A modified toxicity probability interval method for dose-finding trials

Yuan Ji, Ping Liu, Yisheng Li and B. Nebiyou Bekele

Clin Trials 2010 7: 653 originally published online 8 October 2010
DOI: 10.1177/1740774510382799

CLINICAL TRIALS *ARTICLE* *Clinical Trials* 2010; **7:** 653–663

A modified toxicity probability interval method for dose-finding trials

Yuan Ji[a], *Ping Liu*[b], *Yisheng Li*[b] *and B Nebiyou Bekele*[b]

Background Building on earlier work, the toxicity probability interval (TPI) method, we present a modified TPI (mTPI) design that is calibration-free for phase I trials.
Purpose Our goal is to improve the trial conduct and provide more effective designs while maintaining the simplicity of the original TPI design.
Methods Like the TPI method, the mTPI consists of a practical dose-finding scheme guided by the posterior inference for a simple Bayesian model. However, the new method proposes improved dose-finding decision rules based on a new statistic, the unit probability mass (UPM). For a given interval and a probability distribution, the UPM is defined as the ratio of the probability mass of the interval to the length of the interval.
Results The improvement through the use of the UPM for dose finding is three-fold: (1) the mTPI method appears to be safer than the TPI method in that it puts fewer patients on toxic doses; (2) the mTPI method eliminates the need for calibrating two key parameters, which is required in the TPI method and is a known difficult issue; and (3) the mTPI method corresponds to the Bayes rule under a decision theoretic framework and possesses additional desirable large- and small-sample properties.
Limitation The proposed method is applicable to dose-finding trials with a binary toxicity endpoint.
Conclusion The new method mTPI is essentially calibration free and exhibits improved performance over the TPI method. These features make the mTPI a desirable choice for the design of practical trials. *Clinical Trials* 2010; **7**: 653–663. http://ctj.sagepub.com

Introduction

Dose-finding trials in oncology aim to find the maximum tolerated dose (MTD), the highest dose at which the toxicity probability is less than a target probability, denoted by p_T (e.g., $p_T = 0.30$). Usually a grid of dose levels is predetermined before the trial starts. Patients are enrolled sequentially and adaptively treated at a given dose based on the observed dose-limiting toxicity (DLT).

Currently, the most widely used method in practice is the 3+3 method. Since the introduction of the 3+3 method, a large number of statistical methods have been proposed that attempt to improve the design of phase I dose-finding oncology trials. The most prominent work among these is perhaps the continual reassessment method (CRM) by O'Quigley *et al.* [1] and its many extensions [2–6], among others. The basic idea of the CRM is to sequentially update the estimates of

[a]Department of Bioinformatics and Computational Biology, The University of Texas M.D. Anderson Cancer Center, Houston, TX, USA, [b]Department of Biostatistics, The University of Texas M.D. Anderson Cancer Center, Houston, TX, USA
Author for correspondence: Yuan Ji, Department of Bioinformatics and Computational Biology, The University of Texas M.D. Anderson Cancer Center, Houston, TX, USA. E-mail: yuanji@mdanderson.org

10.1177/1740774510382799

dose toxicity probabilities and assign future patients adaptively to the dose deemed closest to the MTD.

Although model-based methods have shown to be superior to the algorithm-based $3+3$ method [7], the simplicity of the $3+3$ method still makes it by far the most popular method chosen for practical trials. There are two reasons for this. First, using model-based methods to monitor a practical trial usually requires a computer program that allows investigators to obtain dose assignments in real time. This is an elegant but tedious matter that calls for close collaboration between physicians, statisticians, and computer programmers. Second, most model-based designs require the conduct of computer simulations before the trial starts, in which statisticians must calibrate the design parameters to achieve desirable operating characteristics. For example, the CRM requires users to provide a set of prior estimates of the toxicity probabilities for the candidate doses to be used in the trial. This set should be elicited from the investigators. However, the performance of the CRM under the computer simulations will depend on the agreement of these prior estimates and the true toxicity probabilities specified in the simulations. A poorly elicited set of prior estimates will lead to poor operating characteristics. Therefore, tuning of the prior estimates becomes an important but sometimes challenging process. Recent attempts [8] have been made to alleviate this issue, although nontrivial calibration is still required.

Ji *et al*. [9] provided a framework different from the CRM. They proposed using simple Bayesian models to describe the observed toxicity data, and introduced a set of decision rules based on toxicity posterior intervals (TPI). Their method is implemented in an Excel spreadsheet, which reduces the burden in making dose assignments during an ongoing trial. However, the performance of the TPI method in certain cases is sensitive to two key parameters, namely K_1 and K_2, which are two weights that define three toxicity intervals, $(0, p_T - K_1\sigma_1)$, $[p_T - k_1\sigma_1, p_T + k_2\sigma_2]$, and $(p_T + K_2\sigma_2, 1)$, where σ_1 and σ_2 are larger values of K_1 and K_2, means more deviations from the target p_T. The posterior probabilities of these three intervals are used for making dose-assignment decisions. A default set of values for K_1 and K_2 was proposed in [9], but it is unclear how sensitive their method is to changes in these values. We have conducted simulations (results not shown) and found that different values of K_1 and K_2 could lead to different results.

In this article, we propose a *calibration-free* modified TPI (mTPI) design. The implementation of the method is essentially effortless, although the underlying statistical theory is not trivial. Implementation of the mTPI method does not require tuning of model parameters, and computer simulations can be carried out in an Excel macro. In addition, similar to the TPI design in [9], the mTPI design is transparent in the sense that physicians can see all the possible dose-finding decisions before the trial starts. Consequently, patients enrolled into the trial can be allocated to appropriate doses without conducting additional computations.

