Paper AS02

Dose-Finding Schemes in Clinical Trials

Michael Brockmann, Chrestos Concept GmbH & Co.KG, Essen, Germany

ABSTRACT

The higher the dose, the higher the efficacy – and also the higher the risk of undesired side effects. Estimating the maximum tolerable dose (MTD) is often an essential part of early phase clinical trials. A major challenge is to find the MTD by using as small patient groups as possible without sacrificing too much accuracy in estimating the MTD. This can be achieved by selecting the right dose-finding scheme. In this paper, we compare the commonly used and very simple 3+3 design to more sophisticated model-based designs such as the modified toxicity probability interval method or the modified continual reassessment model. For all dose-finding schemes presented, we provide algorithms as well as snippets of SAS code. They can be used to successfully implement the corresponding dose-finding scheme into the study. Finally, we compare the different methods by estimating the MTD based on sample toxicity data.

1 INTRODUCTION

Almost every drug has an increasing dose-response relationship: the higher the dose, the higher the efficacy. However, a higher dose also increases the risk of undesired side effects. In early phase clinical trials, especially in oncology, the safety of the patients is of utmost importance. One goal of such a trial –usually the first-in-human (FiH) study– is to estimate the maximum tolerable dose (MTD) in order to ensure the best possible safety for patients of the FiH trial and for patients of all follow-up trials. Therefore, a dose-escalation design is often used for the FiH study. One result of a dose-escalation study should be the recommended dose (RD) which is used in follow-up trials and which should be the dose less than and closest to the MTD.

Common dose-escalation designs consist of different cohorts. Whether the trial continues with the next cohort depends on the toxicity data of previous cohorts. In many oncological trials, the number of patients with dose-limiting toxicities (DLT) is just counted, without considering the degree of toxicity.

One simple and frequently used dose-finding scheme is the so-called 3+3 design. The FDA guidance "Clinical Considerations for Therapeutic Cancer Vaccines" states [1]: "The traditional standard dose-escalation schedule in the development of cancer therapeutics uses the so-called '3+3 design' to avoid selection of a phase 2 clinical trial dose that causes a treatment-limiting toxicity in more than 17% of subjects, a standard considered acceptable as an outpatient therapeutic for patients with limited options and life-threatening diseases. […] Many cancer vaccine trials have used the '3+3 design', and the results show that, except in very rare situations, an MTD for a cancer vaccine may not be identified. In these trials, the dose-toxicity curve may be so flat that the highest dose that can be administered is limited by manufacturing or anatomic issues rather than toxicity. Therefore, this '3+3 design' may not be the most suitable approach to gathering information from early phase trials of cancer vaccines, and alternative trial designs should be considered." That statement of the FDA was one of our reasons to investigate and compare different dose-finding schemes.

Whereas the 3+3 design is a rule-based design, the so-called *Continual Reassessment Model (CRM)* suggested and introduced by Storer [2] and by O'Quigley et al. [3], is a model-based design. It uses a Bayesian approach where the probability of a DLT is modeled by a particular function that monotonically increases with dose. Well-known examples of such a dose-response model are the one- or two-parameter logistic models or the hyperbolic tangent model [4]. Later, the CRM was improved by various authors [5,6], in order to ensure even greater safety for patients. This is called *modified Continual Reassessment Model (mCRM)*.

Another model-based dose-finding scheme is the *modified Toxicity Probability Interval (mTPI)* method [7,8]. One of the advantages of this method over the mCRM is its simple implementation. All calculations, including some pre-defined safety rules, can be performed before the actual trial starts. The suggestion of the model either to escalate, to stay at the current dose level, or to de-escalate can be simply read from a pre-calculated table. This greatly simplifies communication between the different study team members. One disadvantage, however, is that

the mTPI method is less flexible and should only be used if there are few possible dose levels. Otherwise, the number of patients with whom the MTD is to be found could be too high.

In this paper, we provide basic ideas of above-mentioned dose-finding schemes. We further present some algorithms and snippets of SAS code that might help to implement those dose-finding schemes [9]. In particular, we show how they work by applying them to sample toxicity data that are typical for phase I oncological trials. We subsequently compare the different schemes in terms of the number of patients needed to find a reliable result, i.e. to estimate an accurate RD, which is the maximum dose less than the true MTD. Three different scenarios will be discussed: two extreme cases, where the first dose is toxic or where the largest available dose is safe, and one more realistic scenario. We would like to stress that the aim of this paper is, besides a simple comparison of the three methods for those three special scenarios, to show how the methods can be applied to toxicity data and, in particular, how they can be implemented by means of SAS® 9.4 (see Appendix, Sec. 8). More sophisticated comparisons of the 3+3 design to the other two methods can be found, e.g., in Refs. [4,10] for the mCRM and Refs. [7,8] for the mTPI method.

2 THEORETICAL BACKGROUND

In this section, we provide basic ideas of the dose-finding schemes that are examined in this paper, namely the commonly used *3+3 design*, the *modified Continual Reassessment Model (mCRM)*, and the *modified Toxicity Probability Interval (mTPI)* method. They will be used in Sec. 4 to estimate the maximum tolerable dose (MTD) for three different scenarios.

2.1 THE 3+3 DESIGN

The dose-finding scheme that is most commonly used in early phase oncological trials (>80%) is the so-called *3+3 design* and is based on toxicity data of patients. It is applicable to study designs where each treated patient belongs to a cohort of three patients and where only discrete doses of study drug are administered. The number of dose-limiting toxicities (DLTs) is counted for each cohort and a simple algorithm determines the dose of the next cohort.

The FDA proposes the following algorithm [1]: “In a ‘3+3 design’, three patients are initially enrolled into a given dose cohort. If there is no DLT observed in any of these subjects, the trial proceeds to enroll [three] additional subjects into the next higher dose cohort. If one subject develops a DLT at a specific dose, an additional three subjects are enrolled into that same dose cohort. Development of DLTs in more than 1 of 6 subjects in a specific dose cohort suggests that the MTD has been exceeded, and further dose escalation is not pursued.” This traditional algorithm can be graphically depicted as below:

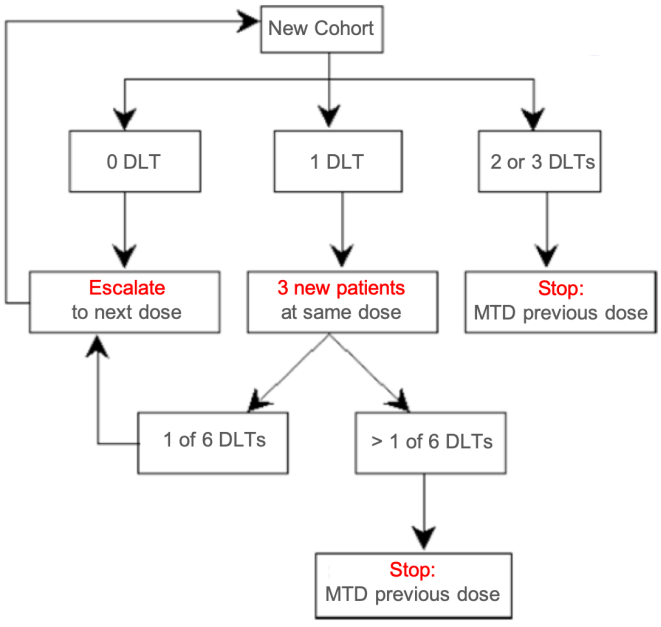

Figure 1: Graphical representation of the traditional 3+3 design. The algorithm starts with a cohort of 3 patients (top). Depending on the number of DLTs, the algorithm stops (right), suggests to enroll three more patients at the same dose level (middle), or to escalate to the next dose level (left) and to continue with a new cohort (top).

