A modified toxicity probability interval method for dose-finding trials

Yuan Ji[1], Ping Liu[2], Yisheng Li[2], and B. Nebiyou Bekele[2]

1 Department of Bioinformatics and Computational Biology

2 Department of Biostatistics

University of Texas M. D. Anderson Cancer Center, Houston, TX 77030

Email: *yuanji@mdanderson.org*

July 1, 2009

Abstract. Building on the strength of an early work, the toxicity posterior interval (TPI) method (Ji, Li, and Bekele, 2007a), we present a modified TPI design which is calibration-free for phase I trials. Like the TPI method, the mTPI consists of a practical dose-finding scheme guided by simple Bayesian models and posterior inference. However, the new method proposes improve dose-finding decision rules based on a new statistics, 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 and the length of the interval. The improvement through the use of the UPM for dose finding is three folds: 1) the mTPI method is 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 known difficult issue; and 3) the mTPI method corresponds to the Bayes rule under a decision theoretic framework and possesses desirable large- and small- sample properties. We provide user-friendly software in both Excel and R to accompany the new method.

Keywords: Bayes rule; Maximum tolerated dose; Oncology; Toxicity; Unit probability mass.

1 Introduction

Dose-finding trials in oncology aim to find the maximum tolerated dose (MTD). Usually a grid of dose levels is predetermined before the trial starts. Patients are enrolled in a sequential manner and treated at a given dose level. Based on the observed dose-limiting toxicity (DLT) responses from enrolled patients, a decision is made to determine the dose level for treating future patients. The trial proceeds until the maximum sample size is reached, at which time the MTD is estimated using the DLT responses from all the enrolled patients.

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, Pepe, and Fisher (1990) and its many variations and extensions (Babb, Rogatko, and Zacks, 1998; Cheung and Chappell, 2000; Piantadosi, Fisher, and Grossman, 1998; Shen and O'Quigley, 1996; Yuan, Chappell, and Bailey, 2007; among others). The basic idea of the CRM is to sequentially update the estimates of 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 (Thall and Lee, 2003), the simplicity of the 3+3 method still makes it by far the most popular method chosen for practical trials. There are two main reasons for this. First, using model-based methods to monitor a practical trial usually requires an electronic and automatic procedure that allows nurses and physicians to obtain dose assignments in real time when a patient needs to be treated. This is an elegant and tedious matter that calls for a close collaboration between physicians, statisticians, and computer scientists. Second, for most model-based designs, a common requirement is to conduct computer simulations

before the trial starts, in which statisticians must calibrate the design parameters in order to achieve desirable operating characteristics for the trial. For example, the classic CRM requires the users to provide a set of prior estimates of the dose toxicity probabilities for the candidate doses to be used in the trial. This set should be elicited from collaborating physicians. 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 and tedious process. Recent attempts (Shing and Cheung, 2009) have been made to alleviate this issue although nontrivial calibration is still required.

Ji, Li, and Bekele (2007a) 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). Due to the simplicity of their models, their method is implemented in an Excel spreadsheet, which greatly 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 define the toxicity intervals that are crucial elements of their method. A default set of values was proposed in their paper, 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 for certain cases.

In this paper, 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 real time. In

addition, similar to the TPI design in Ji et al. (2007a), 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 a statistic named 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. 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. The definition of the EI is solely obtained through consultation with the physicians. We will show that the performance of the UPM method is robust to the definition of 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.

On the surface, one may see similarities between the TPI and mTPI methods. For example, both methods propose decision rules involving the evaluation of posterior probabilities of three intervals. Also, both methods have been implemented in Excel and R with similar presentations. Despite these similarities, we argue that the two methods differ significantly in two major aspects.

- First, the TPI method requires users to calibrate two parameters K_1 and K_2 that affect the performance of the method. In practice, the authors of TPI have received inquiries on how to calibrate the parameters for a real trial when the default values do not work. The authors of TPI have yet provided a satisfactory solution on how to

properly and efficiently calibrate these parameters for a given trial. In contrast, the mTPI method is essentially calibration free, which means that one does not have to calibrate the designs for different trials. Of course, this does not mean that one may not 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.

- 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 onset of 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, which are not probability measures. Instead, we show that they correspond to the Bayes rule under a decision-theoretic framework for a given set of losses.

