Early Phase Trial Designs and Estimand

Yuan Ji, PhD

Program for Computational Genomics & Medicine
NorthShore University HealthSystem
Department of Public Health Sciences
The University of Chicago

SBF 2017

December 4, 2017

COI Disclosure

- Consultants for Seattle Genetics and Biogen.
- Co-Founder of Laiya Consulting, Inc., a statistical solution and consulting company

ESTIMAND

An ESTIMAND includes

1. Target population
2. Endpoint (the variable to be measured)
3. Strategies to account for intercurrent events (potential confounders)
4. Population summaries for statistical inference and comparison

ESTIMAND for confirmatory (late-phase) trials depends on results from early-phase trials, such as

- target population

There could be discrepancies between ESTIMAND for confirmatory trials and important components for early-phase trials, such as

- endpoint

Phase I dose-finding (Oncology)

Consider trials with fixed doses.

Setup Climb up and down a sequence of D ordered doses of a new drug to determine the maximum tolerated dose (MTD).

Data At each dose i, n_i patients are tested, y_i patients experienced toxicity outcome (DLT).

Parameters Dose i has a toxicity probability of p_i (unknown).

Sampling Model Binomial $y_i \mid p_i \sim Bin(n_i, p_i)$

Assumption Toxicity Monotonicity : $p_i \leq p_{i+1}$.

Hidden assumption Efficacy Monotonicity : $q_i \leq q_{i+1}$ – if not, why escalate when the dose is safe?

Goal to find the MTD, defined as the highest dose with toxcity rate lower (or close to) a target rate, p_T, e.g., $p_T = 0.30$.

Intercurrent event Death (how to account for death in dose-finding trials)

ESTIMAND for dose-finding trials

1. Target Population
   - Traditionally, all comers (e.g., all patients with solid tumors)
   - Recently, in Immune-Oncology (e.g., CAR-T trials), targeted cancer types (subtypes). Often the same as late-phase population
2. Endpoint
   - Binary DLT: yes/no
   - Different grades; multiple cycles; time-to-event
3. Intercurrent Event
   - Death: due to toxicity vs. due to disease
   - Side effects from other medications
4. Population Summary
   - DLT rates at different doses
   - Total toxicity risk score (for toxicity grades)
   - Average toxicity rates over cycles

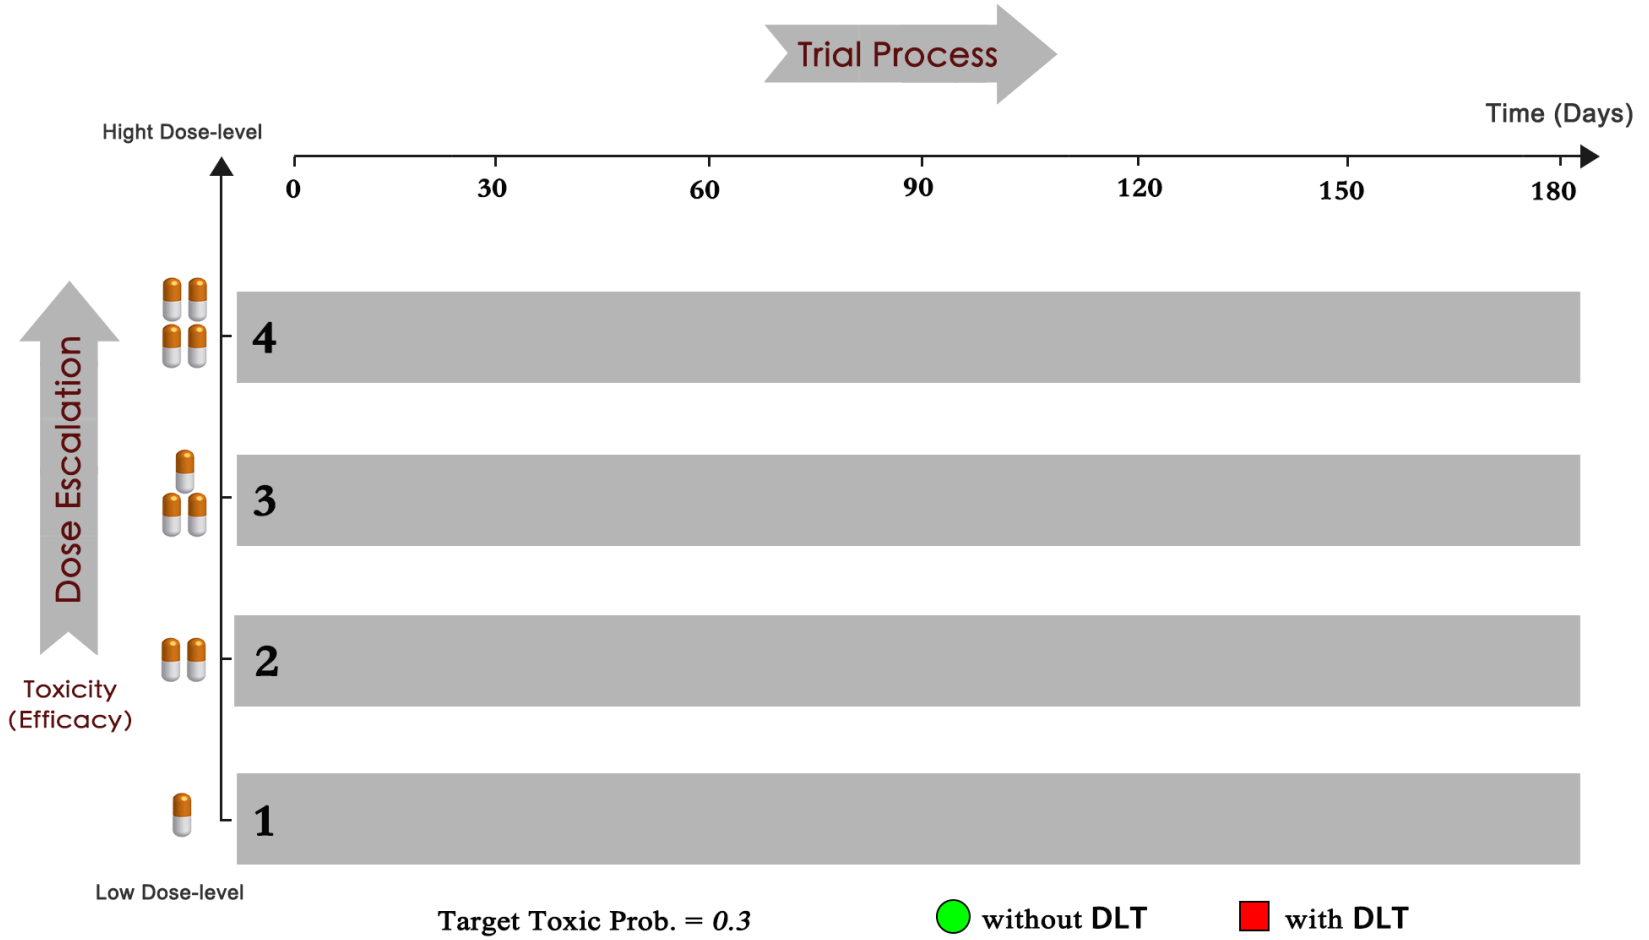
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Dose Escalation
4
3
2
1
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

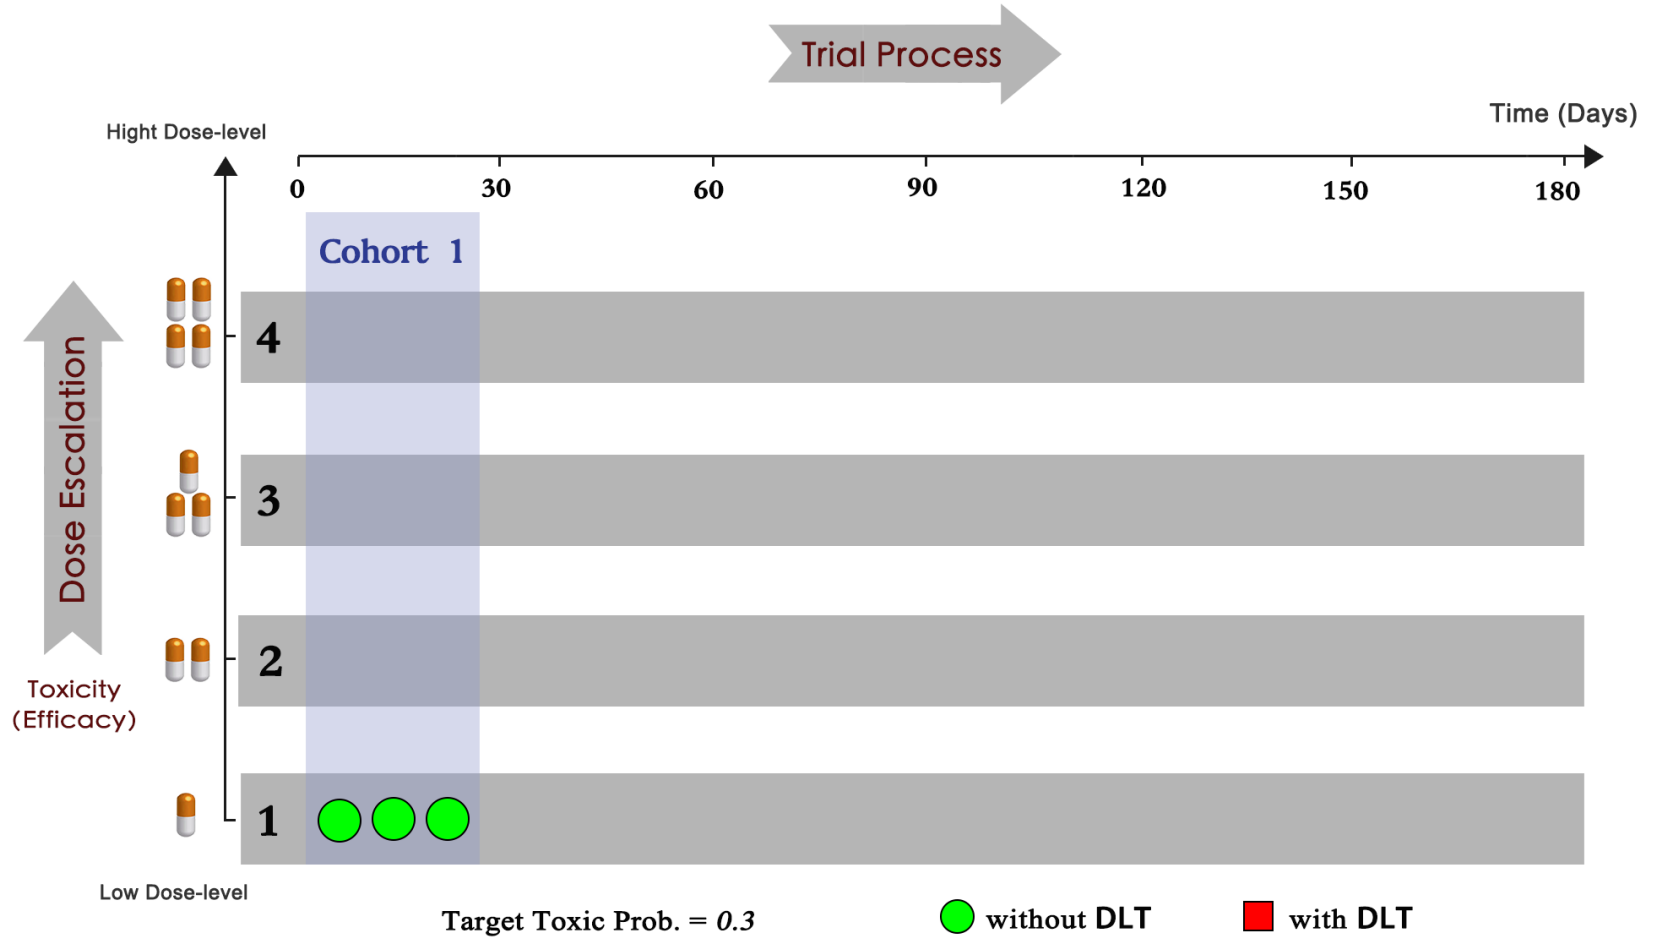
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Dose Escalation
4
3
2
1
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