There are slightly modified versions of the traditional 3+3 algorithm, in particular, those which use a different cohort size or those which allow to de-escalate dose levels [11]. In Sec. 4, we will follow the logic of one of latter. The algorithm is shown in Fig. 2:

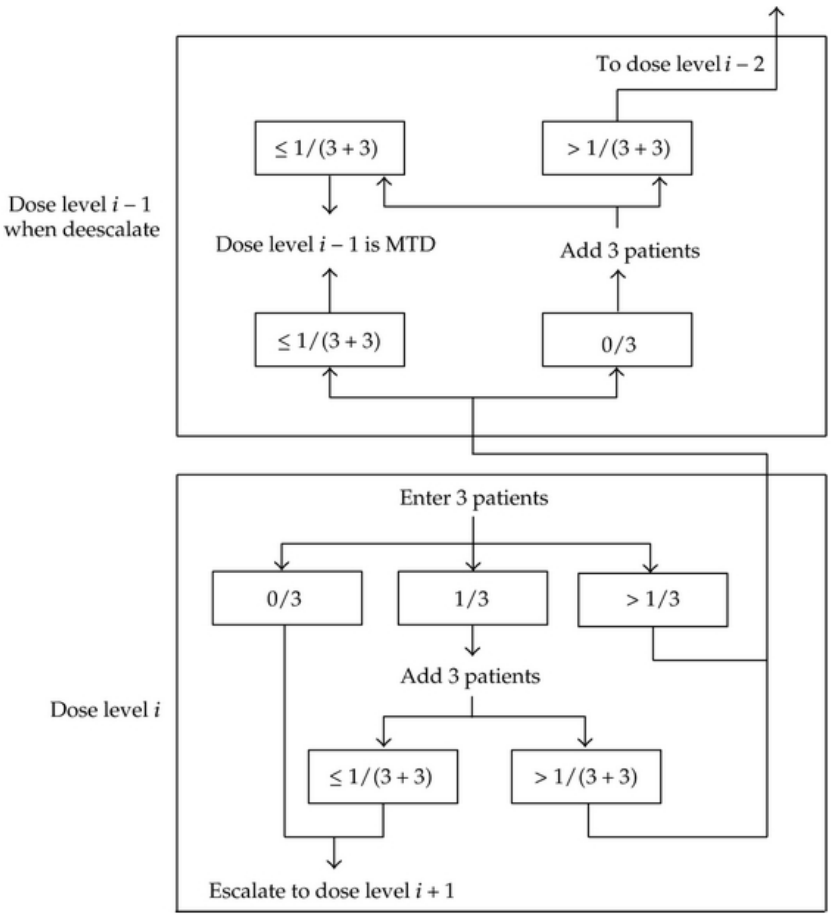

Figure 2: Graphical representation of a modified 3+3 design that allows to de-escalate dose levels (from [12]). The algorithm starts with a cohort of 3 patients (top of the lower panel). Depending on the number of DLTs, the algorithm suggests to escalate to the next dose level (bottom left of the lower panel), to enroll three more patients at the same dose level (middle), or to de-escalate to the next lower dose level and to continue with the algorithm depicted in the upper panel.

2.2 MODIFIED CONTINUAL REASSESSMENT MODEL

The modified continual reassessment model (mCRM) is a model-based design that is based on a Bayesian approach [2-4], where the decision making is data-driven. Whereas the traditional 3+3 design performs well only in a small range of DLT rates (< 20%), the mCRM is not limited to a certain range of DLT rates. On the contrary, the target rate of DLTs, p_t, serves as input to the model.

The mCRM always takes all observed data into account, in contrast to the 3+3-design and the mTPI method. If the current cohort is treated at one of the available dose levels $d_i, i = 1, \ldots, l$, the decision which dose level should be used for the next cohort depends on the data of all cohorts that have been observed until the time of decision, and is not only based on the outcomes of patients at dose level d_i. Furthermore, the model allows to skip dose levels and to jump to dose levels $d_{i\pm2}$, $d_{i\pm3}$, etc. This property limits safety issues if, for instance, the starting dose is too high, and optimizes convergence to the MTD if, for instance, the starting dose is too low.

In this paper, we use a logistic model to simulate the dose-toxicity probability relationship:

$$\text{logit}(p) = \ln\left(\frac{p}{1-p}\right) = \alpha + \beta d, \qquad \text{i.e.} \quad p = p(d) = (1 + e^{-\alpha-\beta d})^{-1} \tag{1}$$

Another realistic and commonly used model is the hyperbolic tangent model $p(d) = \tanh(\alpha + \beta d)$. Both are monotonically increasing functions $p(d) = p_{\alpha,\beta}(d)$ for all real α and positive β. We call p the *(dose limiting) toxicity probability* or, in our context synonymously, the *DLT rate*.

Furthermore, the likelihood of observing x DLTs out of n is given by a binomial distribution. Therefore, the Bayesian approach for the probability distribution of (α, β), given an observation (n, x), reads:

$$f_{\text{posterior}}\big((\alpha,\beta)\big|(n,x)\big) = \frac{p_{\alpha,\beta}^{x}(1-p_{\alpha,\beta})^{n-x} f_{\text{prior}}(\alpha,\beta)}{\int_{0}^{\infty}\int_{-\infty}^{\infty} p_{\alpha',\beta'}^{x}(1-p_{\alpha',\beta'}^{x})^{n-x} f_{\text{prior}}(\alpha',\beta')\, d\alpha'\, d\beta'} \tag{2}$$

In this paper, we choose a normal distribution for α and a positive normal distribution for β as prior distributions. The values of their parameters are given in Sec. 8.1 (numbers are inside the SAS macro code).

The idea is now to simulate the posterior distributions of α and β, which can be achieved by creating posterior samples of pairs (α, β). The PROC MCMC step in the SAS code of Sec. 8.1 does this by "burning in" (option `ntu` of the PROC MCMC step) and using, for instance, a Gibbs sampler for the Markov chain Monte Carlo algorithm. The number of iterations after which the sampling stops (option `nmc`) should be such that the samples reflect the posterior distribution in Eq. (2) well.

