Dose-finding in oncology clinical trials based on unit probability mass

Yuan Ji

yuanji@mdanderson.org

Division of Quantitative Science

University of Texas M.D. Anderson Cancer Center

In collaborations with B.N. Bekele and Y. Li

Summary

- In oncology, patients are severely ill and often resort to advanced and aggressive treatments with potential serious toxicity
- The type of toxicity is called the dose-limiting toxicity (DLT), which could be as serious as permanent organ damage
- Oncologists want to find an appropriate dose level that is effective to the disease and yet is not "too toxic"
- For example, $< 30\%$ of the patients will experience DLT
- The highest dose of which the probability of toxicity is less than p_T, say, $p_T = 0.30$, is called the maximum tolerated dose, or the MTD.

Dose-finding in oncology

Fixed-dose setup:

- Align the doses from the lowest to the highest, say dose $1, 2, \ldots, 8$.
- Usually starting from the lowest dose, although not necessary
- Treat the first cohort of patients (cohort size ≥ 1) at the starting dose
- Depending on the observed toxicity from the treated patients, increase/decrease/not change the dose level for treating the next cohort

Variable-dose setup:

- the number of doses and their levels are not fixed
- the increment of the dose level depends on the type/grade of the toxicity

Key criteria – SPST

Key points when developing or evaluating a design for phase I trials are SPST

- SAFE – safety is the top priority and should be the first thing to check with any new design: scientific and legal implication
- Performs well in a variety of different scenarios – a good design must perform reasonably well in various settings: simulation is the tool to examine the performance (more on this later).
- SIMPLE – complicated designs are almost surely going to be declined by the physicians: they can ALWAYS use 3+3.
- Transparent – it is easier for the physicians to understand if they can see how doses are assigned. They can even help to improve the design by pointing out unreasonable rules, e.g., escalate when 2 DLTs are found out of 4 patients.

SP – required ST – preferred

Some common misconceptions

- Dose-escalation vs.estimation of the MTD
- Binary toxicity vs. graded toxicity
- Monotonicity is not a big deal
- Curve free vs. model free
- Tradeoff between percentage of correct selection vs. patients assignment

The 3+3

Patients are enrolled in cohorts of 3. The core algorithm of the 3+3 contains:

- If none of the three patients in a cohort has a DLT, another cohort of three will be treated at the next higher dose.
- If one of the three patients has a DLT, three more will be treated at the same dose.
- The dose escalation continues until at least two patients of three or six patients have DLTs.

People add in additional rules (e.g., allowing for dose de-escalation in the event of more than two DLTs) to make the algorithm more flexible.

Evaluation of the 3+3

Let us look at the SPST:

- Safe? YES. In general, the 3+3 is known to be conservative.
- Performs well? Not necessary, especially in cases with more doses.
- Simple? YES. No computation for trial conduct.
- Transparent? YES. The rules are laid out up front and clear.

So one of the "required" properties is not satisfied – does not always perform well.

Evaluation of the 3+3 – con't

In addition:

- Problem number one: what if my $p_t = 0.1$? Can 3 patients estimate a 10% proportion?
- Problem number two: why would I want to stop at 6 patients regardless?
- Problem number three: A single DLT will forbid any further escalation. But DLT is a random event!

Example: if $p_i = 0.05$, $P(\geq 1$ DLT in 3 patients$) = 0.14$. So even the true toxicity rate for a dose is 5%, one would at least stop escalatioin 14% of the time!

The CRM – General idea

The first notable pulication about the CRM is O'Quigley, Pepe, and Fisher (1990) (Biometrics).

- Key idea: assuming a dose-toxicity response curve, the CRM continuously reassess the toxicity of all the doses and assign future patients to the dose closest to the MTD.
- Key components (assumptions):
  - Need to specify the form of the dose response curve.
  - Need to specify a probability model with appropriate priors.
  - Need to calibrate the model parameters before trial starts so that appropriate operating characterisitics can be achieved.
  - Need to set up a web-interface to allow for real-time dose-assignment during trial conduct.

