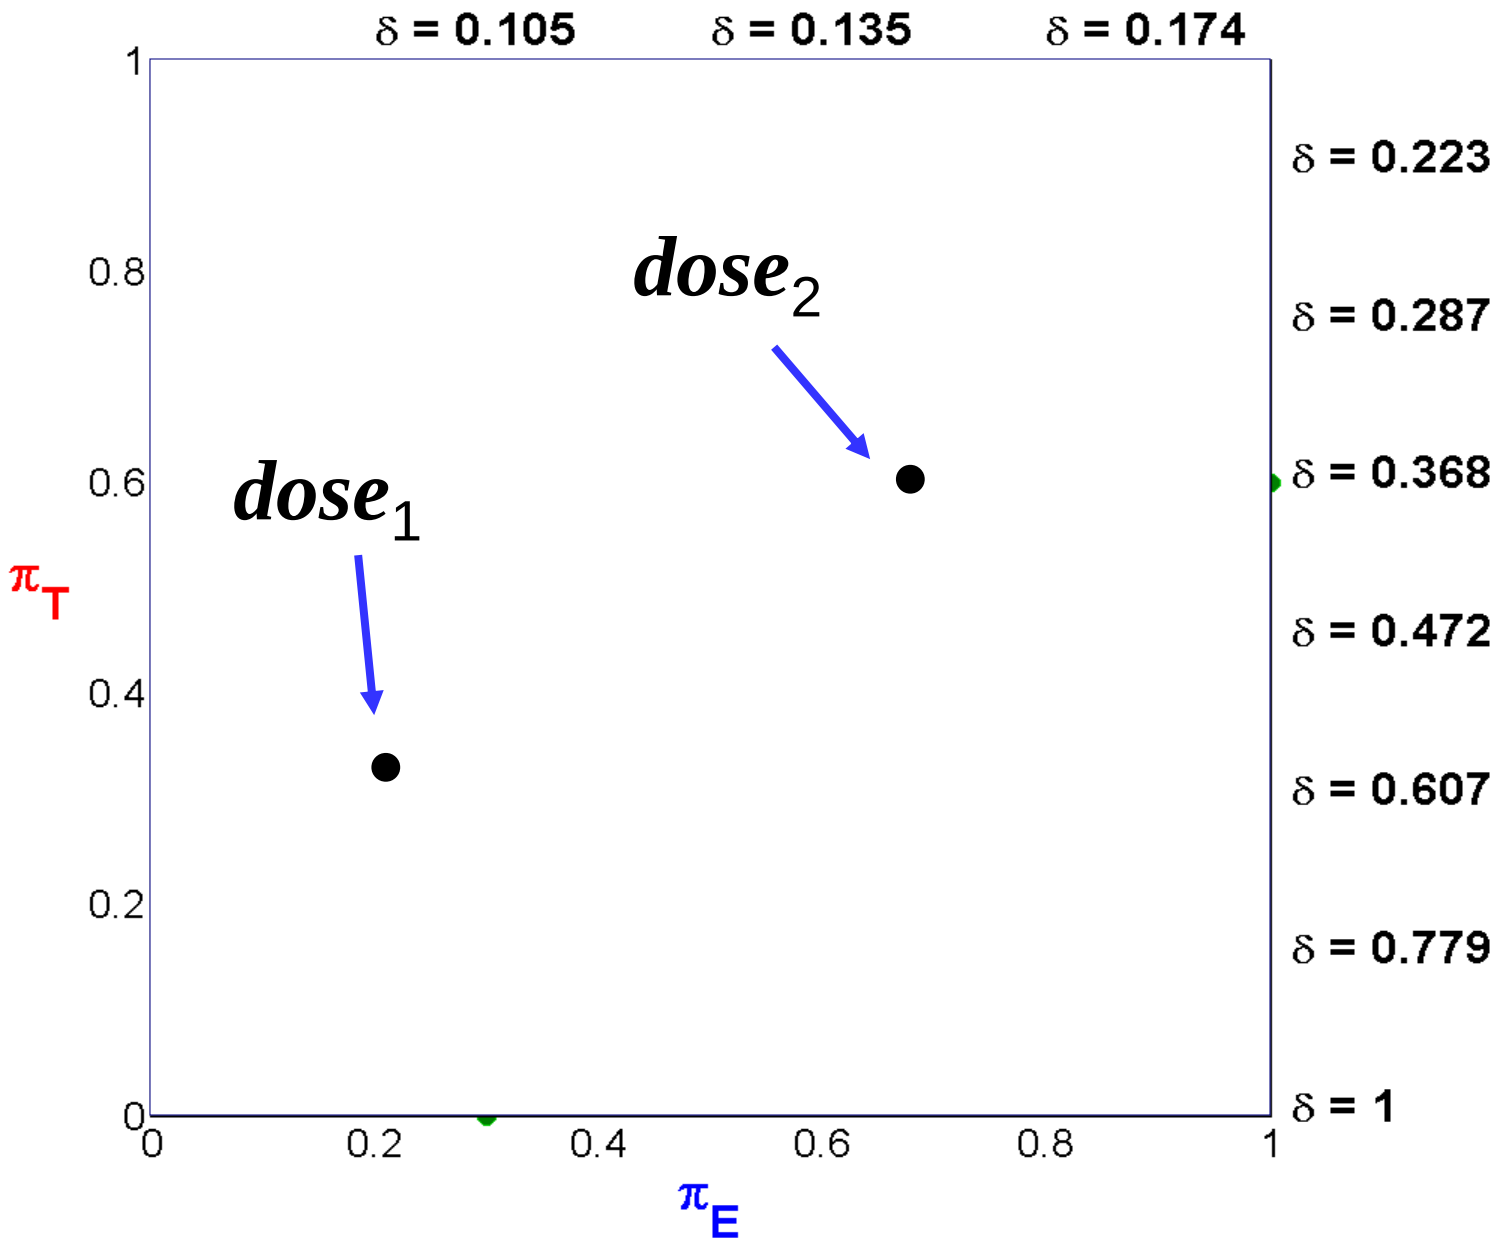

Which of these two π pairs is more desirable?

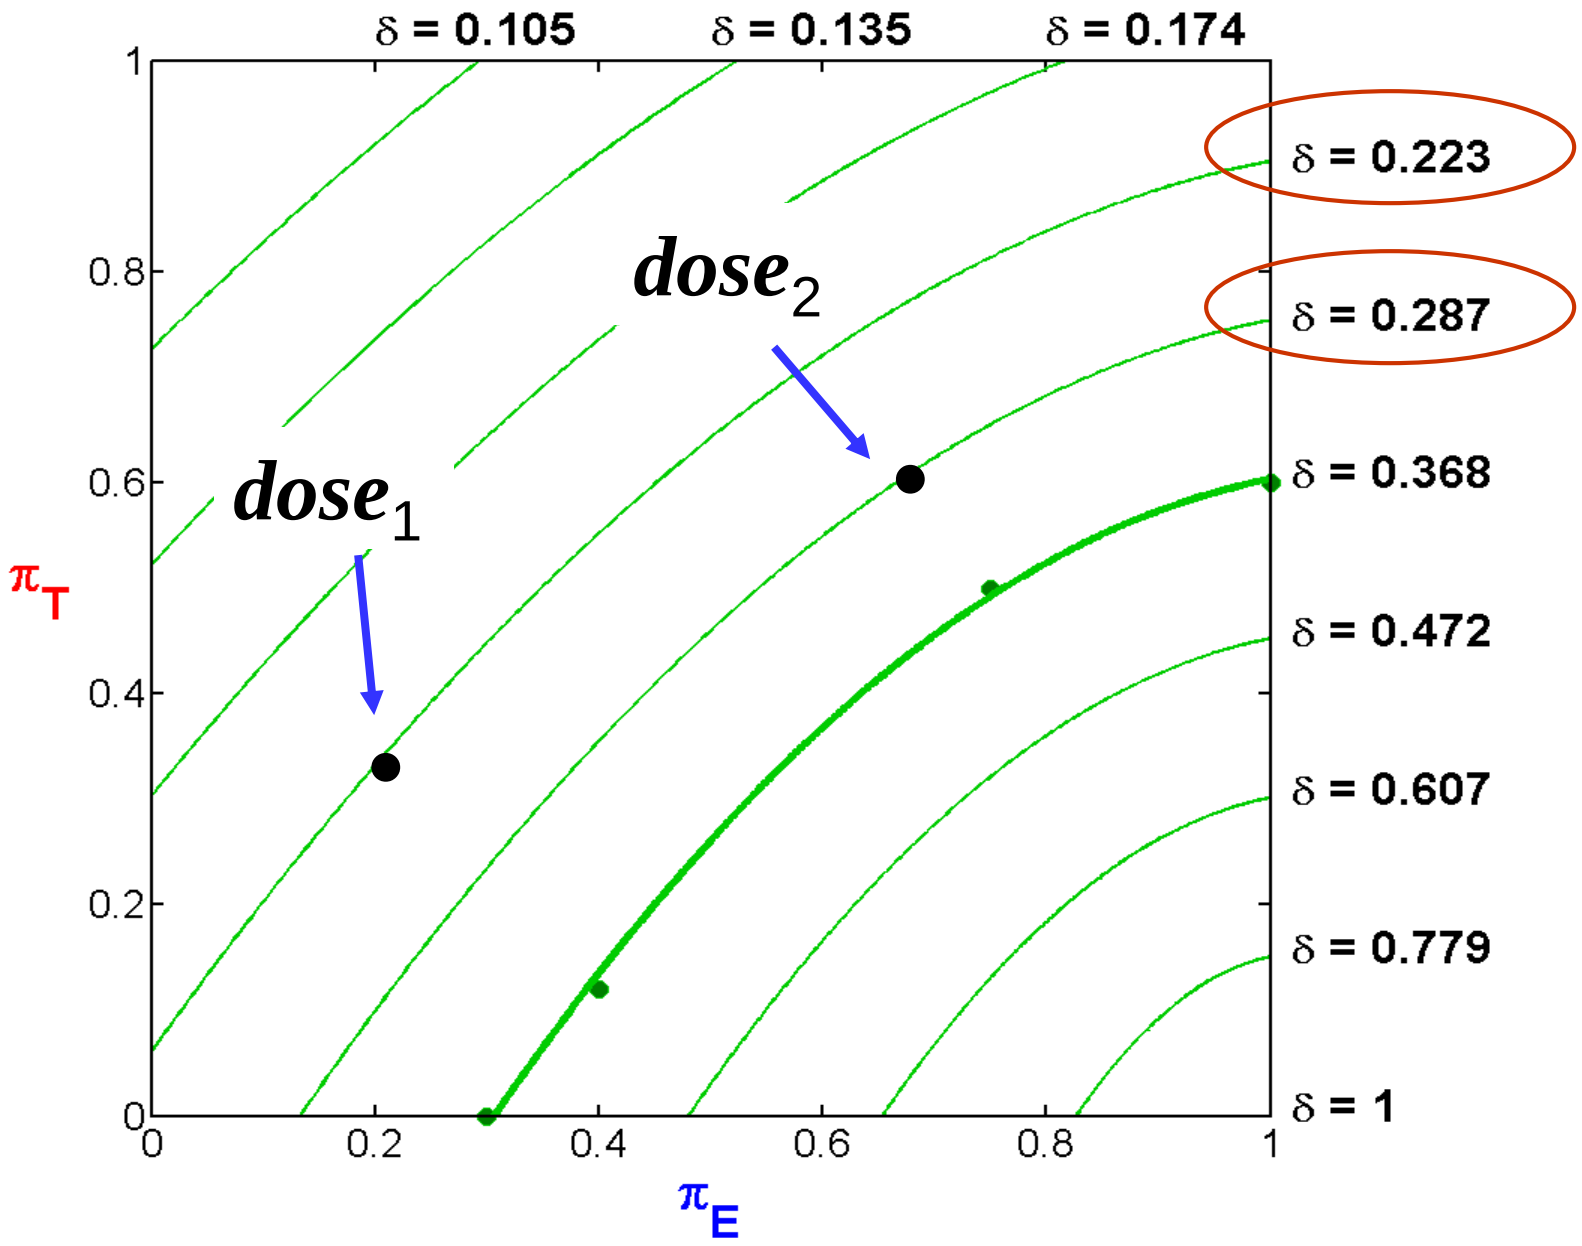

Trial Conduct

1) Physician chooses starting dose

2) A dose is *Acceptable* if
   a) it has both acceptably low toxicity and acceptably high efficacy, or
   b) it is the next higher untried dose and has acceptably low toxicity

3) *Treat each cohort at the most desirable acceptable dose*

4) Do not skip untried doses when escalating

5) If no dose is acceptable ➔ *Stop the trial and do not select any dose (combined futility and safety monitoring)*

6) At the end, select the most desirable acceptable dose

Re-Visiting Doses

The dose chosen for the next cohort may be
higher than, the same as, or lower than
the current dose

After de-escalation due to excessive toxicity or low efficacy, if subsequent outcomes at a lower dose are sufficiently safe and efficacious, then
the algorithm may re-escalate

This is what makes *any* reasonable adaptive dose-finding method less stupid than any 3+3 algorithm

Pathological Trade-Off Contours: A Decade of Dysfunction

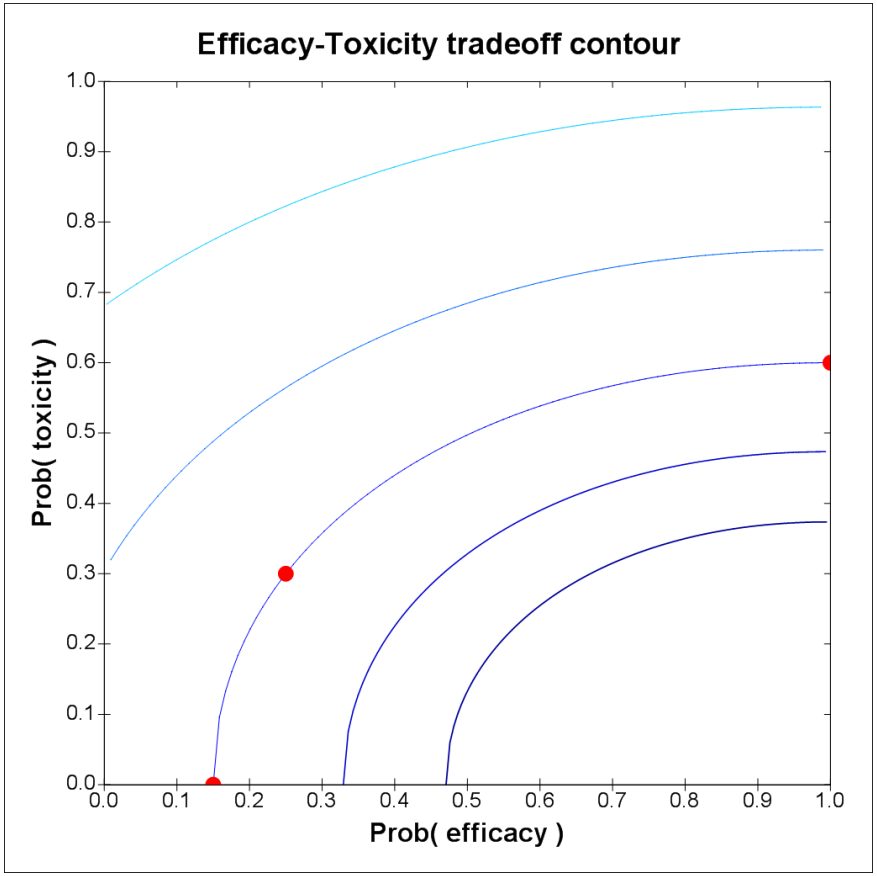

For π_E > .60, this contour requires a HUGE increase in π_E for a small increase in π_T → In scenarios where $\pi_E(d)$ increases steeply with *d* and $\pi_T(d)$ is low, the algorithm gets stuck →

A much steeper contour is needed.

A Phase I-II Trial in Advanced Prostate Cancer

d = 1, 2, 4, 6.6, 10 mcL/kg of Magic Agent (5 dose levels)

Elicited prior means $\mu_E^{(e)}$ = (.20, .40, .60,.80, .90) and
$\mu_T^{(e)}$ = (.02, .04, .06, .08, .10), prior ESS = .90

N_{max} =39, cohort size = 3, first cohort treated at d = 1

.30 = Upper Limit on $\pi_T(d)$, .50 = Lower Limit on $\pi_E(d)$

Target Contour Trade-off pairs giving **pathological contours**

$(\pi_E, \pi_T)^*$ = (.50, 0), (.70, .10), (1, .30)

Target Trade-off pairs giving **non-pathological contours**

$(\pi_E, \pi_T)^*$ = (.50, 0), (.70, .25), (1, .64)

Bad Contour :
Not Steep Enough

Good Contour :
Steep Enough

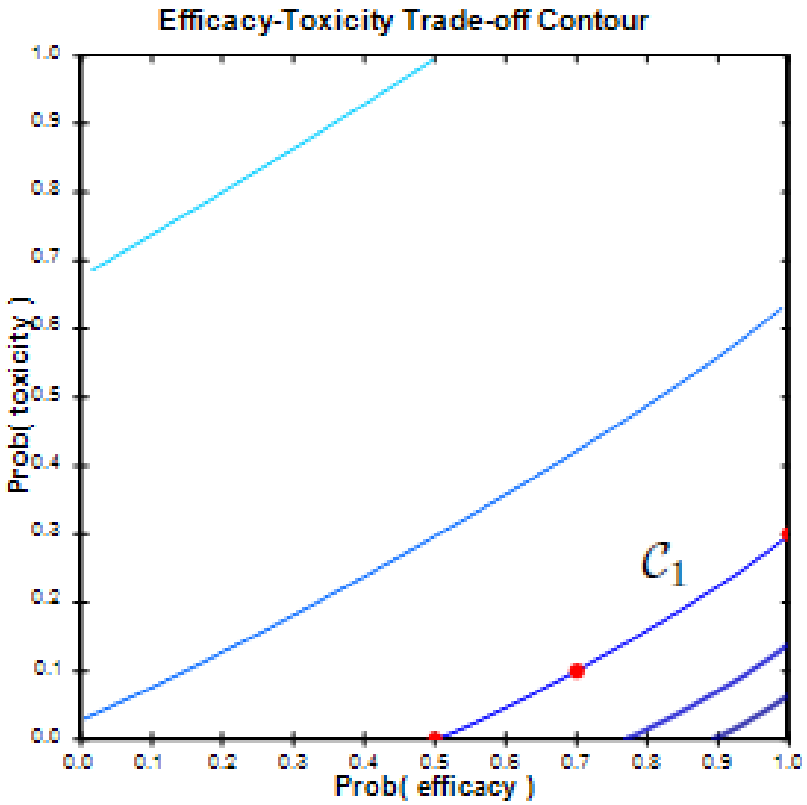

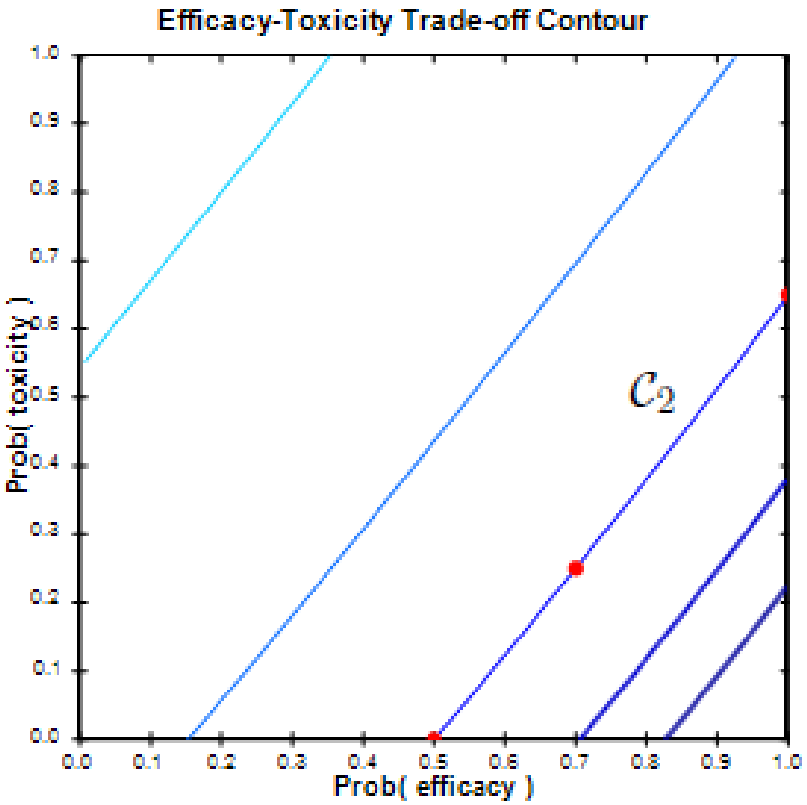

Bad Contour: Not Steep Enough **The algorithm gets stuck at the middle dose = 4, and is unlikely to escalate to a more desirable higher dose**