From the posterior distributions of α and β interesting quantities can be computed, for example,

- Probability of dose d being the MTD, i.e. probability that d is the maximum dose where $p < p_t$
- Probability that the DLT rate is below p_t: $T(d) = \sum_{(\alpha,\beta),p_{\alpha,\beta}(d)<p_t} 1 \;/\; \sum_{(\alpha,\beta)} 1$
- Expectation value of the DLT rate for dose d: $E(p|d) = \sum_{(\alpha,\beta)} p_{\alpha,\beta}(d) \;/\; \sum_{(\alpha,\beta)} 1$

The first quantity is basically a discretization of the distribution of the MTD, where the discretization scheme should reflect all available doses. In our example it is described by the format `dosefmt` used in the PROC FREQ step at the end of the macro `mtd_estimation` in Sec. 8.1. The MTD distribution can be obtained by using the posterior distribution of pairs (α, β) as well as by choosing a target rate, here $p = p_t = 0.3$, and converting Eq. (1):

$$MTD = \frac{\ln(p/(1-p)) - \alpha}{\beta} = \frac{\ln(3/7) - \alpha}{\beta} \tag{3}$$

The mCRM recommends for the next cohort of patients the dose with the highest probability of being the MTD.

In principle, this would be enough to determine the MTD. However, safety of patients is of utmost importance in early phase clinical trials. This is the reason why quantities like $T(d)$ and $E(d)$ should be monitored in a trial that uses the mCRM. The FDA, for example, states [1]: "[…] When using such designs [mCRM], the protocol should describe acceptable parameters for the dosing endpoint (supported by data). Irrespective of which dose-escalation approach is chosen, the study protocol should clearly define DLTs, the subject "off-treatment" criteria, and the study stopping rules that will ensure subject safety." Therefore, some common safety rules have been established for the mCRM. They read:

- *Safety rule 1 (termination):* The MTD is precisely estimated, i.e. the inter-quartile range of the MTD distribution is less than a certain percentage P_{term} of its median.
- *Safety rule 2 (high toxicity):* Minimum dose tested is toxic, i.e. $T(d_1) < 1 - \xi$.
- *Safety rule 3 (low toxicity)*: Maximum possible dose is safe, i.e. $T(d_l) > \xi$.
- *Safety rule 4 (maximum number of patients)*: Total sample size for next dose is already equal to N.

Of course, study-specific modifications are always possible. For the analysis in Sec. 4, we choose $P_{term} = 0.5$, $\xi = 0.8$, and $N = 9$, i.e. a maximum amount of three cohorts.

2.3 MODIFIED TOXICITY PROBABILITY INTERVAL METHOD

The following description of the modified toxicity interval (mTPI) method is based on Ref. [7]. Let us consider dose levels $d_i, i = 1, \ldots, l$, with $d_1 < \cdots < d_l$ of a toxic drug and a monotonically increasing dose-toxicity probability relationship $p = g(d)$, i.e. $g(d_1) = p_1 < \cdots < p_l = g(d_l)$. Let us suppose that x_i of n_i patients at dose d_i experienced a DLT. One of the following three decisions should be recommended by the mTPI method: "E" = escalate to d_{i+1}, "S" = stay at d_i, or "D" = de-escalate to d_{i-1}. The next cohort of the trial will be treated at that respective dose and the mTPI method will be applied again to the whole group of patients at that next dose level. The starting dose usually is, but does not necessarily have to be the lowest dose d_1.

The algorithm how to decide between "E", "S", and "D" uses the so-called *equivalence interval* (EI) and the *unit probability mass* (UPM). The EI is defined as the interval ranging from the "lowest toxicity probability that [a

physician] would be comfortable using to treat future patients without dose escalation" to the "highest toxicity probability that he/she would be comfortable using to treat future patients without dose de-escalation" [7]. The target DLT rate p_t, which defines the MTD via $p_t = g(MTD)$, lies in the EI: $p_t \in [p_t - \varepsilon_1, p_t + \varepsilon_2]$ where $\varepsilon_1, \varepsilon_2 > 0$. A common choice is, for example, $p_t = 0.3$ and $\varepsilon_1 = \varepsilon_2 = 0.05$. The EI divides the interval $[0,1]$ of toxicity probabilities into three subintervals: $[0, p_t - \varepsilon_1)$, $[p_t - \varepsilon_1, p_t + \varepsilon_2]$, $(p_t + \varepsilon_2, 1]$. Doses obtained from these three intervals via $d = g^{-1}(p)$ are deemed lower, close to, and higher than the MTD, respectively. The UPM at level i is defined as the interval average of the posterior density of p_i given all observations $(\boldsymbol{n}, \boldsymbol{x})$ where $\boldsymbol{n} = (n_1, \dots, n_l)$ and $\boldsymbol{x} = (x_1, \dots, x_l)$. The full density of all probabilities $\boldsymbol{p} = (p_1, \dots, p_l)$ is built from binomial distributions for each p_i multiplied with the prior distribution $\pi(\boldsymbol{p})$:

$$f\left(\boldsymbol{p}\middle|(\boldsymbol{n},\boldsymbol{x})\right) = \frac{\pi(\boldsymbol{p}) \prod_{i=1}^{l} p_i^{x_i} (1-p_i)^{n_i - x_i}}{\int_0^1 \dots \int_0^1 \pi(\boldsymbol{p}) \prod_{i=1}^{l} p_i^{x_i} (1-p_i)^{n_i - x_i} \, dp_1 \dots dp_l} \tag{4}$$

The posterior density f_i of toxicity probabilities p_i is the integral of $f\left(\boldsymbol{p}\middle|(\boldsymbol{n},\boldsymbol{x})\right)$ over all other p_j, $j = 1, \dots, l$, $j \neq i$:

$$f_i\left(p_i\middle|(\boldsymbol{n},\boldsymbol{x})\right) = \frac{\int_0^1 \dots \int_0^1 \pi(\boldsymbol{p}) \, p_i^{x_i} (1-p_i)^{n_i - x_i} \prod_{k=1, k\neq i}^{l} p_k^{x_k} (1-p_k)^{n_k - x_k} dp_k}{\int_0^1 \int_0^1 \dots \int_0^1 \pi(\boldsymbol{p}) \prod_{k=1}^{l} p_k^{x_k} (1-p_k)^{n_k - x_k} \, dp_k} \tag{5}$$

Hence, the UPM of interval (q_1, q_2) at level i reads:

$$UPM_i(q_1, q_2) = \frac{1}{q_2 - q_1} \int_{q_1}^{q_2} f_i\left(p_i\middle|(\boldsymbol{n},\boldsymbol{x})\right) dp_i \tag{6}$$

Having this at hand, the decision recommended by the mTPI method can be easily obtained:

Choose "E", "S", or "D" if the interval $[0, p_t - \varepsilon_1)$, $[p_t - \varepsilon_1, p_t + \varepsilon_2]$, or $(p_t + \varepsilon_2, 1]$ has the largest UPM. (7)

A general prior distribution $\pi(\boldsymbol{p})$ intertwines the integrals in Eq. (5) and enormously complicates their calculation. Therefore, a common choice that extremely simplifies the computation consists of independent priors π_i:

$$\pi(\boldsymbol{p}) = \prod_{i=1}^{l} \pi_i(p_i) \tag{8}$$

Inserting this into Eq. (5), the multiple integral factorizes, many terms cancel each other and a rather simplified expression for the UPM, Eq. (6), remains:

$$UPM_i(q_1, q_2) = \frac{1}{q_2 - q_1} \frac{\int_{q_1}^{q_2} \pi_i(p) \, p^{x_i} (1-p)^{n_i - x_i} \, dp}{\int_0^1 \pi_i(p) \, p^{x_i} (1-p)^{n_i - x_i} \, dp} \tag{9}$$

Another common simplification is the choice of prior beta-distributions: $\pi_i(p) \propto p^{a-1}(1-p)^{b-1}$. With this, the UPM just reads

$$UPM_i(q_1, q_2) = \frac{I_{q_2}(a,b) - I_{q_1}(a,b)}{q_2 - q_1} \tag{10}$$

Here, $I_q(a,b) = B(q;a,b)/B(a,b)$ is the regularized incomplete beta function, which is the cumulative distribution function of the beta distribution. Furthermore, $B(q;a,b) = \int_0^q p^{a-1}(1-p)^{b-1} \, dp$ is the incomplete beta function and $B(a,b) = B(1;a,b)$ is the beta function.