The decision rules of the mTPI design are based on the unit probability mass (UPM). For a given interval on the real line, the UPM is defined as the ratio of the probability of the interval (based on a probability measure) and the length of the interval. For example, for a continuous random variable X with cumulative distribution function $F(x)$, the UPM of an interval $(a, b]$ under the distribution $F(x)$ is $\{F(b) - F(a)\}/(b-a)$. The mTPI method only requires a definition of an equivalence interval (EI), $[p_T - \epsilon_1, p_T + \epsilon_2]$, in which any dose is considered as a potential candidate for the true MTD. We will show that the performance of the mTPI method is robust to the definition of the EI. Therefore, one does not need to calibrate the EI for different trials and physicians. Upon determination of the EI, we compute the mTPI of the three resulting toxicity intervals and choose one of three actions, escalating to the higher dose, staying at the same dose, or de-escalating to the lower dose, depending on which corresponds to the interval with the lowest UPM.

In the section 'Dose-finding method', we introduce the basic idea of the mTPI method as well as the dose-finding algorithm. In the sections 'Probability model' and 'Large-sample properties', we present the UMP method and its theoretical properties. We evaluate the performance of the new method under small sample sizes and conduct sensitivity analyses in the section 'Simulation study'. We describe software in the section 'Software' and provide discussions in the last section.

Dose-finding method

The derivation of the dose-finding rules for the mTPI method involves two steps. In the first step, we introduce an EI, which leads to three toxicity probability intervals that partition (0, 1). Building upon the EI, we set up a decision-theoretic framework and derive a Bayes rule. We show that the Bayes rule is equivalent to computing the UPM for the toxicity probability intervals.

Consider d dose levels of a certain cytotoxic drug in a phase I trial. Let p_i be the unknown probability of toxicity associated with the i-th dose, $i = 1, \ldots, d$.

The toxicity probability usually increases with the dose level, so we assume $p_1 < p_2 < \ldots < p_d$.

Suppose that dose i is currently used for treating patients and n_i ($n_i \geq 1$) patients have been treated at this dose. Suppose x_i ($x_i \leq n_i$) patients experienced toxicity. Based on the observed values of x_i and n_i, we assume that physicians choose one of the following three decisions: de-escalate (D) to the next lower dose ($i-1$); stay (S) at the same dose i; or escalate (E) to the next higher dose ($i+1$). Depending on the decision, the next cohort is treated at dose $j \in \{i-1, i, i+1\}$; the values of x_j and n_j are then observed for the new cohort, and an appropriate decision is chosen once again. The trial thus proceeds with the next cohort.

Equivalence interval

The EI is defined as $[p_T-\epsilon_1, p_T+\epsilon_2]$, $\epsilon_1, \epsilon_2 \geq 0$. It contains those doses considered so close to the true MTD that physicians would agree to select them as the estimated MTD. An EI for the trial is elicited from collaborating physicians. For example, if the true MTD has a toxicity probability $p_T = 0.3$, then a physician may agree to select any dose between [0.25, 0.35] as the estimated MTD. In a different trial, another physician may agree on an EI of [0.2, 0.4]. Compared with other model-based designs, the definition of EI does not increase the complexity of the design for phase I trials. The reason is that for almost all the model-based methods, the probability of actually finding a dose with a toxicity probability equal to p_T is zero. Hence, an implicit criterion usually is defined as the distance between the toxicity probability of a given dose and the target p_T. For example, in the CRM, such a measure is an L_1 norm, that is, $|p_i - p_T|$. Here, we explicitly ask the physicians to express their intrinsic measure of such a distance in the form of EIs. A simple guide for elicitation of the EI is provided below.

- Ask the physician to indicate the lowest toxicity probability that he/she would be comfortable using to treat future patients without dose escalation. This determines $p_T - \epsilon_1$.
- Ask the physician to indicate the highest toxicity probability that he/she would be comfortable using to treat future patients without dose de-escalation. This determines $p_T + \epsilon_2$.

Dose-finding algorithm

Defining an EI results in the partition of the unit interval (0, 1) into three subintervals; (0, $p_T-\epsilon_1$), $[p_T-\epsilon_1, p_T+\epsilon_2]$, and $(p_T+\epsilon_2, 1)$. Doses in these three intervals are deemed lower, close to, and higher than the MTD, respectively. With this clarification, we propose a dose-finding algorithm.

A dose-finding algorithm with additional safety rules is given below. These additional rules are important for practical concerns [9].

- Suppose that the current tried dose is i, $i \in \{1, \ldots, d\}$. After the toxicity outcomes of the last cohort are observed, choose E, S, or D if the interval $(0, p_T-\epsilon_1)$, $[p_T-\epsilon_1, p_T+\epsilon_2]$, or $(p_T+\epsilon_2, 1)$ has the largest UPM, respectively. Figure 1 provides an illustration of how the UPM is associated with dose-finding decisions for an EI.
- *Safety rule 1 (early termination)*: Suppose that dose 1 has been used to treat patients. If $\Pr(p_1 > p_T|\text{data}) > \xi$ for a ξ close to 1 (say, $\xi = 0.95$), then terminate the trial due to excessive toxicity. Otherwise, terminate the trial when the maximum sample size is reached.
- *Safety rule 2 (dose exclusion)*: Suppose that the decision is E, to escalate from dose i to $(i+1)$. If $\Pr(p_{i+1} > p_T|\text{data}) > \xi$, for a ξ close to 1 (say, $\xi = 0.95$), then treat the next cohort of patients at dose i and exclude doses $(i+1)$ and higher from the trial, that is, these doses will never be used again in the trial.