The CRM – 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 known 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 "skeleton").
- The probability model is a binomial likelihood: $\prod_{i=1}^{d} p_i^{y_i}(1-p_i)^{n_i-y_i}$ where y_i and n_i are the number of DLTs and patients at dose i, respectively.
- 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 – con't

- The dose-response curve: $p_i = p_{i0}^{\exp(\alpha)}$.
- The prior for α is $N(0, 2)$.
- Solving $E[p_i^{\exp(\alpha)}] = s_i$ for p_i (need to decide s_i).
- The next dose is $\arg\min_i \mid \hat{p}_i - p_T \mid$.

One problem is to choose s_i. So far, one paper (to appear) claims it has an automatic way of choosing s_i. Most times, people do it arbitrarily. For example, with five doses, s_i's are (0.1, 0.2, 0.3, 0.4, 0.5).

The CRM – trial conduct

- The software only provides simulation results.
- How does one actually assign patients to doses when a trial starts?
- Answer:
  - Need to write another software involving numerical integration at real time.
  - Need to set up an interface between statisticians and nurses at clinics.
  - The interface allows the nurses to input the toxicity data, with which statisticians can compute the next dose to assign.

The CRM – brief evaluation

Let us look at the SPST.

- Safe? Somewhat, need additional rules.
- Performs well? Depending on the priors.
- Simple? NO.
- Transparent? No.

A typical set of observation

Notation: E = Escalation; S = Stay; D = De-escalation; U = Unacceptable (too toxic).

Current dose	observed toxicity	Decision
1	0/3	E
2	0/3	E
3	2/3	D
2	2/6	S
2	...	...

At the end of the trial, one dose is selected as the estimated MTD.

In the above case, probably dose 2 will be selected.

Toxicity probability interval method (TPI)

- The method provides a spread sheet for a given number of p_T with an embedded macro
- Suppose patients are treated at dose i
- Identify the number of patients treated at this dose and go to the corresponding column in the table;
- Identify the number patients experienced toxicity and go to the corresponding row in the table;
- the corresponding entry in the cell provides the dose-assignment decision that one needs to take.

Idea

Main idea:

- Suppose patients are currently being assigned to dose i, with toxicity probability p_i.
- Dose-assignment decisions are based on posterior toxicity probability that p_i is in one of the following three intervals: $\{(0, p_T - K_1\sigma_i),\ [p_T - K_1\sigma_i, p_T + K_2\sigma_i],\ (p_T + K_2\sigma_i, 1)\}$, where σ_i is the posterior standard deviation of p_i, and K_1 and K_2 are some small positive constants (default value $K_1 = 1.5$, $K_2 = 1$).
- If the posterior probability of the first, second, or the third interval is the largest, then the toxicity probability p_i is likely to be smaller to p_T, close to p_T, or larger than p_T, respectively, implying the dose level should be escalated, kept unchanged, or de-escalated.

Probability model

- Likelihood function is a product of binomial densities: $l(\boldsymbol{p}) \propto \prod_{i=1}^{d} p_i^{x_i}(1-p_i)^{n_i - x_i}$, where n_i and x_i are the numbers of patients treated at dose i and experienced DLT, respectively.
- The priors of p_i are i.i.d. Beta(α, α), where α takes a small value, e.g., $\alpha = 0.005$, resulting in a U-shaped prior.
- Posteriors are beta with known parameter values.