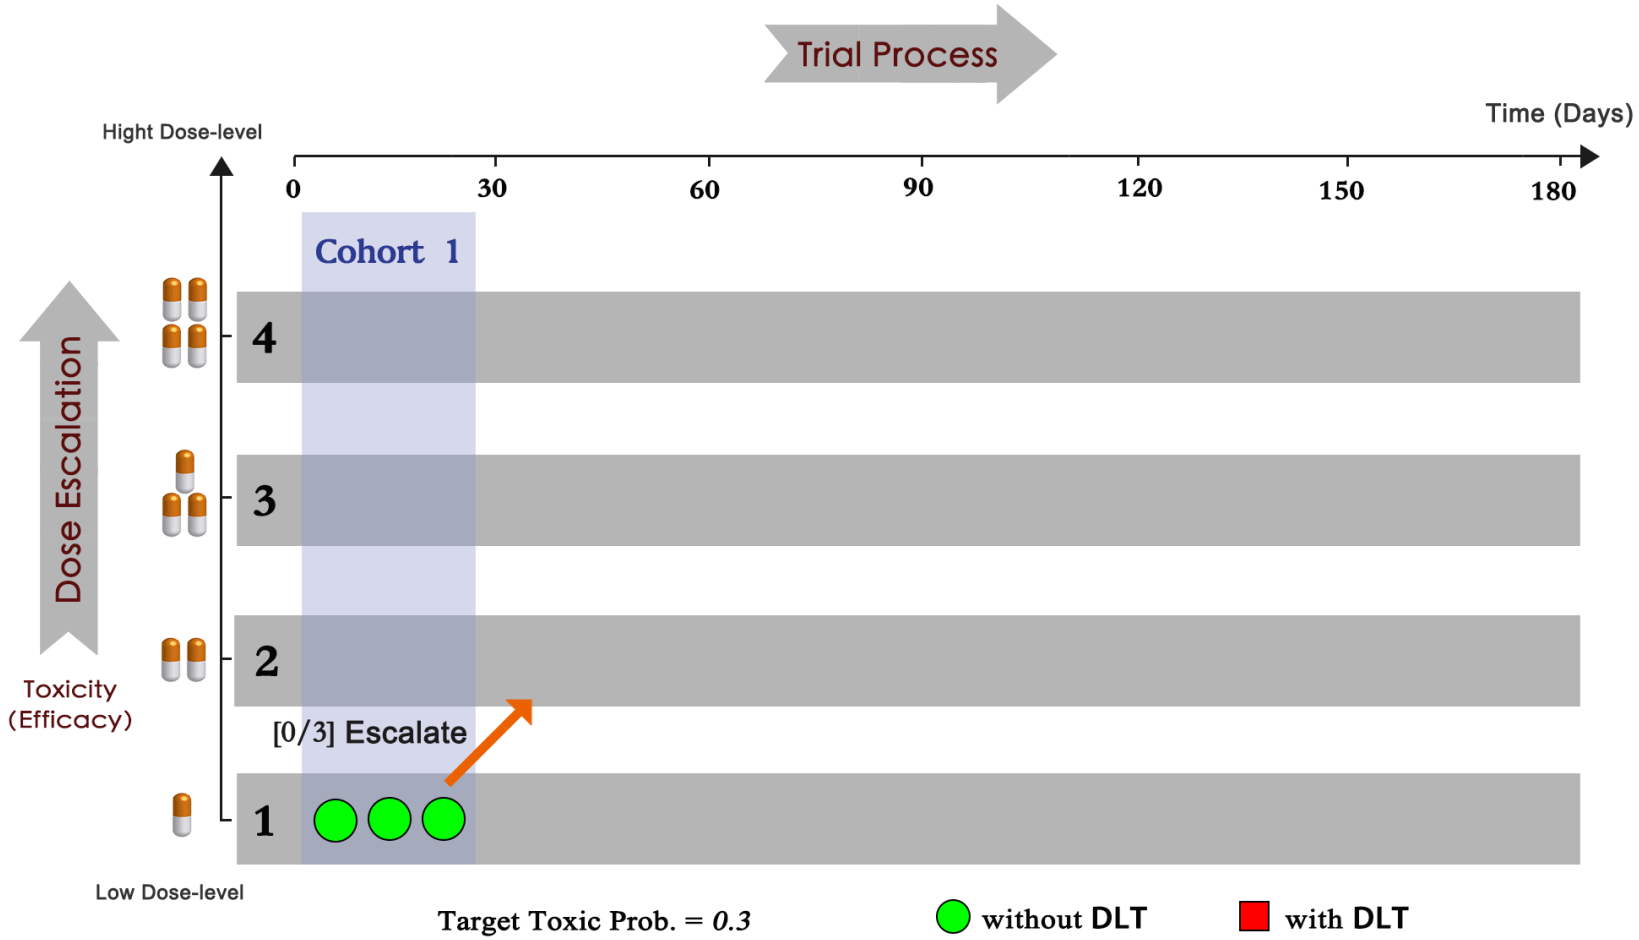
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Dose Escalation
4
3
2
1
Toxicity
(Efficacy)
[0/3] Escalate
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

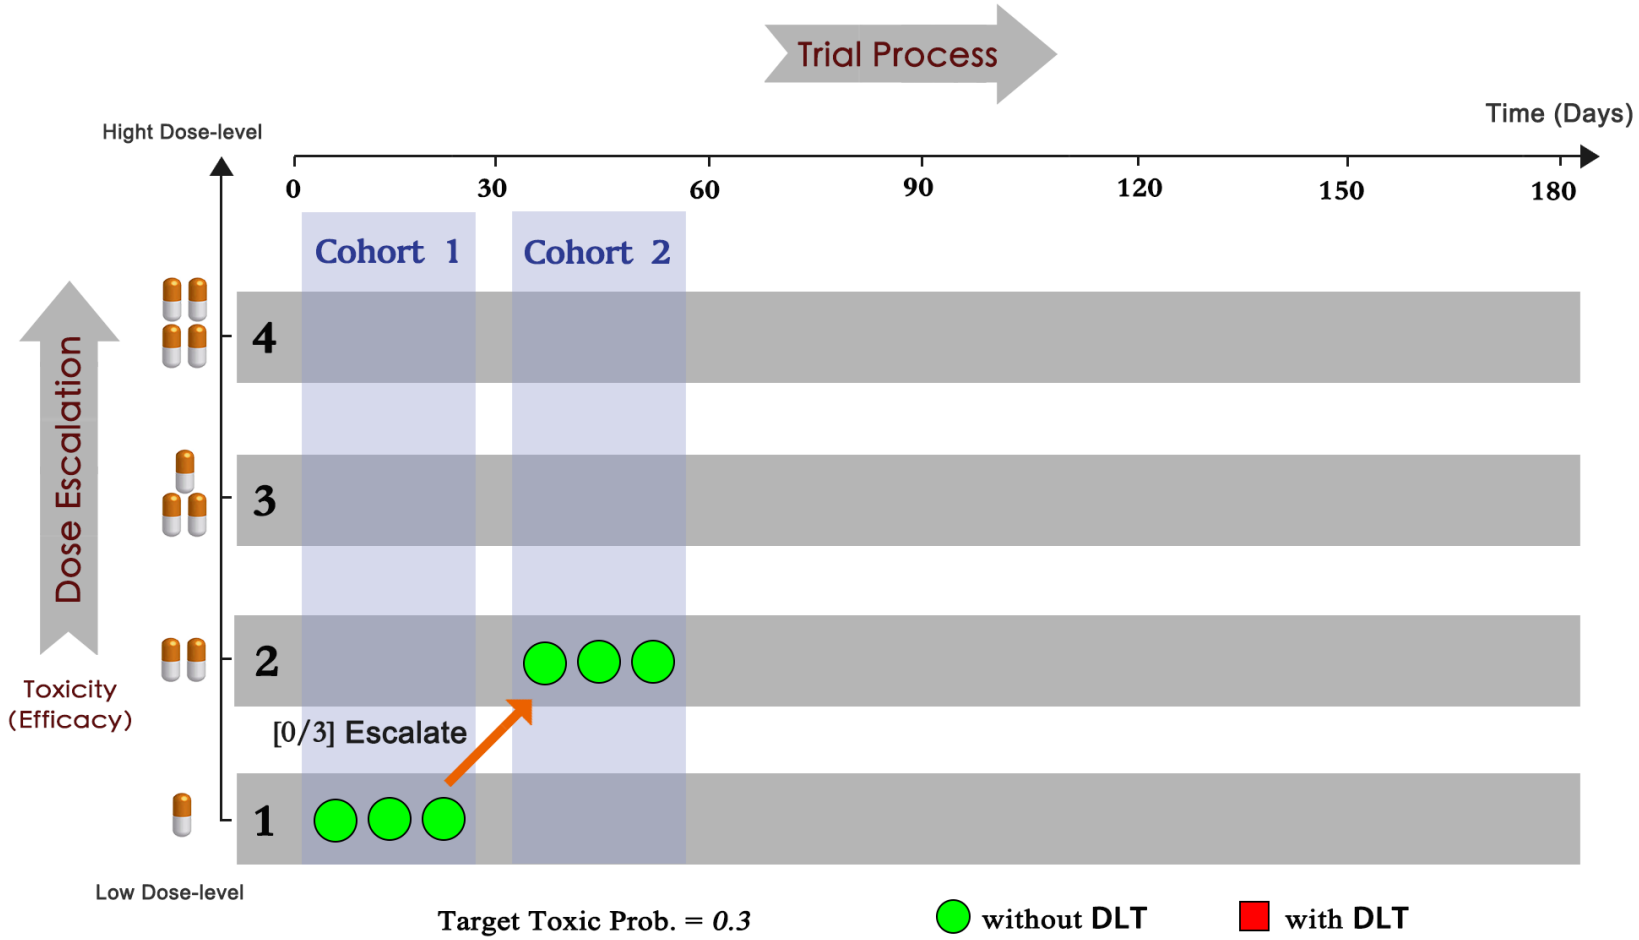
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Dose Escalation
4
3
2
1
Toxicity
(Efficacy)
[0/3] Escalate
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