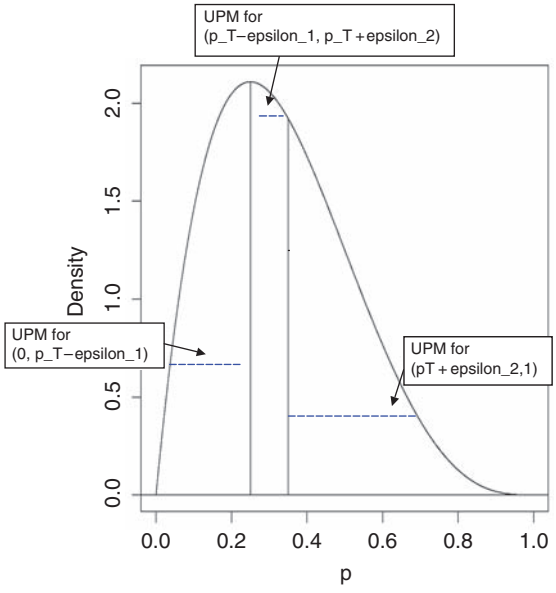

Figure 1 A demonstration of the UPMs for three intervals as defined in the 'Introduction' section. The two vertical lines result in three intervals on the *X*-axis. The UPM for each of the three intervals is indicated by the dashed horizontal line. The equivalence interval in the middle has the highest UPM under the distribution defined by the density curve

- At the end of the trial, select as the estimated MTD the dose with the smallest difference $|\hat{p}_i - p_T|$ among all the tried doses i for which $\Pr(p_i > p_T|\text{data}) <= \xi$. Here $\hat{p}_i$ is a sensible estimate of p_i, for example, the isotonically transformed posterior mean [9]. If two or more doses tie for the smallest difference, the toxicity probabilities of the tied doses can only be $(p_T + q)$ or $(p_T - q)$ for some $q \in (0,\ p_T)$. Perform the following rule:

  – If there is at least one dose at which $\hat{p}_i$ equals $(p_T - q)$, choose the highest dose among those at which $\hat{p}_i$ equal $(p_T - q)$;
  – Otherwise, choose the lowest dose among the tied doses.

Safety rule 2 will terminate a trial early due to high toxicity at the starting dose. This rule is based on a posterior probability, thus taking into consideration all the data at the starting dose. However, due to the thresholding, there is a small probability that a trial could be terminated even when the starting dose is safe, which is the type I error associated with this rule.

Unlike other rules in the algorithm, which concern dose finding, the last rule deals with the selection of the MTD when the trial ends, which is an issue related to statistical inference. We propose to use data from all the patients across doses in this step so that we can borrow strength when estimating the final MTD. We believe that more investigation into this step can be carried out to refine the procedure, perhaps in a similar fashion as Stylianou and Flournoy [10] did for a different dose-finding design.

Probability model

We now focus our attention on technical details of the method. In this section, we provide a statistical framework that leads to the above dose-finding algorithm.

First, we propose a set of penalty functions for choosing a proper decision from among D, S, or E, which is similar to the one in Ji *et al.* [11]. For dose i, define the penalty functions

$$L(D, p_i) = \begin{cases} N_D, & \text{if } p_i - p_T < -\epsilon_1; \\ K_D, & \text{if } -\epsilon_1 \le p_i - p_T \le \epsilon_2; \\ 0, & \text{if } p_i - p_T > \epsilon_2; \end{cases}$$

$$L(S, p_i) = \begin{cases} N_S, & \text{if } p_i - p_T < -\epsilon_1; \\ 0, & \text{if } -\epsilon_1 \le p_i - p_T \le \epsilon_2; \\ M_S, & \text{if } p_i - p_T > \epsilon_2; \end{cases}$$

$$L(E, p_i) = \begin{cases} 0, & \text{if } p_i - p_T < -\epsilon_1; \\ K_E, & \text{if } -\epsilon_1 \le p_i - p_T \le \epsilon_2; \\ M_E, & \text{if } p_i - p_T > \epsilon_2. \end{cases}$$

The six penalties K_D, K_E, M_S, M_E, N_S, and N_D are positive real numbers. For example, quantities K_D or N_D are the penalties for choosing decision D (de-escalate) when dose i is either within the EI or lower than $(p_T - \epsilon_1)$. The values of M_S, N_S, K_E, and M_E can be interpreted similarly. We assign a zero penalty for choosing the right decision.

We derive a straightforward decision rule for dose finding based on posterior expected penalties. Let $\mathcal{X} = \{(x_1,\ n_1), \ldots, (x_d,\ n_d)\}$ be the accumulated data in which n_i patients have been treated at dose i and x_i of them have experienced toxicities, for $i = 1, \ldots, d$. The information set corresponding to $\mathcal{X}$ is a σ-algebra, $\mathcal{F} = \sigma(\mathcal{X})$. The likelihood is the product of the d binomial probability mass functions defined by $(x_1,\ n_1), \ldots, (x_d,\ n_d)$. Suppose that the prior distribution for the vector $\boldsymbol{p} = (p_1, \ldots, p_d)'$ has a density $\pi(\boldsymbol{p})$. Define

$$\begin{aligned} R(\mathrm{D}, p_i) &= E\{L(\mathrm{D}, p_i)|\mathcal{F}\}, \\ R(S, p_i) &= E\{L(S, p_i)|\mathcal{F}\}, \quad \text{and} \\ R(E, p_i) &= E\{L(E, p_i)|\mathcal{F}\} \end{aligned}$$

as the three posterior expected penalties corresponding to $\pi(\boldsymbol{p})$. Let

$$\begin{aligned} q_{Di} &= \Pr(p_i - p_T > \epsilon_2|\mathcal{F}), \\ q_{Si} &= \Pr(-\epsilon_1 \le p_i - p_T \le \epsilon_2|\mathcal{F}), \quad \text{and} \\ q_{Ei} &= \Pr(p_i - p_T < -\epsilon_1|\mathcal{F}); \end{aligned}$$

then

$$\begin{aligned} R(\mathrm{D}, p_i) &= K_D q_{Si} + N_D q_{Ei}; \\ R(S, p_i) &= M_S q_{Di} + N_S q_{Ei}; \quad \text{and} \\ R(E, p_i) &= K_E q_{Si} + M_E q_{Di}. \end{aligned} \tag{1}$$

The Bayes rule that achieves the minimum posterior expected penalty is given by

$$\mathcal{B}_i = \arg \min_{m \in \{\mathrm{D}, S, E\}} R(m, p_i). \tag{2}$$

The performance of the Bayes rule depends on the prior distribution of the p_i and the six penalties. Our approach is to use a simple prior and then specify a sensible set of penalties for that prior. To start, we assume an independent uniform prior for p_i. Note that this is a special case of the proposed beta priors in [9], which contains a detailed

discussion of the priors for phase I designs. Under the uniform prior, we set the six penalties at

$$K_D = K_E = \frac{1}{\epsilon_2 + \epsilon_1}, \quad M_S = M_E = \frac{1}{1 - p_T - \epsilon_2},$$
$$N_D = N_S = \frac{1}{p_T - \epsilon_1}, \tag{3}$$

which possess the following property.