In Section 2, we introduce the basic idea of the mTPI method as well as the dose-finding algorithm. In Sections 3 and 4, we present the probability models, decision rules, and asymptotic properties of the UMP method. We comprehensively evaluate the performance of the new method under small sample sizes and conduct sensitivity analyses in Section 5. We describe related software for Excel and R in Section 6 and provide final remarks in Section 7.

2 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 pseudo 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.

In this section, we first introduce the EI and then describe a general dose-finding algorithm without going into technical details such as the probability model and decision rules. In Section 3, we develop technical details including the probability models and the decision rules.

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 ith dose, $i = 1, \ldots, d$. The toxicity probability usually increases with the dose level, so we assume $p_1 < p_2 < \cdots < p_d$.

Suppose currently that dose i is used for treating patients and n_i ($n_i \geq 1$) patients have been treated. 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 previous 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.

2.1 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 is elicited from collaborating physicians for the trial. 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 $[0.2, 0.4]$. We argue that, 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 real dose with an unknown toxicity probability equal to p_T is zero. Hence, an implicit measure usually is defined for the distance in toxicity probability between a given dose and the target p_T. For example, in the CRM, such a measure is in the form of an L_1 norm, i.e., $|p_i - p_T|$. Here, we explicitly ask the physicians to express their intrinsic measure of such a distance in the form of EIs, which seems to be more intuitive. 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 to continue treating 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 to continue treating future patients without dose de-escaltion. This determines $p_T + \epsilon_2$.

2.2 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 next. **The core of the algorithm** is simple: Suppose dose i is

the dose currently used to treat patients. Compute the UPM for the three subintervals above. Treat the next (cohort of) patients at dose $(i-1)$, i, or $(i+1)$ if interval $(0, p_T - \epsilon_1)$, $[p_T - \epsilon_1, p_T + \epsilon_2]$, or $(p_T + \epsilon_2, 1)$ has the largest UMP, respectively. Continue until the sample size limit is reached. Figure 1 provides an illustration of how the UPM is associated with dose-finding decisions for an EI.

A dose-finding algorithm with additional safety rules is given below. These additional rules are important for practical concerns. See a full discussion of these rules in Ji et al. (2007a).

- Suppose that the current tried dose is i, $i \in \{1, \cdots, d\}$. 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.

- *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 instead of $(i+1)$ and exclude doses $(i+1)$ and higher from the trial, i.e., these doses will never be used again in the trial.

- 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, e.g., the isotonically transformed posterior mean (Ji et al.,

2007a). If two or more doses tie for the smallest difference, perform the following rule. Let $\hat{p}^*$ denote the $\hat{p}_i$ of the tied doses.

- If $\hat{p}^* < p_T$, choose the highest dose among the tied doses.
- If $\hat{p}^* > p_T$, choose the lowest dose among the tied doses.

3 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. (2007b). For dose i, define the penalty functions

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

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

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

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

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

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

then

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

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

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

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 Ji et al. (2007a), which contains a detailed discussion of the priors for phase I designs. Under the uniform prior, we set the six penalties at

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

which possess the following property.

Proposition 1. Under the uniform prior of 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.

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 unit probability masses) 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 penalties given in (3).*

4 Asymptotic properties

The following propositions ensure that with a large sample size the mTPI design will always make the right decision. While this does not guarantee any good properties in the case of small sizes (such as in dose-finding trials), it ensures that the mTPI design is theoretical sound.

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 $n_i > N$, $\mathcal{B}_i = D$ $a.s$;

The proof of Proposition 2 is given in Appendix A. This proposition states that when the sample size is large, the mTPI method will choose the correct dose-finding action. For example, when the true toxicity probability of dose i is in the EI, asymptotically the mTPI method will always choose to stay at dose i.

Theorem 1. 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, all the future patients will be treated at dose i.

Proof of Theorem 1 is immediate based on Proposition 2. The results of Theorem 1 are self-explanatory.

5 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 Appendix B and Table 4.

5.1 Overall performance

Based on a previous clinical trial described in Goodman (1995), we examined the overall performance of the mTPI design. The trial had eight doses with a maximum sample size of 30 patients. We compared our method to their method, as well as the 3+3 method and the CRM. The performance of the 3+3 method is much worse than the other methods. For simplicity, we will not present the results for the 3+3 methods.