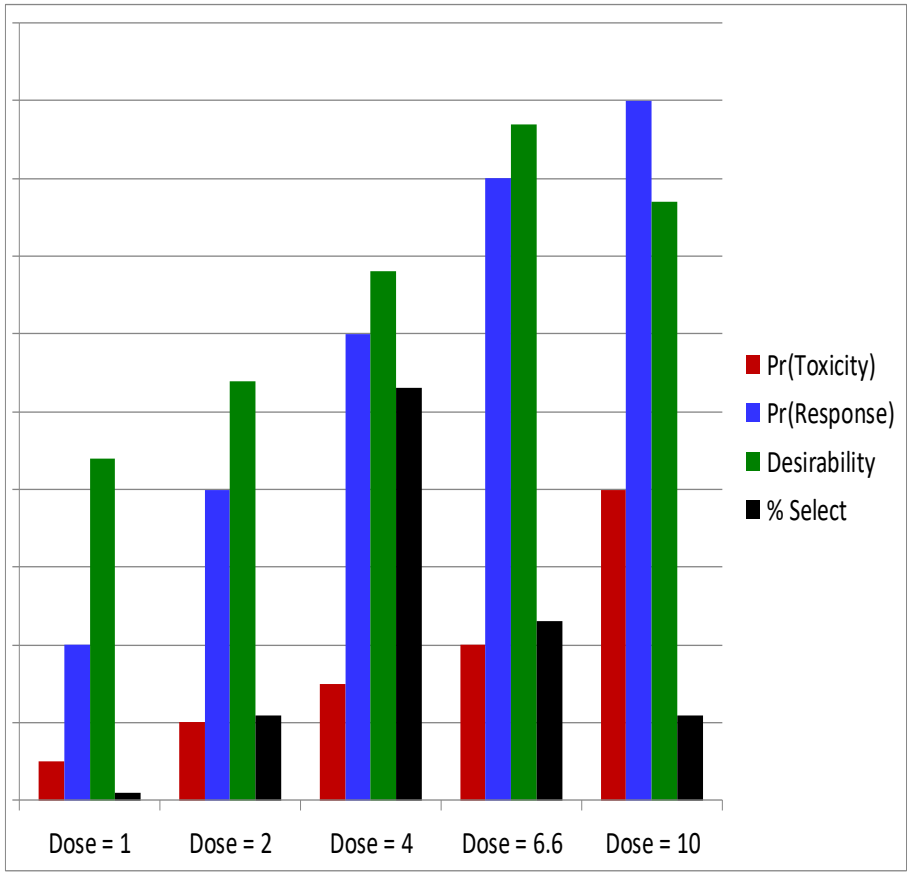

Good Contour : Steep Enough

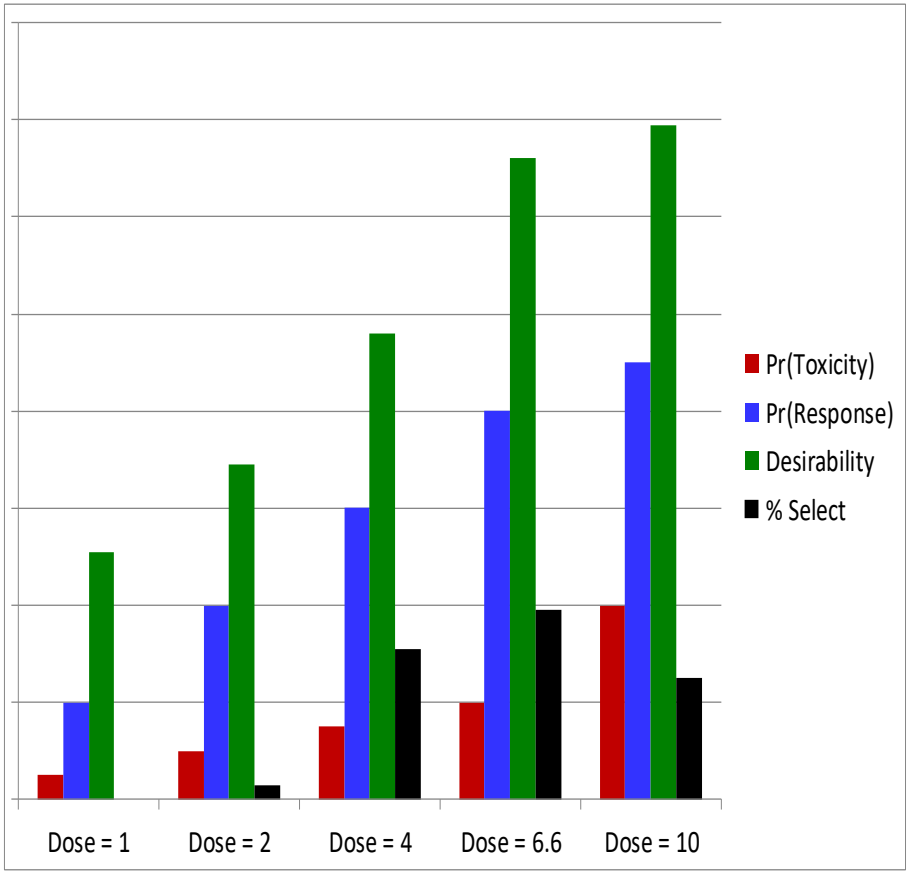

Prior Effective Sample Size (ESS)
(Morita, Thall and Mueller, 2008)

A fundamental question in Bayesian statistics:

How much information is contained in the prior?

Intuitive Motivation for ESS : Saying *Beta*(*a*, *b*) has ESS = *a* + *b* implicitly refers to the well-known fact that

$\theta \sim Beta(a, b)$ and $Y \mid \theta \sim Binom(n, \theta)$ →

$\theta \mid Y, n \sim Beta(a + Y, b + n - Y)$ which has $ESS = a + b + n$

But for many commonly used parametric Bayesian models

***it is not obvious how to determine the ESS of the prior*.**

Prior Effective Sample Size (ESS)
(Morita, Thall and Mueller 2008)

Example:

Usual normal linear regression model,

$$E(Y \mid X) = \beta_0 + \beta_1 X$$

$\text{var}(Y) = \sigma^2 \rightarrow \quad \theta = (\beta_0, \beta_1, \sigma^2)$ with

$(\beta_0, \beta_1) \sim$ Biv Normal, $\quad \sigma^2 \sim$ Inverse χ^2

For prior $p(\theta \mid \xi)$, the hyperparameter has

$\dim(\xi) = 2+3+1 = 6.$

What is the prior ESS for given ξ ?

Prior Effective Sample Size (ESS)
(Morita, Thall and Mueller 2008)

A simple algorithm: Use the method of moments

For each of several probabilities

$\pi_1(\boldsymbol{d}, \theta), \ldots, \pi_k(\boldsymbol{d}, \theta)$, approximate prior$\{\pi_j(\boldsymbol{d}, \theta)\}$

by a *Beta(a, b)* so

$$E\{\pi_j(\boldsymbol{d}, \theta) \mid \xi\} = \mu = a/(a+b)$$

$$\text{var}\{\pi_j(\boldsymbol{d}, \theta) \mid \xi\} = \sigma^2 = \mu(1-\mu)/(a+b+1) \rightarrow$$

$$\text{ESS} \sim a+b = \mu(1-\mu)/\sigma^2 - 1$$

This gives k ESS values. Just use their mean.

An Overly Informative Prior

For a,b = 0,1, x = dose,

$\pi_{a,b}(x, \theta) = \Pr(Y_E = a, Y_T = b \mid x, \theta)$

$= \pi_E^a(1-\pi_E)^{1-a}\,\pi_T^b(1-\pi_T)^{1-b} + (-1)^{a+b}\,\pi_E(1-\pi_E)\pi_T(1-\pi_T)\,(e^{\psi}-1)/(e^{\psi}+1)$

with logit $\pi_T(x,\theta) = \mu_T + x\beta_T$, logit $\pi_E(x,\theta) = \mu_E + x\beta_{E,1} + x^2\beta_{E,2}$

Model parameters: $\theta = (\mu_T, \beta_T, \mu_E, \beta_{E,1}, \beta_{E,2}, \psi)$, **p = 6**

The model for the allogeneic stem cell transplant trial in Thall and Cook (2004) has prior **ESS = 8.9**, **equivalent to 3 cohorts of patients!! This prior was far too informative**

A Strategy for Determining Priors in the Eff-Tox Regression Model

Fix the prior means $\tilde{\mu}_{\mu_E}, \tilde{\mu}_{\beta_{E,1}}, \tilde{\mu}_{\beta_{E,2}}, \tilde{\mu}_{\mu_T}, \tilde{\mu}_{\beta_T}$, and $\tilde{\mu}_{\psi}$

The new EffTox V5.0.1 GUI asks you to input :

1. Prior means of $\pi_E(d,\theta)$ and $\pi_T(d,\theta)$ for each d

2. The prior mean ESS that you desire for the $\pi_E(d,\theta)$ marginal and also for the $\pi_T(d,\theta)$ marginal

The EffTox program computes the hyperparameters that give your desired ESS values.

Using ESS to Calibrate the Prior: Applying the **Three Bears Criterion** (N_{max}=39, c=3)

			Dose				
		1	2	4.4	6	10	
Prior ESS	Desira-bility	**41**	**57**	**96**	**78**	**66**	None
10	% Sel	0	0	**15**	**68**	17	0
.90	% Sel	0	1	**45**	**38**	11	5
.02	% Sel	0	4	26	18	34	**17**

Too Big →

Just Right ! →

Too Small →

Where to find the latest version
EffTox V5.0.1 of the program

https://biostatistics.mdanderson.org/SoftwareDownload

Example of the GUI
Applying the Three Bears Criterion for N_{max} = 39 and c = 3

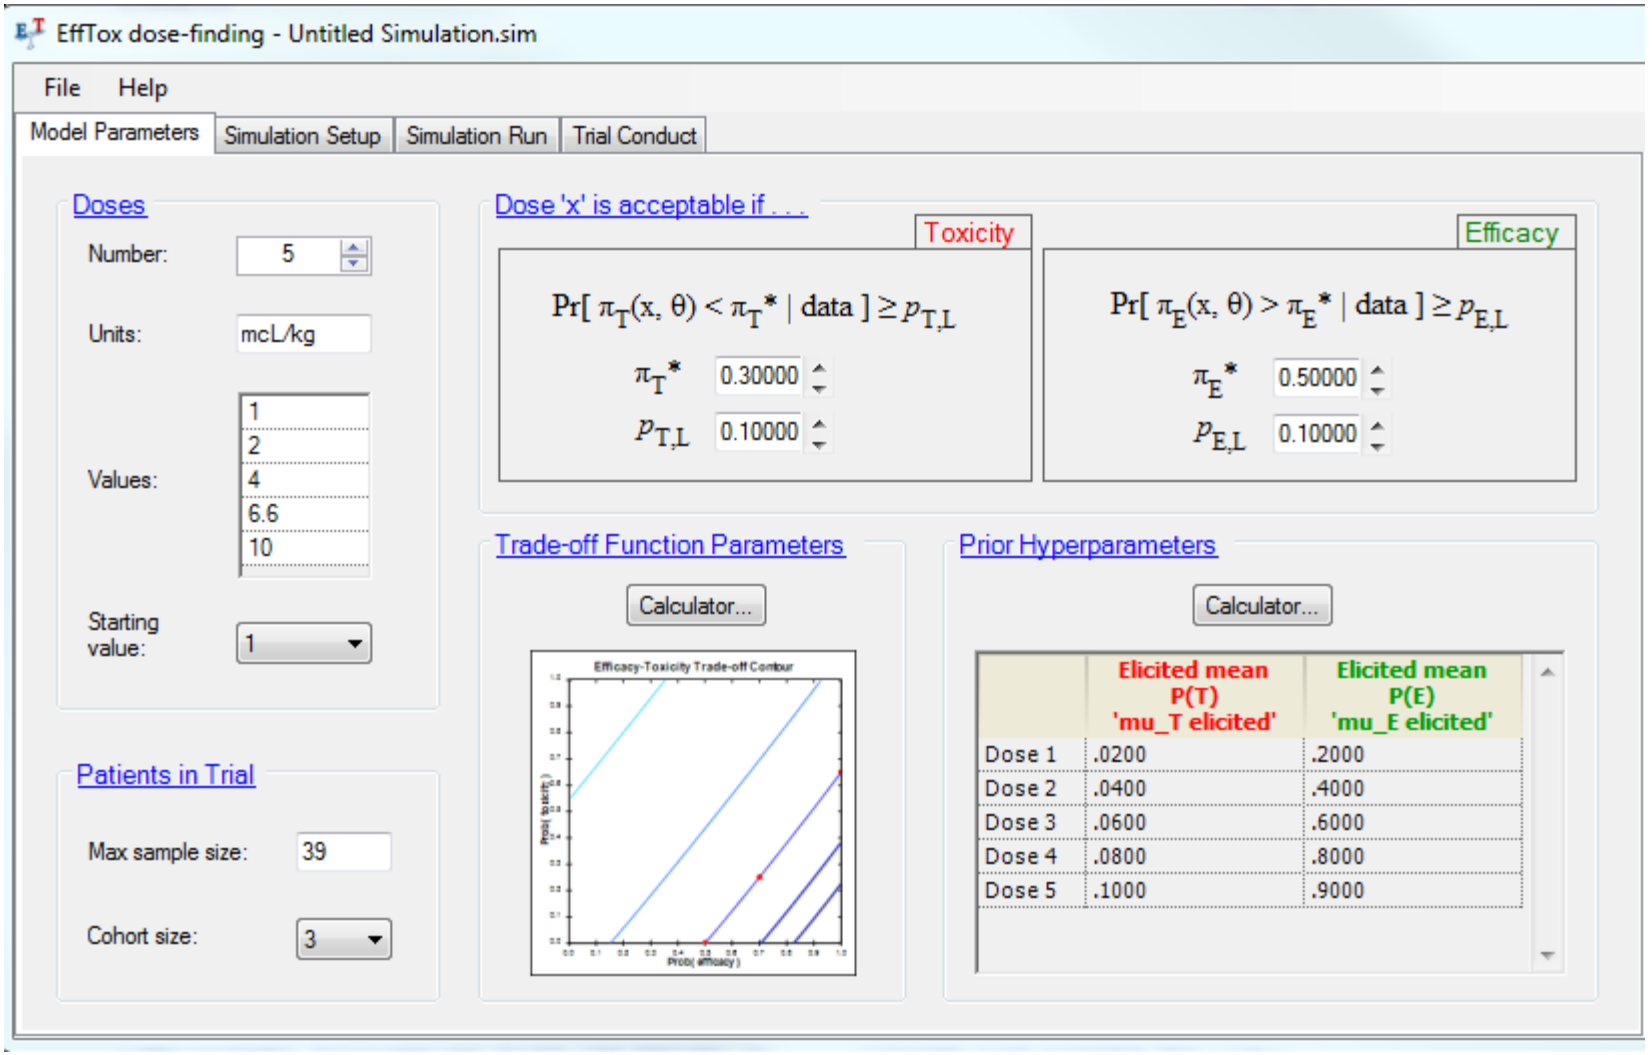

An EffTox Phase I-II Trial of Lenalidomide for Myeloma Patients Undergoing Autologous Stem Cell Transplant

Preparative regimen = fixed dose of IV melphalan + oral Lenalidomide at one of the doses { 25, 50, 75,100 } mg/m^2 on each of days -8, -7, …, -2 before transplant

Toxicity = Regimen-related death, graft failure, or grade 3,4 atrial fibrillation, deep venous thrombosis, or pulmonary embolism within 30 days post transplant

Efficacy = Alive and in CR at day 30 post transplant

.20 = Upper Limit on $\pi_T(x)$, .15 = Lower Limit on $\pi_E(x)$

Trade-off pairs (π_E, π_T) = (.15, 0), (.30, .15), (1, .50)