Proposition 1.
Under the uniform prior for p_i and the penalties in (3), the prior expected penalties for D, E, and S are the same.

The proof is straightforward and omitted. Proposition 1 implies that the uniform prior and the set of penalties in (3) are 'unbiased' *a priori* in that one does not prefer any of the three actions over the others before the trial starts. We note that similar results can be obtained with an arbitrary beta prior although the performance of such a setup is part of our ongoing research [12].

Linking to the UPM: It is immediate that given the penalties in (3), the three posterior expected penalties $R(D, p_i)$, $R(S, p_i)$, and $R(E, p_i)$ equal (1− the UPMs) for the intervals $(0, p_T-\epsilon_1)$, $[p_T-\epsilon_1, p_T+\epsilon_2]$, and $(p_T+\epsilon_2, 1)$. Therefore, the Bayes rule $\mathcal{B}_i$ chooses E, S, or D if the interval $(0, p_T-\epsilon_1)$, $[p_T-\epsilon_1, p_T+\epsilon_2]$, or $(p_T+\epsilon_2, 1)$ has the largest UPM. *In words, the mTPI design proposed in Section 2 is equivalent to the Bayes rule $\mathcal{B}_i$ under the decision-theoretic framework above with the penalties given in (3).*

Large-sample properties

We briefly report some large-sample properties of the mTPI design. Although they do not imply that the design would perform well in the case of small sample sizes (such as in dose-finding trials), they ensure that the mTPI design is theoretically sound. In the next section, we will evaluate small-sample properties of the mTPI method based on computer simulations. For ease of exposition, we place the theoretical derivations in the Appendix and only discuss their practical implication below.

The first large-sample property is that the mTPI method will choose the correct dose-finding action from among D, S, and E in large samples. That is, if a dose is too toxic, the mTPI method will always de-escalate from this dose given enough information. Similarly, if a dose is below the MTD, then the mTPI method will always escalate. These results are summarized as Proposition 2 in the Appendix. Building upon these results, our second large-sample property states that when enough patients are treated in a practical trial, the mTPI design will always choose a dose in the equivalence interval to treat all future patients, given that such a dose is one of the target candidates in the trial. The property is summarized as Theorem 1 in the Appendix.

Simulation study

We conducted extensive simulation studies and sensitivity analyses with comparisons to established methods. We present the simulation results in subsections, with each focusing on one aspect of the mTPI design. In addition, we performed simulation studies for a second trial with different setups. The results are summarized in the Supplementary Material.

Overall performance

Based on a clinical trial design described in Goodman [13], we examined the overall performance of the mTPI design. The trial had eight doses with a maximum sample size of 30 patients. The target toxicity $p_T = 0.25$. The starting dose was the lowest dose and the cohort size was three. We simulated 1000 trials on computer. For each trial, unless the safety rule 1 in our proposed algorithm is invoked, that is, dose level 1 is deemed too toxic, patients will be enrolled and assigned to appropriate doses based on each design until the maximum sample size is reached. We recorded the patient assignments, toxicity responses, and final doses selected as the MTD for all the trials. The results are tabulated in Table 1. We compared the proposed mTPI to the TPI method, the CRM [5], and the 3+3 [7]. We used the CRM model in [5] that defines the CRM skeleton $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_d)$. The skeleton is a set of prespecified and fixed toxicity probabilities with the constraint $\phi_1 < \phi_2 < \ldots < \phi_d$. According to [5], the CRM models the probability of toxicity at the i-th dose as $p_i = \phi_i^{\exp(-\beta)}$; we assume that β follows a normal prior distribution with mean 0 and standard deviation 2. The final MTD is the dose with minimal $|\hat{p}_i - p_T|$ where $\hat{p}_i$ is the posterior mean of p_i. Note that there are many versions of the CRM (e.g., proposed in [4,14–16]) and that some of these models may perform better than the one-parameter model used here.

We used $\epsilon_1 = \epsilon_2 = 0.05$, which was arbitrary. A sensitivity analysis in the next section will

Table 1 Simulation results comparing the proposed mTPI method, the TPI in [9], the CRM in [1], and the 3+3 design in [7]. The true probabilities of toxicity are presented as percentages for each scenario (first row of each scenario). The selection percentages for the true MTDs are in bold face

		Recommendation percentage at dose level $p_T = 0.25$								Toxicity percentage*	Average number of patients
	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		
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		
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			
3+3	% MTD	0	0	1	2	25	**56**	11	0	0	13	24
	# pt	3.1	3.2	3.2	3.3	3.9	4.8	2.3	0.3			
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			
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			
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			
3+3	% 38	9	1	0	0	0	0	0	0	**52**	43	6
	# pt	4.7	0.5	0.6	0.7	0	0	0	0			
Scenario 5		15	25	35	45	55	65	75	85	None		
mTPI	% MTD	29	**45**	20	4	0	0	0	0	0	24 (20, 27)	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			
3+3	% MTD	29	**37**	20	7	1	0	0	0	0	26	12
	# pt	4.4	3.9	2.4	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			
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			

*Overall % of patients with DLTs over all the simulated trials. **In all scenarios other than Scenario 4, the mTPI, TPI, and CRM will stop the trial when the maximum number of patients (30) is reached. For Scenario 4, the trial will stop early since all the doses are too toxic.

demonstrate the robustness of the method to the choices of ϵ's. For the CRM we specified the prior toxicity probability for dose i to be $\phi_i = 0.05 * i$. For the TPI method in [9], we used the recommended parameter values $K_1 = 1$ and $K_2 = 1.5$, which were calibrated by the authors of the article.