Table 1 shows that the mTPI method exhibits desirable operating characteristics compared to its major competitors. Keep in mind that the mTPI method is much easier to implement. We used $\epsilon_1 = \epsilon_2 = 0.05$, which was arbitrary. A sensitivity analysis in the next section will demonstrate the robustness of the method to the choices of ϵ's. In working with other methods we had to manually calibrate some design parameters to achieve desirable performance. For example, we had to choose a set of eight prior toxicity probabilities using the CRM. Because we did not know where to start, we tried several arbitrary sets of toxicity probabilities and chose the one that gave a reasonable performance (in our opinion). In particular, for the CRM we specified the prior toxicity probability for dose i to be $q_i = 0.05 * i$. For the TPI method in Ji et al (2007a), we used the recommended parameter values $K_1 = 1$ and $K_2 = 1.5$, which were calibrated by the original authors of the paper.

We now summarize key findings from Table 1. First, the CRM is inferior 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. Second, the mTPI method treats on average fewer patients at doses higher than the MTD than the other two methods, while maintaining about the same or higher numbers of patients at the MTD (Figure 2). The exceptions are Scenarios 5 and 6, in which the CRM is slightly better since the set of prior probabilities we

used is close to the true toxicity probabilities. Third, in all but one scenarios, the toxicity percentage of the mTPI is the lowest among the three methods, indicating that it is the safest design. Considering that the mTPI method is the simplest method operationally, these results are particularly encouraging.

5.2 Sensitivity analyses

We conducted additional sensitivity analysis for the mTPI method. First, we varied the values of the ϵ's and reran the computer simulations for all the 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; however it is reassuring to observe the simulation results seen in the table.

Second, fixing $\epsilon_1 = \epsilon_2 = 0.05$ we tried different beta prior distributions for the mTPI. 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. However, the performance of the mTPI method under other priors is not terrible, although it is worse than the uniform prior. Again, for simplicity, we only presented the results for Scenario 1 in Table 3. We note that when strong prior information is available, one can use a more informative prior distribution for the 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.

6 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 addin file. The Excel program contains a dose-finding table that consists of all the possible dose-finding actions for a given trial. Figure 3 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 3 will be generated. Using this table, one can carry out all the dose assignments throughout a trial without needing to conduct additional computations again. For example, suppose patients are being treated at dose i with x_i DLTs 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 have 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, 5,000 simulations are usually obtained in a few seconds.

We also provide R functions with the same capabilities. Both Excel and R programs are available to download at http://odin.mdacc.tmc.edu/∼yuanj/.

7 Final remarks

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 must be treated. At M.D. Anderson Cancer Center, the biostatistical group usually sets up a web interface allowing physicians to input patients characteristics and trial data. In return, the website outputs the dose assignment for the next patient. This is viewed as a black box, which sometimes make physicians skeptical and uncomfortable. Perhaps the most attractive feature of the TPI and the mTPI is the availability of the Excel spreadsheet that frees physicians from additional burdens needed to make dose assignment decisions.

Under the mTPI, the specification of ϵ_1 and ϵ_2 will not be determined by computer simulations. It is a fully subjective decision made by physicians. We highly recommend the use of this interval in practice because it is almost never the case that a prespecified dose in the trial happens to have a toxicity probability exactly equal to p_T. Usually, a dose close enough is selected and the specification of this interval formally defines "how close is close."

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 there is not 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.

References

Babb, J., Rogatko, A. & Zacks, S. (1998), 'Cancer phase I clinical trials: efficient dose escalation with overdose control', *Statistics in Medicine* **17**, 1103–1120.

Carlin, B. & Louis, T. (2000), *Bayes and empirical Bayes methods for data analysis*, 2nd edn, Chapman & Hall/CRC, Boca Raton, Florida.

Cheung, Y. & Chappell, R. (2000), 'Sequential designs for phase i clinical trials with late-onset toxicities', *Biometrics* **56**, 1177–1182.

Gelman, A., Carlin, J., Stern, H. & Rubin, D. (1995), *Bayesian Data Analysis*, Chapman & Hall/CRC, Boca Raton, Florida.

Goodman, S., Zahurak, M. & Piantadosi, S. (1995), 'Some practical improvements in the continual reassessment method for phase I studies', *Statistics in Medicine* **14**, 1149–1161.

Ji, Y., Li, Y. & Bekele, B. (2007a), 'Dose-finding in phase I clinical trials based on toxicity probability intervals', *Clinical Trials* **4**, 235–244.