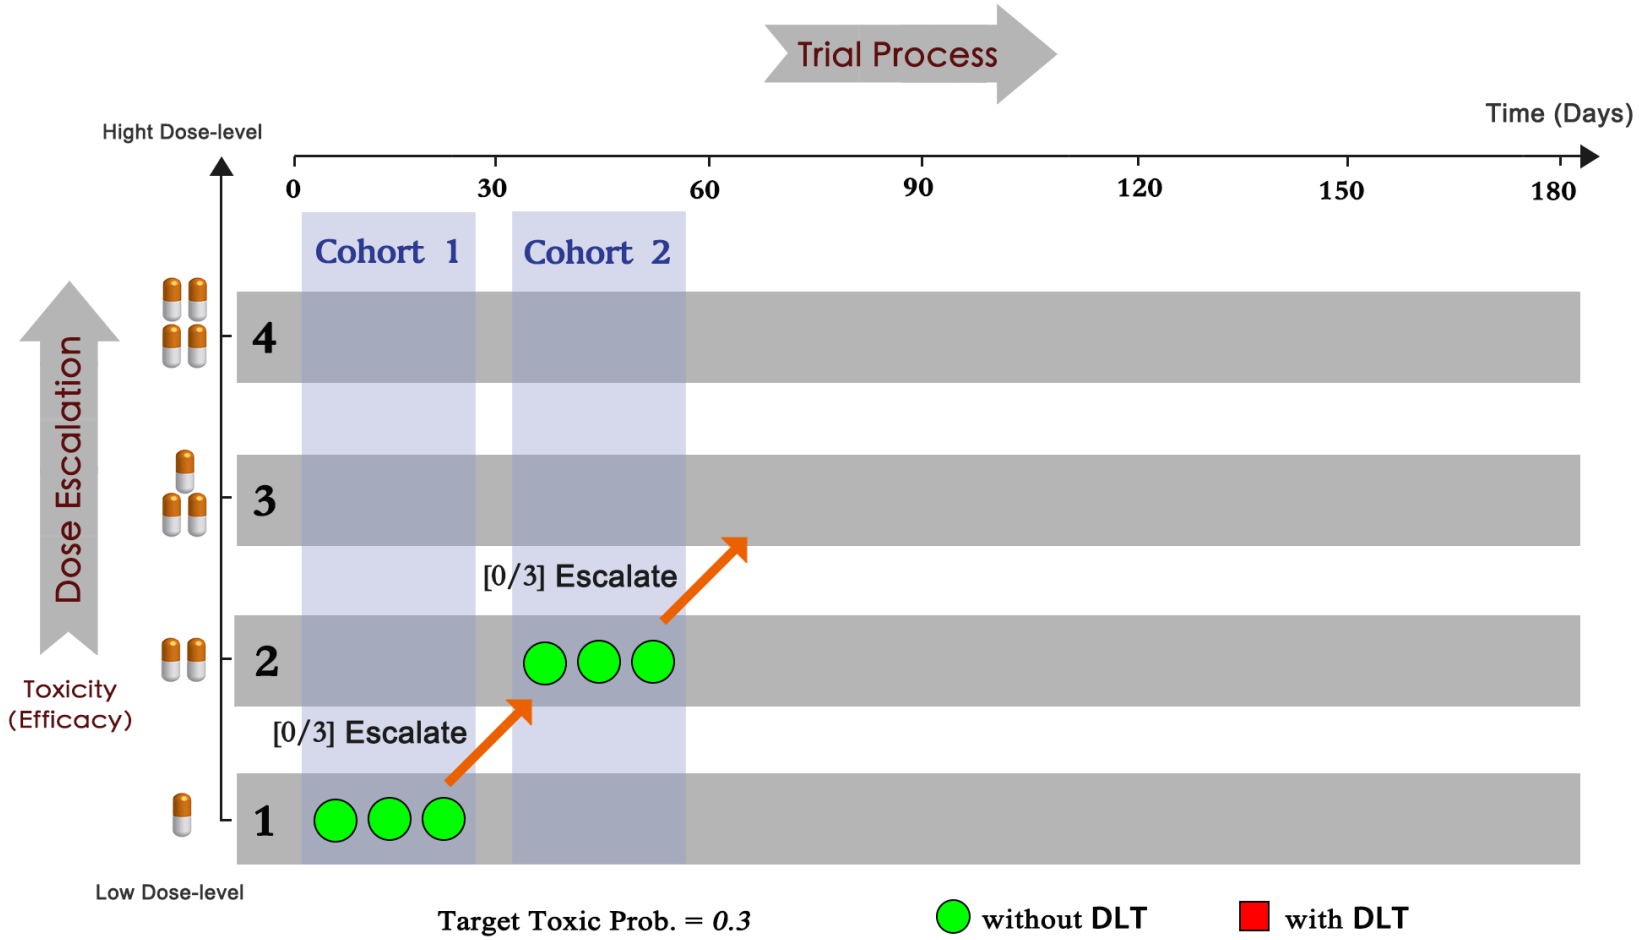
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Dose Escalation
4
3
2
1
[0/3] Escalate
[0/3] Escalate
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

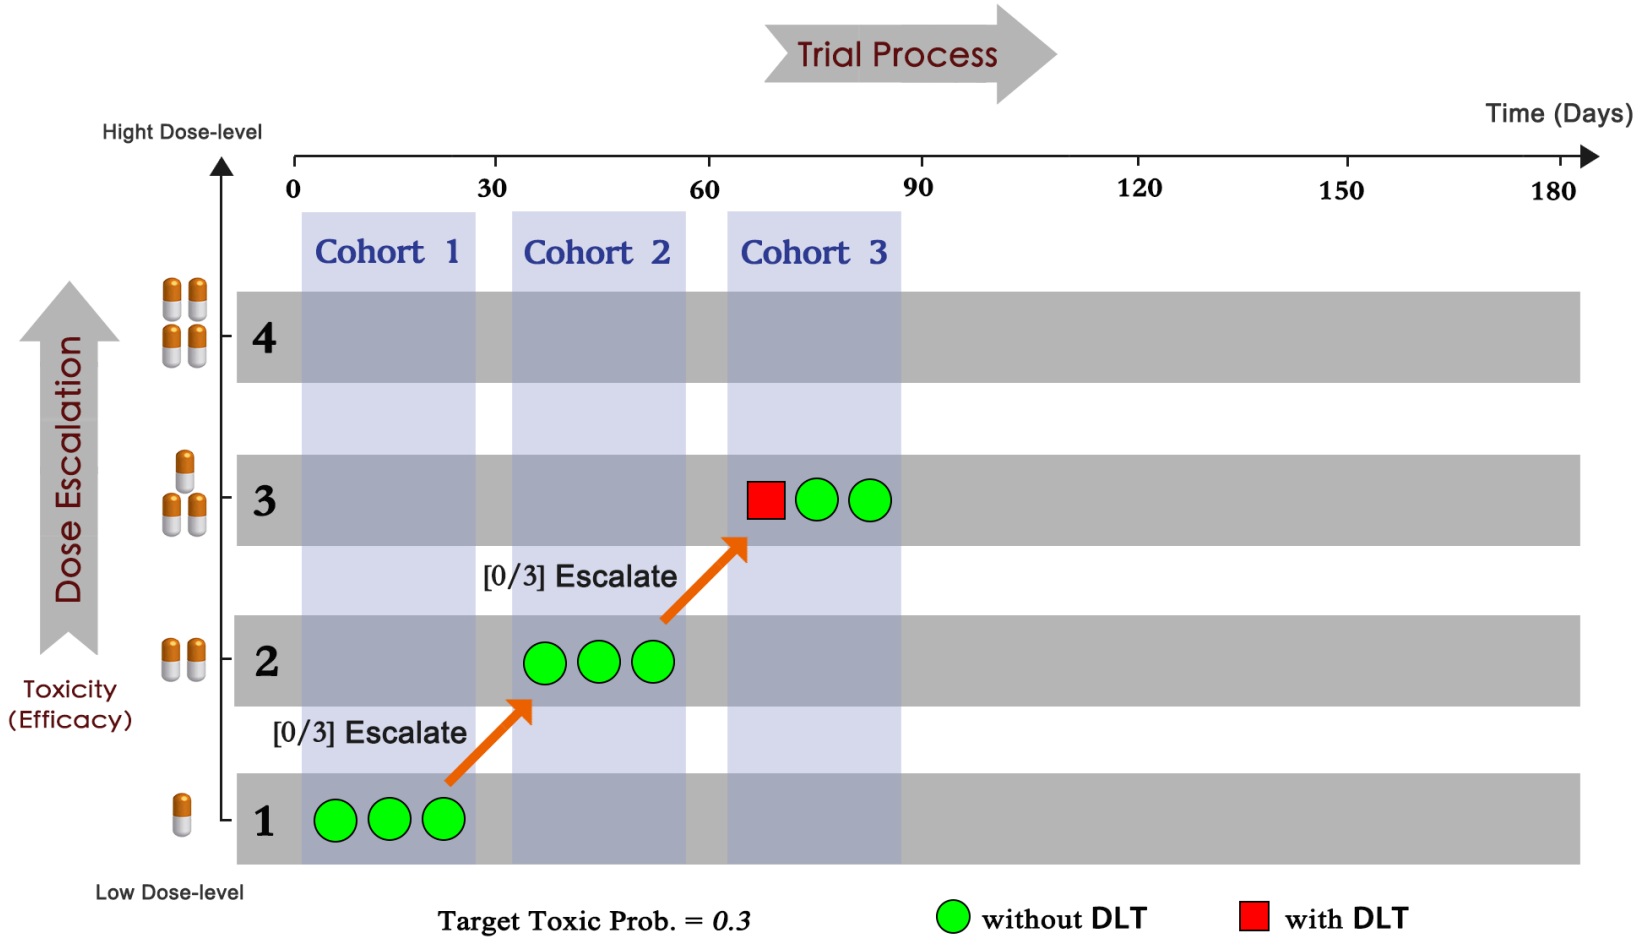
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Dose Escalation
4
3
2
1
[0/3] Escalate
[0/3] Escalate
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

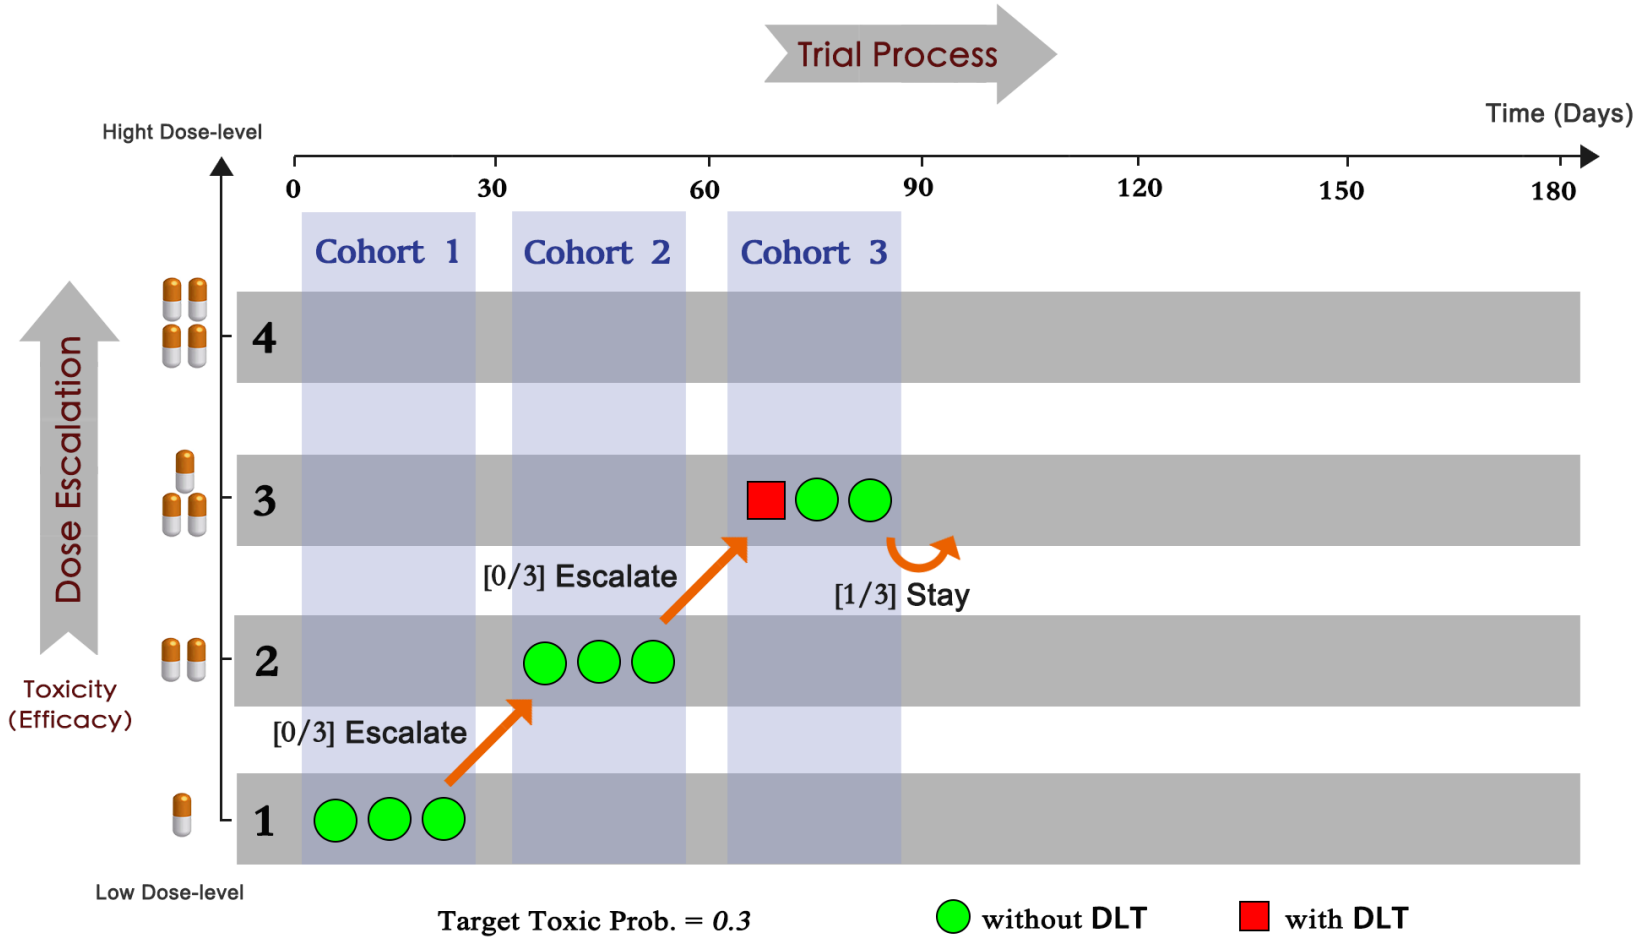
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Dose Escalation
4
3
2
1
[0/3] Escalate
[0/3] Escalate
[1/3] Stay
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