N_{max} = 60, cohort size = 3, first cohort treated at 25 mg/m^2

Lenalidomide Autologous SCT Trial: Simulation Scenario 1

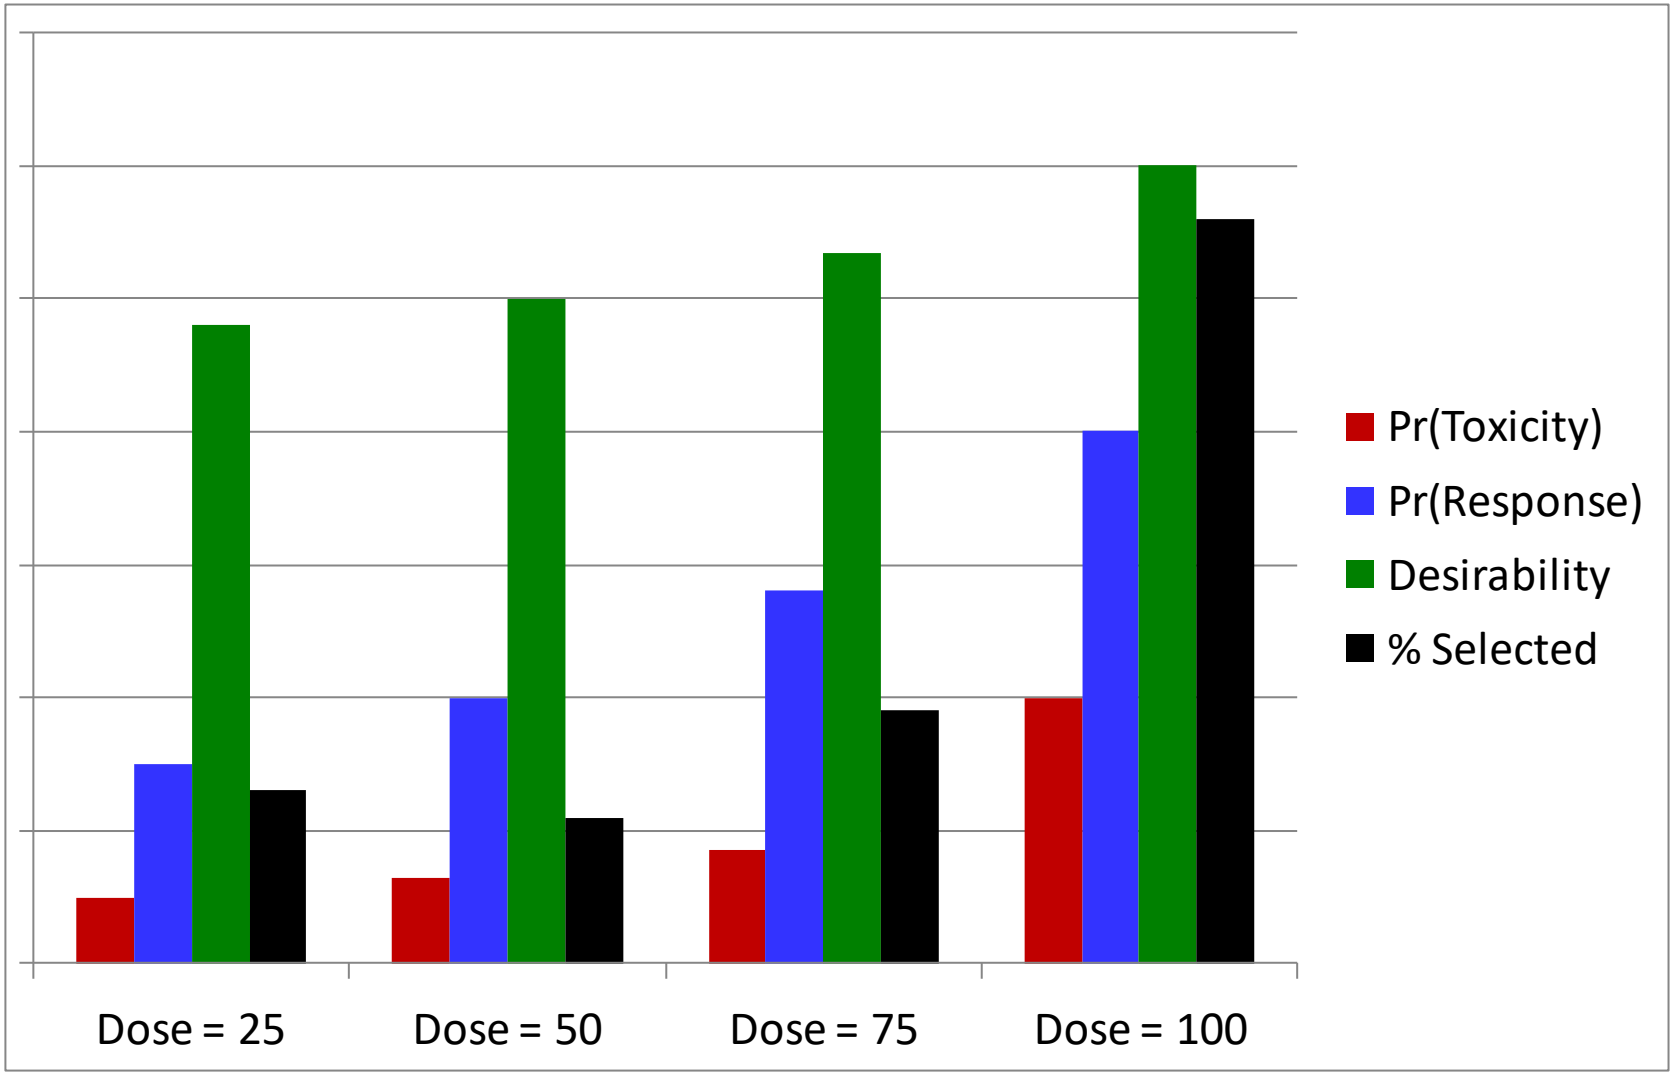

Lenalidomide Autologous SCT Trial: Simulation Scenario 2

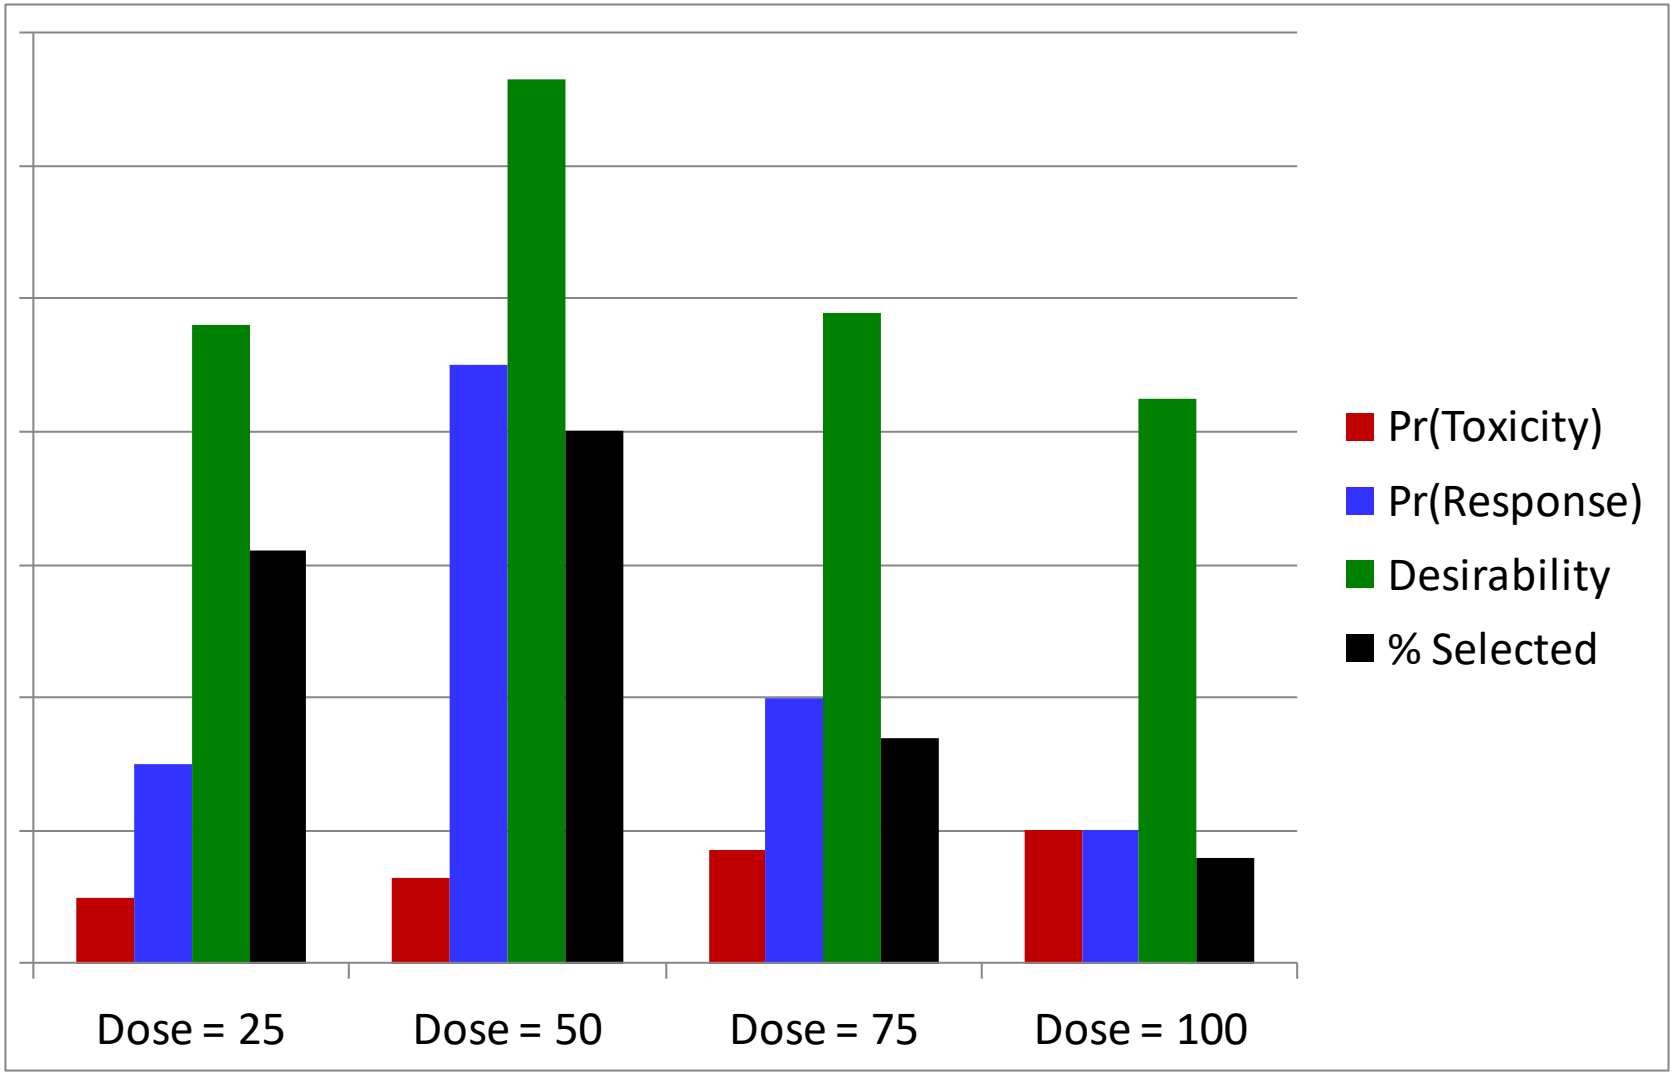

Lenalidomide Autologous SCT Trial: Simulation Scenario 3

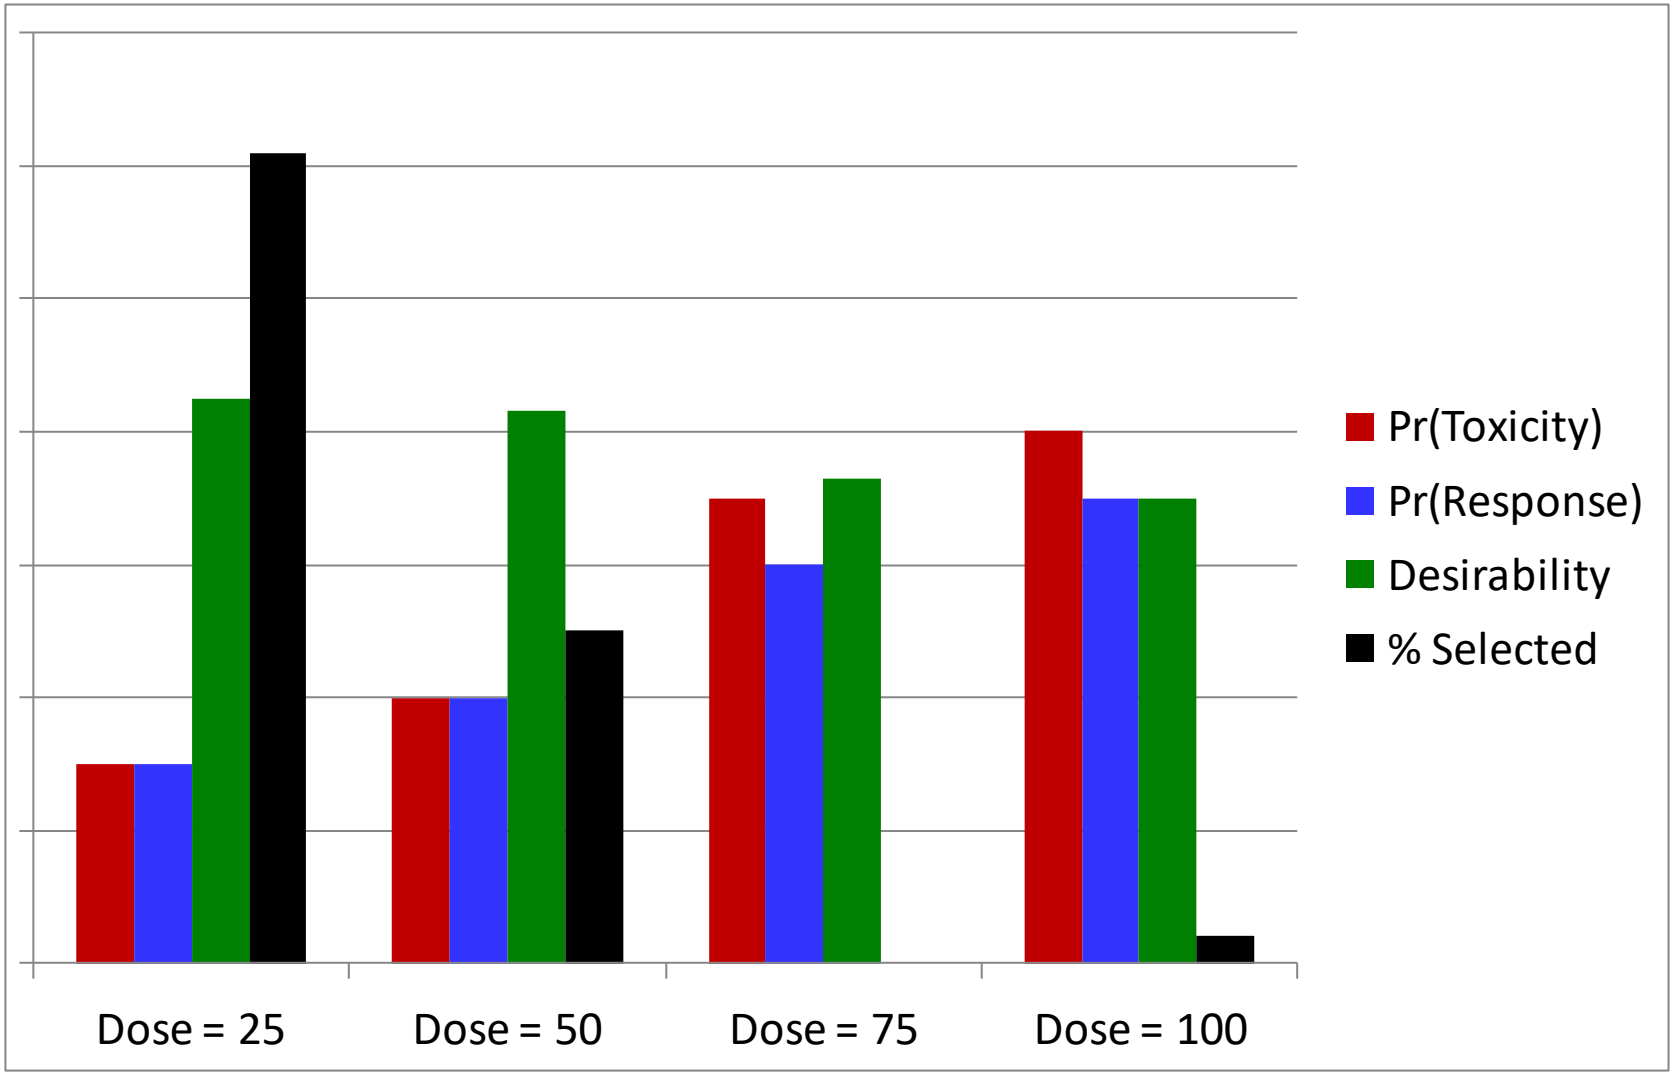

Simulation Comparisons of 3+3, CRM, and EffTox
Dose Selection %

			Dose Level					
			d_1	d_2	d_3	d_4	d_5	None
Scenario 1	True $\pi_E(d_j)$		.05	.10	.30	.55	.60	
	True $\pi_T(d_j)$		.10	.20	.35	.40	.60	
	Trade-off		34	32	38	58	48	
	N=30	3+3 A	29	39	16	6	0	11
		3+3 B	19	39	25	11	1	5
		CRM	2	38	45	13	2	0
		EffTox	2	11	25	28	25	9
	N=45	CRM	0	36	52	10	1	0
		EffTox	1	8	28	37	16	10
	N=60	CRM	0	34	58	8	0	0
		EffTox	0	7	30	42	10	12

- CRM has target .30 for π_T
- EffTox has upper limit .40 on π_T and lower limit .20 on π_E

Scenario 2	True $\pi_E(d_j)$		.30	.60	.65	.70	.75	
	True $\pi_T(d_j)$		.20	.25	.45	.60	.70	
	Trade-off		48	81	66	59	57	
	N=60	3+3 A	30	31	7	0	0	31
		3+3 B	31	38	15	1	0	15
		CRM	9	70	17	0	0	4
		EffTox	14	54	28	3	1	1
Scenario 3	True $\pi_E(d_j)$		.05	.10	.30	.50	.70	
	True $\pi_T(d_j)$		.05	.10	.20	.30	.35	
	Trade-off		37	38	48	61	86	
	N=60	3+3 A	10	28	31	16	0	15
		3+3 B	5	19	33	24	6	13
		CRM	0	0	21	40	38	0
		EffTox	0	3	14	13	68	1
Scenario 4	True $\pi_E(d_j)$		.01	.02	.04	.06	.07	
	True $\pi_T(d_j)$		.12	.15	.30	.50	.55	
	Trade-off		30	30	24	18	17	
	N=60	3+3 A	18	38	27	3	0	15
		3+3 B	15	32	38	9	1	6
		CRM	0	13	77	9	0	1
		EffTox	0	0	1	3	3	93

Comparing Dose Selection % for Four Methods:
Simulation Results in One Scenario (N= 60 patients)

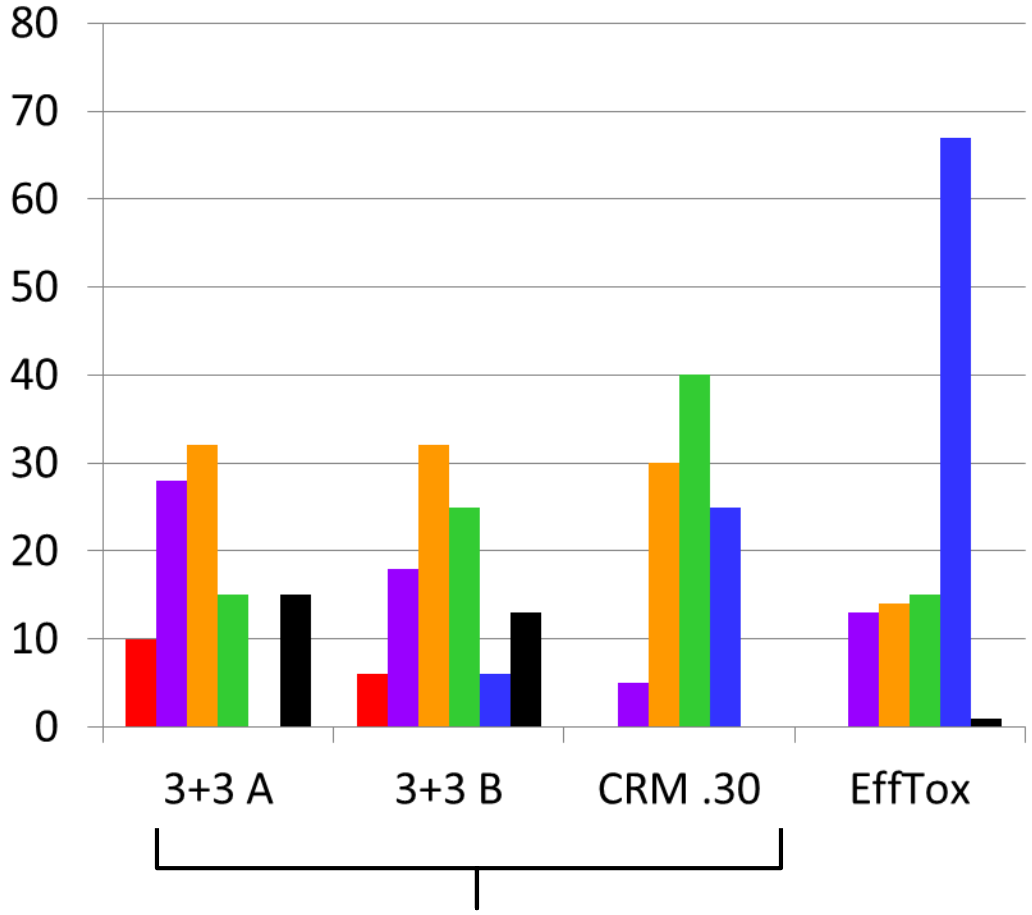

	π_E^{true}	π_T^{true}
Dose 1	.05	.05
Dose 2	.10	.10
Dose 3	.30	.20
Dose 4	.50	.30
Dose 5	.70	.35
None		

Utility-Based Phase I-II Trials

Medical Practice

- Efficacy and Toxicity both matter for the patient →

  Any reasonable statistical method should use both.

- Utilities and trade-offs underlie all medical decision-making → They are natural tools for statistical decision-making.

Advantages of Using Utilities

1. U(Toxicity, Efficacy), or U(Tox, Eff_1, Eff_2) maps a multidimensional outcome to a 1-dimensional criterion that

- quantifies risk-benefit trade-offs
- can be used to make decisions about doses, or more generally about treatments.

2. In practice, physicians do not write down their utilities, unless they are elicited by a statistician designing a clinical trial. ***Physicians LOVE to give their numerical utilities. So, their utilities now are made EXPLICIT.***

Computing Mean Utility

		Toxicity	
		No	Yes
Efficacy	No	40 $\pi_{00} = .40$	0 $\pi_{01} = .10$
	Yes	100 $\pi_{10} = .30$	70 $\pi_{11} = .20$

Mean Utility

$= U(0,0)\,\pi_{00} + U(0,1)\,\pi_{01} + U(1,0,)\,\pi_{1,0} + U(1,1)\,\pi_{1,1}$

$= 40 \times .40 + 0 \times .10 + 100 \times .30 + 70 \times .20$

$= 16 + 0 + 30 + 14 = 60$

		Toxicity	
		No	Yes
Efficacy	No	40 π_{00} = ?	0 π_{01} = ?
	Yes	100 π_{10} = ?	70 π_{11} = ?