Ji, Y., Li, Y. & Yin, G. (2007b), 'Bayesian dose-finding in phase i clinical trials based on a new statistical framework', *Statistica Sinica* **17**, 531–547.

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

Piantadosi, S., Fisher, J. & Grossman, S. (1998), 'Practical implementation of a modified continual reassessment method for dose-finding trials', *Cancer Chemotherapy and Pharmacology* **41**, 429–436.

Sen, P. & Singer, J. (1993), *Large Sample Methods in Statistics*, Chapman & Hall/CRC, Boca Raton, Florida.

Shen, L. & O'Quigley, J. (1996), 'Consistency of continual reassessment method under model misspecification', *Biometrika* **83**, 395–405.

Shing, L. & Cheung, Y. (2009), Model calibration in the continual reassessment method. to appear *Clinical Trials*.

Thall, P. & Lee, S.-J. (2003), 'Practical model-based dose-finding in phase i clinical trials: Methods based on toxicity', *International Journal of Gynecological Cancer* **13**, 251–261 (invited).

Yuan, Z., Chappell, R. & Bailey, H. (2007), 'The continual reassessment method for multiple toxicity grades: A Bayesian quasi-likelihood approach', *Biometrics* **63**, 173–179.

Appendix A

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

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

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 2000, 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 MVN_d denotes a d-dimensional multivariate normal distribution. By the Cramér-Wold Device (Sen and Singer 1993, P. 106),

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

where $\tilde{p}_i$ is the ith 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., 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 $i \neq j$

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

are bounded, and

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

goes to ∞ as n_i goes to ∞. Let λ_j be the jth 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 $\epsilon_2 > 0$.

Appendix B

Additional simulation results are summarized in Table 4 for a trial with five doses. The target rate $p_t = 0.1$.

Table 1: Simulation results comparing the proposed mTPI method, the TPI, and the CRM. 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			
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			

Table 1 (continued)

		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 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			

Table 1 (continued)

		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 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			

$*$ Overall % toxicity out of all the simulated trials.

Table 2: Simulation results of the mTPI using different ϵ_1 and ϵ_2 values. 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	
$\epsilon_1 = \epsilon_2 = .05$	% 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			

Table 3: Simulation results comparing the results of the mTPI method with $\epsilon_1 = \epsilon_2 = 0.05$ using different beta prior distributions $Beta(a, b)$. 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			

Table 4: Simulation results comparing the results of the mTPI with $\epsilon_1 = \epsilon_2 = 0.05$ for a trial with five dose levels and a maximum sample size of 21 patients. The cohort size is three.

	Recommendation percentage at dose level $p_T = 0.10$						Toxicity percentage*	Average number of patients
Dose	1	2	3	4	5			
Scenario 1	1	10	20	40	50	none		
% MTD	20	**49**	28	3	0	0	12	21
# Pts	5.4	8.5	5.2	1.6	0.3			
Scenario 2	1	1	10	30	50	none		
% MTD	4	18	**64**	14	0	0	12	21
# Pts	3.8	4.9	7.5	3.8	0.8			
Scenario 3	10	30	50	60	60	none		
% MTD	**75**	8	0	0	0	17	18	19
# Pts	12.4	5.5	1.0	0.1	0			
Scenario 4	40	50	60	70	80	none		
% MTD	5	0	0	0	0	**95**	42	7
# Pts	12.4	5.5	1.0	0.1	0			