In addition to the decision algorithm (7), the mTPI method contains stopping (safety) rules (see Ref. [7]):

- *Safety rule 1 (early termination):* Suppose that d_1 has been used to treat patients. If $\Pr\left(p_1 > p_t \,\middle|\, (\boldsymbol{n},\boldsymbol{x})\right) > \xi$ for a ξ close to 1 (say, $\xi = 0.95$), then terminate the trial due to excessive toxicity.
- *Safety rule 2 (dose exclusion)*: Suppose that the decision is "E", i.e. to escalate from d_i to d_{i+1}. If the probability $\Pr\left(p_{i+1} > p_t \,\middle|\, (\boldsymbol{n},\boldsymbol{x})\right) > \xi$, then treat the next cohort of patients at d_i and exclude all d_j with $j > i$ from the trial, i.e. these doses will never be used again in the trial (label: "DU").
- *Safety rule 3 (maximum number of patients)*: Total sample size for next dose is already equal to N.

Choices of a prior distribution $\pi(\boldsymbol{p})$, the EI defined by p_t, ε_1, and ε_2, the safety parameter ξ, etc. are study specific. For the scenarios discussed below in Sec. 4 we will choose a uniform prior (i.e. a beta distribution with $a = b = 1$), $EI = [0.25, 0.35]$, $p_t = 0.3$, $\xi = 0.95$, and a maximum number of three cohorts at the same dose level (in order to make it somehow comparable to the 3+3 design).

3 TOXICITY DATA

The sample data considered in this paper consists of patients who have received the same drug, but in different doses: 100 mg, 200 mg, … up to 1000 mg (per day). For the sake of simplicity and for comparability only discrete doses are allowed. This is indeed often the case, for instance, if the drug is taken orally as a pill. If the drug is available in continuous doses, like a liquid administered by an intravenous infusion, the 3+3 design and the mTPI method cannot be applied and only the mCRM makes sense as a possible dose-finding scheme. Furthermore, all cohorts consist of 3 patients.

The structure of toxicity data of early phase clinical trials usually looks as follows (see Table 1). Each patient has a unique identification number ("Subject identifier"). There is one record for each patient. The dose received by each patient is encoded in the variable "Treatment arm". The variable "Cohort" (= 1, 2, …) represents the order of administered doses. The group of patients treated with the first dose is cohort 1, the group of patients treated with the next recommended dose is cohort 2 and so on. The number of different values and the order of variable "Cohort" does not necessarily have to coincide with the number and order of available doses. Treatment arm and cohort are defined by patient's starting dose. If a patient reduces the dose during the study due to other toxicities or safety issues, treatment arm and cohort remain nonetheless unchanged. In addition to the subject identifier, the treatment arm, and the cohort number, there is a variable named "DLT Flag" that indicates whether the patient had a dose-limiting toxicity or not ("Y" or "N").

Toxicity data - Scenario 3

Subject identifier	Cohort	Treatment arm	DLT flag
1001	1	100 mg	N
1002	1	100 mg	N
1003	1	100 mg	N
1004	2		
1005	2		
1006	2		

Table 2: Part of the sample toxicity data of Scenario 3. Starting cohort 1 consists of three patients who can be distinguished by their "Subject identifier" and who received a dose of 100 mg ("Treatment arm"). In this example, there is no DLT observed for the first cohort (DLT flag = "N"). As we shall see, in this case, all dose-finding schemes presented in this paper recommend to escalate to a higher dose. Therefore, there will be a second cohort of patients, whose dose depends on the applied method.

As mentioned in the introduction of this paper, we examine three different scenarios. The two extreme cases, where the lowest available dose is toxic (Scenario 1) or where the largest available dose is safe (Scenario 2), are somehow trivial. We therefore show in Tab. 1 the starting point of a more realistic scenario (Scenario 3).

Since the total number of treated patients depends on the applied dose-finding method, the sample data of Tab. 1 is not complete yet and only the first part is shown (cohorts 1 and 2). Later, this data gets more and more complete, depending on the applied dose-finding scheme, until the maximum tolerable dose (MTD) is found. We start with a cohort of 3 patients who have received the lowest possible dose available (100 mg). The number of DLTs can be easily obtained by counting the records where the DLT flag is "Y". In the example of Tab. 1 the number of DLTs for the first cohort is zero. Dose escalation is recommended by all dose-finding schemes presented below. The next three patients enrolled to the study belong to a different cohort 2. In cases of the 3+3 design or the mTPI method these patients will receive the next higher dose of 200 mg. The recommended dose of the mCRM depends on the parameters chosen (prior distribution, MTD rate, etc). Therefore, the treatment arm information as well as the DLT outcome are left open in Tab. 1 and will be filled in later on.

For the sake of simplicity, we transform the data set described above to one in which we summarize the results "per cohort". The subject identifier gets replaced by the number of subjects per cohort and the DLT flag by the number of DLTs per cohort. The transformed data set of the example shown in Tab. 1 looks as follows.

Summarized toxicity data - Scenario 3

Cohort	Number of subjects	Treatment arm	Number of DLTs
1	3	100 mg	0
2	3		.

Table 2: Part of the summarized sample toxicity data of Scenario 3. Subject identifier and DLT flag are omitted. Instead, there are two new variables which count the number of subjects and DLTs per cohort, respectively.

For the analysis of the next section and, in particular, for the SAS macro presented in Sec. 8.1, we shall use this kind of (summarized) data structure. The summarized data shown in Tab. 2 serves as starting point of Scenario 3 for all dose-finding schemes.

4 APPLICATION OF DOSE-FINDING SCHEMES

In this section we present three different scenarios. For each dose-finding scheme, we first discuss two extreme cases, where the first dose is toxic (Scenario 1) or where the largest available dose is safe (Scenario 2), in order to show that stopping rules of the algorithms, in particular for the mCRM, work well. Subsequently, we examine a more realistic scenario (Scenario 3) in order to show how the dose-finding algorithms work in detail.

4.1 THE 3+3 DESIGN

In Scenario 1 all three patients of the first cohort have a DLT. The 3+3 design as depicted in Fig. 2 clearly states to de-escalate to the next lower dose. Since there is no lower dose, the study stops. No dose is recommended as MTD, due to further safety rules of the study.