Prior-posterior

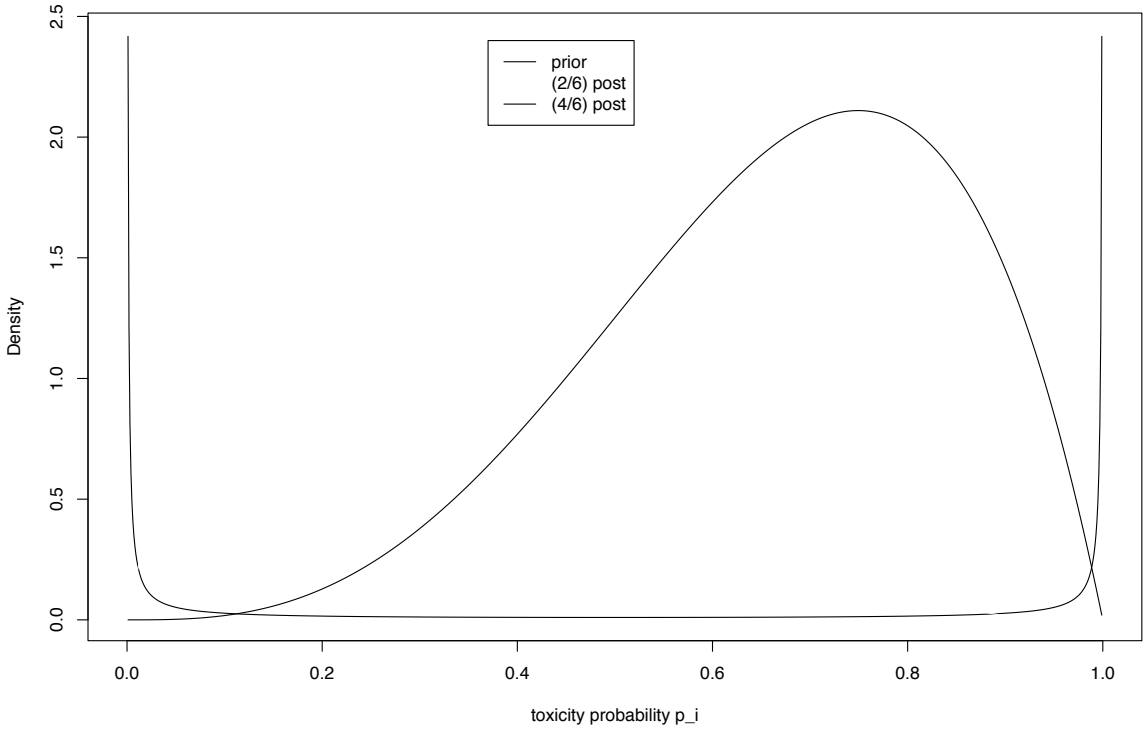

Decision rules

Let D, S, E denote the decision to de-escalate to dose $(i-1)$, stay at dose i, and escalate to dose $(i+1)$, respectively. Following the main idea, define the posterior probabilities for the three intervals:

$$q(D,i) = P(p_i - p_T > K_1\sigma_i|\text{data}),$$

$$q(S,i) = P(-K_2\sigma_i \leq p_i - p_T \leq K_1\sigma_i|\text{data}),$$

$$q(E,i) = P(p_i - p_T < -K_2\sigma_i|\text{data}).$$

The dose-assignment rule

$$\mathcal{B}_i = \arg \max_{m\in\{D,S,E\}} q(m,i),$$

i.e. take the decision that has the maximum posterior probability.

Two issues

- What if the first dose is very toxic?
- What if dose $i-1$ is safe, but dose i is very toxic?

Exclusion rule

Define

$$\mathcal{T}_i = 1\left\{P(p_i > p_T|\text{data}) > \xi\right\},$$

where $1\{\}$ is the indicator function and $\xi \in (0,1)$ is a cutoff value (e.g., $\xi = 0.95$). For a large value of ξ, $\mathcal{T}_i = 1$ implies that dose i is very likely to be highly toxic, and escalation to this dose should be permanently prohibited. To incorporate this rule, modified decision rule is given by

$$\mathcal{B}_i^{(e)} = \arg \max_{m \in \{D,S,\tilde{E}\}} q(m,i),$$

where $q(\tilde{E},i) = q(E,i)(1-\mathcal{T}_{i+1})$.Therefore, if $\mathcal{T}_{i+1} = 1$, the probability $q_{\tilde{E},i}$ equals zero and the assignment rule $\mathcal{B}_i^{(e)}$ can be only D, to de-escalate, or S, to stay.