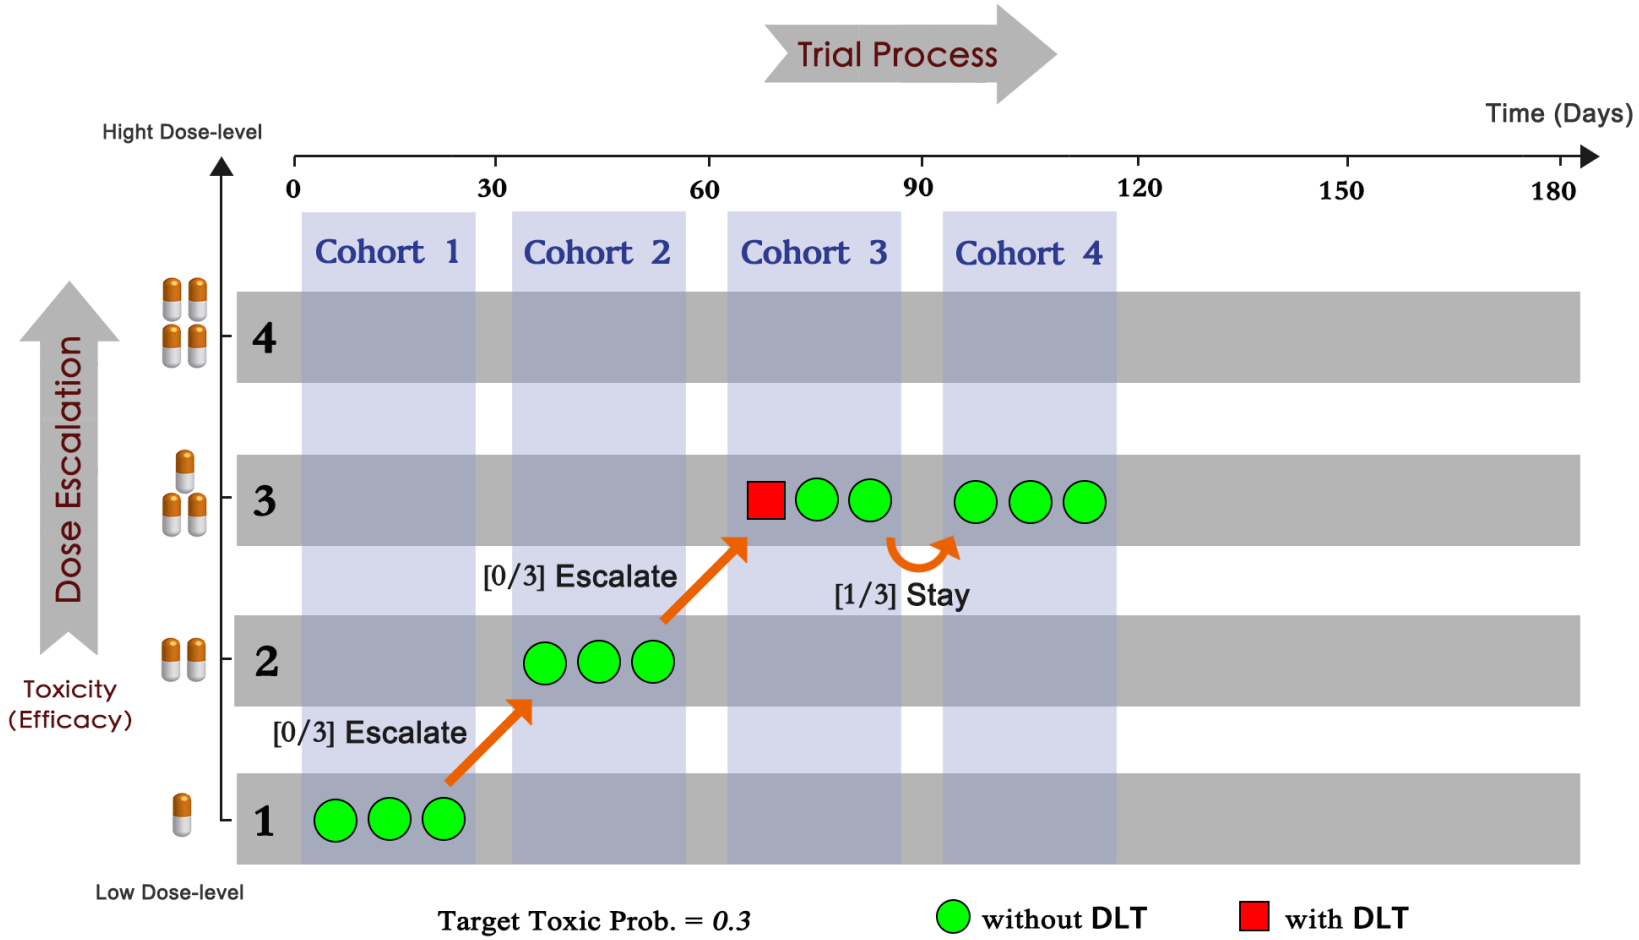
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Cohort 4
Dose Escalation
4
3
2
1
[0/3] Escalate
[0/3] Escalate
[1/3] Stay
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

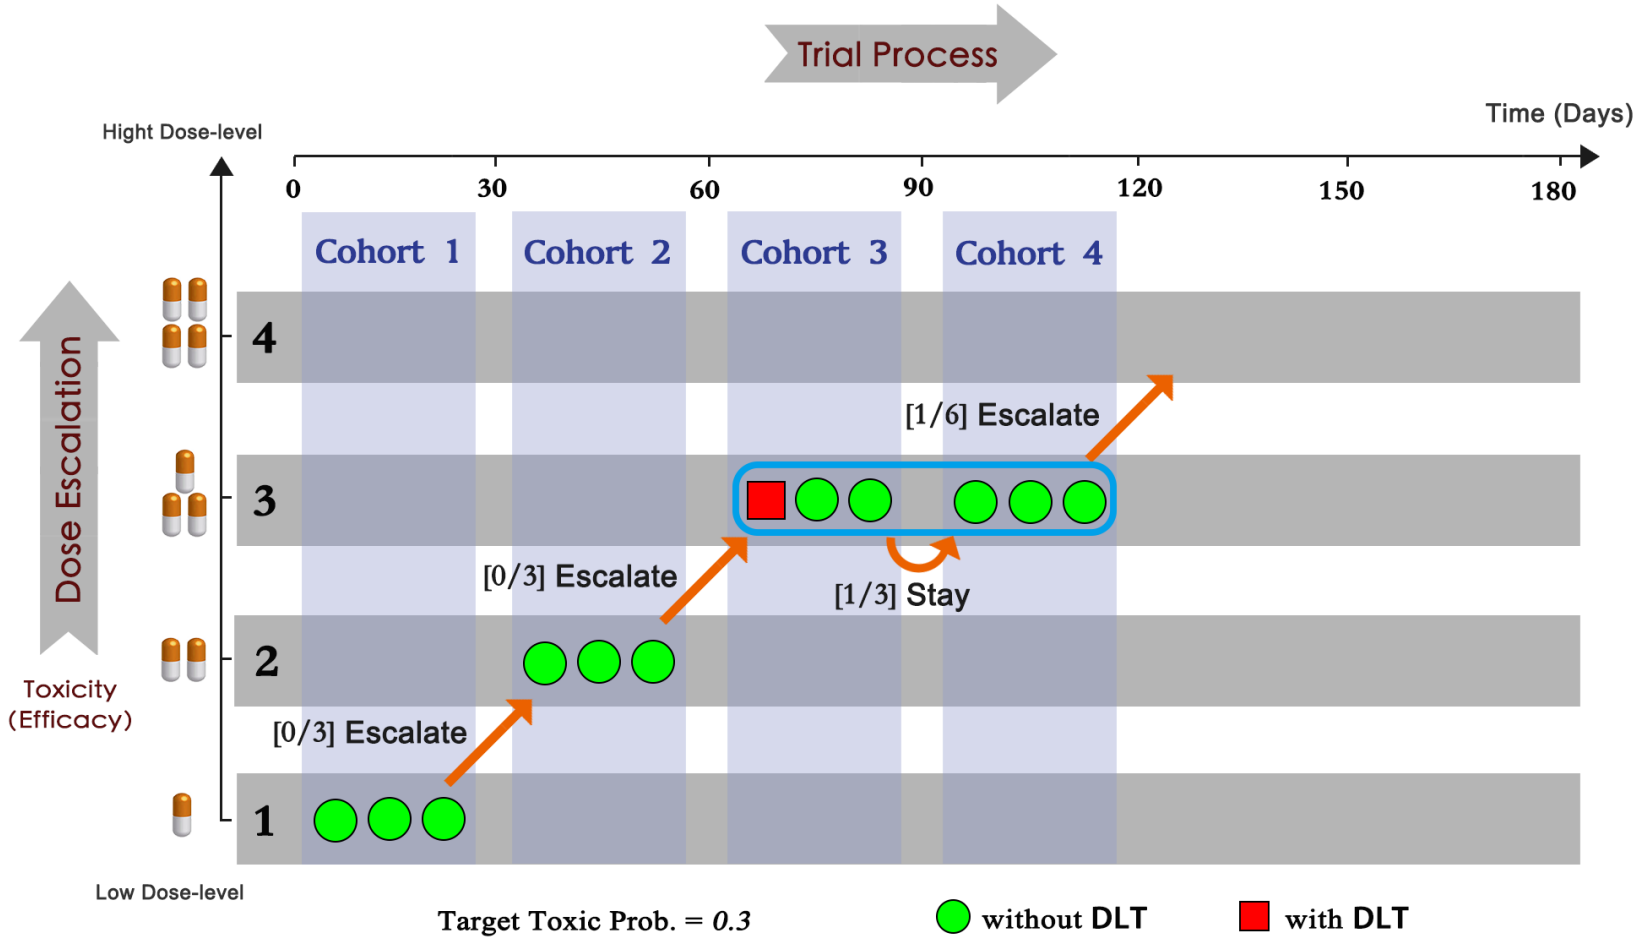
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Cohort 4
Dose Escalation
4
3
2
1
[1/6] Escalate
[0/3] Escalate
[1/3] Stay
[0/3] Escalate
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