In Scenario 2 no patient experiences a DLT up to a dose of 600 mg. For the cohort at 700 mg one DLT is observed. The 3+3 design (see Fig. 2) recommends to enroll a second cohort at 700 mg. Let us suppose there is no DLT observed for that cohort and neither it is for all following cohorts. Escalation then continues up to the highest available dose of 1000 mg. The total number of enrolled cohorts is 6 (100 mg, …, 600 mg) plus 2 at 700 mg plus 3 (800 mg, …, 1000 mg). This is a total of 33 patients, where one of them had a DLT at a dose level of 700 mg. The MTD is the highest available dose of 1000 mg.

Let us now focus on a more realistic scenario (Scenario 3). Let us assume that the first observed DLTs are two DLTs at a dose of 700 mg. The 3+3 design suggests to de-escalate to the next lower dose (upper panel of Fig. 2). Let us further assume that there is 1 DLT observed for the next cohort at 600 mg. From the upper panel of Fig. 2 we see that no further patients will be enrolled and the MTD is 600 mg. In total, there are 24 patients and 3 DLTs.

4.2 MODIFIED CONTINUAL REASSESSMENT MODEL

Let us now apply the modified continual reassessment model (mCRM) to the three scenarios introduced in the previous paragraph. In order to make everything comparable to the results of the 3+3 design where, vaguely put, one DLT out of three is considered to be acceptable, we choose a target DLT rate of $p_t = 0.3$, see Eq. (3). The main task is now to compute the posterior distribution of parameters α and β in Eq. (2), and hence the posterior distributions of the MTD and of rates p given the doses $d \in \{100,\ 200, \ldots, 1000\}$ (doses in mg).

The SAS code [9] that simulates the posterior distribution by means of the Markov chain Monte Carlo (MCMC) algorithm as well as computes some parameters to test the stopping rules as described in Sec. 2.2 is presented in the Appendix (Sec. 8.1). It needs as input the toxicity data set (see e.g. Tab. 2), the DLT rate p_t, and a list of all available (discrete) doses. From the output data set (`&indat.out_final.sas7bdat`) we read the dose that is recommended to be used for the next cohort of patients, i.e. the dose with the highest probability being the MTD (see Sec. 2.2).

Scenario 1 is the case when for all three patients treated at a dose of 100 mg a DLT is observed:

```
DATA toxic;
    INPUT n dlts dose @@;
    DATALINES;
    3 3 100
    ;
RUN;
%mtd_estimation(indat = toxic);
```

Table 3: Output data TOCIXOUT_FINAL.sas7bdat. The number of rates equal to $p_t = 0.3$ in the distribution of dose-rate functions, Eq. (1) simulated by the PROC MCMC step, is counted for each dose interval. The percentage is the probability of the lower interval limit being the MTD.

Dose d	Frequeny count	Prob. (%) of dose d being the MTD
0	1167	58.35
100	315	15.75
200	174	8.70
300	109	5.45
400	48	2.40
500	43	2.15
600	21	1.05

From the output data set `TOXICOUT_FINAL.sas7bdat` generated by the PROC FREC step at the end of the macro `%mtd_estimation` in Sec. 8.1 we see that the category `low-100='0'` contains most of the counts (about 58%) and the algorithm strongly suggests to lower the dose from 100 mg to zero. Suppose we were able to enroll three more patients at 100 mg and suppose two of them had a DLT, the macro code would write the following message to the log "`WARNING: Min. dose tested is toxic: T_1 = 19.0% < 20%`", which reflects the second stopping rule of Sec. 2.2.

Scenario 2 is a bit more complex. Let us assume that the first DLT appears at 700 mg. Therefore, we start with an input data set that shows zero DLTs at 100 mg:

```
DATA safe;
    INPUT n dlts dose @@;
    DATALINES;
    3 0 100
    ;
RUN;
%mtd_estimation(indat = safe);
```

Table 4: First output data set SAFEOUT_FINAL.sas7bdat of macro %mtd_estimation for Scenario 2. Input data set was the data set SAFE.sas7bdat generated by the data step from above. Columns are the same as in Tab. 3.

Dose d	Frequeny count	Prob. (%) of dose d being the MTD
0	22	1.10
100	46	2.30
200	110	5.50
300	150	7.50
400	169	8.45
500	162	8.10
600	131	6.55
700	119	5.95
800	108	5.40

From the output data set `SAFEOUT_FINAL.sas7bdat` (see Tab. 4) we conclude that the probability of dose d being the MTD (see Sec. 2.2) has its maximum for d = 400 mg (8.45%). So, we add a cohort with zero DLTs at 400 mg. The next maximum value is about 7.7% at a dose of 500 mg. After adding a cohort with zero DLTs at 500 mg, we find a maximum of 6.55% at a dose of 700 mg. For the next cohort of three patients (at a dose of 700 mg) one DLT is observed:

```
DATA safe;
    INPUT n dlts dose @@;
    DATALINES;
    3 0 100
    3 0 400
    3 0 500
    3 1 700
    ;
RUN;
%mtd_estimation(indat = safe);
```

Table 5: Output data set SAFEOUT_FINAL.sas7bdat of the fourth run of macro %mtd_estimation, after adding the cohort at 700 mg to the input data set SAFE.sas7bdat. Columns are the same as in Tab. 3.

Dose d	Frequeny count	Prob. (%) of dose d being the MTD
200	1	0.05
300	15	0.75
400	64	3.20
500	131	6.55
600	183	9.15
700	232	11.60
800	188	9.40
900	127	6.35
1000	136	6.80

Again, the maximum probability (11.6%) is at 700 mg. The mCRM therefore recommends to continue with a second cohort at 700 mg.

Let us suppose there is no DLT observed for that cohort. The mCRM then recommends a dose of 900 mg (8.6%) for the next cohort. If there are again no DLTs at 900 mg, the next macro call yields:

```
DATA safe;
    INPUT n dlts dose @@;
    DATALINES;
    3 0 100
    3 0 400
    3 0 500
    3 1 700
    3 0 700
    3 0 900
    ;
RUN;
%mtd_estimation(indat = safe);
```

Table 6: Last output data set SAFEOUT_FINAL.sas7bdat of macro %mtd_estimation for Scenario 2. Columns are the same as in Tab. 3. The maximum of the MTD distribution, see Eq. (3) in Sec. 2.2, is at the maximum available dose of 1000 mg.

Dose d	Frequeny count	Prob. (%) of dose d being the MTD
100	1	0.05
400	5	0.25
500	17	0.85
600	37	1.85
700	77	3.85
800	134	6.70
900	131	6.55
1000	137	6.85
1100	122	6.10
1200	119	5.95

Further, the following message is written to the log: "`WARNING: Max dose is safe: T_10=85.5% > 80%`". It means that the third stopping rule of Sec. 2.2 (for the maximum dose of 1000 mg) is met. There is a total number of 18 patients that was needed to achieve this, where one of them had a DLT (at a dose of 700 mg).