Dose-finding algorithm

- Suppose that the current tried dose is i for $i = 1, \ldots, d$. After the toxicity outcomes of the last cohort are observed, select the dose for treating the next cohort among $\{(i-1), i, (i+1)\}$ based on the assignment rule $\mathcal{B}_i^{(e)}$. There are two exceptions: if $i = 1$, the next available doses are $\{i, (i+1)\}$; if $i = d$, the next available doses are $\{(i-1), i\}$.

Dose-finding algorithm con't

- Suppose that dose 1 is a dose that has been tried previously. If $\mathcal{T}_1 = 1$, terminate the trial due to excessive toxicity. Otherwise, terminate the trial when the maximum sample size is reached.
- In the special case of cohort of size 1, by convention, do not apply the exclusion rule $\mathcal{T}_i$ until two or more patients have been evaluated at a dose.

Dose-finding algorithm con't

- At the end of the trial, select the dose as the estimated MTD with the smallest difference $|\hat{p}_i^* - p_T|$ among all the tried doses i for which $\mathcal{T}_i = 0$.
- Quantity $\hat{p}_i^*$ is the isotonic transformation estimator of the posterior mean $\hat{p}_i$ so that $\hat{p}_j^* \leq \hat{p}_i^*$ for $j > i$.
- If two or more doses tie for the smallest difference, perform the following rule. Let p^* denote the transformed posterior mean of the tied doses.
  - If $p^* < p_T$, choose the highest dose among the tied doses.
  - If $p^* > p_T$, choose the lowest dose among the tied doses.