But we do not know (π_{00}, π_{01}, π_{10}, π_{11}), so we write down a Bayesian model with parameters θ

$\pi(a,b \mid d, \theta) = Pr(Y_T = a, Y_E = b \mid d, \theta)$ for a, b = 0 or 1,

and use it to compute the mean utility of each d :

$$U(d, \theta) = \pi(1,1 \mid d, \theta)\, U(1,1) + \pi(1,0 \mid d, \theta)\, U(1,0) + \pi(0,1 \mid d, \theta)\, U(0,1) + \pi(0,0 \mid d, \theta)\, U(0,0)$$

As dose-outcome data are observed during a clinical trial, θ and $U(d, \theta)$ estimated, for making decisions.

The Meanings of "Mean Utility"

But we do not know θ → We apply ***Bayesian Statistics*** to decide which dose is "optimal," as follows:

The *mean utility of dose d given parameters* θ is

$$u(d, \theta) = E\{ U(Y) \mid d, \theta \} = \sum_y U(y)\, P(Y=y \mid d, \theta)$$

The *posterior mean utility* of dose *d, given the data*, is

$$\phi(d, \text{data}) = E_\theta \{ u(d, \theta) \mid \text{data} \}$$

The dose with largest posterior mean utility is given to the next cohort of patients.

As new data are obtained during the trial, we repeat this computation, learn sequentially about the dose-utility function, and the "optimal" dose may change.

Adaptive Randomization to Deal With Stickiness

Well Known Fact

Any sequentially adaptive statistical decision rule based on an optimality criterion may risk getting stuck at a sub-optimal (locally optimal) action. If so, one fails to adequately explore the action space and identify a truly optimal action: "Stickiness", or the "Exploitation versus Exploration" problem.

Practical Solution

After an in initial burn-in, use sequential Adaptive Randomization (AR): Treat each new cohort at dose *d* chosen randomly with probability proportional to $\phi(d, data)$

Utility Based Sequential Decision Making

1) Goal: Given a set of experimental treatment regimes $\{\rho_1, \ldots, \rho_k\}$ (doses, schedules, (d,s) combinations, etc.), sequentially choose a "best" regime for each successive cohort of patients, in real time, based on 2 or 3 or more Efficacy and Toxicity outcomes

2) Utilities : Use elicited utility U(Efficacy,Toxicity), to choose each cohort's regime

3) Bayesian Computations: Map (ρ,***data***) to the *posterior mean utility* $\phi(\rho,$***data***) of each treatment regime ρ , **or** find the regime ρ that has largest

$$\pi(\rho, \textbf{\textit{data}}) = \Pr[\, u(\rho, \theta) = max\{u(\rho*, \theta)\} \mid \textbf{\textit{data}}\,]$$

Utility Based Sequential Decision Making

4) Maximize either $\phi(\rho,\textbf{\textit{data}})$ or $\pi(\rho,\textbf{\textit{data}})$ to choose the best ρ

5) Acceptability : Restrict selections to τ that are acceptable, in terms of safety and efficacy. If all τ are "unacceptable" then stop the trial.

6) Sequential Adaptive Randomization (AR) :

After an initial burn-in, **repeatedly randomize** among doses with $u(\rho,\textbf{\textit{data}})$ close to the maximum, to avoid getting stuck at a suboptimal regime.

A Phase I-II Pediatric Radiation Therapy Trial

Diffuse Intrinsic Pontine Gliomas (DIPGs)

- Very aggressive brain tumors
- No treatment with substantive anti-disease activity exists
- **Radiation Therapy (RT) is standard treatment, but is mainly palliative**
- RT dose-toxicity & dose-efficacy profiles are not well understood

Subjects: Children, median age = 5 years, with DIPGs

Three RT dose levels: "Biologically Equivalent Doses" in Gy, given serially per a fractionation schedule

A Phase I-II Pediatric Radiation Therapy Trial

Efficacy = # improvements in :

- Clinical Symptoms
- Radiographic Appearance of the Tumor
- Quality of Life

→ Y_E = **0, 1, 2, or 3**

Toxicity Defined in terms of fatigue, nausea/vomiting, headache, skin toxicities, blindness, brain edema or necrosis with Y_T = **Low, Moderate, High, or Severe**

Both Efficacy (Y_E) and Toxicity (Y_T) are scored by day 42

Number of (Efficacy, Toxicity) outcomes = 3x4 = 12

Elicited Numerical Joint Outcome Utilities of 16 possible outcomes

		Toxicity Severity			
		Low	**Moderate**	**High**	**Severe**
Efficacy Score	0	**50**	**25**	**10**	0
	1	**85**	**50**	**15**	**5**
	2	**92**	**60**	**20**	**7**
	3	100	**75**	**25**	**10**

U(Toxicity, Efficacy) is used to make decisions adaptively in the trial ("learn-as-you go")

1) Decide which radiation does are acceptable

2) Choose best dose for each cohort of 3 children: **"Best" means "Has the highest posterior (mean) utility"**

Some Properties of the Utilities

		Toxicity Severity			
		Low	**Moderate**	**High**	**Severe**
Efficacy Score	0	50	25	10	0
	1	85	50	15	5
	2	92	60	20	7
	3	100	75	25	10

Question:

Why not just use "DLT" = {High, Severe} and apply a simple dose finding method (e.g. "3+3" or "CRM") ?

Answer:

U(0,Moderate) = U(3, High) = 25 → Scoring these two outcomes as "No DLT" and "DLT" makes no sense!

Bivariate Ordinal Dose-Outcome Model

Y_1 = Efficacy index {*0, 1, 2, 3*}

Y_2 = Toxicity index {*0, 1, 2, 3*} (low,mod,high, severe)

→ 16 possible (Efficacy, Toxicity) outcomes

x = dose, indexed by *1,2, …, J,*

$$\pi_{k,y,x} = \Pr(Y_k = y \mid x, \theta) \quad \text{for } k = 1, 2$$

Bivariate Ordinal Dose-Outcome Model

$$\lambda_{k,y,x} = e^{\theta_{k,y,x}} / (1 + e^{\theta_{k,y,x}})$$

$$\begin{aligned}
\pi_{k,0,x} &= 1 - \lambda_{k,1,x} \\
\pi_{k,1,x} &= \lambda_{k,1,x} - \lambda_{k,1,x}\lambda_{k,2,x} \\
\pi_{k,2,x} &= \lambda_{k,1,x}\lambda_{k,2,x} - \lambda_{k,1,x}\lambda_{k,2,x}\lambda_{k,3,x} \\
\pi_{k,3,x} &= \lambda_{k,1,x}\lambda_{k,2,x}\lambda_{k,3,x}.
\end{aligned}$$

Establishing Priors

Elicited prior mean outcome probabilities for the RT trial

		Y_1 = Toxicity Severity				Y_2 = Efficacy Score			
x	BED	Low	Moderate	High	Severe	0	1	2	3
1	40.00	0.65	0.20	0.12	0.03	0.20	0.40	0.35	0.05
2	45.76	0.55	0.25	0.15	0.05	0.10	0.30	0.45	0.15
3	53.39	0.40	0.30	0.23	0.07	0.10	0.20	0.50	0.20

Computing Prior Hyperparameters

24 elicited probabilities, p=19 hyperparameters :

1. Estimate19 prior means from the elicited probabilities
2. Calibrate the hyper-variances to ensure small overall prior ESS

Radiation Therapy Trial Conduct

Approximating each prior($\pi_{k,x,y}$) as a beta →

Prior ESS values were 0.31 to 0.70, with mean 0.42.

A 10% limit was imposed on Pr(High or Severe toxicity) →

x unacceptably toxic if $\Pr(\pi_{1,x,3} \geq .10 \mid \mathcal{D}_n) > .80$

- N_{max} = 30. Treat the first 3 pats. at ***x*** =1, then adapt, do not skip dose level ***x*** =2 when escalating at the start.
- AR applied for doses with posterior mean utility close to maximum
- Posteriors computed using MCMC with Gibbs sampling.

Operating Characteristics of the RT Trial Design

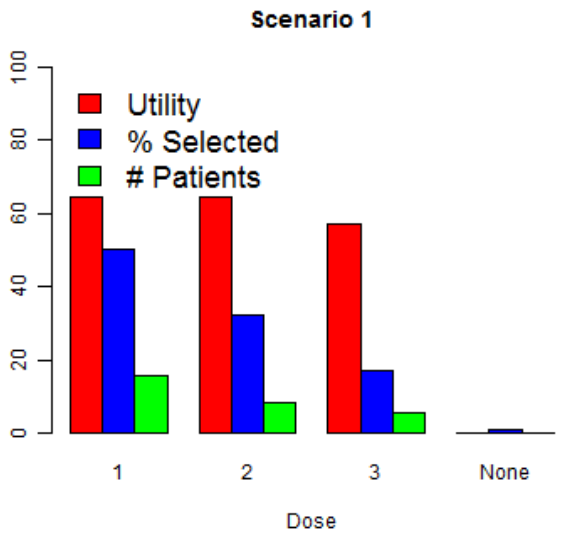

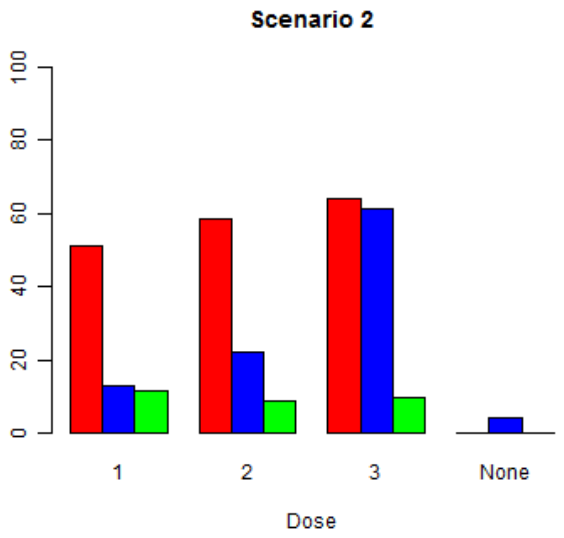

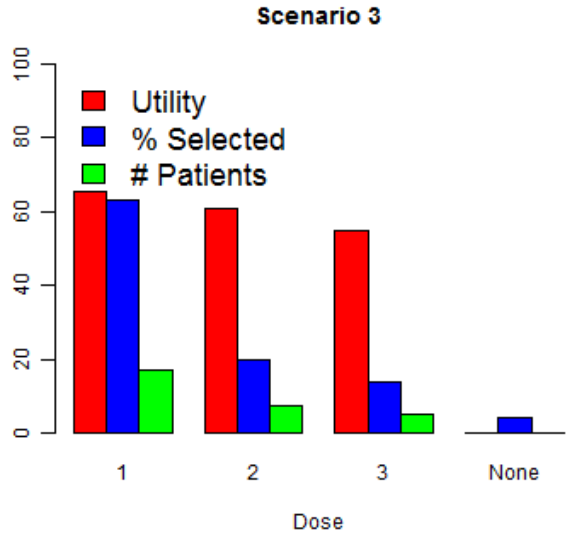

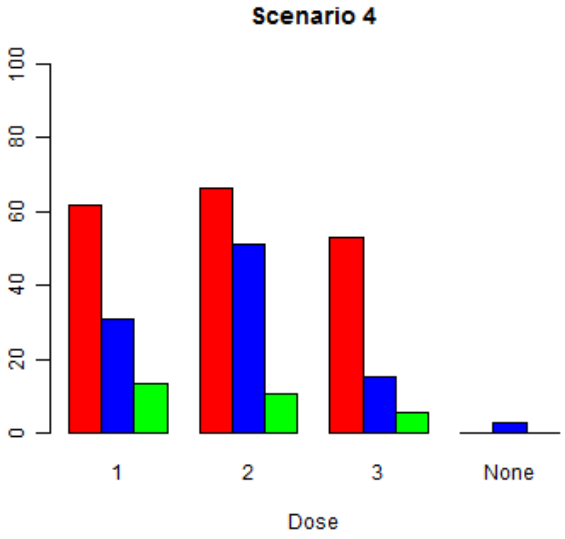

Operating Characteristics of the RT Trial Design

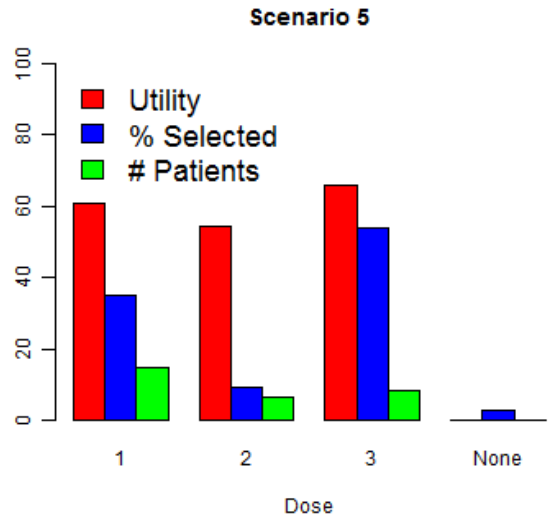

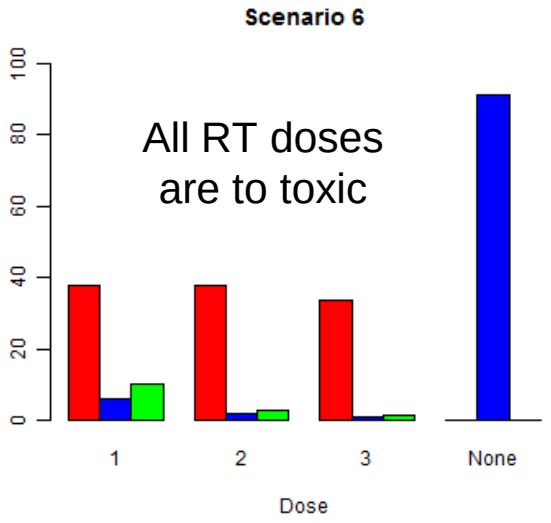

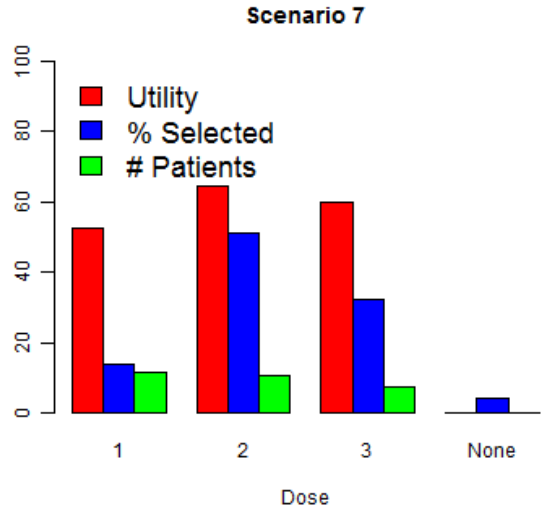

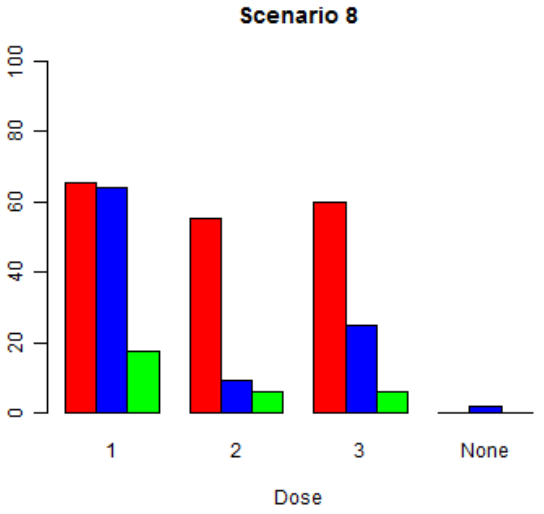

Conclusions About Utility-Based Designs

Utilities

Using joint utilities of (Efficacy, Toxicity) is vastly superior to using Toxicity only and ignoring Efficacy.

Adaptive Randomization

Randomizing among doses with posterior mean utility close to the maximum is insurance against cases where the greedy algorithm gets stuck at an inferior dose.

Safety and Futility

The marginal probability rules work extremely well to screen out unsafe or ineffective doses

A Clinical Trial Treatment Development Design Hierarchy

Statistical Design	Characterization	IQ
Phase I-II Using Efficacy-Toxicity Trade-Offs or Utilities	Honest, Sensible, and Useful, But Not A Panacea	120
Phase I Using A Model-Based Toxicity-Only Method	Reasonably Intelligent, But Flawed Because Efficacy is Ignored	100
Phase I Using Any "3+3" Algorithm	Dumb As a Sack of Hammers. Should Be Illegal.	80

Bayesian Designs for Phase I-II Clinical Trials

A Half-Day Short Course
2020 ASA Biopharmaceutical Section Regulatory-Industry Statistics Workshop

Peter F. Thall and Ying Yuan
Department of Biostatistics
M.D. Anderson Cancer Center

Textbook: Yuan, Nguyen, and Thall
Bayesian Designs for Phase I-II Clinical Trials
Chapman & Hall/CRC Biostatistics Series, 2016

Lecture Schedule (Eastern Standard Time)

Lecture 1: 2:00 – 2:50 **Thall**
- Problems with the Phase I → Phase II paradigm
- Phase I-II designs

(5 minute break)

Lecture 2: 2:55 – 3:40 **Thall**
- The EffTox phase I-II design
- Utility based phase I-II designs

(10 minute break)

Lecture 3: 3:50 – 4:45 **Yuan**
- Model assisted phase I-II designs
- Phase I-II designs for immunotherapies

(5 minute break)

Lecture 4: 4:50 – 5:30 **Yuan**
- Tissue agnostic phase I-II designs
- Handling late onset toxicity and efficacy

OBD vs. MTD

- For targeted and immune therapies, the conventional assumption that efficacy increases with the dose may not hold
  - Efficacy often plateaus or even decreases at high doses
  - To obtain optimal treatment effect, immunotherapy and targeted agents are not necessarily administered at the maximum tolerated dose (MTD)
- The appropriate objective of dose finding trials for immunotherapy and targeted therapy is to find the Optimal Biological Dose (OBD)

OBD vs. MTD

Require a paradigm shift !

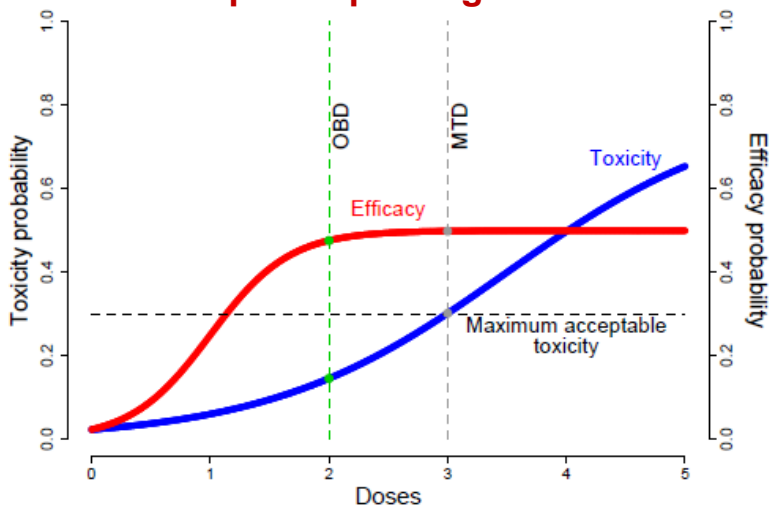

Elements of phase I-II trial design

1. **Toxicity and efficacy outcomes** that characterize potential risks and benefits of the treatment being studied
2. **Risk-benefit trade-off criterion** that characterizes and quantifies the trade-off between efficacy and toxicity for each dose
3. **Statistical model** describing the dose-toxicity and dose-efficacy relationships
4. **Adaptive decision rule** that determines the best dose for the next cohort, based on the (dose, toxicity, efficacy) data from all previous patients
5. **Admissibility rules** that protect patients in the trial from unacceptably toxic or inefficacious doses
6. **Stopping rule** that terminates the trial early if the all doses being considered are unacceptably toxic or inefficacious

Yan, F., Thall PF, Lu KH, Gilbert MR and Yuan, Y. (2018) Phase I-II clinical trial design: A state-of-the-art paradigm for dose finding. *Annals of Oncology* , 29, 694-699.