Let us now consider Scenario 3. Up to a dose of 700 mg the procedure is the same as for Scenario 2. The difference is that now two DLTs are observed instead of only one DLT. The mCRM algorithm suggests to go back to 500 mg (18.3%). Let us assume that there is one DLT observed for the next cohort. The final input data are as follows:

```
DATA real;
    INPUT n dlts dose @@;
    DATALINES;
    3 0 100
    3 0 400
    3 0 500
    3 2 700
    3 1 500
    ;
RUN;
%mtd_estimation(indat = real);
```

Table 7: The final input data set of Scenario 3. Up to now, there are in total 15 patients and 3 observed DLTs.

Summarized data: mCRM - Scenario 3

Cohort	Number of subjects	Treatment arm	Number of DLTs
1	3	100 mg	0
2	3	400 mg	0
3	3	500 mg	0
4	3	700 mg	2
5	3	500 mg	1

Calling the macro `%mtd_estimation` with this input data yields a maximum percentage of 22.5% at 500 mg. Furthermore, the following message is written to the log: "`WARNING: Limit for coefficient of variation reached: CV(MTD) = 0.425 < 1/2, i.e. IQR less than half of median => MTD precisely estimated!`". It means that the first stopping rule of Sec. 2.2 is met. The recommended dose is 500 mg. In total, there were 15 patients treated and 3 DLTs observed.

4.3 MODIFIED TOXICITY PROBABILITY INTERVAL METHOD

First of all, let us compute a table of recommendations for dose (de-)escalation by means of the modified toxicity probability (mTPI) method (see Sec. 2.3). For our sample study design, we choose $p_t = 0.3$ as target DLT rate. As an acceptable result for the final DLT rate we choose a symmetric toxicity probability interval, which is the equivalence interval (EI) introduced in Sec. 2.3, around p_t: [0.25, 0.35]. It means that the MTD which is found after determination of the mTPI algorithm is such that the DLT rate lies (at least with a high probability) in that interval. Furthermore, we choose a flat prior distribution (a = b = 1).

The SAS code [9] that we used to compute the table of mTPI's recommendations is written in form of a macro and is presented in the Appendix (Sec. 8.2). If we call that macro by `%create_mTPI_table(nmin=3, nmax=12, dltsmax=7)`, we obtain the following table:

ESD(U)-Table of the mTPI method

# DLTs	n=3	n=4	n=5	n=6	n=7	n=8	n=9	n=10	n=11	n=12
0	E	E	E	E	E	E	E	E	E	E
1	S	S	S	E	E	E	E	E	E	E
2	D	S	S	S	S	S	S	S	E	E
3	DU	DU	D	S	S	S	S	S	S	S
4		DU	DU	DU	D	D	S	S	S	S
5			DU	DU	DU	DU	DU	D	S	S
6				DU	DU	DU	DU	DU	DU	D
7					DU	DU	DU	DU	DU	DU

Table 8: Entries represent recommendations of the mTPI method for dose (de-)escalation ("E" = escalate, "S" = stay, "D" = de-escalate, "DU" = de-escalate and never use again). The derivation of the table is based on a symmetric toxicity probability interval of width 0.1 around $p_t = 0.3$, a flat prior distribution (a=b=1) and an exclusion probability of 95% (see safety rules 1 and 2 in Sec. 2.3).

Before we apply the modified toxicity interval (mTPI) method to our three scenarios we would like to stress that the mTPI method is not well-suited for a study design with so many different dose levels available (100 mg, 200 mg, …, 1000 mg) and with such small cohort sizes of three patients. Instead, it shows its real superiority over the 3+3 design for studies with few dose levels available by enrolling patients into larger cohorts (n = 4, 5, etc.) [7,8]. However, for the sake of comparability, we consider the study design introduced in Sec. 3. Therefore, we only need columns 1, 2, 5, and 8 of Tab. 8 (i.e. #DLTs, n=3, n=6, and n=9).

For Scenario 1, i.e. three DLTs at the lowest dose level, it can be seen from the 2nd column (n=3) and 4th row of Tab. 8 (i.e. 3 DLTs; row of headers does not count) that the mTPI method suggests to de-escalate to the next lower dose and to never use the dose again ("DU"). It means that the first stopping rule of Sec. 2.3 is met and the trial stops. Therefore, Scenario 1 proceeds for the mTPI in the same way as for the 3+3 design.

The enrollment of cohorts in Scenario 2 is also exactly the same as for the 3+3 design. At dose 700 mg there is one DLT. The suggestion of the mTPI method is to stay at the dose, see 2nd column (n=3), 2nd row (1 DLT): "S". A second cohort at 700 mg is enrolled. Since there is no additional DLT at 700 mg, the 5th column (n=6) and 2nd row of Tab. 8 (still 1 DLT) gives the suggestion to escalate to the next higher dose ("E"). Since there is no further DLT observed, the dose is increased step by step until the maximum dose of 1000 mg is reached. From the 2nd column (n=3 at 1000mg) and 1st row (no DLT) we see that the mTPI method suggests to escalate to the next higher dose, which is not possible. Therefore, the trial stops and the highest available dose of 1000 mg is also the recommended dose.

The procedure of dose escalation of Scenario 3 by means of the mTPI method is again very similar to the one using the 3+3 design. The only difference happens when the second cohort at 600 mg is enrolled after de-escalation from 700 mg (2nd column (n=3), 3rd row (2 DLTs): "D"). There is one DLT observed for that cohort at 600 mg. We conclude from the 5th column (n=6) and 2nd row of Tab. 8 (one DLT) that escalation is suggested. Since the dose of 700 mg is still allowed (only "D" suggested for the previous de-escalation step, not "DU"), a second cohort will be enrolled at 700 mg. How the trial proceeds, depends on the outcome of that cohort. For instance, if there is no DLT or only one, a dose of 700 mg will be recommended as dose of the next cohort. In case of two (or even three) additional DLTs the dose of 700 mg will be never used again: n = 6 at 700 mg, and more than three DLTs. Let us assume the latter case is true. In principle, there should be a next cohort at 600 mg. Since there are already two cohorts at 600 mg with a total of only one DLT, the next suggestion of the mTPI method would be to continue at 600 mg, independent from the outcome as can be seen from column 8 (n=9) and rows 2-5 of Tab. 8. This would result in a total number of four cohorts at 600 mg. This is not allowed according to the third safety rule with N = 9 defined in Sec. 2.3. Therefore, enrollment of a third cohort at 600 mg is not necessary. The recommended dose is 600 mg and there is a total of 27 treated patients with five DLTs.

5 CONCLUSIONS

In clinical trials, especially in oncology, toxicity data like dose-limiting toxicities (DLTs) are used to estimate the maximum tolerable dose (MTD). In this paper, we analyzed sample toxicity data (Scenario 1-3) in order to compare the most common methods to find the recommended dose (RD) for all follow-up trials, i.e. the 3+3 design, the modified continual reassessment model (mCRM), and the modified toxicity probability interval (mTPI) method. Due to its simplicity, the 3+3 design is still the "Gold standard" in oncology. The main advantages of the other two methods are that decision making is data-driven and that they are more flexible. As we have seen in Sec 4.3, the traditional 3+3 design can be somehow covered by the mTPI method, whereby the latter has not shown its full strength.