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 1		5	25	50	60	70	80	90	95	none		
Bayes	% MTD	13	**79**	8	0	0	0	0	0	0	25	30
	# Pts	7.7	16.1	5.8	0.5	0	0	0	0			
3+3	% MTD	24	**58**	16	2	0	0	0	0	0	25	12
	# Pts	4.0	5.0	2.6	0.4	0	0	0	0			
BCD	% MTD	10	**78**	11	1	0	0	0	0	0	24	30
	# Pts	11.4	11.5	5.2	1.4	0.3	0.1	0	0			
CFM	% MTD	6	**80**	14	0	0	0	0	0	0	29	30
	# Pts	5.2	16.3	7.5	0.9	0	0	0	0			
CRM	% MTD	6	**83**	11	0	0	0	0	0	0	27	30
	# Pts	5.7	18.6	4.9	1.0	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 2		1	2	3	4	5	6	25	50	none		
Bayes	% MTD	0	0	0	0	2	22	**62**	14	0	12	30
	# pt	3.2	3.2	3.4	3.5	3.7	4.5	5.9	2.6			
3+3	% MTD	0	0	0	2	3	21	**46**	8	0	11	27
	# pt	3.1	3.2	3.3	3.4	3.3	3.7	4.5	2.2			
BCD	% MTD	0	0	1	2	7	24	**56**	10	0	10	30
	# pt	3.2	3.6	3.6	3.5	3.8	5.4	4.8	2.1			
CFM	% MTD	0	0	0	0	1	22	**61**	16	0	12	30
	# pt	3.1	3.0	3.1	3.5	3.7	5.1	6.3	2.1			
CRM	% MTD	0	0	1	1	5	22	**50**	21	0	13	30
	# pt	3.1	3.4	3.3	3.7	3.6	4.4	5.1	3.4			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 3		1	5	50	60	70	80	90	95	none		
Bayes	% MTD	0	**79**	21	0	0	0	0	0	0	22	30
	# pt	5.5	13.2	10.2	1.0	0	0	0	0			
3+3	% MTD	0	**70**	28	2	0	0	0	0	0	22	13
	# pt	3.1	5.2	4.4	0.7	0.1	0	0	0			
BCD	% MTD	0	**60**	39	1	0	0	0	0	0	22	30
	# pt	4.9	14.3	8.2	2.2	0.4	0	0	0			
CFM	% MTD	0	**56**	44	0	0	0	0	0	0	28	30
	# pt	3.1	11.7	13.1	2.0	0.1	0	0	0			
CRM	% MTD	0	**49**	51	0	0	0	0	0	0	26	30
	# pt	3.1	13.0	12.0	1.8	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 4**		40	50	60	70	80	90	95	99	none		
Bayes	% MTD	31	2	0	0	0	0	0	0	**67**	41	19
	# pt	16.8	1.8	0.2	0	0	0	0	0			
3+3	% MTD	38	9	1	0	0	0	0	0	**52**	43	6
	# pt	4.7	0.5	0.6	0.7	0	0	0	0			
BCD	% MTD	38	2	0	0	0	0	0	0	**60**	45	18
	# pt	12.6	4.6	1.2	0.2	0	0	0	0			
CFM	% MTD	38	3	1	0	0	0	0	0	**58**	42	14
	# pt	11.7	1.9	0.5	0.1	0	0	0	0			
CRM	% MTD	47	2	0	0	0	0	0	0	**51**	42	23
	# pt	20.2	2.5	0.2	0	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 5		15	25	35	45	55	65	75	85	none		
Bayes	% MTD	31	**41**	21	7	0	0	0	0	0	24	30
	# pt	12.4	9.5	5.5	1.9	0.3	0	0	0			
3+3	% MTD	29	**37**	20	7	1	0	0	0	8	26	12
	# pt	4.4	3.9	2.4	0.9	0.2	0	0	0			
BCD	% MTD	21	**46**	22	6	1	0	0	0	5	26	29
	# pt	10.6	9.2	5.7	2.5	0.8	0.1	0	0			
CFM	% MTD	15	**44**	32	7	0	0	0	0	0	27	30
	# pt	8.0	10.6	8.0	2.6	0.4	0	0	0			
CRM	% MTD	36	**47**	14	2	0	0	0	0	0	23	30
	# pt	13.8	11.4	3.6	0.9	0.2	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox %	Ave. n
	Dose	1	2	3	4	5	6	7	8			
Scenario 6		5	15	25	35	45	55	65	75	none		
Bayes	% MTD	2	24	**42**	24	7	0	0	0	0	22	30
	# pt	5.1	8.2	9.2	5.7	1.6	0.3	0	0			
3+3	% MTD	9	28	**34**	22	5	0	0	0	0	21	15
	# pt	3.6	4.3	3.8	2.3	0.8	0.2	0	0			
BCD	% MTD	1	29	**44**	19	6	1	0	0	0	21	30
	# pt	6.8	8.7	7.5	4.4	1.9	0.6	0.1	0			
CFM	% MTD	0	14	**49**	29	6	0	0	0	0	24	30
	# pt	3.9	6.2	10.7	7.1	1.8	0.3	0	0			
CRM	% MTD	4	37	**45**	12	2	0	0	0	0	20	30
	# pt	5.5	11.5	8.9	3.4	0.7	0.1	0	0			

Evaluation of the TPI

Let us look at the SSTP:

- Simple? YES.
- Safe? YES.
- Transparent? YES.
- Performs well? YES.

Looks like we have got a perfect design. But...

TPI – Not there yet

Two issues:

- The interval rule is based on $(p_T - K_1\sigma_i, p_T + K_2\sigma_i)$.
  - σ_i is the posterior standard deviation of p_i – easy to compute (in closed form).
  - K_1 and K_2 are fixed and must be given. Default values are $K_1 = 1.5$ and $K_2 = 1$. However, they need to be calibrated for certain trials.
- The number of patients treated at doses above the MTD is a bit higher (compared to the CRM, for example) – room for improvement.

The modified TPI (mTPI) method

The mTPI is based on a new statistics called the unit probability mass. The mTPI improves the TPI on two aspects:

- mTPI is calibration free – does not require tuning of parameters
- mTPI is safer – treats fewer patients at over-toxic doses while maintaining other good performance properties of the TPI.