Phase I-II trial design

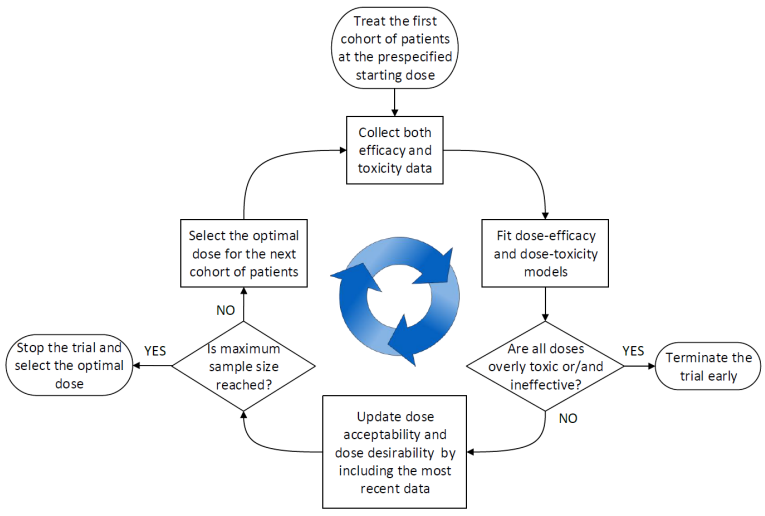

Yan, F., Thall PF, Lu KH, Gilbert MR and Yuan, Y. (2018) Phase I-II clinical trial design: A state-of-the-art paradigm for dose finding. *Annals of Oncology* , 29, 694-699.

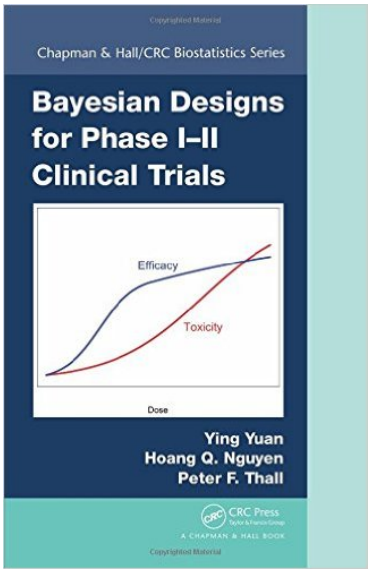

Model-Assisted Phase I-II Trial Designs

- U-BOIN design
- BOIN12 design

Zhou, Y., Lee. J.J. and Yuan, Y. (2019) A Utility-based Bayesian Optimal Interval (U-BOIN) Phase I/II Design to Identify the Optimal Biological Dose for Targeted and Immune Therapies. *Statistics in Medicine*, 38(28):5299-5316.
Lin R, Zhou Y, Yan F, Li D and Yuan Y (2020) BOIN12: Bayesian Optimal Interval Phase I/II Trial Design for Utility-Based Dose Finding in Immunotherapy and Targeted Therapies, JCO precision oncology, minor revision invited

JCO® Precision Oncology
An American Society of Clinical Oncology Journal

SPECIAL SERIES: STATISTICAL METHODS IN PRECESION ONCOLOGY

review article

Model-Assisted Designs for Early-Phase Clinical Trials: Simplicity Meets Superiority

Ying Yuan, PhD[1]; J. Jack Lee, PhD[1]; and Susan G. Hilsenbeck, PhD[2]

Model-assisted designs refer to a class of design that uses a statistical model to derive the design for efficient decision making, similar to model-based design; but like the algorithm-based design, its dose escalation and de-escalation rule can be pre-determined before the onset of the trial, and thus can be implemented as simple a way as the algorithm-based designs.

U-BOIN design

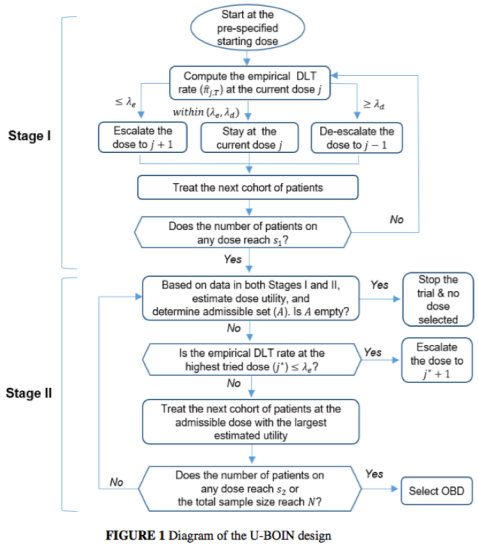

FIGURE 1 Diagram of the U-BOIN design

Zhou, Y., Lee. J.J. and Yuan, Y. (2019) A Utility-based Bayesian Optimal Interval (U-BOIN) Phase I/II Design to Identify the Optimal Biological Dose for Targeted and Immune Therapies. *Statistics in Medicine*, 38(28):5299-5316.

Toxicity and efficacy outcomes

- Consider a phase I-II trial with J prespecified doses
- Let $Y_T = 0, \cdots, Q-1$ denote a Q-level toxicity endpoint, where a higher level represents a more severe toxicity
- Let $Y_E = 0, \cdots, R-1$ denote the R-level efficacy endpoint, where a higher level represents a more desirable response
- (Y_T, Y_E) can be equivalently represented by a single variable Y with $k = R \times Q$ levels
  - (no tox, eff), (no tox, no eff), (tox, eff), (tox, no eff),

Efficacy-toxicity model

- Define $\pi_{jk} = \Pr(Y = k | d = j)$, where d is the dose level
- We assume that Y follows a Dirichlet-multinomial model

$$Y = k | d = j \sim Multinomial(\pi_{j1}, \cdots, \pi_{jK})$$
$$(\pi_{j1}, \cdots, \pi_{jK}) \sim Dirichlet(a_1, \cdots, a_K)$$

where $\sum_{k=1}^{K} a_k = 1$ to have a vague prior with prior sample size of 1.

Efficacy-toxicity model

- At an interim decision time, assume that n_j patients have been treated at dose j, among which n_{jk} patients had outcome $Y = k$.
- Given the interim data $D_j = (n_{j1}, \cdots, n_{jK})$, the posterior is

$$(\pi_{j1}, \cdots, \pi_{jK})|D_j \sim Dirichlet(a_1 + n_{j1}, \cdots, a_K + n_{jK})$$

Risk-benefit tradeoff

- Use utility to measure the desirability (i.e., risk-benefit tradeoff) of the doses

Tox	Toxicity	Response		
		No ($Y_E = 0$)	Yes ($Y_E = 1$)	/PR = 2)
No (Y	No ($Y_T = 0$)	ψ_2 =30	ψ_1 =100	=100
Yes (	Yes ($Y_T = 1$)	ψ_4 =0	ψ_3 =50	=60

Utility

- Let ψ_k denote the utility ascribed to outcome $Y = k$, with ψ_1=100 (most desirable) and ψ_K=0 (least desirable)
- The true mean utility (i.e., desirability) for dose j is given by

$$U_j = \sum_{k=1}^{K} \psi_k \pi_{jk}$$

- The OBD is the dose with the highest desirability
- The estimate of desirability is given by

$$\hat{U}_j = \sum_{k=1}^{K} \psi_k E(\pi_{jk} | D_j)$$

Why utility?

Theorem An alternative approach of quantifying the desirability of dose j based on the marginal probability of toxicity $p_{T,j}$ and the marginal probability of efficacy $p_{E,j}$, i.e.,

$$U_j' = p_{E,j} - w p_{T,j},$$

is a special case of the utility approach with $\psi_2 + \psi_3 = 100$ and $w = \psi_2/\psi_3$.

Toxicity	Response	
	No (Y_E = 0)	Yes (Y_E = 1)
No (Y_T = 0)	50	100
Yes (Y_T = 1)	0	50

$$U_j' = p_{E,j} - p_{T,j}$$

Lin R, Zhou Y, Yan F, Li D and Yuan Y (2020) BOIN12: Bayesian Optimal Interval Phase I/II Trial Design for Utility-Based Dose Finding in Immunotherapy and Targeted Therapies, JCO precision oncology, revision invited

Steps to elicit utility

1. Fix the value of the utility for the most desirable outcome $Y = 1$ as $\psi_1 = 100$, and the least desirable outcome $Y = K$ as $\psi_K = 0$.
2. Ask clinicians to use these two utilities as a reference to score the utility values $\psi_2, \cdots, \psi_{K-1}$ for the other $K - 1$ possible outcomes $Y = 2, \cdots, K - 1$ to quantify the risk-benefit trade-off under each outcome.

Toxicity	Response	
	No ($Y_E = 0$)	Yes ($Y_E = 1$)
No ($Y_T = 0$)	30	100
Yes ($Y_T = 1$)	0	50

Elicit utility (cont.)

- One possible criticism for using the utility values is that they require subjective input. However, we are inclined to view this as a strength rather than a weakness.
- This is because the utilities must be elicited from the physicians planning the trial, and thus their numerical values are based on the physician's experience in treating the disease and observing the good and bad effects.

Elicit utility (cont.)

- The process of specifying the utility requires physicians to carefully consider the potential risks and benefits of the treatment that underlie their clinical decision making in a more formal way and incorporate that into the trial.
- In addition, our simulation study and previous studies show that the design is generally not sensitive to the numerical values of the utility as long as it reflects a similar trend.

Dose admissibility criteria

- Let ϕ_T denote the upper limit of the toxicity rate, and ϕ_E denote the lower limit of the efficacy rate, specified by physicians.
- Dose j is defined as **admissible** if the following two criteria are satisfied

$$\text{(Safety)} \quad \Pr(\pi_T < \phi_T | D) > C_T$$
$$\text{(Efficacy)} \quad \Pr(\pi_E > \phi_E | D) > C_E$$

where C_T (e.g., = 0.05) and C_E (e.g., = 0.1) are prespecified probability cutoffs.

U-BOIN design

- U-BOIN is consisted of two seamlessly connected stages
- Stage I: dose escalation based on toxicity. The objective is to quickly explore the dose space to identify a set of admissible doses that are reasonably efficacious and safe
- Stage II: adaptively allocate patients to the estimated OBD based on the toxicity and efficacy data accrued from both stages I and II

Stage I: dose escalation using BOIN

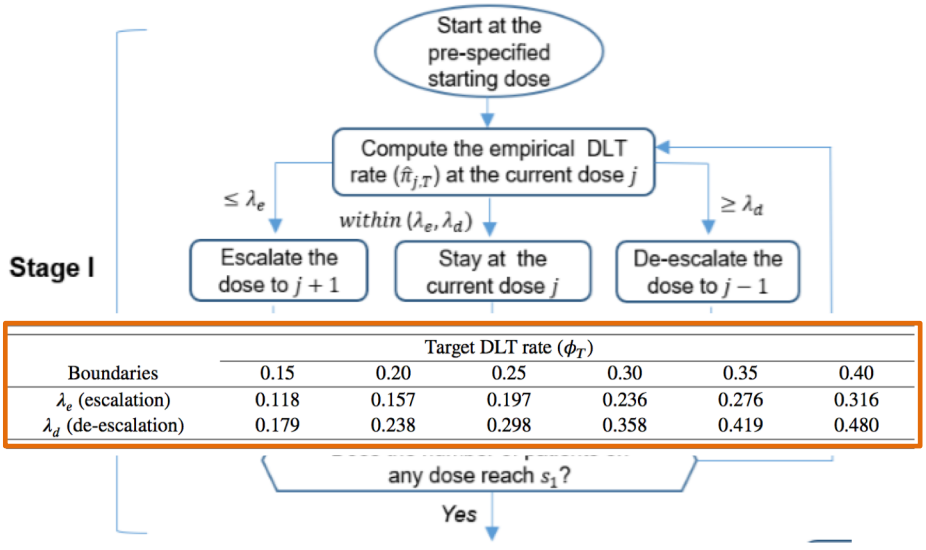

Boundaries	Target DLT rate (ϕ_T)					
	0.15	0.20	0.25	0.30	0.35	0.40
λ_e (escalation)	0.118	0.157	0.197	0.236	0.276	0.316
λ_d (de-escalation)	0.179	0.238	0.298	0.358	0.419	0.480

Stage II: identify OBD

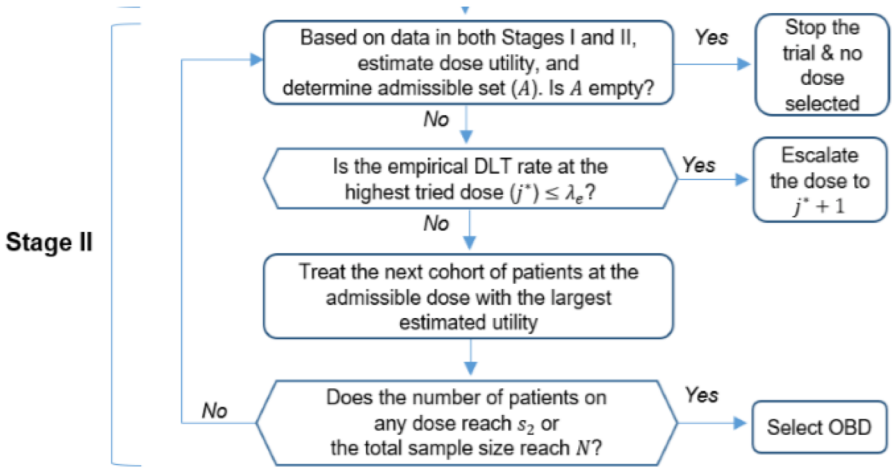

Two-stage dose-finding algorithm

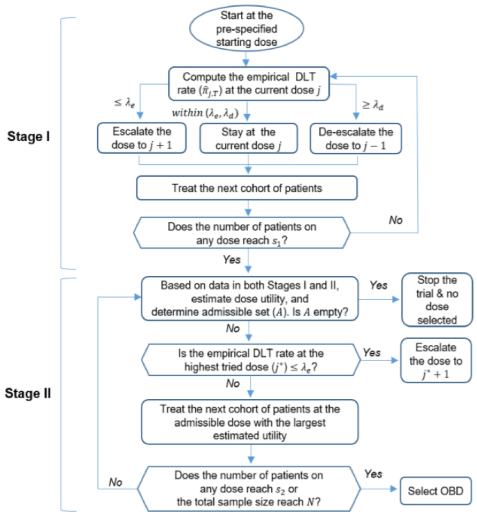

FIGURE 1 Diagram of the U-BOIN design

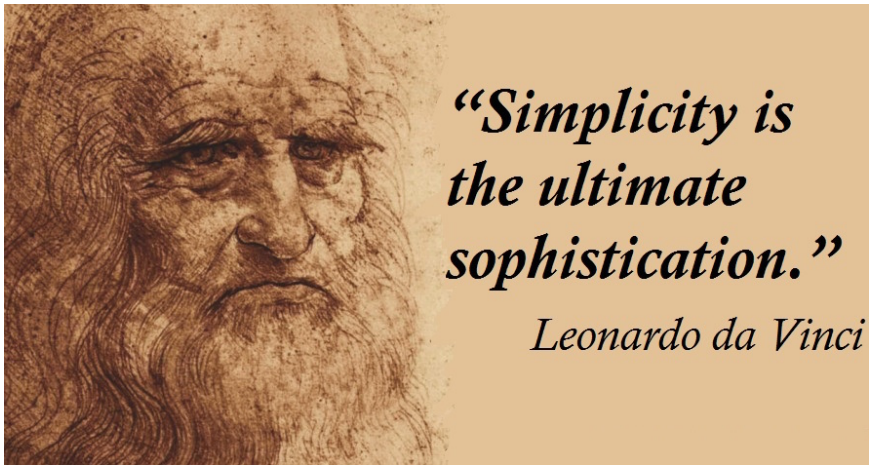

Stage I decision table

Decision	The number of patients treated at the current dose			
	3	**6**	**9**	**12**
Escalate if No. of DLT <=	0	1	1	2
De-escalate if No. of DLT <=	1	2	3	4
Eliminate if No. of DLT >=	3	4	5	6

Stage II decision table

Table 2. Utility table when 3 patients are treated on a dose

#Eff	#Tox	#(Eff=1,Tox=0)	Utility	#Eff	#Tox	#(Eff=1,Tox=0)	Utility
<1	Any	Any	0	2	1	2	61.2
1	0	1	51.2	2	2	0	43.8
1	1	0	38.8	2	2	1	48.8
1	1	1	43.8	2	> 2	Any	0
1	2	0	31.2	3	0	3	86.2
1	2	1	36.2	3	1	2	73.8
1	> 2	Any	0	3	2	1	61.2
2	0	2	68.8	3	> 2	Any	0
2	1	1	56.2				

Delayed response

- In some trials, Y_E may require a long time to be ascertained
- Consequence: some Y_E are unavailable at the interim time, making adaptive decisions difficult
- Approach
  - Use multiple imputation to impute unobserved Y_E
  - Leverage the measure of biological activity (e.g., immune response) to impute Y_E

BOIN12 design

Another model-assisted phase I-II design

BOIN12 design

Compared to U-BOIN

- BOIN12 is a single-stage design targeting the OBD from the beginning of the trial
- BOIN12 generally requires a small sample size, thus is particularly suitable when the number of dose is larger (e.g., >3)
- U-BOIN, however, is a good choice when
  - The MTD is of substantial interest
  - The number of doses is small (≤ 3)
  - Interested in collecting some data (e.g., PK/PD) over multiple doses

Lin R, Zhou Y, Yan F, Li D and Yuan Y (2020) BOIN12: Bayesian Optimal Interval Phase I/II Trial Design for Utility-Based Dose Finding in Immunotherapy and Targeted Therapies, JCO precision oncology, minor revision invited

Toxicity and efficacy outcomes

- Let $Y_T = 0, \cdots, Q-1$ denote a Q-level toxicity endpoint, where a higher level represents a more severe toxicity
- Let $Y_E = 0, \cdots, R-1$ denote the R-level efficacy endpoint, where a higher level represents a more desirable response
- (Y_T, Y_E) can be equivalently represented by a multinomial variable Y with $K = R \times Q$ levels

$$Y = k|d \sim Multinomial\big(\pi_1(d), \cdots, \pi_K(d)\big), \qquad k = 1, \ldots, K$$

Lin R, Zhou Y, Yan F, Li D and Yuan Y (2020) BOIN12: Bayesian Optimal Interval Phase I/II Trial Design for Utility-Based Dose Finding in Immunotherapy and Targeted Therapies, JCO precision oncology, minor revision invited

Risk-benefit tradeoff

- Use utility to measure the desirability of the doses

Toxicity	Response	
	No ($Y_E = 0$)	Yes ($Y_E = 1$)
No ($Y_T = 0$)	$u_2 = 30$	$u_1 = 100$
Yes ($Y_T = 1$)	$u_4 = 0$	$u_3 = 50$

Quasi-binomial model for utility

- Let $u_1, \ldots, u_4$ denote the utility ascribed to the four possible outcomes
- The mean utility for dose d is given by

$$u(d) = u_1\pi_1(d) + u_2\pi_2(d) + u_3\pi_3(d) + u_4\pi_4(d)$$

- Define standardized utility $u^*(d) = u(d)/100$, such that $u^*(d) \in [0, 1]$ and is a weighted average of ($\pi_1(d)$, $\pi_2(d)$, $\pi_3(d)$, $\pi_4(d)$)
- $u^*(d)$ thus can be viewed as a probability and modelled using the binomial distribution with "quasi-binomial" data $(x(d), n(d))$, where

$$x(d) = \frac{u_1 y_1(d) + u_2 y_2(d) + u_3 y_3(d) + u_4 y_4(d)}{100}$$

and $n(d)$ is the number of patients treated at d.