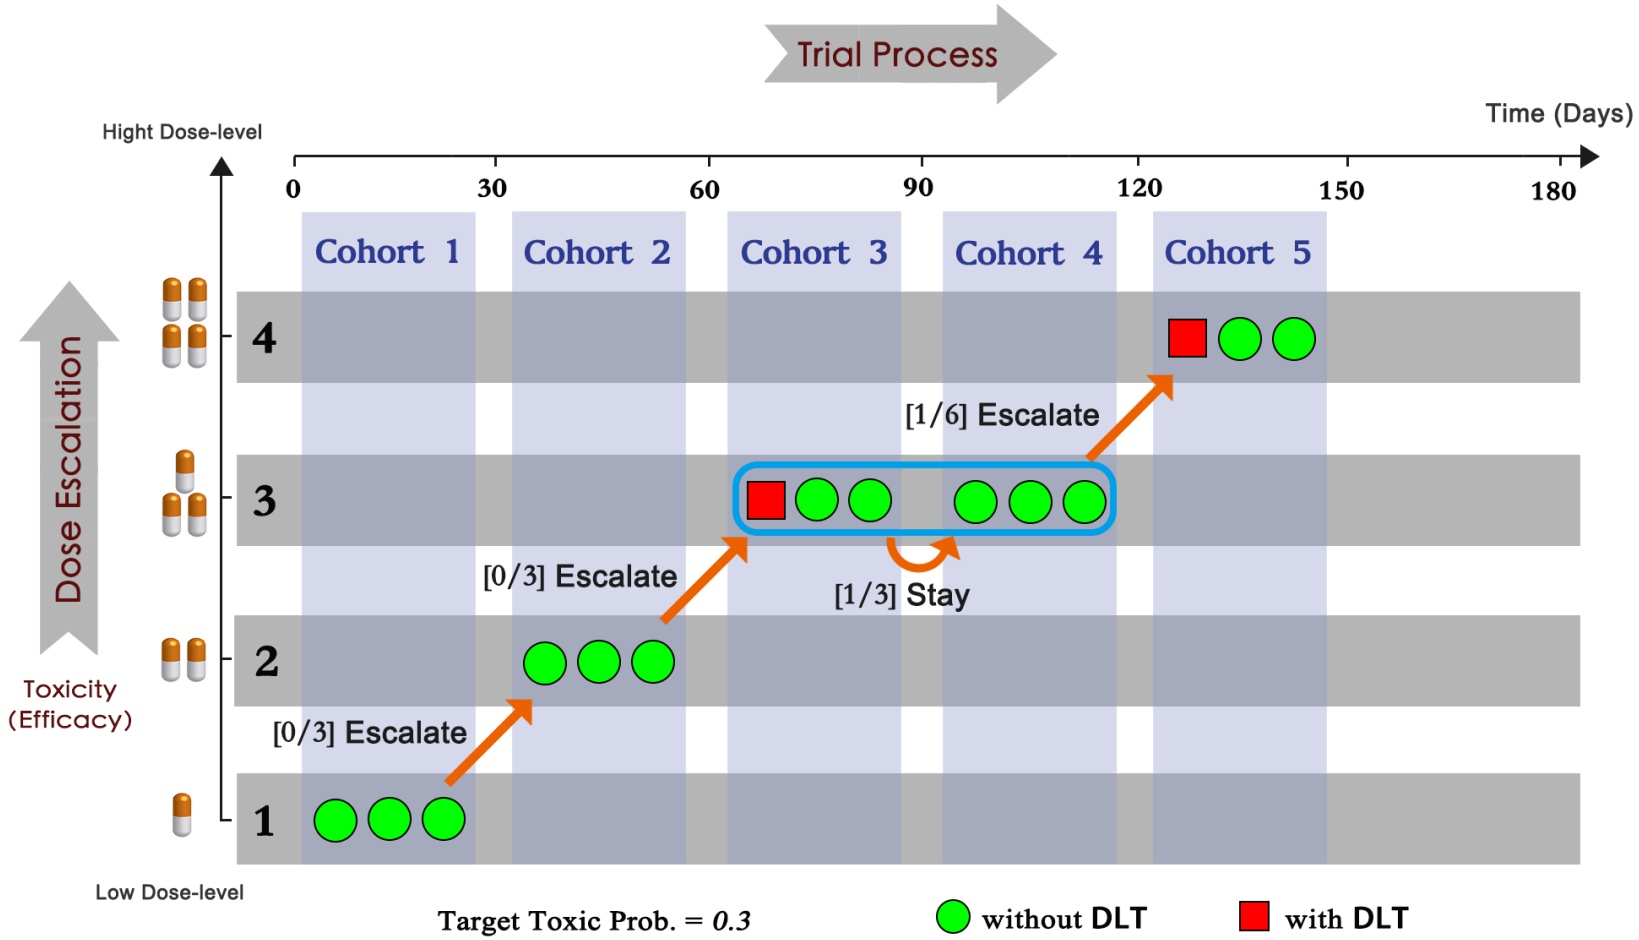
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Cohort 4
Cohort 5
Dose Escalation
4
3
2
1
[1/6] Escalate
[0/3] Escalate
[1/3] Stay
[0/3] Escalate
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

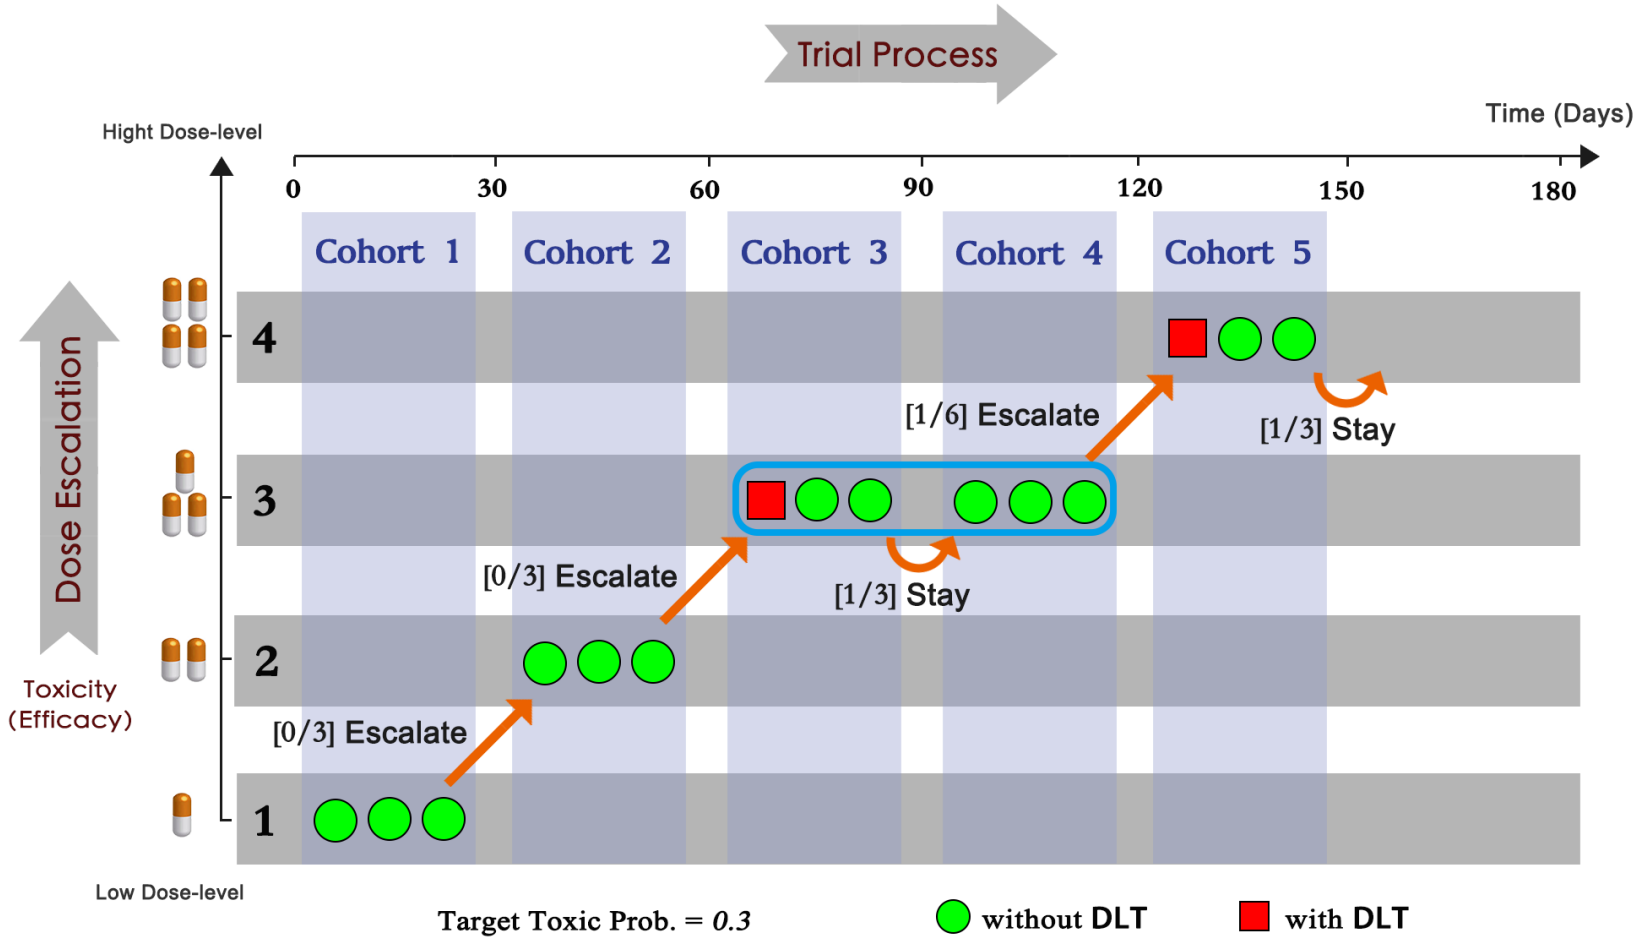
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Cohort 4
Cohort 5
Dose Escalation
4
3
2
1
[1/6] Escalate
[1/3] Stay
[0/3] Escalate
[1/3] Stay
[0/3] Escalate
Toxicity
(Efficacy)
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