The mTPI – introduction

For interval (a, b), the UPM is defined as

$$\frac{F(b) - F(a)}{b - a},$$

where $F(\cdot)$ is a probability distribution function. In the mTPI method, the framework for dose-findin is the following:

- Define an equivalence interval (EI) $I_{EI} = (p_T - \epsilon_1, p_T + \epsilon_2)$, where any dose with probability of toxicity within the interval can be considered as a MTD.
- Compute the UPM of three intervals, $I_{safe} = (0, p_T - \epsilon_1)$, $I_{toxic} = (p_T + \epsilon_2, 1)$ and I_{EI} under the posterior distribution of p_i.
- Choose to escalate, de-escalate, or stay at dose i if the UPM of I_{safe}, I_{toxic}, or I_{EI} is the largest, respectively.

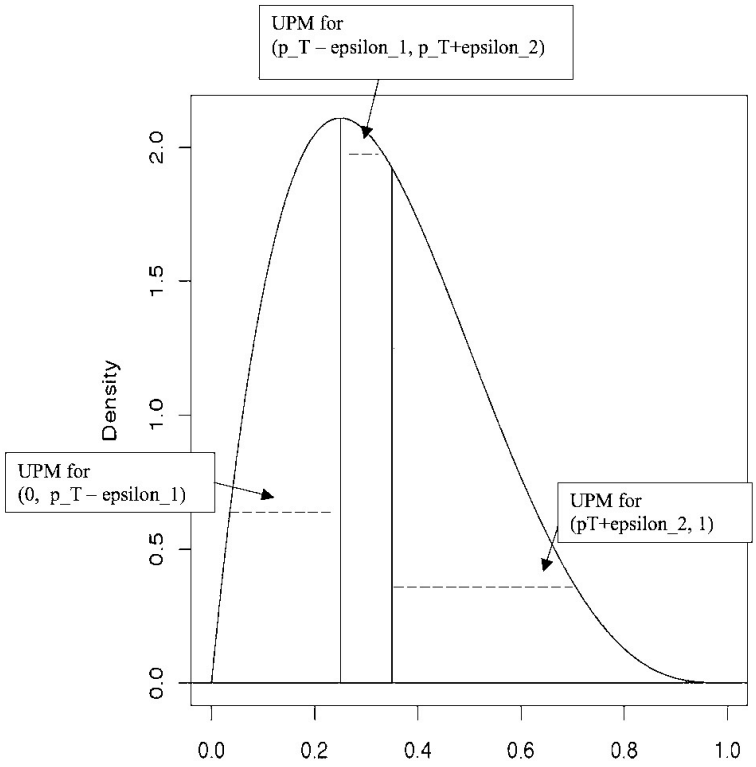

The mTPI – EI

The equivalence interval $I_{EI} = (p_T - \epsilon_1, p_T + \epsilon_2)$ is the only thing that one needs to decide.

- We will show that the choices of ϵ_1 and ϵ_2 do not affect the performance of the mTPI.
- We should elicit the EI as a general practice: since we never get a dose with exact toxicity probability p_T.

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox pctg*	Ave n
	Dose	1	2	3	4	5	6	7	8			
Scenario 1		5	25	50	60	70	80	90	95	none		
mTPI	% MTD	14	**78**	8	0	0	0	0	0	0	24	30
	# Pts	7.1	18.3	4.4	0.2	0	0	0	0			
TPI	% MTD	13	**79**	8	0	0	0	0	0	0	25	30
	# Pts	7.7	16.1	5.8	0.5	0	0	0	0			
CRM	% MTD	6	**83**	11	0	0	0	0	0	0	27	30
	# Pts	5.7	18.6	4.9	1.0	0	0	0	0			
Scenario 2		1	2	3	4	5	25	50	60	none		
mTPI	% MTD	0	0	0	2	16	**71**	10	1	0	16	30
	# Pts	3.2	3.5	3.5	4.0	5.2	8.1	2.3	0.1			
TPI	% MTD	0	0	0	0	19	**70**	11	0	0	15	30
	# pt	3.2	3.2	3.3	3.6	5.0	8.0	3.3	0.3			
CRM	% MTD	0	0	1	1	20	**61**	16	2	0	16	30
	# pt	3.1	3.4	3.3	3.7	4.7	7.0	3.8	0.9			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox pctg*	Ave n
	Dose	1	2	3	4	5	6	7	8			
Scenario 3		1	5	50	60	70	80	90	95	none		
mTPI	% MTD	0	**82**	17	0	0	0	0	0	0	21	30
	# pt	3.2	15.9	10.3	0.6	0	0	0	0			
TPI	% MTD	0	**79**	21	0	0	0	0	0	0	22	30
	# pt	5.5	13.2	10.2	1.0	0	0	0	0			
CRM	% MTD	0	**49**	51	0	0	0	0	0	0	26	30
	# pt	3.1	13.0	12.0	1.8	0	0	0	0			
Scenario 4		40	50	60	70	80	90	95	99	none		
mTPI	% MTD	31	2	0	0	0	0	0	0	**67**	41	19
	# pt	16.8	2.0	0.2	0	0	0	0	0			
TPI	% MTD	31	2	0	0	0	0	0	0	**67**	41	19
	# pt	16.8	1.8	0.2	0	0	0	0	0			
CRM	% MTD	47	2	0	0	0	0	0	0	**51**	42	23
	# pt	20.2	2.5	0.2	0	0	0	0	0			