Quasi-binomial model for utility

- Thus, the quasi-binomial likelihood of the observe data $D(d)$ is

$$L(\, D(d) \mid u^*(d)\,) \propto \big(u^*(d)\big)^{x(d)}\big(1 - u^*(d)\big)^{n(d)-x(d)}$$

- Under the Bayesian framework, assign $u^*(d)$ a Beta prior, i.e., $u^*(d) \sim Beta(\alpha, \beta)$, the posterior distribution of $u^*(d)$ arises as,

$$u^*(d) \mid D(d) \sim Beta\big(\alpha + x(d), \beta + n(d) - x(d)\big)$$

Adaptive decision rule

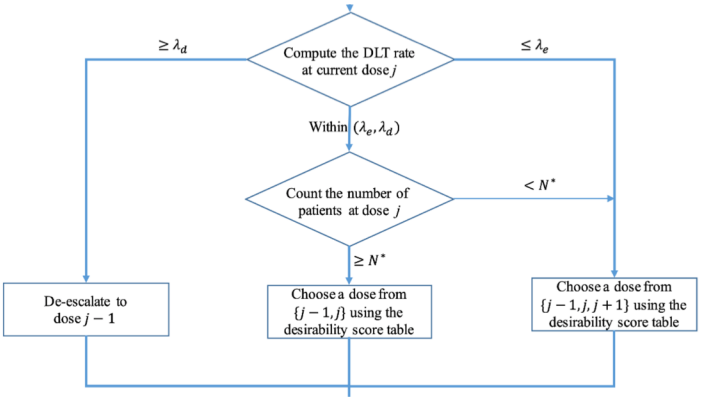

Table 1. Optimal dose escalation and de-escalation boundaries

	Target toxicity rate for the MTD						
Boundary	0.1	0.15	0.2	0.25	0.3	0.35	0.4
λ_e (escalation)	0.078	0.118	0.157	0.197	0.236	0.276	0.316
λ_d (de-escalation)	0.119	0.179	0.238	0.298	0.358	0.419	0.479

Adaptive decision rule

Table S2. Decision table for the BOIN12 design with the target toxicity rate $\phi_T = 0.35$, the target efficacy rate $\phi_E = 0.25$, and the utility specification $u_1 = 100$, $u_2 = 30$, $u_3 = 50$, and $u_4 = 0$, up to six patients.

No. Patients	No. DLT	No. Eff	No. (No DLT, Eff)	Desirability Score	No. Patients	No. DLT	No. Eff	No. (No DLT, Eff)	Desirability Score
0	0	0	0	59	6	2	1	1	31
3	0	0	0	30	6	2	2	0	39
3	0	1	1	53	6	2	2	1	45
3	0	2	2	71	6	2	2	2	49
3	0	3	3	79	6	2	3	1	56
3	1	0	0	20	6	2	3	2	60
3	1	1	0	36	6	2	3	3	65
3	1	1	1	43	6	2	4	2	69
3	1	2	1	58	6	2	4	3	72
3	1	2	2	63	6	2	4	4	77
3	1	3	2	74	6	2	5	3	2
3	2	0	0	11	6	2	5	4	6
3	2	1	0	25	6	2	6	4	8

Simulation

- Compare BOIN12, EffTox, and 3+3+CE designs
- 3+3+CE design: dose escalation using the 3+3 design, followed by a cohort expansion at the identified MTD using Simon's two-stage design
- Five dose levels
- Patients are treated in cohorts of 3 with N=36

Efficacy / Toxicity	Yes	No
No	$u_1 = 100$	$u_2 = 40$
Yes	$u_3 = 60$	$u_4 = 0$

Scenarios

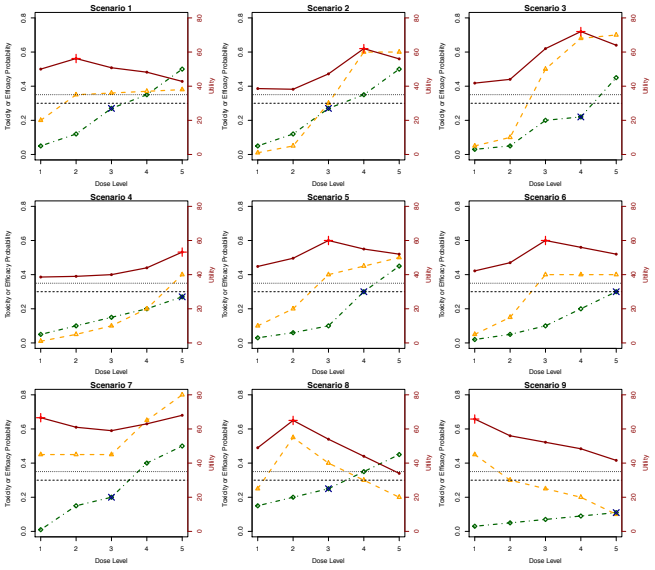

Correct selection percentage

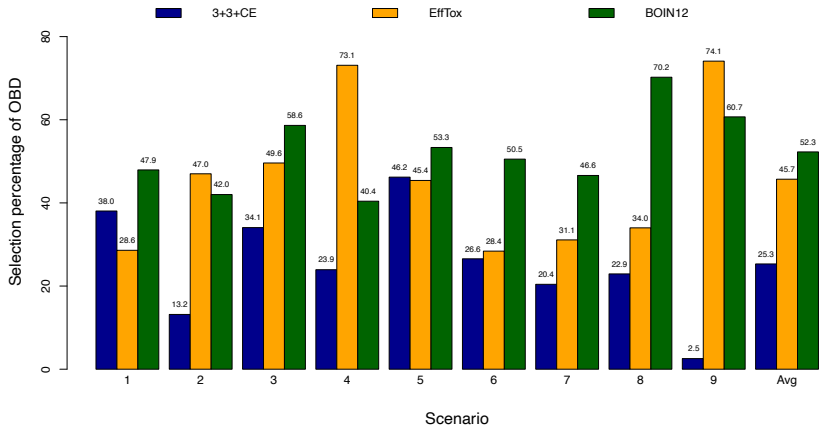

Number of patients at OBD

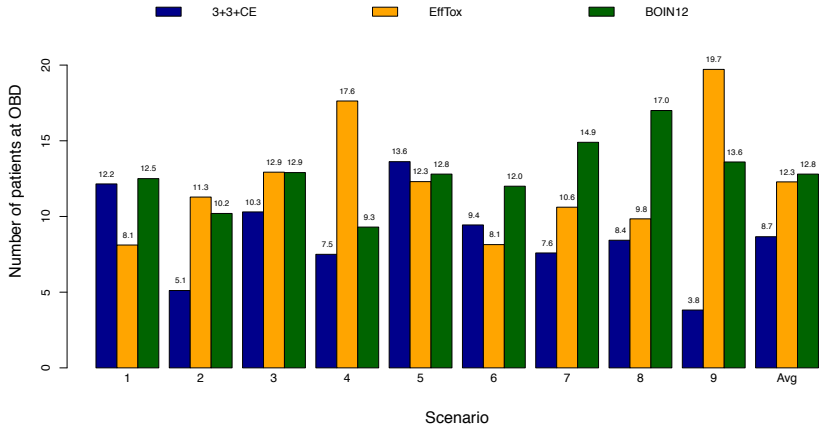

Number of patients overdosed

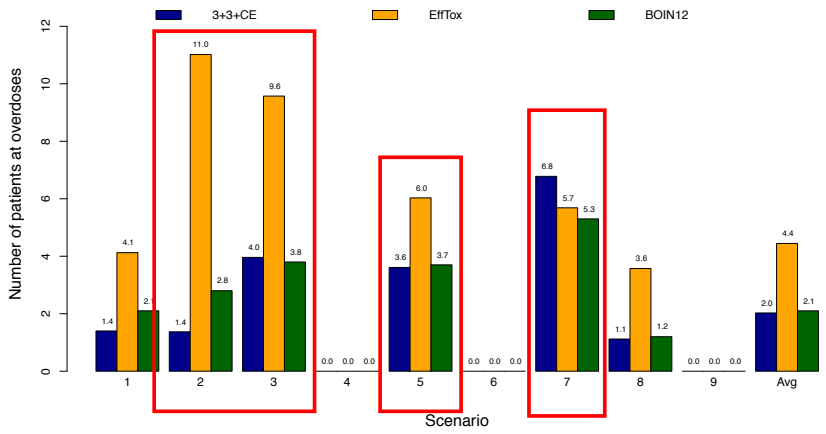

www.trialdesign.org

www.trialdesign.org

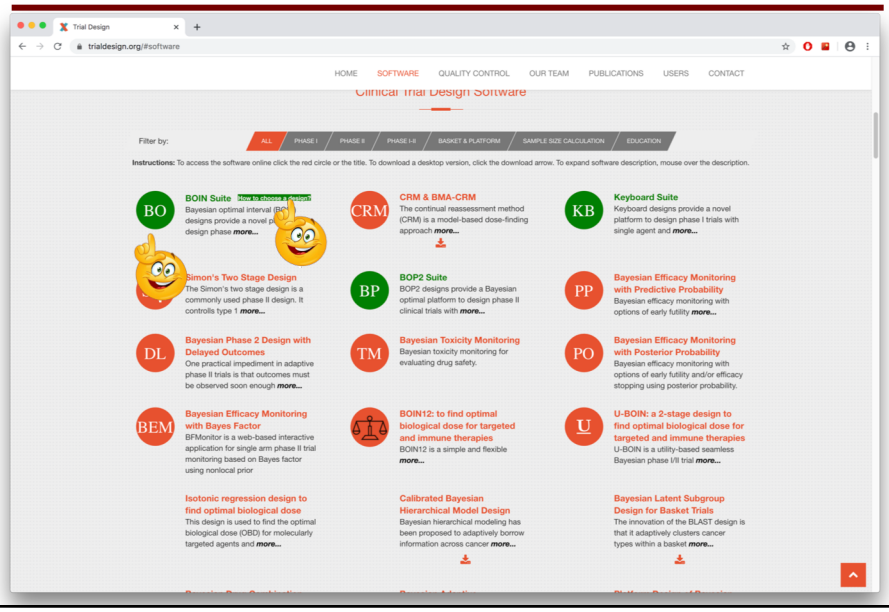

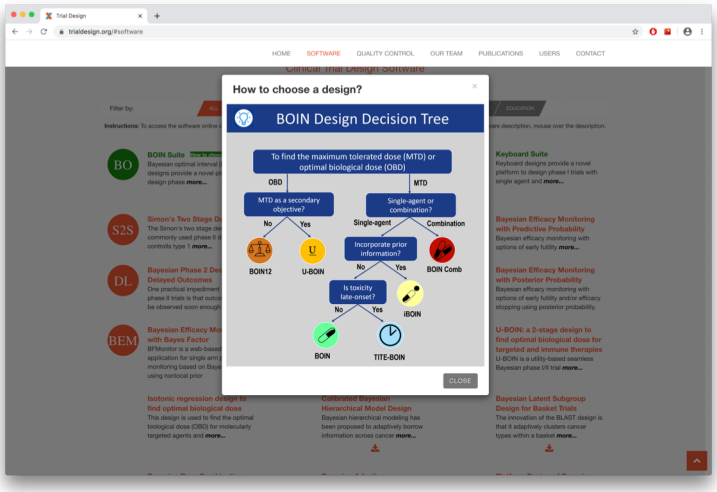

How to choose a design?

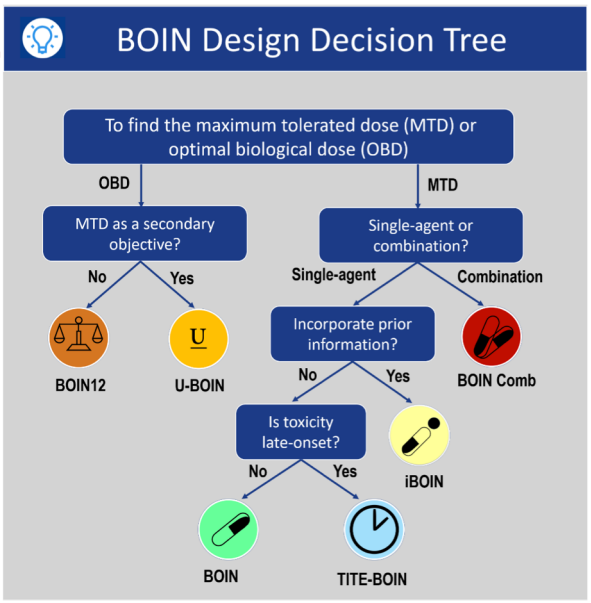

CLOSE

www.trialdesign.org

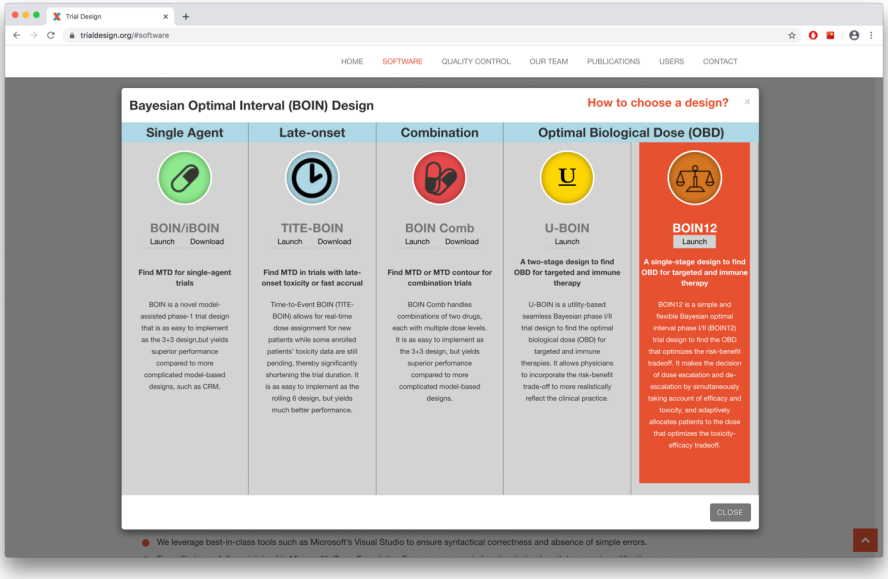

Using multiple endpoints to improve the efficiency of phase I-II design for immunotherapy

Liu, S., Guo, B. and Yuan, Y. (2018) A Bayesian Phase I/II Trial Design for Immunotherapy. *Journal of the American Statistical Association*, 113, 1016-1027.

Notation

- Consider a phase I-II trial with J prespecified doses, $d_1 < \cdots < d_j$, under investigation.
- Let Y_T denote the binary toxicity outcome, with $Y_T = 1$ indicating toxicity, and $= 0$ otherwise.
- Let Y_E denote the tumor response, with $Y_E = 0, 1,$ and 2 indicating PD, SD and PR/CR, respectively.
- Let Y_I denote a measure of the immune response (e.g., the count of CD8+ T-cells or the concentration of cytokine), which takes a real value after appropriate transformation.
- Adaptive decisions in the trial (e.g., dose assignment and selection) are based on the behavior of the trinary vector (Y_I, Y_T, Y_E) as a function of dose d.

Factorization

- To reflect the fact that in immunotherapy, clinical responses rely on the activation of the immune system, the joint distribution $[Y_I, Y_T, Y_E \,|d]$ is factorized as

$$[Y_I, Y_T, Y_E \,|d, \boldsymbol{\theta}] = [Y_I|d, \boldsymbol{\theta}_1][Y_T, Y_E|d, Y_I, \boldsymbol{\theta}_2]$$

where $\boldsymbol{\theta}$ is the vector of the parameters, and $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are subvectors of $\boldsymbol{\theta}$.

Model for immune response Y_I

- Model immune response $[Y_I|d, \boldsymbol{\theta}_1]$ using an Emax model,

$$Y_I|d, \theta_1 = \alpha_0 + \frac{\alpha_1 d^{\alpha_3}}{\alpha_2^{\alpha_3} + d^{\alpha_3}} + \varepsilon$$

where α_0 is the baseline immune activity in the absence of the IT; α_1 is the maximum immune activity that is possibly achieved by the IT above the baseline activity (i.e., E_{max}); α_2 is the dose that produces half of the maximum immune activity (i.e., ED_{50}); α_3 is the Hill factor that controls the steepness of the dose-response curve; and ε is the random error, which is normally distributed with a mean of 0 and variance σ^2, i.e., $\varepsilon \sim N(0, \sigma^2)$.

Latent variable model for Y_T and Y_E

- Let Z_T and Z_E denote two continuous latent variables that are related to Y_T and Y_E, respectively, as follows,

$$Y_T = \begin{cases} 0 & if\ Z_T < \zeta_1 \\ 1 & if\ Z_T \geq \zeta_1 \end{cases} \quad \text{and} \quad Y_E = \begin{cases} 0 & if\ Z_E < \xi_1 \\ 1 & if\ \xi_1 \leq Z_E < \xi_2 \\ 2 & if\ Z_E \geq \xi_2 \end{cases}$$

  where ζ_1, ξ_1 and ξ_2 are unknown cut points.
- Z_T and Z_E can be interpreted as the patient's latent traits, and Y_T and Y_E are the clinical manifestations of unobserved Z_T and Z_E.
- When Z_T and Z_E pass certain thresholds, certain clinical outcomes (e. g., toxicity, CR/PR) are observed.

Toxicity model

- $[Z_T, Z_E | d, Y_I]$ follows a bivariate normal distribution

$$\begin{pmatrix} Z_T \\ Z_E \end{pmatrix} \Big| Y_I, d \sim N_2 \left(\begin{pmatrix} \mu_T(Y_I, d) \\ \mu_E(Y_I, d) \end{pmatrix}, \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix} \right)$$

  where $\mu_k(Y_I, d) = E(Z_k | Y_I, d)$, $k = E$ or T, is the conditional mean of Z_k.
- $\mu_T(Y_I, d)$ follows a threshold model,

$$\mu_T(Y_I, d) = \beta_0 + \beta_1 d + I(Y_I > \beta_3)\beta_2 Y_I$$

  Where β_0, β_1, β_2, β_3 are unknown parameters, and the indicator $I(Y_I > \beta_3) = 1$ when $Y_I > \beta_3$, and 0 otherwise.
- Immune response Y_I induces toxicity only when it passes threshold β_3.

Efficacy model

- Model the mean structure $\mu_E(Y_I, d)$ for efficacy as

$$\mu_E(Y_I, d) = \gamma_0 + \gamma_1 Y_I + \gamma_2 Y_I^2$$

- Although the quadratic model cannot directly take an increase-then-plateau shape, it works reasonably well in that case.
- In addition, as the Emax model allows Y_I to plateau with the dose d, the above model indeed accommodates the case that efficacy Y_E plateaus with d.

Dose admissibility criteria

- To safeguard patients from overly toxic or ineffective doses.
- Let $\pi_T = \Pr(Y_T = 1|d)$ denote the toxicity rate and $\pi_E = \Pr(Y_E > 0|d)$ denote the response rate of SD/PR/CR.
- Let ϕ_T denote the upper limit of the toxicity rate, and ϕ_E denote the lower limit of the response rate, specified by physicians.

Dose admissibility criteria

- A dose d is defined as **admissible** if it satisfies both the safety requirement
$$\Pr(\pi_T > \phi_T | \mathrm{D}_n) < C_T$$
and the efficacy requirement
$$\Pr(\pi_E < \phi_E | \mathrm{D}_n) < C_E$$
where C_T and C_E are prespecified cutoffs, and D_n is the observed data from n treated patients.
- We can also add immune response to define admissible $\Pr(\pi_I > \phi_I | \mathrm{D}_n) < C_I$
- Let $\mathcal{A}$ denote all admissible doses
- Dose assignment and selection are restricted to $\mathcal{A}$

Desirability and optimal biological dose

- Use a utility $U(Y_I, Y_T, Y_E)$ to map multi-dimensional outcomes into a single index to measure the desirability of a dose in terms of the risk-benefit tradeoff.