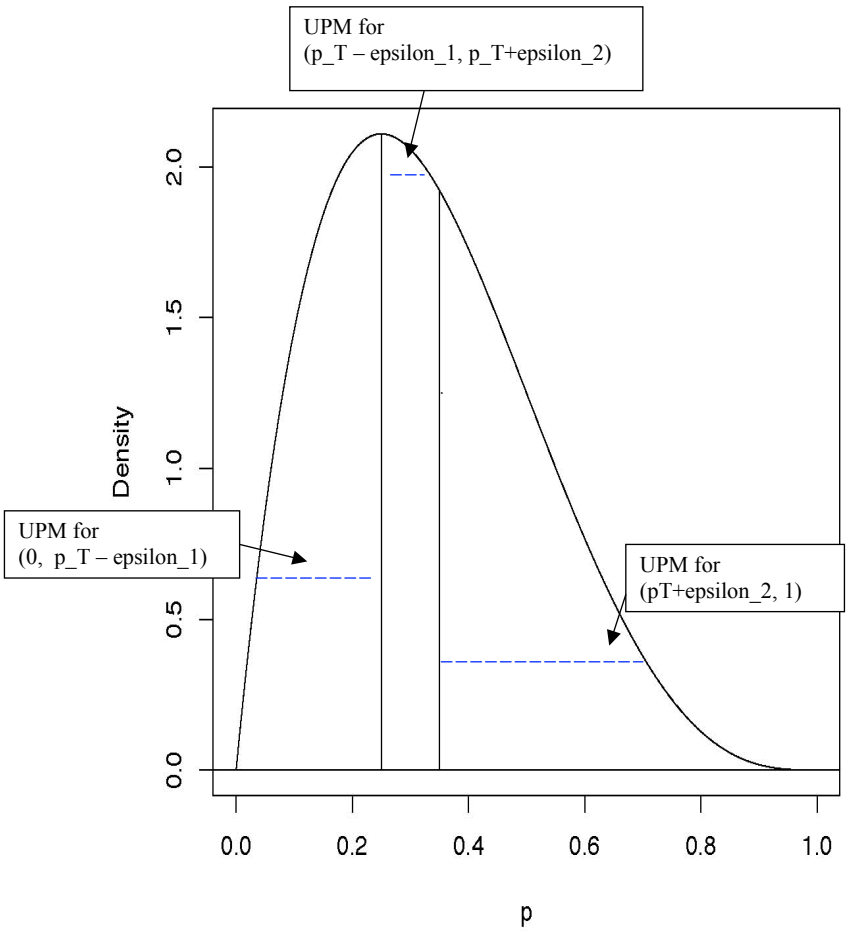

Figure 1. A demonstration of the UPMs for three intervals. The two vertical lines results in three intervals on the X-axis. The UPM for each of the three intervals is indicated by the dashed horizontal line. The equivalence interval in the middle has the highest UPM under the distribution defined by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

1. At the MTD

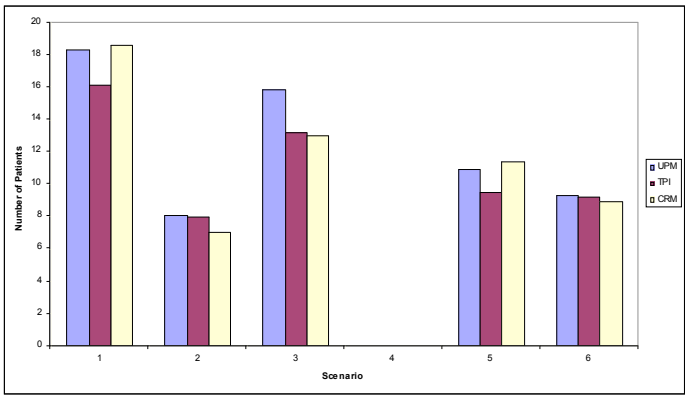

2. All doses higher than the MTD

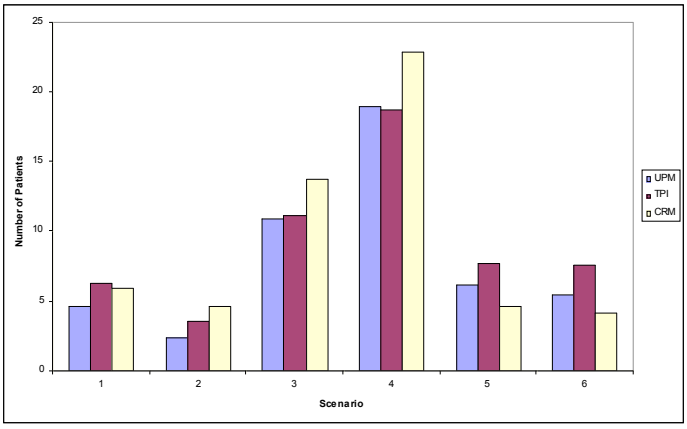

Figure 2 (colored). Summary of the number of patients treated at the MTD (upper panel) and at doses higher than the MTD (lower panel).