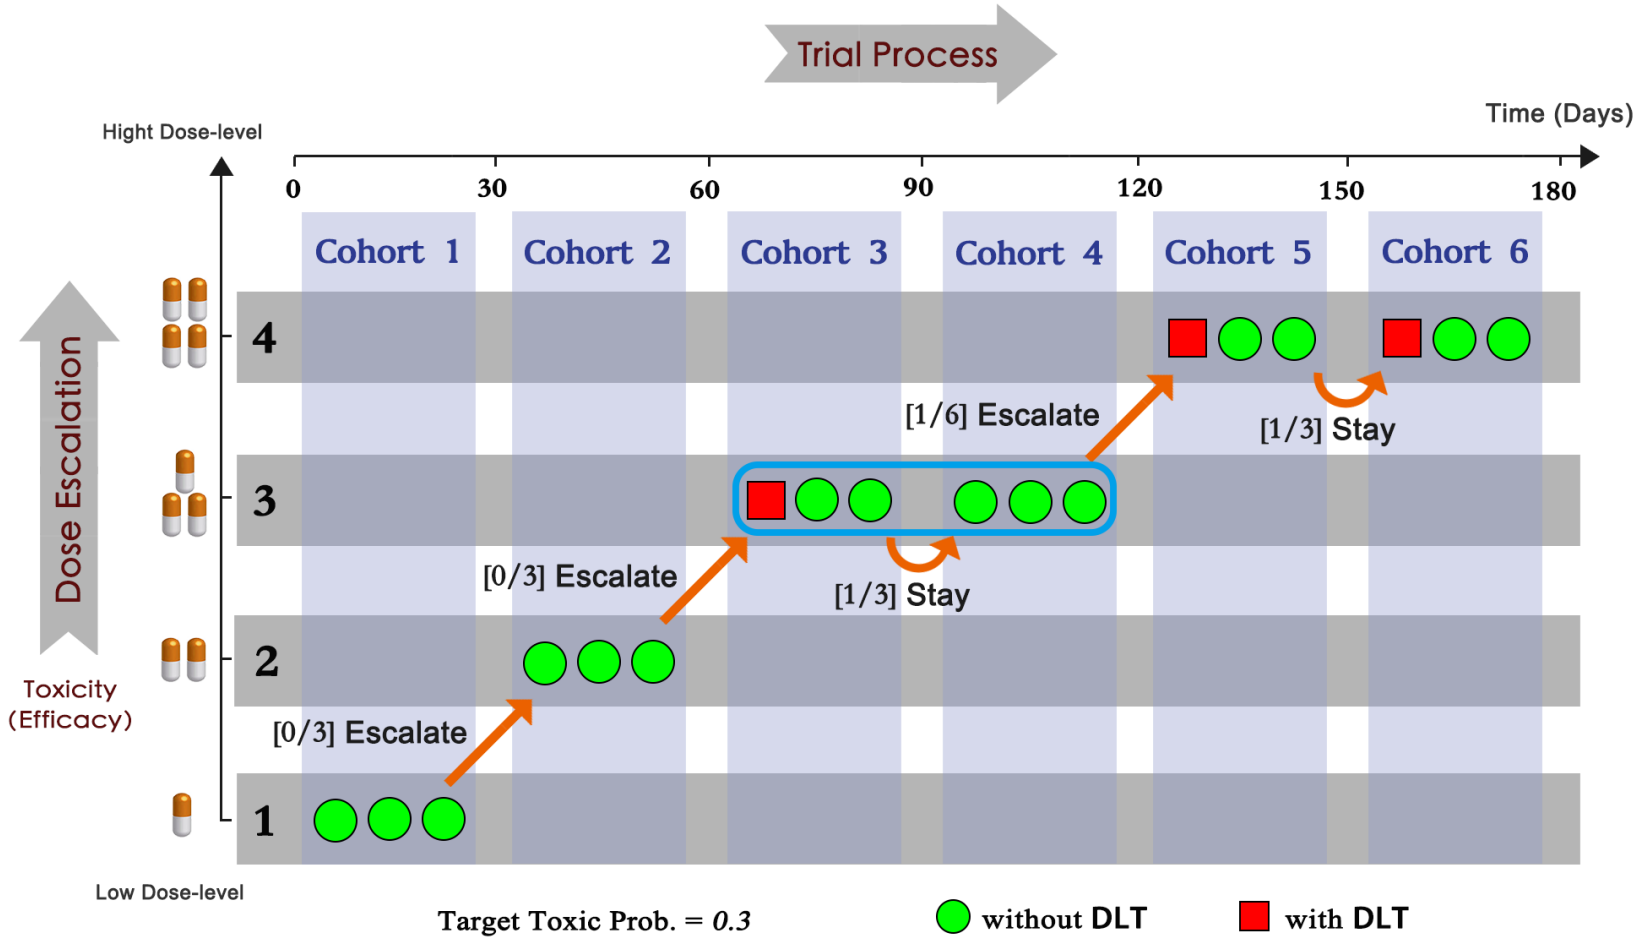
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Cohort 4
Cohort 5
Cohort 6
Dose Escalation
4
3
2
1
Toxicity
(Efficacy)
[0/3] Escalate
[0/3] Escalate
[1/3] Stay
[1/6] Escalate
[1/3] Stay
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

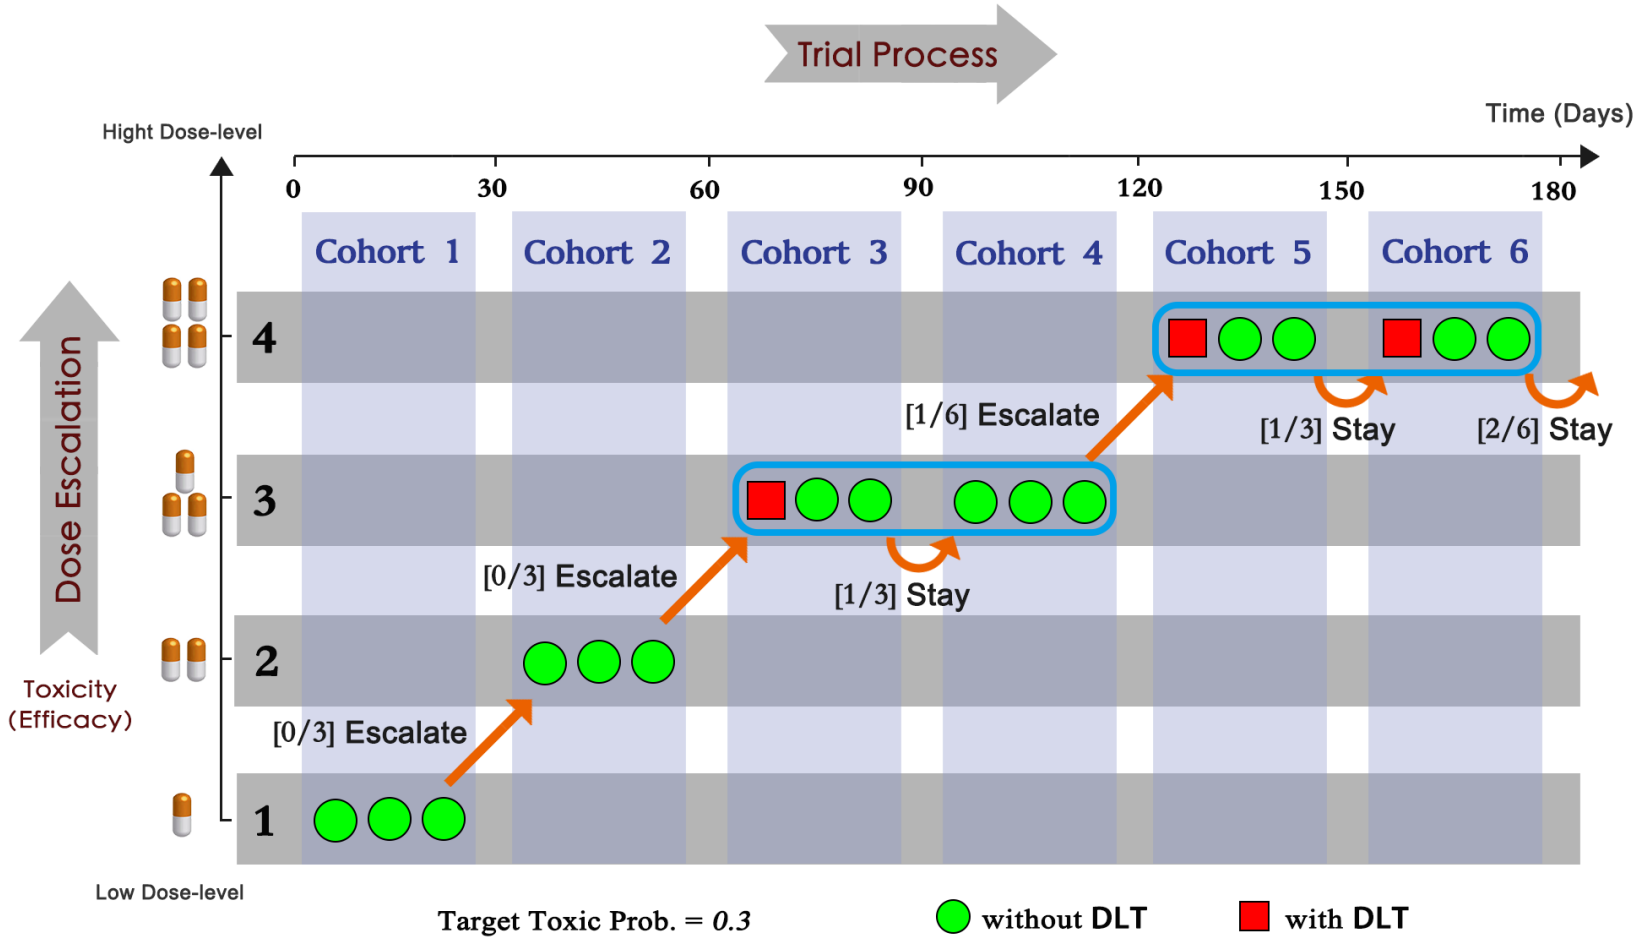
Trial Process
Time (Days)
Hight Dose-level
0
30
60
90
120
150
180
Cohort 1
Cohort 2
Cohort 3
Cohort 4
Cohort 5
Cohort 6
Dose Escalation
4
3
2
1
Toxicity
(Efficacy)
[0/3] Escalate
[0/3] Escalate
[1/3] Stay
[1/6] Escalate
[1/3] Stay
[2/6] Stay
Low Dose-level
Target Toxic Prob. = 0.3
without DLT
with DLT

Existing Designs

3+3

- Storer (1989). Algorithmic design; simple and transparent; popular among physicians
- Lacks a standard program; performs worse than model-based designs

CRM

- The first model-based design. First publication in 1990.
- Many different versions; black-box to physicians

mTPI and mTPI-2

- Ji et al. (2010); Guo et al. (2017). Model-based interval design in an algorithmic presentation
- Getting popular in the community; user-friendly software; simple and transparent

CCD and BOIN

- CCD (Ivanova et al. 2007), extension of CCD – BOIN (Liu and Ying 2015); also model-based interval designs but with a different framework
- Simple inference based on point estimate of toxicity probabilities; BOIN asymptotic behavior is strange

The Industry-Standard 3+3 Design

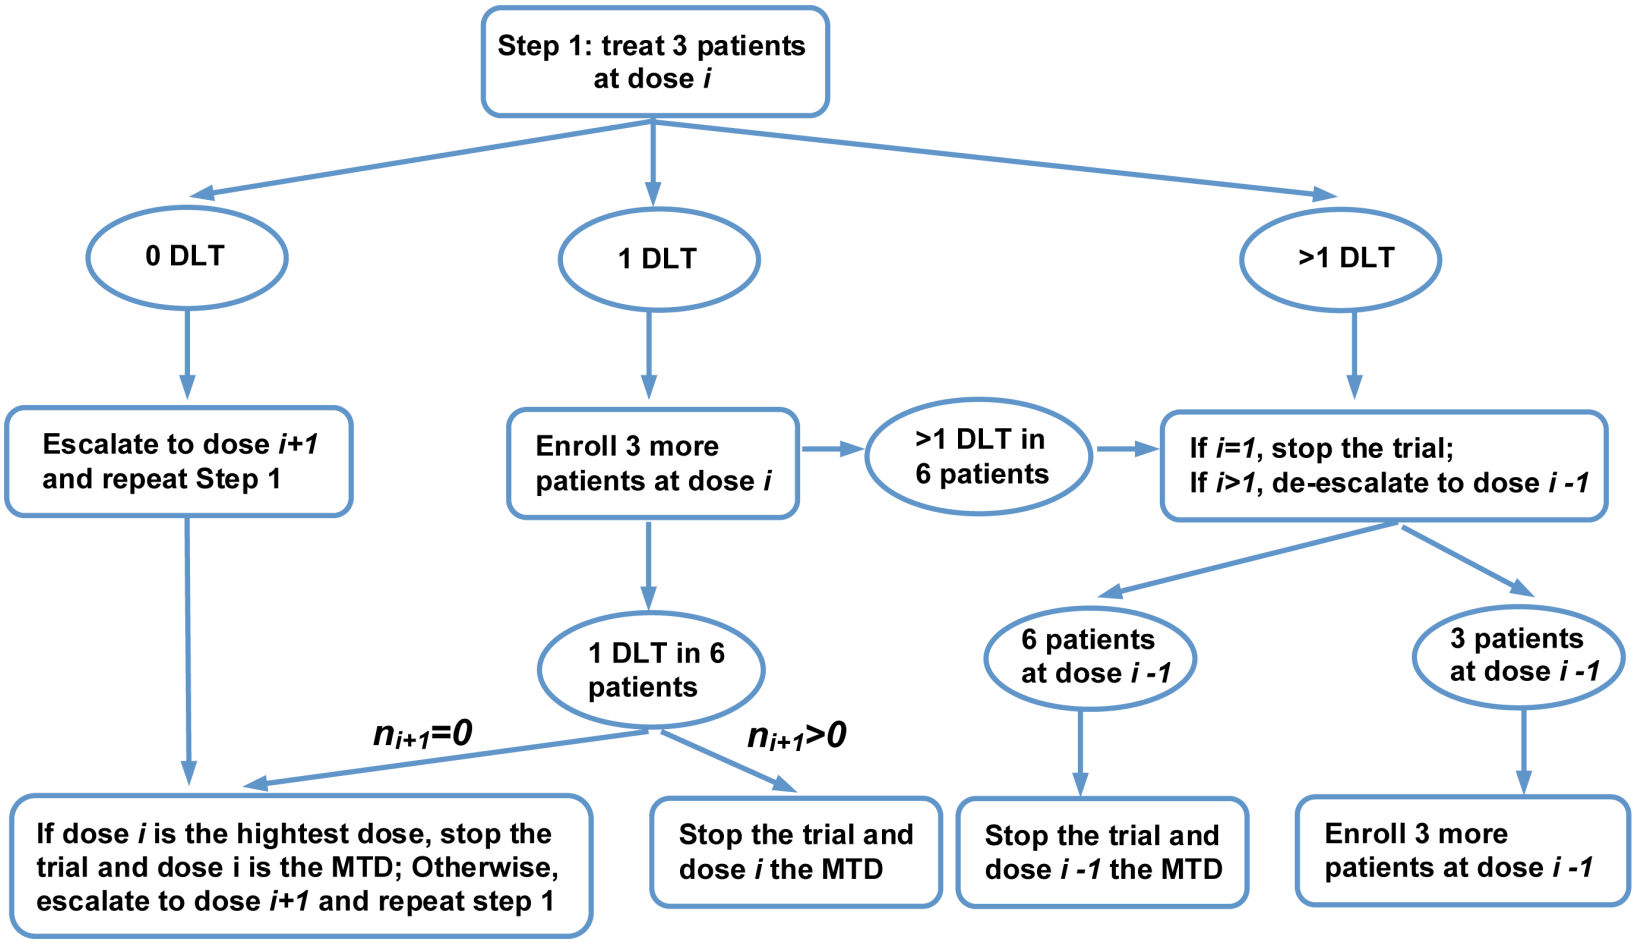

– Yang et al. (2015)