		Efficacy		
Toxicity	Immune response	PD ($Y_E = 0$)	SD ($Y_E = 1$)	CR/PR ($Y_E = 2$)
No ($Y_T = 0$)	Desirable ($\tilde{Y}_I = 1$)	5	70	**100**
	Undesirable ($\tilde{Y}_I = 0$)	0	50	80
Yes ($Y_T = 1$)	Desirable ($\tilde{Y}_I = 1$)	0	20	45
	Undesirable ($\tilde{Y}_I = 0$)	**0**	10	35

Calculation of utility

- For a given dose d, its true utility is given by

$$E(U(d)|\boldsymbol{\theta}) = \int U(\widetilde{Y}_I, Y_T, Y_E) f(\widetilde{Y}_I, Y_T, Y_E | d, \boldsymbol{\theta})\, d\widetilde{Y}_I dY_T dY_E$$

- Since $\boldsymbol{\theta}$ is not known, the utility of dose d must be estimated.
- Given interim data D_n collected from the first n patients at a decision-making point in the trial, the utility of dose d is estimated by its posterior mean

$$E(U(d)|\mathsf{D}_n) = \int E(U(d)|\theta) p(\theta|\mathsf{D}_n)\, d\theta$$

Dose-finding algorithm

- The first cohort of patients is treated at the lowest dose d_1.
- Assume that r cohort(s) of patients have been enrolled in the trial, where $r = 1, \cdots, R-1$, and let d_h denote the current highest tried dose. To assign a dose to the $(r+1)$th cohort of patients:
  - If the posterior probability of toxicity at d_h satisfies $\Pr(\pi_T(d_h) < \phi_T|\mathsf{D}_n) > C_{es}$ and $d_h \neq d_J$, then we treat the $(r+1)$th cohort of patients at d_{h+1} where C_{es} denote the probability for escalation based on toxicity, and $n = m \times r$.

Dose-finding algorithm (cont.)

- Otherwise, we identify the admissible set $\mathcal{A}$ and adaptively randomize the $(r+1)$th patient or cohort of patients to dose $d_j \in \mathcal{A}$ with probability

$$\psi_{j,n} = \Pr[U(d_j) = \max\{U(d_{j'}),\ j' \in \mathcal{A}\} \mid \mathsf{D}_n]$$

  If A is empty, the trial is terminated.
- Once the maximum sample size of N is exhausted, the dose in $\mathcal{A}$ with the largest posterior mean utility $E(U(d)|\mathsf{D}_n)$ is recommended.

Shotgun: A Tissue-Agnostic Phase I-II Design

Tissue-agnostic trials

- Tissue-agnostic clinical trials have become increasingly important in the development of targeted therapy and immunotherapy for cancer
- Tissue-agnostic drugs approved by FDA
  - Entrectinib for treating NTRK gene fusion cancer patients, regardless of cancer types
  - Pembrolizumab for tumors with microsatellite instability-high (MSI-H) or mismatch repair decifient (dMMR) tumors
  - Larotrectinib for NTRK gene fusion tumors

Conventional design

- One indication at a time, grossly inefficient !

Shotgun design

Shotgun design

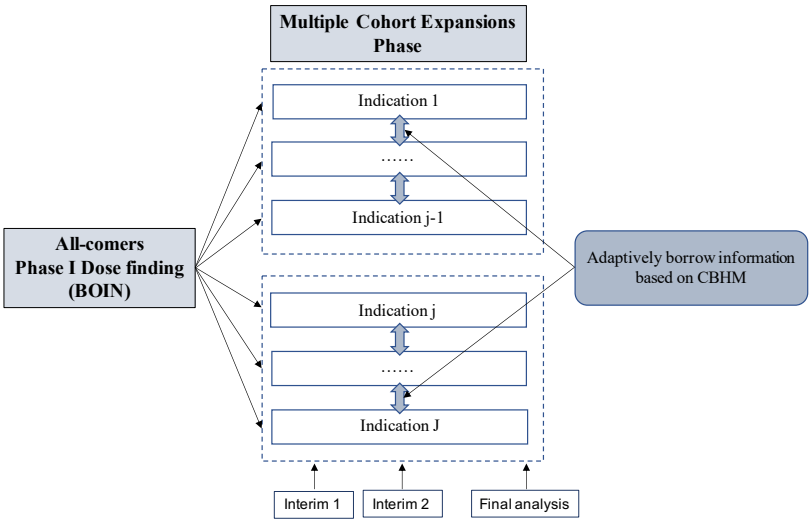

Stage I: dose escalation using BOIN

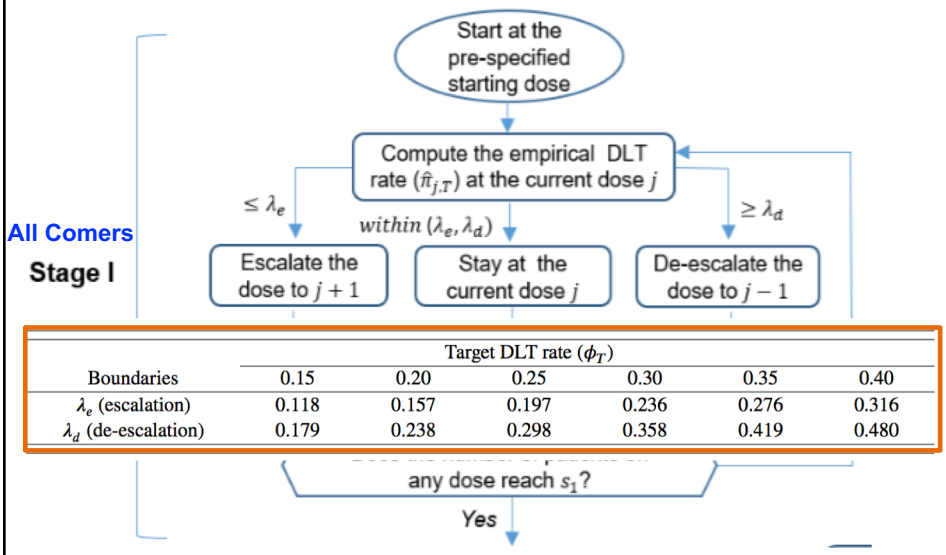

Boundaries	Target DLT rate (ϕ_T)					
	0.15	0.20	0.25	0.30	0.35	0.40
λ_e (escalation)	0.118	0.157	0.197	0.236	0.276	0.316
λ_d (de-escalation)	0.179	0.238	0.298	0.358	0.419	0.480

Stage II: multiple cohort expansion

- After the RP2D is determined, indication-specific cohort expansions are initiated to assess the efficacy of the drug in each of the G indications.
- The toxicity and efficacy data from the patients who were treated at the RP2D in stage I will be included in stage II to improve the efficiency of the design.
- Borrowing information, e.g., using the Bayesian hierarchical model (BHM, Thall, et al., 2003), across cohorts is attractive

Bayesian hierarchical model

- Let q_g denote the response rate in indication g, $q_{0,g}$ denote null response rate that is deemed futile, and $q_{1,g}$ denote the target response rate that is deemed promising.
- BHM

$$x_g|q_g \sim Binomial(q_g)$$

$$\theta_g = \log\left(\frac{q_g}{1-q_g}\right) - \log\left(\frac{q_{0,g}}{1-q_{0,g}}\right)$$

$$\theta_g|\theta,\sigma^2 \sim N(\theta,\sigma^2)$$

where x_g is the number of responses in indication g, $g = 1,\cdots,G$

Clustered BHM (CBHM)

- The exchangeable assumption underlying the BHM, however, is not always appropriate, resulting in inflated type I error or low power
  - It is not uncommon that that some tumors are responsive to the treatment, while others are not
- CBHM is to address this issue: based on the interim data, cluster the cohorts into responsive and non-responsive subgroups; and then use the BHM to borrow information within each subgroup.

Clustered BHM (CBHM)

- **Clustering Step**
  - Bayesian clustering rule: an indication is allocated to the responsive cluster R if it satisfies

$$\Pr\left(q_g > \frac{q_{0,g}+q_{1,g}}{2} \middle| data\right) > 0.5\left(\frac{n_g}{N_{g,2}}\right)^w;$$

    otherwise allocated to the non-responsive cluster $\bar{R}$, where n_g is the interim sample size, $N_{g,2}$ is the prespecified maximum sample size of indication g for phase II, and w is a tuning parameter (recommended w = 2 or 3)
  - The interim-sample-size-dependent adaptive cutoff improves the performance, borrowing the idea from the BOP2 design

Clustered BHM (CBHM)

- **Borrowing Step**
  - Apply the BHM to responsive cluster R and non-responsive cluster $\bar{R}$ independently

$$x_g|q_g \sim Binomial(q_g), \qquad g \in R \text{ or } \bar{R}$$

$$\theta_g = \log\left(\frac{q_g}{1-q_g}\right) - \log\left(\frac{q_{0,g}}{1-q_{0,g}}\right)$$

$$\theta_g|\theta,\sigma^2 \sim N(\theta,\sigma^2)$$

$$\theta \sim N(0,10^6), \sigma^2 \sim N(10^{-6},10^{-6})$$

  - When the number of indication in a cluster is 1, use beta-binomial model.

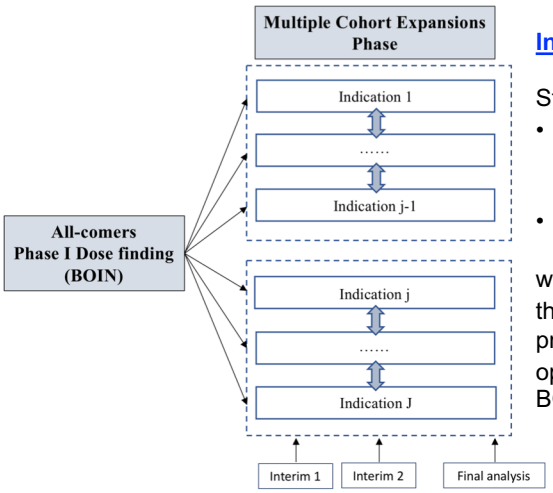

Interim go/no-go rule:

Stop the accrual to indication g if

- $\Pr(q_g \le q_{0,g}|data) < C(n_g)$ (futility)
- $\Pr(p_g < \phi_T|data) > C_T$ (toxicity)

where p_g is the toxicity rate, ϕ_T is the toxicity upper limit, C_T is a probability cutoff, and $C(n_g)$ is the optimal probability cutoff taken from BOP2 design (www.trialdesign.org)

Simulation

- 4 doses, and 5 indications with accrual rate of 3, 2.5, 2, 1.5 and 2 per month
- n=24 for stage I, and maximum n=21 per indication
- Null response rate $q_{0,1} = \cdots = q_{0,5} = 0.05$, and alternative response rate $q_{1,1} = \cdots = q_{1,5} = 0.3$
- Compare shotgun with the convention design (3+3 design followed by Simon's two-stage design)

Performance metrics

- Correct discovery rate (CDR)

$$CDR = \frac{Number\ of\ correct\ discoveries}{Number\ of\ promising\ indications}$$

- False discovery rate (FDR)

$$FDR = \frac{Number\ of\ false\ discoveries}{Number\ of\ indications\ claimined\ as\ promising}$$

- Adjusted discovery rate (ADR)

$$ADR = \frac{Number\ of\ correct\ discoveries - Number\ of\ false\ discoveries}{Number\ of\ promising\ indications}$$

Correct discovery rate (CDR)

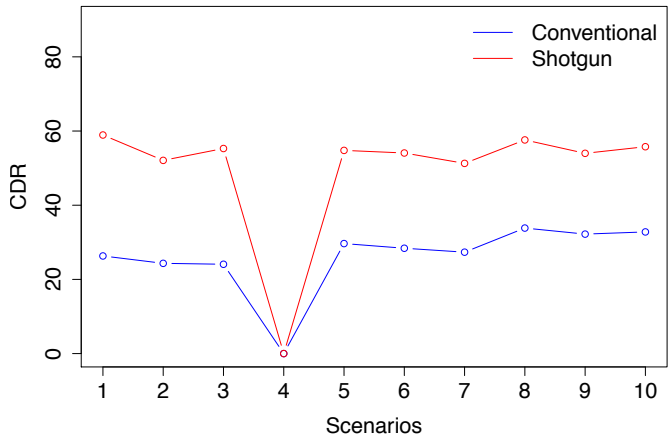

False discovery rate (FDR)

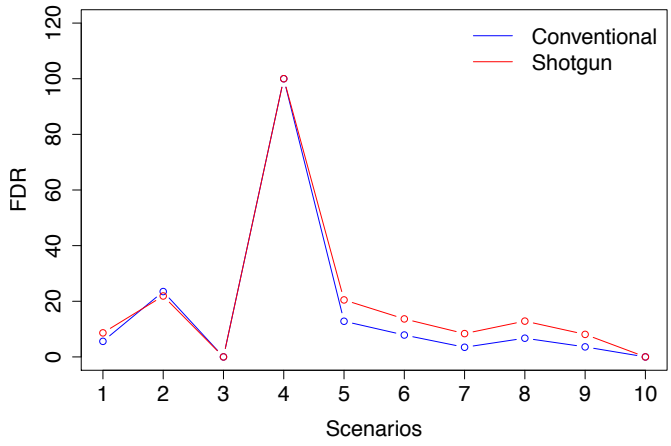

Adjusted discovery rate (ADR)

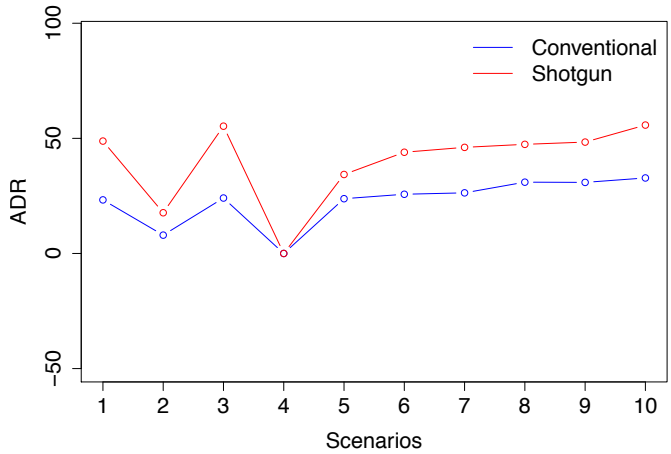

Conclusion

- Shotgun design provides a highly efficient and flexible Bayesian phase I-II method to develop tumor-agnostic drugs
  - Eliminates the white space between phase I and phase II
  - Allows the total number of interim analyses and interim times to vary from one indication to another
  - Borrows information efficiently and accurately across indications using the novel CBHM

Handle Delayed Toxicity and Efficacy

Late-onset efficacy and toxicity

- A major impediment to implement adaptive designs in practice is that outcomes must be observed soon enough to apply the adaptive decision rules to choose treatments or doses for newly accrued patients.
- This is true for all outcome-dependent adaptive trial designs, regardless of the phase and type.
- In phase I-II trials, this problem arises if either toxicity or efficacy is not scored quickly, relative to the trial's accrual rate.

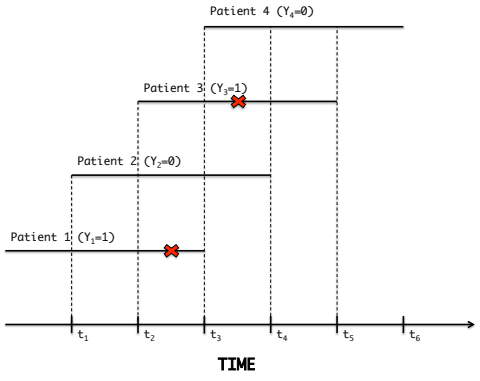

Challenge: how to treat the new patients while waiting to evaluate the outcomes of the previous patients?

Trial example

- A phase I–II dose-finding to determine the optimal dose of three fractionated stereotactic radiation therapy (SBRT) doses, given either with or without a novel radiomodulating agent.
- Toxicity is defined as a grade 3 or 4 gastro-intestinal (GI) toxicity, occurring within 90 days from the start of therapy.
- Efficacy is defined as stable disease (SD) or better, compared to baseline, as evaluated at day 90 from the start of therapy.
- Expected accrual rate: 2 patients/month in each arm
- LO-EffTox design was used, and the accrual has been completed!

Jin, I.H., Liu, S., Thall, P. and Yuan, Y. (2014) Using Data Augmentation to Facilitate Conduct of Phase I/II Clinical Trials with Delayed Outcomes. Journal of American Statistical Association, 109, 525-536.

Three general approaches

- Weight the observation with incompleted follow-up with a “partial credit” (Cheung and Chappell, 2000)
  - Weight the observation with its follow-up time
- Model the toxicity and efficacy outcomes as time-to-event endpoints (Yuan and Yin, 2009)
- Regard unobserved toxicity and efficacy outcomes as missing data and handle them using missing data methodology

Three general approaches

- Weight the observation with incompleted follow-up with a “partial credit” (Cheung and Chappell, 2000)
  - Weight non-DLT with its follow-up time
- Model the toxicity and efficacy outcomes as time-to-event endpoints (Yuan and Yin, 2009)
- **Regard unobserved toxicity and efficacy outcomes as missing data and handle them using missing data methodology**

Jin, I.H., Liu, S., Thall, P. and Yuan, Y. (2014) Using Data Augmentation to Facilitate Conduct of Phase I/II Clinical Trials with Delayed Outcomes. Journal of American Statistical Association, 109, 525-536.

Delayed outcomes as missing data

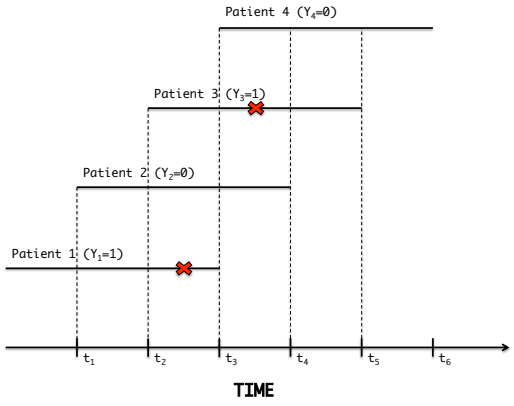

Delayed outcomes as missing data

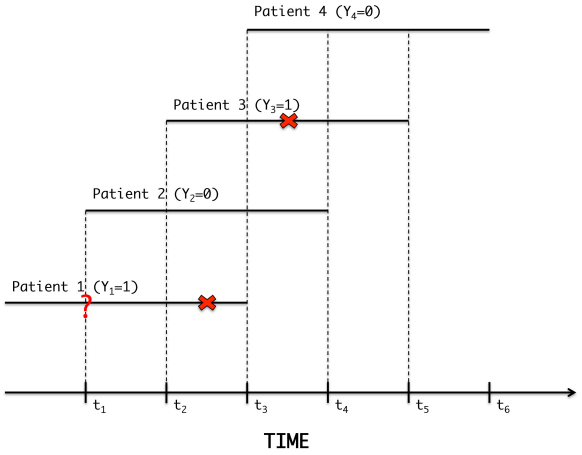

Delayed outcomes as missing data

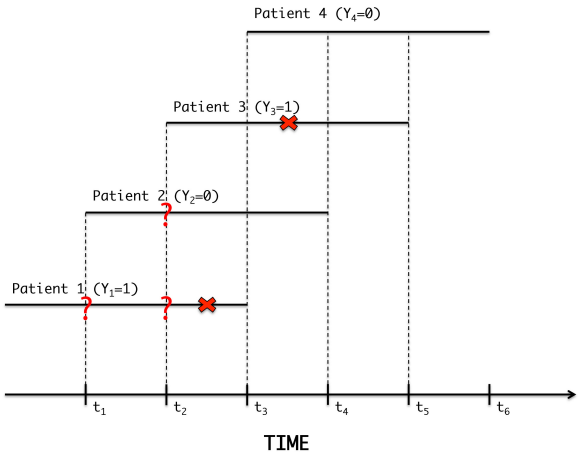

Delayed outcomes as missing data

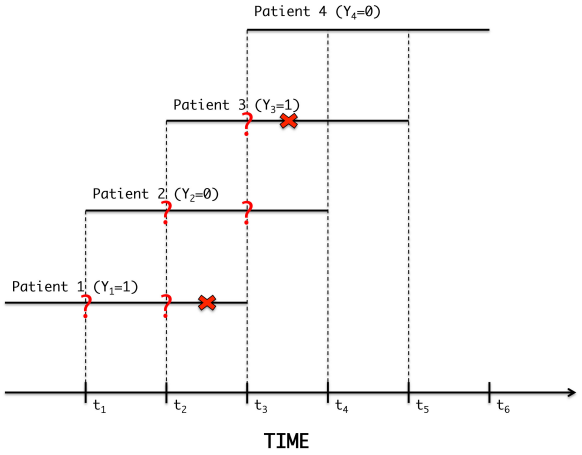

Advantages of missing-data approach

- Intuitive and natural
- Likelihood based, thus efficient and rigorous
- Very general, can be used with almost any adaptive designs to handle delayed outcomes
  - phase I (Liu, Yin and Yuan, 2013), phase II (Cai, Liu and Yuan, 2014), phase I-II (Jin et al., 2014), ...
- Include the original design without delayed outcomes as a special case