Table 1 shows that overall the mTPI method exhibits comparable operating characteristics to those of the other methods. We observe that the 3+3 does not perform as well as other methods in most scenarios. In particular, it usually yields smaller percentages of selecting the correct MTD. It appears that the 3+3 is too conservative in that it is unable to escalate quickly even when the doses are safe. The CRM, while generally performing well for most scenarios, has worse performance for Scenario 3, in which there is a large gap between the true toxicity probabilities of adjacent doses. This is due to the mismatch between the underlying dose–response model for the CRM and the true toxicity probabilities specified in the scenario. Third, we summarize the numbers of patients treated at the MTD and at doses above the MTD below, for the mTPI, TPI, and CRM methods across the six scenarios. The '*na*' in the vectors V's corresponds to Scenario 4, in which all the doses are above the MTD.

# patients at the MTD	*# patients at doses above the MTD*
(18.3, 8.1, 15.9, *na*, 10.9, 9.3)	(4.6, 2.4, 10.9, 19.0, 6.2, 5.5)
(16.1, 8.0, 13.2, *na*, 9.5, 9.2)	(6.3, 3.6, 11.2, 19.0, 7.7, 7.6)
(18.6, 7.0, 13.0, *na*, 11.4, 8.9)	(5.9, 4.7, 13.8, 23.0, 4.7, 4.2)

For Scenarios 1–4, the mTPI method treats, on average, fewer patients at doses above the MTD than the other two methods while maintaining about the same or higher numbers of patients at the MTD. For Scenarios 5 and 6, the CRM is better partly because the set of prior probabilities we used was close to the true toxicity probabilities. Last, in all but one scenario, the toxicity percentage of the mTPI is the lowest among the three model-based methods (mTPI, TPI, CRM), indicating that it is the safest model-based design. The 3+3 design almost always yields the lowest overall toxicity percentage due to its conservative nature. Considering that the mTPI method is the simplest model-based method operationally, these results are particularly encouraging.

Sensitivity analyses

We conducted additional sensitivity analyses for the mTPI method. First, we varied the values of the ϵ's and reran the computer simulations for all five scenarios. For simplicity, we arbitrarily choose Scenario 1 and present the simulation results in Table 2. The results demonstrate the robustness of the method to different values of ϵ's at two extremes, one with large ϵ's and the other with small values. This is not surprising because, by definition, the method compares per-unit probability mass for an interval and is therefore robust to how wide the interval is.

Second, after fixing $\epsilon_1 = \epsilon_2 = 0.05$, we tried different beta prior distributions for the mTPI method (Table 3). Since the penalties of the mTPI method are calibrated for the uniform prior (or beta(1,1)), we obtained the best performance under this prior. The performance of the mTPI method under other priors is comparable, although worse than the uniform prior. We note that when strong prior information is available, one can use a more informative prior distribution for the

Table 2 Simulation results of the mTPI using different ϵ_1 and ϵ_2 values that define the equivalent interval $[p_T - \epsilon_1, p_T + \epsilon_2]$ (Section 'Overall performance'). The selection percentages for the true MTDs are in bold face

		Recommendation percentage at dose level $p_T = 0.25$								Toxicity percentage*		Average number of patients
	Dose	1	2	3	4	5	6	7	8			
Scenario 1		5	25	50	60	70	80	90	95	none		
$\epsilon_1 = \epsilon_2 = 0.05$	% 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			
$\epsilon_1 = \epsilon_2 = 0.2$	% 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			
$\epsilon_1 = \epsilon_2 = 0.001$	% 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			

*Overall % of patients with DLTs over all the simulated trials.

Table 3 Simulation results comparing the results of the mTPI method with $\epsilon_1 = \epsilon_2 = 0.05$ using different beta prior distributions Beta(a, b) (Section 'Overall performance'). The selection percentages for the true MTDs are in bold face

		Recommendation percentage at dose level $p_T = 0.25$									Toxicity percentage*	Average number of patients
	Dose	1	2	3	4	5	6	7	8			
Scenario 1		5	25	50	60	70	80	90	95	none		
$a = b = 1$ (default)	% 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			
$a = b = .05$	% MTD	15	**76**	8	1	0	0	0	0	0	27	30
	# Pts	6.3	17.1	6.0	0.5	0	0	0	0			
$a = 1$, $b = 3$	% MTD	9	**77**	12	1	0	0	0	0	0	29	30
	# Pts	4.9	16.9	7.5	0.6	0	0	0	0			
$a = .1$, $b = .3$	% MTD	16	**73**	10	0	0	0	0	0	0	31	30
	# Pts	5.1	14.2	7.4	2.4	0.7	0	0	0			

*Overall % of patients with DLTs over all the simulated trials.

mTPI method. However, usually little prior information is known about the toxicity of the treatments in phase I trials, as they typically involve novel therapies.

Lastly, in [12] we investigated an alternative scheme of specifying and calibrating different beta priors by the same approach as stated in Proposition 1. That is, we proposed alternative beta priors and modified our penalty functions so that the prior expected penalties for D, E, and S are the same. We examined the performance of these priors under a variety of simulated scenarios as part of ongoing research. Initial results (not shown) suggest that the performance of the beta(1,1) prior in this article are comparable to that of the other beta priors.