The 3+3 Design

Standard in clinical community (simple, transparent, "make sense", perform reasonably well if $p_T \sim 17\%$ or 30%.)

Favorite Target to statisticians

- Not model-based
- No more than 6 patients per dose
- Arbitrary choice of "3"
- Conservative – slow escalation

Dominant in practice (e.g., Rogatko et al. 2007)

- $>$98% of all phase I trials between 1991-2006 were based on 3+3 or its variations
- Out of 1,235 trials during the period, 20 were based on CRM; 3 based on EWOC (a variation of CRM)
- getting less popular and more frequently criticized recently (Nie et al., 2016)

The 3+3 Design: Is it conservative?

- Ji and Wang (2013, *JCO*) showed that with matched sample size, 3+3 is less safe and reliable when compared to the mTPI design , a model-based design.

- The 2015 FDA/AACR Dose-Finding Symposium concluded that (Nie et al., 2016, *Clinical Cancer Research*)

  "*The MTD/3+3 approach is not optimal and may result in recommended doses that are unacceptably toxic for many patients and in dose reduction/interruptions that might have an impact on effectiveness.*"

The CRM Design – a specific model

Perhaps the most popular version of the CRM is the power model:

- The dose-response curve : $p_i = p_{i0}^{\exp(\alpha)}$, where p_{i0} are fixed and prespecified constants, and α is a parameter that describes the dose response curve.
- The prior for α is $N(0, 2)$.
- The p_{i0}'s are decided by solving $E[p_i^{\exp(\alpha)}] = s_i$, where s_i's are a set of prior probabilities that one must determine (called "skeletons").
- A binomial likelihood: $\prod_{i=1}^{d} p_i^{y_i}(1-p_i)^{n_i - y_i}$.
- Posterior of α is obtained by numerical integration.
- The next dose is $\arg\min_i \mid \hat{p}_i - p_T \mid$, where $\hat{p}_i$ is the posterior mean.

The CRM Design – trial conduct

- Challenging to implement in practice (logistic and people issues)
- How does one actually conduct a practical trial using CRM?:
  - Need to set up an infrastructure between statisticians and nurses at clinics (potentially multiple sites).
  - Anti-CRM ad-hoc rules: Coherence and Over-dose control (e.g., no skipping dose when escalation)
  - Team meetings are needed for every patient allocation – CRM decisions may be overruled.

The mTPI and mTPI-2 designs are interval designs. The CCD and BOIN designs are a different type of interval designs.

Hallmark of "Interval Designs"

The decision of dose finding involves inference based on toxicity probability intervals.

- Interval designs : up-and-down decisions based on intervals (mTPI, mTPI-2, CCD, BOIN)

Stay	Escalate	De-escalate
$p_i \in (p_T - \epsilon_1, p_T + \epsilon_2)$	$p_i \in (0, p_T - \epsilon_1)$	$p_i \in (p_T + \epsilon_2, 1)$

- Non-interval designs : CRM chooses the dose

$$\arg\min_i |\hat{p}_i - p_T|,$$

3+3 uses up-and-down decisions based on

$$\frac{y_i}{n_i},$$

The mTPI Design

Interval designs are a simple, transparant, intuitive and model-based .

Sample size = 21 ; Target probability = 30% ; epsilon 1 = 0.05 ; epsilon 2 = 0.05 ; OLD

E: Escalate to the next higher dose; **S**: Stay at the same dose; **D**: De-escalate to the previous lower dose; **DU**: De-escalate to the previous lower dose and the current dose will never be used again in the trial;

Number of Patients

Number of Toxicity	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
0	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E
1	D	S	S	S	S	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E
2		DU	D	S	S	S	S	S	S	S	E	E	E	E	E	E	E	E	E	E	E
3			DU	DU	D	S	S	S	S	S	S	S	S	S	S	E	E	E	E	E	E
4				DU	DU	DU	D	D	S	S	S	S	S	S	S	S	S	S	S	S	E
5					DU	DU	DU	DU	DU	D	S	S	S	S	S	S	S	S	S	S	S
6						DU	DU	DU	DU	DU	DU	D	S	S	S	S	S	S	S	S	S
7							DU	DU	DU	DU	DU	DU	DU	D	S	S	S	S	S	S	S
8								DU	DU	DU	DU	DU	DU	DU	DU	DU	D	S	S	S	S
9									DU	DU	DU	DU	DU	DU	DU	DU	DU	DU	DU	S	S
10										DU	DU	DU	DU	DU	DU	DU	DU	DU	DU	DU	DU
11											DU	DU	DU	DU	DU	DU	DU	DU	DU	DU	DU
12												DU	DU	DU	DU	DU	DU	DU	DU	DU	DU
13													DU	DU	DU	DU	DU	DU	DU	DU	DU
14														DU	DU	DU	DU	DU	DU	DU	DU
15															DU	DU	DU	DU	DU	DU	DU
16																DU	DU	DU	DU	DU	DU
17																	DU	DU	DU	DU	DU
18																		DU	DU	DU	DU
19																			DU	DU	DU
20																				DU	DU
21																					DU

– An mTPI decision table

but in an algorithmic form

Interval-Based Decision Rules

Divid $(0, 1)$ into three intervals:

$$(0, p_T - \epsilon_1), \quad \underbrace{(p_T - \epsilon_1, p_T + \epsilon_2)}_{Equivalence\ Interval}, \quad (p_T + \epsilon_2, 1)$$

.

Associate with

$$E, \qquad S, \qquad D$$

Measure the **unit probability mass (UPM)** of each interval under the posterior of p_i.

Decide the action corresponding to the largest UPM (Bayes rule – more on this next).

UPM and Bayes rule

$$\text{UPM (interval)} = \frac{\text{post. prob}\ \{\ p_i \in (\text{interval})\}}{\text{length (interval)}}$$

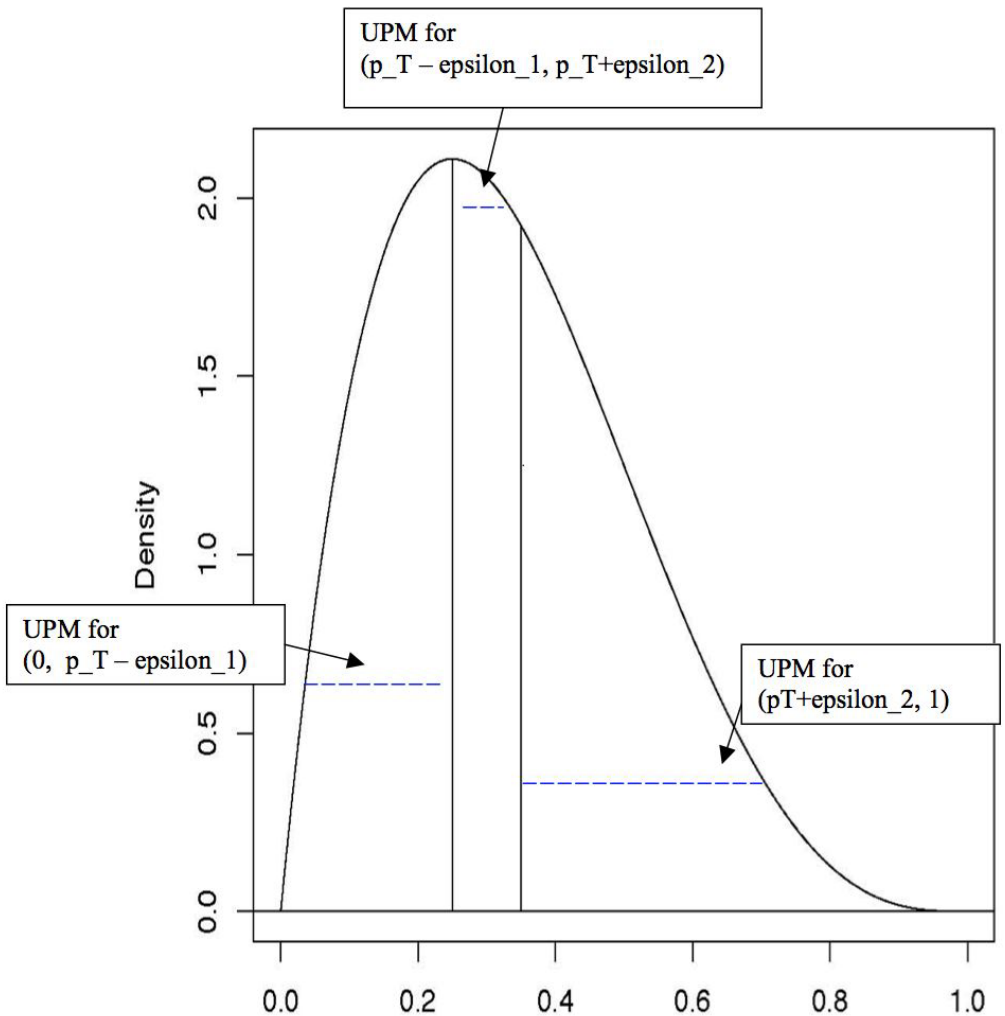

Ockham's razor and interval length

In mTPI, when 3 out of 6 patients have DLT and if $p_T = 0.3$, the decision based on UPM is S, to stay at the current dose. Why?

Ockham's razor states the principle that an explanation of the facts should be no more complicated than necessary

Bayesian model selection requires a prior $p(M_k)$ for the candidate model k and a prior $p(\theta_k \mid M_k)$. Models are selected based on $p(M_k \mid data)$ and automatically applies the Ockham's razor: when two models fit the data equally well, the smaller one wins.

The mTPI design considers three intervals that partition the sample space $(0, 1)$ for the probability of toxicity p_d at a given dose d:

$$\begin{aligned} M_E: &\quad p_d \in (0, p_T - \epsilon_1) \\ M_S: &\quad p_d \in (p_T - \epsilon_1, p_T + \epsilon_2) \\ M_D: &\quad p_d \in (p_T + \epsilon_2, 1) \end{aligned} \tag{1}$$

Typically, p_T ranges from 0.1 to 0.3 in phase I trials, and ϵ's are usually small, say ≤ 0.05. So M_S the middle interval is the smallest (shortest).