General approach

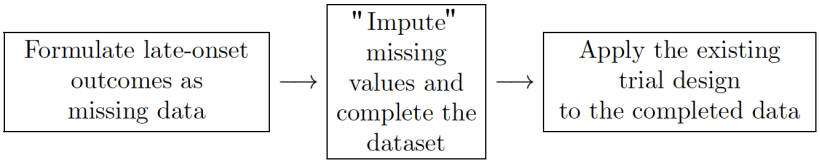

Notations

- Y_T and Y_E denote binary toxicity and efficacy outcomes
- U_T and U_E denote fixed follow-up times for assessing toxicity and efficacy
- X_T and X_E denote times to toxicity and efficacy
- V denote the patient's follow up time, by design $V \leq U$
- $X_T^O = V \wedge X_T$ and $X_E^O = V \wedge X_E$

Notations

- Missing data

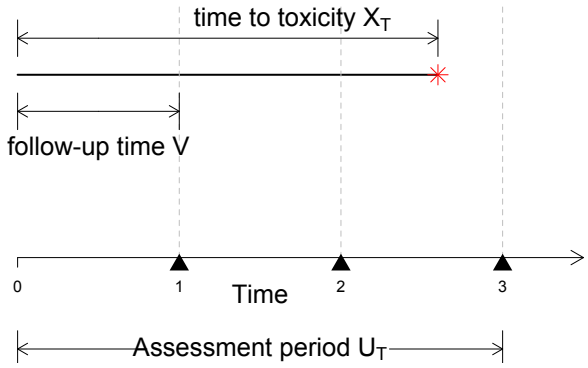

Notations

- No missing data

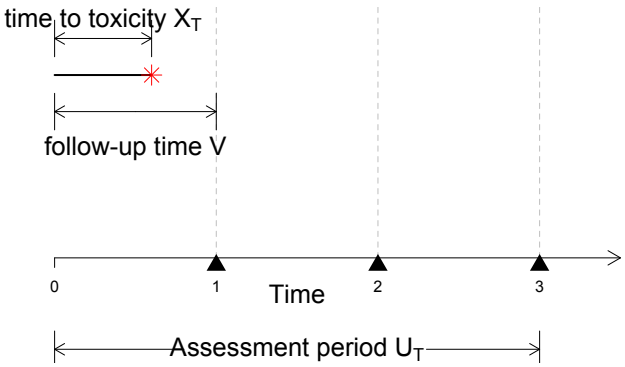

Missingness of outcomes

- For $k = E$ or T,

$$Y_k = \begin{cases} missing & if\ X_k > V and\ V < U_k\ (i.e., X_k^O = V), \\ 1 & if\ X_k \le V \le U_k\ (i.e., X_k^O = X_k), \\ 0 & if\ X_k > V = U_k\ (i.e., X_k^O = U_k). \end{cases}$$

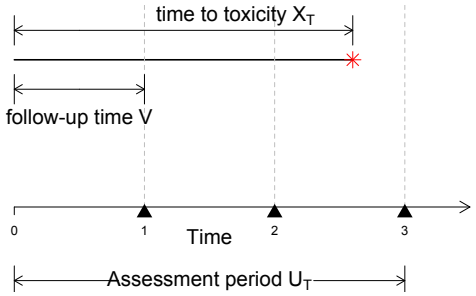

Missing data mechanism

- Let $M_k = I(Y_k = missing)$ denote the missingness indicator for $k = E, T$
- Three types of missing data mechanism (Little and Rubin, 2002)
  - Missing completely at random (MCAR)} if
    $$f(M_k|Y_k, \theta) = f(M_k|\theta)$$
  - Missing at random (MAR)} if
    $$f(M_k|Y_k, \theta) = f\left(M_k|Y_{k,obs}, \theta\right)$$
  - Nonignorable or not missing at random (NMAR)} if
    $$f(M_k|Y_k, \theta) = f\left(M_k|Y_{k,mis}, \theta\right)$$

Nonignorable missing data

- The missing data induced by delayed outcomes are non-ignorable (Yuan and Yin, 2010).
- Formally, $\Pr(M_k = 1|Y_k = 0) > \Pr(M_k = 1|Y_k = 1)$.
- By Bayes' rule,

$$\frac{\Pr(Y_k = 1|M_k = 0)}{\Pr(Y_k = 0|M_k = 0)} > \frac{\Pr(Y_k = 1|M_k = 1)}{\Pr(Y_k = 0|M_k = 1)}$$

  That is, the odds of event k decrease if Y_k is missing, so the missing indicator M_k contains information about the future value of Y_k.
- It is necessary to model the missing data mechanism.

Missing data mechanism

- Missing data induced by delayed outcomes are a special type of nonignorable missing data with a known missing data mechanism as follows

$$Y_k = \begin{cases} missing & if\ X_k > V and\ V < U_k\ (i.e., X_k^O = V), \\ 1 & if\ X_k \le V \le U_k\ (i.e., X_k^O = X_k), \\ 0 & if\ X_k > V = U_k\ (i.e., X_k^O = U_k). \end{cases}$$

That is,

$$M_k = \begin{cases} 1 & if\ X_k > V and\ V < U_k\ (i.e., X_k^O = V), \\ 0 & otherwise \end{cases}$$

- As a result, we need to model X_k, the time to toxicity and efficacy.

Model for times to toxicity and efficacy

- Model marginal survival functions $S_E(x_E|Y_E = 1)$ and $S_T(x_T|Y_T = 1)$ using piecewise exponential distributions
  - Partition the follow-up period $[0, U_k]$ into L_k intervals.
  - For dose d, assume hazard $\lambda_{k,l}$ on the l-th subinterval, where all $\lambda_{k,l} > 0$.
  - The marginal survival function for X_k is give by

    $$S_k(x|d, Y_j = 1, \lambda_k, \gamma_k) = exp\left\{-\sum_{l=1}^{L} w_{k,l}(x)\lambda_{k,l}\right\}, x > 0$$

    where the weights $w_{k,l}(x) = h_{k,l} - h_{k,l-1}$ if $x > h_{k,l}$; $w_{k,l}(x) = x - h_{k,l-1}$ if $x \in [h_{k,l-1}, h_{k,l})$; and otherwise $w_{k,l}(x) = 0$.

Bayesian data augmentation

- Imputation step, in which the missing values of Y_T and Y_E are sampled from their full conditional distributions,
- Posterior step, in which a posterior sample of unknown parameters is simulated based on the completed data including the imputed Y_T and Y_E.
  - In what follows, we illustrate the Bayesian data augmentation approach using the EffTox design.
  - The method can be applied to any phase I/II design (e.g., utility-based designs).

Quick review of EffTox design

- Logistic models for marginal distributions of Y_T and Y_E

$$\text{logit}(\pi_E) = \mu_E + d\beta_{1,E} + d^2\beta_{2,E},$$
$$\text{logit}(\pi_T) = \mu_T + d\beta_{1,T} + d^2\beta_{2,T}$$

  where $\pi_E = \Pr(Y_E = 1|d), \pi_T = \Pr(Y_T = 1|d)$ and d is dosage.

- Gumbel model for the joint distribution of Y_T and Y_E

$$\pi_{a,b} \equiv \Pr(Y_E = a, Y_T = b) = (\pi_E)^a(1-\pi_E)^{1-a}(\pi_T)^b(1-\pi_T)^{1-b} + (-1)^{a+b}\pi_E(1-\pi_E)\pi_T(1-\pi_T)\left(\frac{e^\psi - 1}{e^\psi + 1}\right),$$

  where ψ parameterized association between Y_T and Y_E.

Toxicity-efficacy tradeoff

- $\pi_1^* = (0.15,0); \pi_2^* = (1.0,0.55); \pi_3^* = (0.3,0.3)$

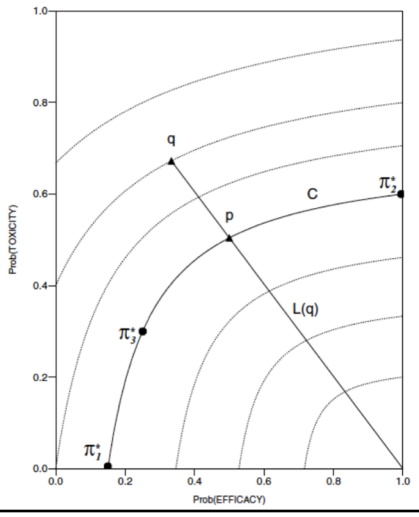

Dose assignment rule

- When a new patient arrives, based on outcomes from patients who have been treated in the trial, update $\hat{\pi}_{a,b}$ and assign that patient to the most desirable dose (according to the tradeoff).

Imputation step

Three possible missing data patterns

- Both Y_E and Y_T are missing,

$$\Pr(Y_E = 1, Y_T = 1|\mathcal{D}_{obs}) = \frac{\pi_{1,1}S_{11}}{\sum_{a=0}^{1}\sum_{b=0}^{1}\pi_{a,b}S_{ab}},$$

$$\Pr(Y_E = 1, Y_T = 0|\mathcal{D}_{obs}) = \frac{\pi_{1,0}S_{10}}{\sum_{a=0}^{1}\sum_{b=0}^{1}\pi_{a,b}S_{ab}},$$

$$\Pr(Y_E = 0, Y_T = 1|\mathcal{D}_{obs}) = \frac{\pi_{0,1}S_{01}}{\sum_{a=0}^{1}\sum_{b=0}^{1}\pi_{a,b}S_{ab}},$$

$$\Pr(Y_E = 0, Y_T = 0|\mathcal{D}_{obs}) = \frac{\pi_{0,0}S_{00}}{\sum_{a=0}^{1}\sum_{b=0}^{1}\pi_{a,b}S_{ab}},$$

where $S_{ab} = \Pr(X_E > V, X_T > V|Y_E = a, Y_T = b)$ and $a, b = \{0, 1\}$.

Imputation step

- Y_E is missing but Y_T is observed, we draw the missing value of Y_E

$$\Pr(Y_E = 1|\mathcal{D}_{obs}) = \begin{cases} \dfrac{\pi_{1,1}S_{10}}{\pi_{1,1}S_{10} + \pi_{0,1}} & if\ Y_T = 1 \\ \dfrac{\pi_{1,0}S_{10}}{\pi_{1,0}S_{10} + \pi_{0,0}} & if\ Y_T = 0 \end{cases}$$

where $S_{ab} = \Pr(X_E > V, X_T > V|Y_E = a, Y_T = b)$ and $a, b = \{0, 1\}$.

- Y_T is missing but Y_E is observed, we draw the missing value of Y_T

$$\Pr(Y_T = 1|\mathcal{D}_{obs}) = \begin{cases} \dfrac{\pi_{1,1}S_{01}}{\pi_{1,1}S_{01} + \pi_{1,0}} & if\ Y_E = 1 \\ \dfrac{\pi_{0,1}S_{01}}{\pi_{0,1}S_{01} + \pi_{0,0}} & if\ Y_E = 0 \end{cases}$$

Posterior step

- Given the imputed data, sequentially sample the unknown model parameters from their full conditional distribution as follows:
  - Sample θ sequentially from its posterior distribution $f\big(\theta|\mathcal{D}(Y)\big)$ where θ is the parameters for the bivariate toxicity-efficacy model.
  - Sample $\lambda_{j,k}$, and ϕ, $k = 1, \cdots, K_j$ and $j = \{E, T\}$ sequentially from its posterior $f\big(\lambda, \phi|\mathcal{D}(Y)\big)$.

Simulation

- Five dose levels with raw doses
  $$\boldsymbol{d} = (2.5, 5.0, 7.5, 10.0, 12.5)$$
- Assume $U_E = U_T = 6$ weeks with accrual rate $\alpha = 1.5$/week.
- $N_{max} = 3{\times}16 = 48$ and $K = 6$.
- The trade-off contour, $\mathcal{C}$,was determined by fitting a quadratic curve to the trade-off target pairs $(\pi_E, \pi_T) = (0.15, 0), (0.45, 0.20), (1, 0.60)$.

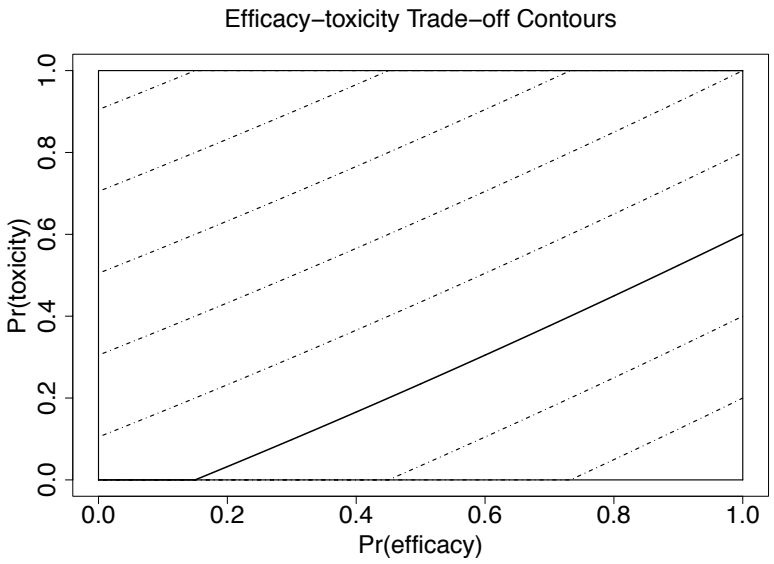

Simulation

- A dose d is considered acceptable if

$$\Pr\{\pi_E(d,\theta) > \underline{\pi}_E | \mathcal{D}(Y)\} > p_E \text{ and}$$

$$\Pr\{\pi_T(d,\theta) < \bar{\pi}_T | \mathcal{D}(Y)\} > p_T$$

  with $\underline{\pi}_E = 0.25$, $\bar{\pi}_T = 0.35$ and $p_E = p_T = 0.10$.

- We considered 4 different true dose-toxicity and -efficacy scenarios.

True toxicity-efficacy scenarios

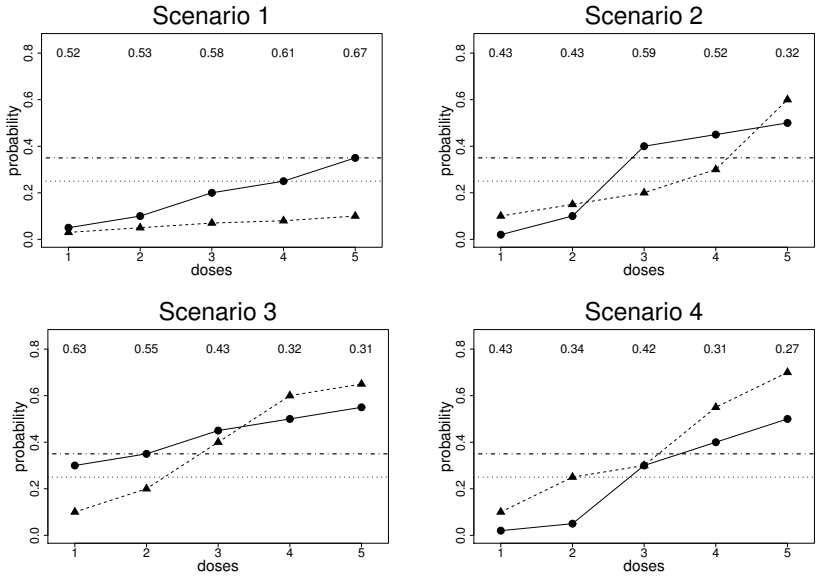

Methods for comparison

- "One Level Down" rule: If some of the patients treated at the current dose d_r have not yet been evaluated fully, i.e. $Y_{i,E,d_r} = missing$ or $Y_{i,T,d_r} = missing$, then any new patient is treated at d_{r-1}.
- "Look Ahead" rule: if all possible values of Y_E or Y_T that currently are missing will not change the dose assignment decision, then treat any new patient based on that dose assignment decision. Otherwise, we make new patients wait to be treated or turn them away and treat them off protocol.
- "Complete Case" rule: Use all complete cases, where both Y_E and Y_T are observed, so compute d^{opt} and treat the next patient immediately.

Two performance metrics

- Desirability-weighted selected percentage:

$$\bar{\delta} = \frac{\sum_{r=1}^{5} \delta_r^{true} \; P(\text{select } d_r) I(d_r \in \mathcal{A}^{true})}{\sum_{r=1}^{5} \delta_r^{true} \; I(d_r \in \mathcal{A}^{true})}$$

  where δ_r^{true} denotes the true desirability of dose d_r and $\mathcal{A}^{true}$ denotes the true set of acceptable doses. This criterion quantifies dose selection reliability and thus potential benefit for future patients.
- The ratio N_E/N_T where N_E and N_T denote the number of patients who experienced efficacy and toxicity, respectively. This criterion quantifies benefit to the patients in the trial, hence may be considered an index of ethical desirability.

Results

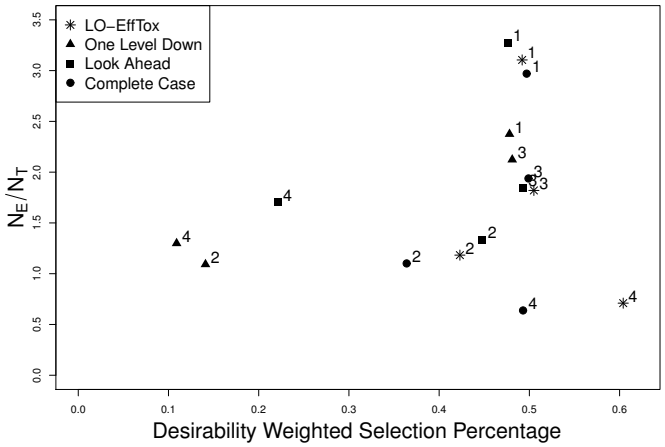

Trial duration

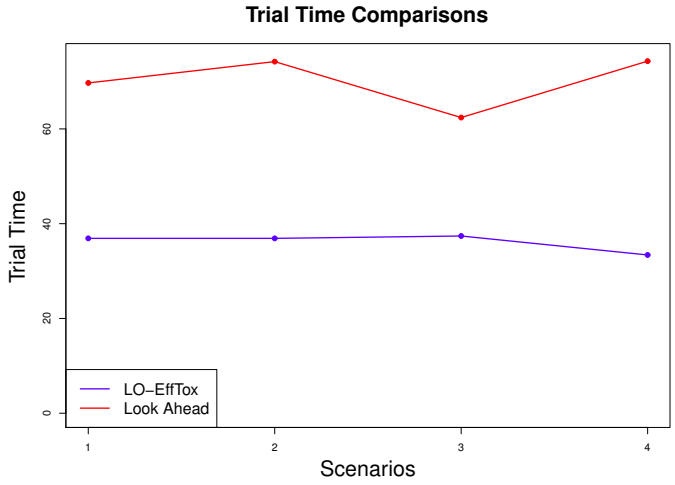

Conclusion

- Delayed outcomes are the major practical impediment of any outcome-adaptive clinical design.
- A general methodology to address this problem is to treat unobserved outcomes as nonignorable missing data.
- Bayesian data augmentation can be conveniently used to handle the resulting missing data.
- Simulation studies show that the methodology outperforms the existing approaches.

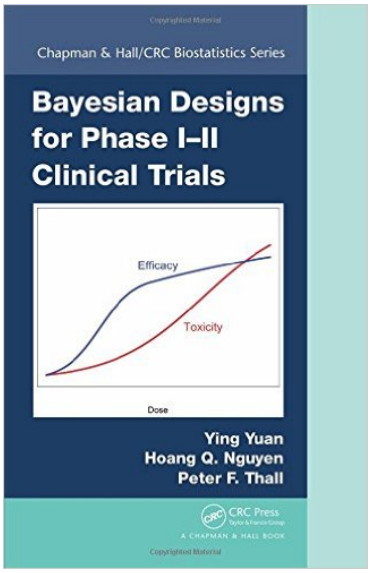


Thank You!