Operating characteristics

		Recommendation percentage at dose level $p_T = 0.25$									Tox pctg*	Ave n
	Dose	1	2	3	4	5	6	7	8			
Scenario 5		15	25	35	45	55	65	75	85	none		
mTPI	% MTD	29	**45**	20	4	0	0	0	0	0	24	30
	# pt	12.4	10.9	5.0	1.1	0.1	0	0	0			
TPI	% MTD	31	**41**	21	7	0	0	0	0	0	24	30
	# pt	12.4	9.5	5.5	1.9	0.3	0	0	0			
CRM	% MTD	36	**47**	14	2	0	0	0	0	0	24	30
	# pt	13.8	11.4	3.6	0.9	0.2	0	0	0			
Scenario 6		5	15	25	35	45	55	65	75	none		
mTPI	% MTD	2	28	**42**	23	4	0	0	0	0	20	30
	# pt	4.9	10.2	9.3	4.5	0.9	0.1	0	0			
TPI	% MTD	2	24	**42**	24	7	0	0	0	0	22	30
	# pt	5.1	8.2	9.2	5.7	1.6	0.3	0	0			
CRM	% MTD	4	37	**45**	12	2	0	0	0	0	20	30
	# pt	5.5	11.5	8.9	3.4	0.7	0.1	0	0			

Sensitivity to the EI

		Recommendation percentage at dose level $p_T = 0.25$									Tox pctg*	Ave n
	Dose	1	2	3	4	5	6	7	8			
Scenario 1		5	25	50	60	70	80	90	95	none		
$\epsilon_1 = \epsilon_2 = .05$	% MTD	14	**78**	8	0	0	0	0	0	0	24	30
	# Pts	7.1	18.3	4.4	0.2	0	0	0	0			
$\epsilon_1 = \epsilon_2 = .2$	% MTD	15	**76**	9	0	0	0	0	0	0	24	30
	# Pts	7.7	18.4	3.7	0.2	0	0	0	0			
$\epsilon_1 = \epsilon_2 = .001$	% MTD	14	**78**	8	0	0	0	0	0	0	24	30
	# Pts	7.1	18.3	4.3	0.2	0	0	0	0			

Discussion

The mTPI and TPI are SSTP.

- The mTPI is safer.
- The mTPI is calibration free.

I would now only use the mTPI. So do we have the same Excel software?

Software

Yes. Both mTPI amd TPI are available to download for free.

- Paper and software http://odin.mdacc.tmc.edu/∼yuanj/

Simulation-based evaluation

The choice of scenarios is rarely discussed in the literature. However, it is critical as failure to include important scenarios may lead to undesirable consequences in practice.

- Early-stopping scenarios (0.50, 0.60, ...)
- Fast-escalating scenarios (0.01, 0.03, 0.05, 0.10, 0.15, 0.3)
- Big-gap scenarios (0.05, 0.10, 0.50, 0.60, ...)
- Regular scenarios (0.10, 0.20, 0.30, 0.40, ...)

Happy dose-finding!

Thank you!