Ockham's razor and interval length – Con't

So mTPI is based on the decision rule for dose i

$$\mathcal{D}_{\mathsf{mTPI}} = \arg \max_{k \in \{E,S,D\}} UPM(k,i) \tag{2}$$

Ockham's razor and interval length – Con't

So mTPI is based on the decision rule for dose i

$$\mathcal{D}_{\text{mTPI}} = \arg \max_{k \in \{E,S,D\}} UPM(k,i) \tag{2}$$

Turns out this rule is the Bayes rule for the following decision framework:

Theorem 1. *Given the sampling model* $y_i \mid p_i \sim Bin(n_i, p_i)$ *and priors*

$$\begin{aligned} f(p_i \mid M_k) &\sim \frac{1}{S(M_k)} I(p_i \in M_k) \\ p(M_k) &= \frac{1}{3} \end{aligned}$$

independently for all doses, and given the 0-1 loss function $\ell(i, M_j)$ *in* (3) *for three decisions, where* $i, j \in \{E, S, D\}$, *decision rule* $\mathcal{D}_{mTPI}$ *in* (2) *is optimal in the sense that it minimizes the posterior expected loss.*

$$\ell(a = i, m_d = M_j) = \begin{cases} 1, & \text{if } i \neq j; \\ 0, & \text{if } i = j, \end{cases} \qquad \text{for } i, j \in \{E, S, D\}. \tag{3}$$

UPM is the marginal posterior probability of each model (interval)

Turns out

$$UPM(k,i) = Pr(M_k \text{ is true} \mid \{x_i, n_i\})$$

This is the marginal posterior probability of model k.
The size of the model is the length of the interval. Ockham's razor picks model $M_S : p_i \in (0.25, 0.35)$; mTPI selects decision "S" when $x_i = 3$ out of $n_i = 6$ patients experience the DLT.

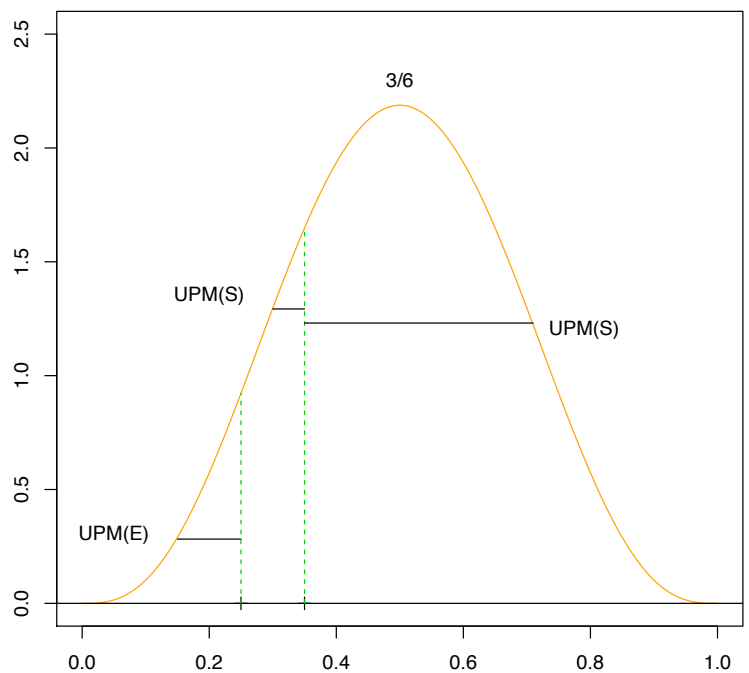

The mTPI-2 design (Guo et al., 2017): Blunt the Ockham's Razor

Divid $(0, 1)$ into subintervals with equal length , same as that of $(p_T - \epsilon_1, p_T + \epsilon_2)$.

Pick the decision $\{D, S, E\}$ corresponding to the interval with the largest UPM.

Still the Bayes (optimal) rule

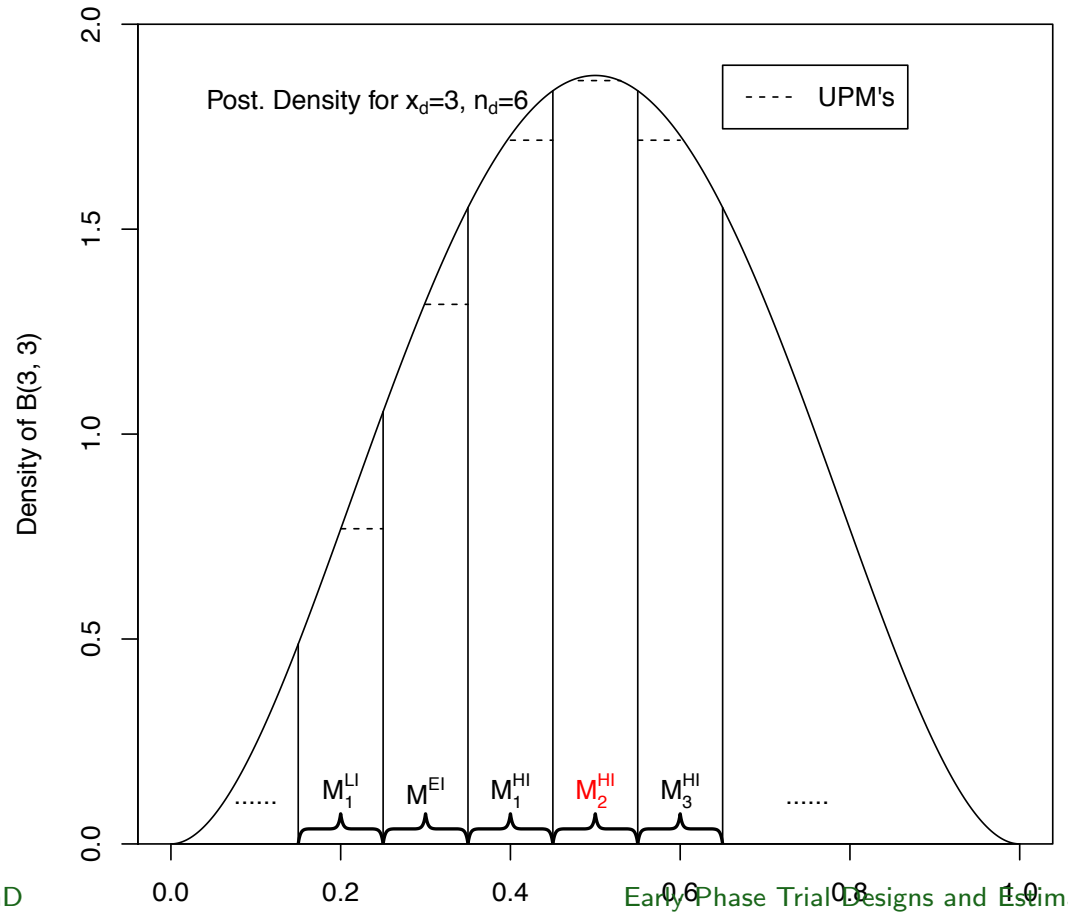

The mTPI-2 Design Compares Favorable to Other Designs

Number of Patients

Number of DLTs	1		2		3		4		5		6		7		8		9		10	
	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2	mTPI	mTPI2
0	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E	E
1	D	D	S	D	S	S	S	S	S	E	E	E	E	E	E	E	E	E	E	E
2			DU	DU	D	D	S	D	S	D	S	S	S	S	S	S	S	E	S	E
3					DU	DU	DU	DU	D	D	S	D	S	D	S	D	S	S	S	S
4							DU	DU	DU	DU	DU	DU	D	D	D	D	S	D	S	D
5									DU	DU	DU	DU	DU	DU	DU	DU	DU	DU	D	D
6											DU	DU	DU	DU	DU	DU	DU	DU	DU	DU
7													DU	DU	DU	DU	DU	DU	DU	DU
8															DU	DU	DU	DU	DU	DU
9																	DU	DU	DU	DU
10																			DU	DU

E: Escalate to the next higher dose; **S**: Stay at the same dose; **D**: De-escalate to the previous lower dose; **DU**: De-escalate to the previous lower dose and the current dose will never be used again in the trial;

The mTPI-2 Design Compares Favorable to Other Designs

We compared mTPI, mTPI-2, 3+3, CRM, BOIN in terms of reliability (power of finding MTD) and safety.

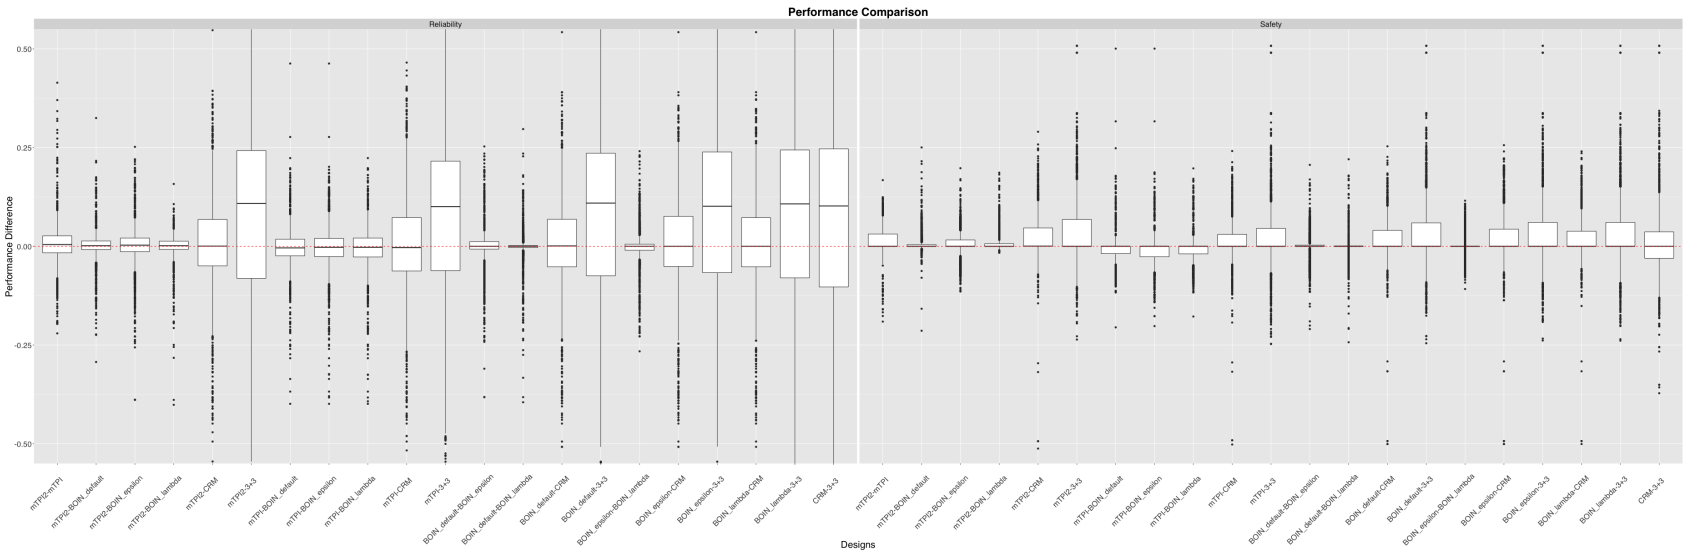

Ji and Yang (2017; https://arxiv.org/abs/1706.03277)

A New Web-based Integrative Dose-Finding Tool

http://udesign.laiyaconsulting.com

A New Web-based Integrative Dose-Finding Tool

http://udesign.laiyaconsulting.com

Web Based: No need to download any software; works on MAC, PC, iPAD, and smart phones – just need an internet browser (e.g., Chrome, FireFox)

Integrative: Offers up to six designs, 3+3, CRM, mTPI-2, BLRM, etc.

Many new features: CRM decision table, etc.

User-friendly: Demo...

Go beyond toxicity probability intervals

- Interval designs are transparent, simple, and easy to implement. And they are model based.
- Existing drug development more often require incorporation of efficacy outcome in dose finding
- Interval designs can be effective for some cases, such as the use of toxicity and efficacy probability intervals (TEPI) in Li et al., (2016) for CAR-T therapies
- The posterior probabilities of the intervals can be used to assess the uncertainties of the dose-finding decisions, which can lead to randomized dose-finding designs (ongoing work).

Thank You!