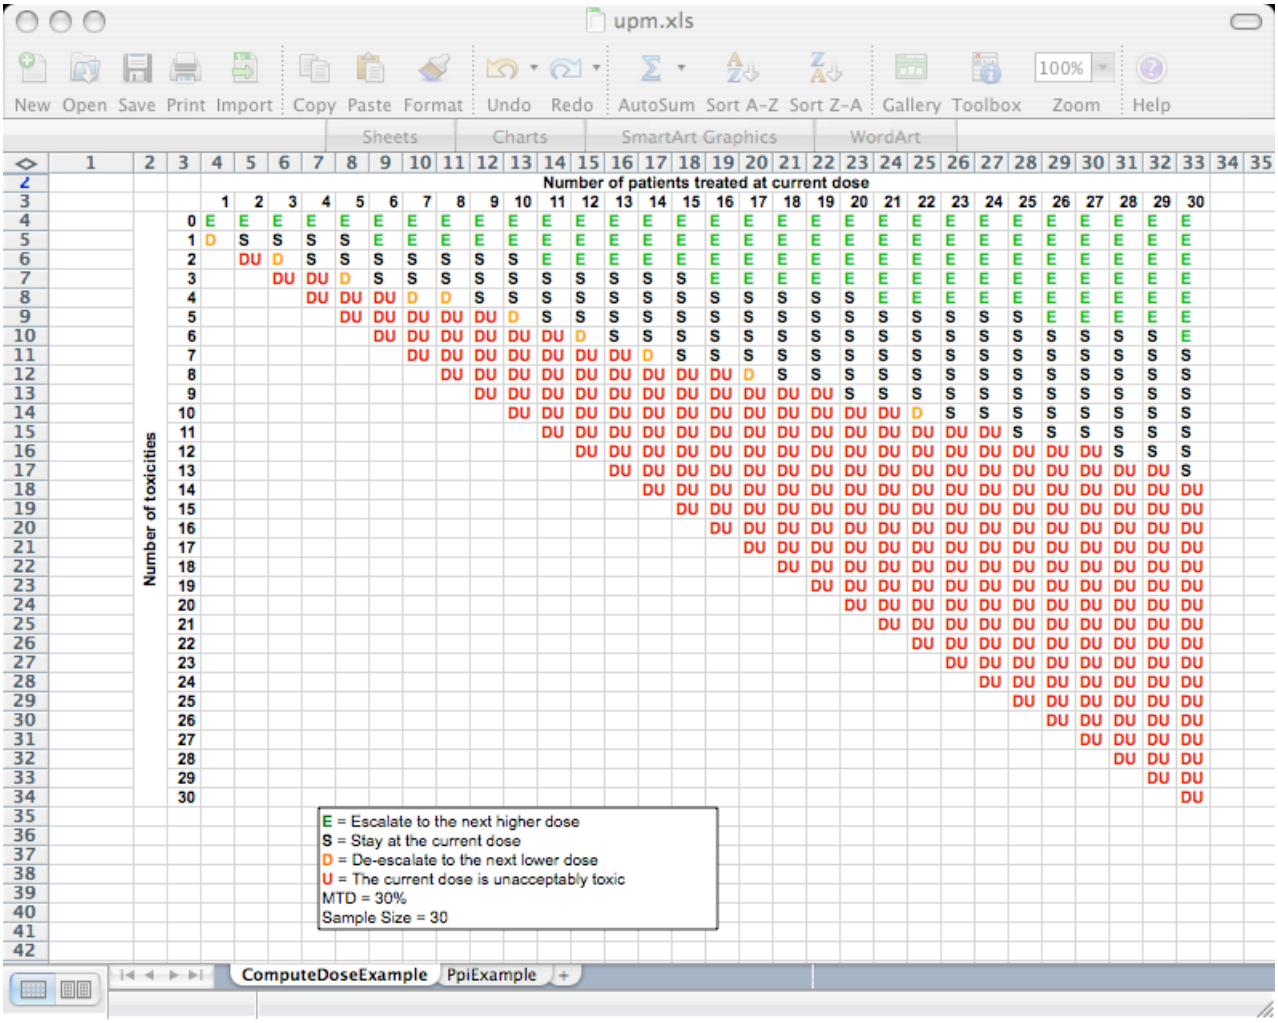

Figure 3 (colored). A screenshot of the Excel macro for the mTPI method. The table is uniquely determined upon specification of the sample size, the EI, and p_T. The letters in different colors are computed based on the decision rules under the mTPI method and represent different dose-finding actions. 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.