Software

The mTPI method is available in both Excel and R programs. In the Excel program, the method is presented in a macro with an add-in file. The Excel program contains a dose-finding table that consists of all the possible dose-finding actions for a given trial. Figure 2 presents a screenshot of the table. To use the macro, one needs to provide the sample size, the EI, and the toxicity probability p_T of the MTD. Then by clicking a button we embedded in Excel, a table in the form shown in Figure 2 will be generated. Using this table, one can carry out all the dose assignments for a trial without needing to conduct additional computations. For example, suppose patients are being treated at dose i with x_i dose-limiting toxicities observed out of n_i patients. In the Excel table, locate the row and column that correspond to x_i and n_i, respectively. The appropriate decision is given by the letter in the corresponding cell of the Excel table. Thus, if $x_i = 1$ out of $n_i = 3$ patients has experienced DLT, then the decision is 'S' to stay at the current dose. Note that these decisions do not depend on dose level i. That is why we only need to provide one table to carry out dose-finding decisions at various doses. In the second page of the Excel macro (not shown), we embedded other buttons to conduct computer simulations similar to those shown in Table 1. Results for, say, 5000 simulations are usually obtained in a few seconds.

We also provide R functions with the same capabilities. In particular, the R functions allow users to fully specify the trial parameters, including the sample size, the EI, the cohort size, the starting dose, the true toxicity probabilities for different scenarios, and the prior distributions. Both Excel and R programs are available to download at http://odin.mdacc.tmc.edu/~ylji/.

Discussion

The mTPI method is an improved version of the TPI. While taking advantage of some desirable features of the TPI (such as simplicity and user-friendly software), the mTPI method extends the TPI in two aspects.

- First, the TPI method requires users to calibrate two parameters K_1 and K_2 (see 'Introduction' section for their definitions) that affect the performance of the method. In contrast, based on our experience the mTPI method does not need to be calibrated for different trials. Of course, one may change the settings of the mTPI method if it is necessary to do so. For example, it is certainly feasible to modify the prior distributions if historical data are available on the toxicity of the treatment, although this is not particularly common when designing a phase I trial.

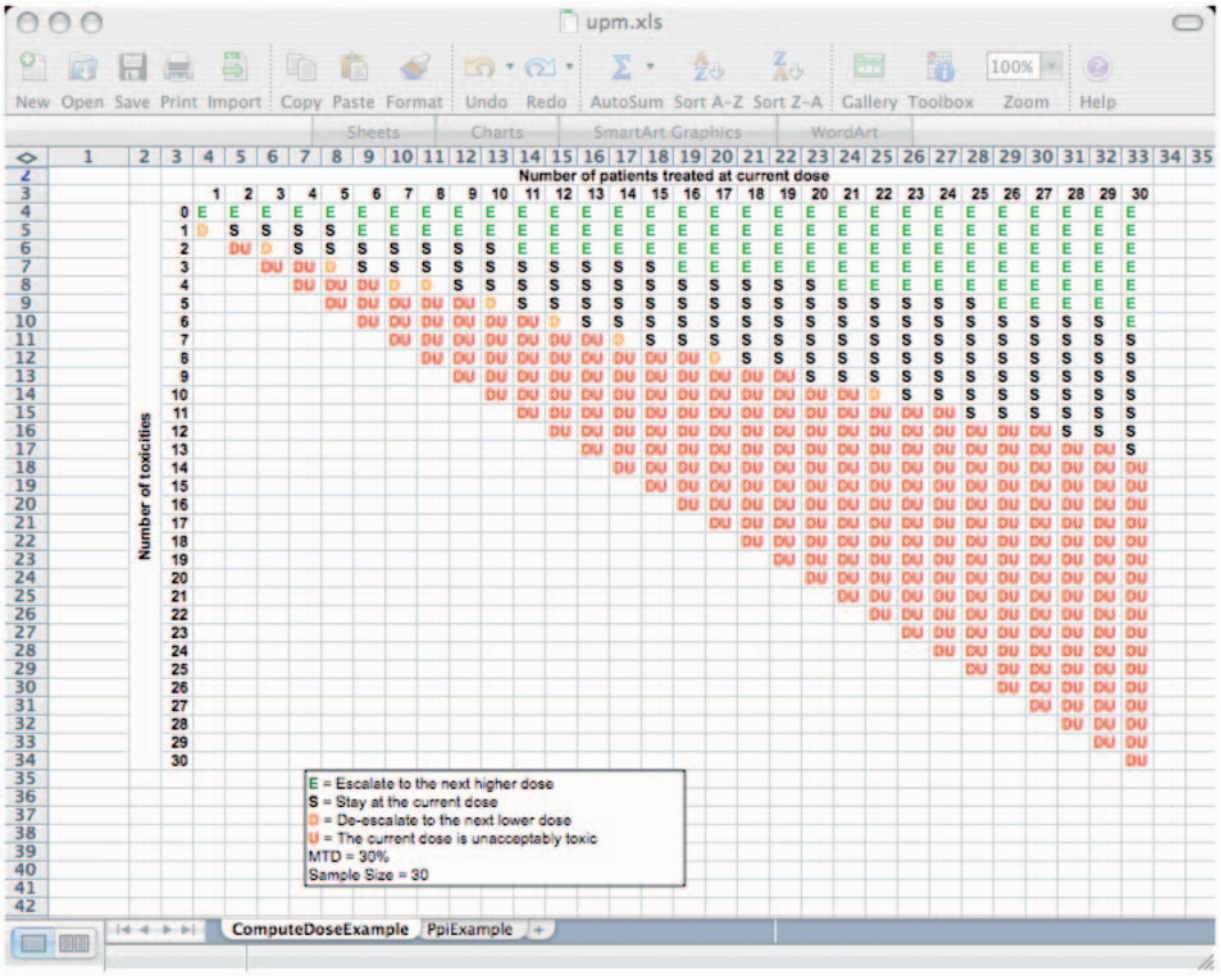

Figure 2 A screenshot of the Excel macro for the mTPI method. The table is uniquely determined upon specification of the sample size, the EI, and p_T (see 'Introduction' section). The letters in different colors are computed based on the decision rules under the mTPI method and represent different dose-finding actions. In addition to actions *D*, *S*, and *E*, the table includes action *U*, which is defined as the execution of the *dose exclusion rule* in the proposed dose-finding algorithm (Section 'Dose-finding algorithm')

- Second and more importantly, the key statistics used for posterior inference are different between the two methods. For the TPI method, the decisions are based on the posterior probabilities of the three intervals defined by K_1, K_2, and the posterior standard deviations of the toxicity probabilities. For the mTPI method, the equivalence intervals are prespecified before the trial and do not depend on any parameters of the probability model. In addition, the decisions of the mTPI method are based on the evaluation of the unit probability masses. We show that they correspond to the Bayes rule under a decision-theoretic type of framework.