In particular, we have shown how the two model-based dose-finding schemes (mCRM and mTPI) can be realized by means of SAS [9]. In the Appendix, Sec. 8, we present two SAS macro codes which can be used to implement the corresponding design into the study. We would like to stress that the mTPI method is well-suited for study designs with only a few discrete dose levels available. It has the advantage that all calculations can be made before the trial actually starts. On the other hand, the mCRM outperforms the mTPI method for study designs where many dose levels are available (as in our sample trial with 10 different dose levels). We could observe this behavior in our Scenario 3. Let us compare the main results of Sec. 4 for this scenario. Tables 9 and 10 summarize these results.

Toxicity data: Scenario 3 - Comparison

Cohort	Number subjects	Dose 3+3 design	# DLTs 3+3 design	Dose mCRM	# DLTs mCRM	Dose mTPI	# DLTs mTPI
1	3	100 mg	0	100 mg	0	100 mg	0
2	3	200 mg	0	400 mg	0	200 mg	0
3	3	300 mg	0	500 mg	0	300 mg	0
4	3	400 mg	0	700 mg	2	400 mg	0
5	3	500 mg	0	500 mg	1	500 mg	0
6	3	600 mg	0		.	600 mg	0
7	3	700 mg	2		.	700 mg	2
8	3	600 mg	1		.	600 mg	1
9	3		.		.	700 mg	2

Table 9: Summary of full trial's toxicity data of Scenario 3 for all three methods discussed in Sec. 4.

Recommended dose: Scenario 3 - Comparison

Category	3+3	mCRM	mTPI
Number of treated patients	24	15	27
Number of observed DLTs	3	3	5
Recommended dose (RD)	600 mg	500 mg	600 mg
Expected DLT rate at RD	0.17	0.22	[0.25, 0.35]
Prob. RD less or equal MTD	67.1 %	80.8 %	74.5 %

Table 10: Summary of Scenario 3 for all three methods (see Sec. 4). Compared are total number of patients and DLTs as well as outcomes for the recommended dose (RD), expected DLT rates at RD, and a percentage for the probability of the RD being less or equal to the true MTD.

We observe that the number of treated patients is much less for the mCRM as compared to the other two dose-finding schemes. One of the main reasons is that the mCRM allows to skip dose levels. Therefore, the convergence to the final recommended dose is much faster.

Furthermore, we compare in Tab. 10 the probability of the recommended dose (RD) being less or equal to the true MTD. It can be obtained, for the mCRM by summing up all probabilities being the MTD of doses larger than or equal to the RD, for the 3+3 design by integrating the likelihood function to find the scenario 1 DLT out of 6 over all DLT rates larger than $p_{tol} = 0.3$, and for the mTPI method by integrating the likelihood function to find the scenario 1.5 DLTs out of 9 over all DLT rates larger than $p_{tol} = 0.3$. The latter is an approximation and takes into account that there actually should be a third cohort of 3 patients at 600 mg (as described in Sec. 4.3) with most likely no or one DLT.

In summary, it can be concluded that model-based designs are better suited to finding a reliable result for the RD. In addition, the mCRM or mTPI method usually need a smaller total amount of patients as compared to the 3+3 design. In our Scenario 3, however, the mTPI method performs similarly to the 3+3 design. This is just because cohort sizes are too small and the number of different available dose levels is too high for the mTPI method. For other study designs, the mTPI method is much better suited and outperforms the 3+3 design, see e.g. Refs. [7,8] for a comprehensive comparison of those two designs. The only argument in favor of the traditional 3+3 design is its simplicity. The FDA, for instance, says [1]: "Given the relatively acceptable safety profile of some classes of cancer vaccines, alternative dose-escalation approaches, such as […] continuous reassessment, may be considered instead of the standard '3+3 design'."

6 REFERENCES

[1] FDA, https://www.fda.gov/; *FDA guidance: Clinical Considerations for Therapeutic Cancer Vaccines,* https://www.fda.gov/downloads/biologicsbloodvaccines.

[2] Storer BE. *Design and Analysis of Phase I Clinical Trials.* Biometrics. 1989 Sep; 45: 925-937.

[3] O'Quigley J, Pepe M, Fisher L. *Continual reassessment method: a practical design for phase 1 clinical trials in cancer.* Biometrics. 1990 Mar; 46(1): 33-48.

[4] Garrett-Mayer E. *The continual reassessment method for dose-finding studies: a tutorial.* Clin Trials. 2006; 3: 57-71.

[5] Goodman SN, Zahurak ML, Piantadosi S. *Some practical improvements in the continual reassessment method for phase I studies.* Stat Med. 1995; 14: 1149-1161.

[6] Faries D. *Practical modifications of the continual reassessment method for phase I cancer clinical trials.* J Biopharm Stat. 1994; 4: 147-164.

[7] Ji Y, Liu P, Li Y, Bekele N. *A modified toxicity probability interval method for dose-finding trials.* Clin Trials. 2010 Dec; 7(6): 653-63.

[8] Ji Y, Wang SJ. *Modified Toxicity Probability Interval Design: A Safer and More Reliable Method Than the 3 + 3 Design for Practical Phase I Trials*. J Clin Oncol. 2013 May; 31(14): 1785-1791.

[9] The output/code for this paper was generated using SAS® software, version 9.4. Copyright © 2018 SAS Institute Inc. SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc., Cary, NC, USA.

[10] Iasonos A 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: 465-477.

[11] Lin Y, Shih WJ. *Statistical properties of the traditional algorithm-based designs for phase I cancer clinical trials.* Biostatistics. 2001 June; 2 (2): 203-215.

[12] Zhao Y. *Escalation scheme for 3+3 design with dose de-escalation*. (Adapted from Lin and Shih [7].) https://www.researchgate.net/figure/Escalation-scheme-for-3-3-design-with-dose-de-escalation-adapted-from-Lin-and-Shih-4_fig1_258385063.

7 CONTACT INFORMATION

Your comments and questions are valued and encouraged. Please feel free to contact the author at

Author Name	Michael Brockmann
Company	CHRESTOS Concept GmbH & Co.KG
Address	Giradetstraße 1-5
City / Postcode	45131 Essen
Work Phone:	+49 151 70239205
Email:	michael.brockmann@chrestos.de
Web:	www.chrestos.de