We believe that the CRM is an excellent method with many advantages over 3+3 like designs. Furthermore, our simulations show that, in certain situations, the CRM can outperform our method. However, we want to re-emphasize the importance of simplicity for a model-based dose-finding design. We believe that the simplicity of the method for early phase trials is the dominating factor that decides whether the method will be embraced by physicians in practice. This not only involves the availability of software, but also the amount of effort required to monitor a trial. For example, the CRM has been implemented by many researchers with available software for conducting simulations. However, physicians still need to work closely with statisticians and computer scientists in order to make dose assignment decisions whenever a patient needs to be treated. Perhaps the most attractive feature of the TPI and the mTPI methods is the availability of the Excel spreadsheet that frees

physicians from additional burdens on making dose assignment decisions.

An arguable point is our choice of the independent prior models for the p_i's. From a statistical inference point of view, the assumption of independence is clearly not optimal as the toxicity probabilities of the doses are assumed to be ordered and hence dependent. So theoretically it is desirable to introduce dependent models for p_i's. However, we believe that for phase I trials with small sample sizes, especially when only a few (up to 6, for example) patients are treated at a given dose, the dependence introduced by prior models will have a strong influence on the operating characteristics of the dose-finding design. Therefore, when the dependence introduced in the prior models does not agree with the true toxicity probabilities, bias may be introduced in a small-sample situation, which subsequently may lead to poor performance under the design. More seriously, there is no way of knowing the level of agreement in practice since for practical trials one does not know the true toxicity probabilities. Therefore, we believe that introducing dependent models for a phase I trial could be tricky and we resorted to a simple approach as in [9]. The independent prior models performs quite well compared to existing approaches. This point has been carefully debated in [9].

A special case in practice involves trials with a cohort size of one. We generally do not recommend making a dose-finding decision on any dose when fewer than two patients are treated. However, our software did not build this rule into the computer code. We do not expect many practical trials with a cohort size of one because it takes a long time to complete such trials and they do not have much of an advantage over using cohort sizes larger than one. If needed, one can simply modify our code and add a rule to only invoke dose finding after at least two patients have been treated at a given dose.

A closely related but separate problem is the selection of the MTD at the end of the trial. In this article, we adopted the same approach as in [9], where the MTD is selected based on the distance $|\hat{p}_i - p_T|$ between the isotonic-transformed posterior mean $\hat{p}_i$ and the target p_T (see the proposed dose-finding algorithm in the section 'Dose-finding method'). In addition, one can easily obtain interval estimates for $\hat{p}_i$ by sampling independently from the posterior beta distribution for each p_i and performing isotonic transformation on the posterior samples. Using the transformed samples, one can easily obtain posterior intervals for $\hat{p}_i$ numerically. Note that the selection of the MTD is statistically different from the design of dose-finding trials in that the former is related to statistical inference based on observed data and the latter is related to design of experiments before data are collected. Even though in practice these are two essential components of a dose-finding trial, they require different statistical thinking and remedies. This has also been noted by other authors (e.g., Stylianou and Flournoy [16]). Further discussion related to the selection of the MTD goes beyond the scope of this article and will be considered in our future work.

Another task for future work is to impose early stopping in a trial when sufficient evidence indicates that one dose corresponds to the MTD and that all future patients will be assigned to that dose. This can be realized based on, for example, posterior predictive inference. However, early stopping is not always required as some institutional review boards and protocol review committees may expect to see the trial all the way to the end, given that no safety or ethical concerns are raised.

References

1. **O'Quigley J, Pepe M, Fisher L.** Continual reassessment method: a practical design for phase i clinical trials in cancer. *Biometrics* 1990; **46**: 33–48.
2. **Babb J, Rogatko A, Zacks S.** Cancer phase I clinical trials: efficient dose escalation with overdose control. *Stat Med* 1998; **17**: 1103–20.
3. **Cheung YK, Chappell R.** Sequential designs for phase i clinical trials with lateonset toxicities. *Biometrics* 2000; **56**: 1177–82.
4. **Piantadosi S, Fisher JD, Grossman S.** Practical implementation of a modified continual reassessment method for dose-finding trials. *Cancer Chemother Pharmacol* 1998; **41**: 429–36.
5. **Shen LZ, O'Quigley J.** Consistency of continual reassessment method under model misspecification. *Biometrika* 1996; **83**: 395–405.
6. **Yuan Z, Chappell R, Bailey H.** The continual reassessment method for multiple toxicity grades: A Bayesian quasi-likelihood approach. *Biometrics* 2007; **63**: 173–79.
7. **Thall PF, Lee S-J.** Practical model-based dose-finding in phase *I* clinical trials: Methods based on toxicity. *Int J Gynecol Cancer* 2003; **13**: 251–61.
8. **Lee SM, Cheung YK.** Model calibration in the continual reassessment method. *Clin Trials* 2009; **6**: 227–38.
9. **Ji Y, Li Y, Bekele BN.** Dose-finding in phase I clinical trials based on toxicity probability intervals. *Clin Trials* 2007; **4**: 235–44.
10. **Stylianou M, Flournoy N.** Dose finding using the biased coin up-and-down design and isotonic regression. *Biometrics* 2002; **58**: 171–77.
11. **Ji Y, Li Y, Yin G.** Bayesian dose-finding in phase I clinical trials based on a new statistical framework. *Stat Sinica* 2007; **17**: 531–47.
12. **Li Q, Tsui K-W, Ji Y.** The choice of prior distribution in dose-finding of phase I clinical trials. Technical Report. Available at: http://odin.mdacc.tmc.edu/~ylji/prior.pdf/, 2010.
13. **Goodman SN, Zahurak ML, Piantadosi S.** Some practical improvements in the continual reassessment method for phase I studies. *Stat Med* 1995; **14**: 1149–61.
14. **Garrett-Mayer E.** The continual reassessment method for dose-finding studies: a tutorial. *Clin Trials* 2006; **3**: 57–71.

15. **Iasonos A, Wilton AS, Riedel ER,** *et al*. A comprehensive comparison of the continual reassessment method to the standard 3 + 3 dose escalation scheme in phase i dose-finding studies. *Clin Trials* 2008; **5**: 465–77.
16. **Zohar S, O'Quigley J.** Sensitivity of dose-finding studies to observation errors. *Contemp Clin Trials* 2009; **30**: 523–30.
17. **Carlin BP, Louis TA.** *Bayes and Empirical Bayes Methods for Data Analysis* (2nd edn). Chapman & Hall/CRC, Boca Raton, Florida, 2000.
18. **Sen PK, Singer JM.** *Large Sample Methods in Statistics*. Chapman & Hall/CRC, Boca Raton, Florida, 1993.
19. **Gelman A, Carlin JB, Stern HS, Rubin DB.** *Bayesian Data Analysis*. Chapman & Hall/CRC, Boca Raton, Florida, 1995.

Appendix

Proposition 2. For any $\epsilon_1 > 0$ and $\epsilon_2 > 0$,

- if $p_{i0} \in [p_T - \epsilon_1,\ p_T + \epsilon_2]$, then there exists $N > 0$, when $n_i > N$, $\mathcal{B}_i = S$ *a.s*;
- if $p_{i0} < p_T - \epsilon_1$, then there exists $N > 0$, when $n_i > N$, $\mathcal{B}_i = E$ *a.s*;
- if $p_{i0} > p_T + \epsilon_2$, then there exists $N > 0$, when $n_i > N$, $\mathcal{B}_i = D$ *a.s*;

Proof of Proposition 2:
Given the prior density $\pi(\boldsymbol{p})$ and the binomial likelihood, the posterior density of $\boldsymbol{p}$ is given by

$$f(\boldsymbol{p}|\text{data}) \propto \prod_{i=1}^{d} p_i^{x_i}(1-p_i)^{n_i - x_i}\pi(\boldsymbol{p}). \qquad (4)$$

Define the ''generalized'' observed Fisher information matrix, denoted by I^{π}, as follows:

$$I^{\pi}_{ij} = -\frac{\partial^2}{\partial p_i \partial p_j}\log f(\boldsymbol{p}|\text{data})|_{\boldsymbol{p}=\tilde{\boldsymbol{p}}}, \quad i, j = 1, \ldots, d,$$

where $\tilde{\boldsymbol{p}}$ is the posterior mode of $\boldsymbol{p}$. By the Bayesian central limit theorem (Carlin and Louis [17], p. 122), when n_i is large for $i = 1, \ldots, d$,

$$\boldsymbol{p}|\text{data} \dot{\sim} MVN_d(\tilde{\boldsymbol{p}}, \{I^{\pi}\}^{-1}),$$

where $MV\,N_d$ denotes a d-dimensional multivariate normal distribution. By the Cramér-Wold Device (Sen and Singer, [18], P. 106),

$$p_i|\text{data} \dot{\sim} N(\tilde{p}_i, \sigma_i^2), \qquad (5)$$

where $\tilde{p}_i$ is the i-th component of $\tilde{\boldsymbol{p}}$ and σ_i^2 is the (i, i)th element of $\{I^{\pi}\}^{-1}$. Under suitable regularity conditions, the posterior mode $\tilde{\boldsymbol{p}}$ is consistent (Gelman *et al*. [19] p. 106). Because the first and second partial derivatives of $\pi(\boldsymbol{p})$ are bounded in the neighborhood of $\boldsymbol{p}_0$, when $\tilde{\boldsymbol{p}}$ is close to $\boldsymbol{p}_0$, it follows that for $i \neq j$

$$I^{\pi}_{ij} = -\frac{\partial^2}{\partial p_i \partial p_j}\log \pi(\boldsymbol{p})|_{\boldsymbol{p}=\tilde{\boldsymbol{p}}}$$

are bounded, and

$$I^{\pi}_{ii} = \frac{(n_i - x_i)\tilde{p}_i^2 + (\tilde{p}_i - 1)^2 x_i^2}{\tilde{p}_i^2(1-\tilde{p}_i)^2} - \frac{\partial^2}{\partial p_i^2}\log \pi(\boldsymbol{p})|_{\boldsymbol{p}=\tilde{\boldsymbol{p}}}$$

goes to ∞ as n_i goes to ∞. Let λ_j be the j-th eigenvalue of I^{π} with associated eigenvector $\mathbf{x_j}$, the fact that $\lambda_j = \mathbf{x_j'}\mathbf{I}^{\pi}\mathbf{x_j}/\mathbf{x_j'}\mathbf{x_j}$ implies that $\lambda_j \to \infty$, for $j = 1, \ldots, d$. Hence $\sigma_j^2 \to 0$, $j = 1, \ldots, d$. Combined with the consistency of $\tilde{\boldsymbol{p}}$, the result in (5), and that $p_{i0} = p_T$, as $n_i \to \infty$,

$$P(p_T - \epsilon_1 \leq p_i \leq p_T + \epsilon_2|\text{data}) \to 1, a.s.,$$

for any $\epsilon_1 > 0$ and $\epsilon_2 > 0$.

Theorem 1.
Let p_{i0} represent the true toxicity probability for dose i. Among all the doses specified in the trial, if there exists a unique dose i such that $p_{i0} \in [p_T - \epsilon_1, p_T + \epsilon_2]$, then there exists $N > 0$ when the number of patients treated in the trial is larger than N and all the future patients will be treated at dose i.

Proof of Theorem 1 is immediate based on Proposition 2. Theorem 1 ensures the consistency of the mTPI design in that when a dose is in the equivalence interval, at some point of the trial, all future patients will be assigned to that dose.