8 APPENDICES

8.1 MACRO CODE TO ESTIMATE THE MTD BY MEANS OF THE MCRM

The following macro `%mtd_estimation` needs as input a (summarized) toxicity data set `&indat` (see for example Tabs. 2 or 7), a target DLT rate p_t (`&DLTrate`), and a list of all available discrete doses separated by a blank (`&doses`). By default, the last two input parameters are set to the values of our example. Note that the heart of the computation, the PROC MCMC step, runs stable under SAS® version 9.4. Previous versions of SAS® software might produce different (unstable) results. The macro ycode reads:

```
%MACRO mtd_estimation(
    indat   =
  , DLTrate = 0.3
  , doses   = %str(100 200 300 400 500 600 700 800 900 1000)
);

%* MCMC algorithm to compute posterior distribution of alpha, beta and MTD;
PROC MCMC data=&indat. ntu=1000 nmc=20000 nthin=10 outpost=&indat.out seed=112358;
    mu_a = -3.0; sigma_a = 1;
    PRIOR alpha ~ normal(mu_a, var=sigma_a**2);
    PARMS (alpha) -3.0;

    mu_b =  0.0015; sigma_b = 0.003;
    f_b = -(beta - mu_b)**2/2.0/sigma_b**2 - 1.0e16*(1.0 - sign(beta));
    PRIOR beta  ~ general(f_b);
    PARMS (beta) 0.0015;

    p = logistic(alpha + beta*dose);
    MODEL dlts ~ binomial(n, p);
RUN;

%* T(d), i.e. tolerable probability that DLT rate < &DLTrate. for a dose d;
%* E(p|d), i.e. expectation value of DLT rate p given dose d;
%LOCAL i;
DATA &indat.out_mtd(DROP=T: E: _:) &indat.out_TE(KEEP = T: E:);
    SET &indat.out END=eof;
    mtd = (log(&DLTrate./(1.0-&DLTrate.)) - alpha)/beta;
    %DO i=1 %TO %sysfunc(countw(&doses., ' '));
        RETAIN _count&i. _sum_p&i. 0;
        _dose&i. = %scan(&doses., &i., ' ');
        IF 1.0/(1.0+exp(-alpha-beta*_dose&i.)) < &DLTrate. THEN _count&i. +1;
        _sum_p&i. + 1.0/(1.0+exp(-alpha-beta*_dose&i.));
        IF eof THEN DO;
            T&i. = _count&i./_N_;
            E&i. = _sum_p&i./_N_;
            IF &i. = %sysfunc(countw(&doses., ' ')) and T&i. > 0.8 THEN DO;
                _T&i._l = put(100*T&i., 4.1);
                PUT "WAR" "NING: Max dose is safe: T_&i.=" _T&i._l "% > 80%";
            END;
        END;
    %END;
    IF . < T1 < 0.2 THEN DO;
        _T1_l = put(100*T1, 4.1);
        PUT "WAR" "NING: Min. dose tested is toxic: T_1 = " _T1_l "% < 20%";
    END;
    OUTPUT &indat.out_mtd;
    IF eof THEN OUTPUT &indat.out_TE;
RUN;

%* Coefficient of variation;
ODS OUTPUT summary = &indat.out_median_iqr_mtd;
PROC MEANS DATA = &indat.out_mtd p50 qrange;
    VAR mtd;
RUN;
ODS OUTPUT CLOSE;
```

```
DATA &indat.out_cv_mtd;
    SET &indat.out_median_iqr_mtd;
    cv_mtd = abs(mtd_qrange/mtd_p50);
    IF cv_mtd < 0.5 THEN PUT "WAR" "NING: Limit for coefficient of variation
reached: CV(MTD) = " cv_mtd " < 1/2, i.e. IQR less than half of median => MTD
precisely estimated!";
RUN;

%* Prob. of dose d being the MTD, i.e. prob. that d is max dose with p < 20%;
PROC FREQ data=&indat.out_mtd noprint;
    TABLES mtd / nocum out=&indat.out_final;
    format mtd dosefmt.;
RUN;

%MEND mtd_estimation;
```

The format `dosefmt.`, which is used in the PROC FREQ step at the end of the macro, reflects the available dose levels (of our sample trial):

```
PROC FORMAT;
    VALUE dosefmt
           low - 100  = '0'
          100 <- 200  = '100'
          200 <- 300  = '200'
          300 <- 400  = '300'
          400 <- 500  = '400'
          500 <- 600  = '500'
          600 <- 700  = '600'
          700 <- 800  = '700'
          800 <- 900  = '800'
          900 <- 1000 = '900'
         1000 <- 1100 = '1000'
         1100 <- 1200 = '1100'
         1200 <- 1300 = '1200'
         1300 <- 1400 = '1300'
         1400 <- 1500 = '1400'
         1500 <- high = '>1500';
RUN;
```

In case of other discrete dose levels, which are available to administer to patients, this format has to be adapted accordingly. Note that we expanded the format above the highest available dose of 1000 mg a little bit (up to fictive doses of 1500 mg) in order to avoid accumulation of probabilities in the last category which actually comes from the (infinitely long) tail of the distribution function.

8.2 MACRO CODE TO CREATE TABLE OF MTPI'S RECOMMENDATIONS

The SAS code [9] which can be used to compute the table of recommendations of the modified toxicity probability interval (mTPI) method looks as follows:

```
%MACRO create_mTPI_table(
    nmin    =
  , nmax    =
  , dltsmax =
  , a       = 1
  , b       = 1
  , ptol    = 0.3
  , eps1    = 0.05
  , eps2    = 0.05
  , excl    = 0.95
);

DATA mtpi_table(drop = a b ptol eps1 eps2 excl pp escal stabl deesc);

    LENGTH n dlts 8 saf $2 idlabel $10;
    LABEL dlts = "# DLTs";
```

```
    DO n = &nmin. TO &nmax.; DO dlts = 0 TO &dltsmax.;
        IF dlts > n THEN saf = " ";
        ELSE DO;
            a = %sysevalf(&a.);
            b = %sysevalf(&b.);
            ptol = %sysevalf(&ptol.);
            eps1 = %sysevalf(&eps1.);
            eps2 = %sysevalf(&eps2.);
            excl = %sysevalf(&excl.);

            * Posterior probability;
            pp = 1.0 - probbeta(ptol, a + dlts, b + n - dlts);

            * Unit probability mass (UPM);
            escal =  probbeta(ptol-eps1, a+dlts, b+n-dlts) / (ptol-eps1);
            stabl = (probbeta(ptol+eps2, a+dlts, b+n-dlts)
                        - probbeta(ptol-eps1, a+dlts, b+n-dlts)) / (eps2+eps1);
            deesc = (1.0 - probbeta(ptol+eps2, a+dlts, b+n-dlts ))/(1.0-ptol-eps2);

            * mTPI's recommendations: "E"= escalate, "S"= Stay, "D"= de-escalate;
            IF      escal > stabl and escal > deesc THEN saf = "E";
            ELSE IF stabl > escal and stabl > deesc THEN saf = "S";
            ELSE IF deesc > stabl and deesc > escal THEN saf = "D";

            * "DU" = De-escalate and never use again;
            IF pp > excl THEN saf = "DU";
        END;

        idlabel = "n=" || strip(put(n, best.));
        OUTPUT;
    END; END;
RUN;

PROC SORT data=mtpi_table;
    BY dlts n;
RUN;

PROC TRANSPOSE data=mtpi_table out=mtpi_table_final(drop=_name_);
    BY dlts;
    ID n;
    IDLABEL idlabel;
    VAR saf;
RUN;

* Table output;
ODS HTML;
TITLE "ESD(U)-Table of the mTPI method";
PROC PRINT data=mtpi_table_final label noobs;
RUN;
ODS HTML CLOSE;

%MEND create_mTPI_table;
```