

Sequential Designs for Phase I Clinical Trials with Late-Onset Toxicities

Ying Kuen Cheung* and Rick Chappell

Department of Statistics and Department of Biostatistics and Medical Informatics,
University of Wisconsin, Madison, Wisconsin 53706, U.S.A.

*email: ken@stat.wisc.edu

SUMMARY. Traditional designs for phase I clinical trials require each patient (or small group of patients) to be completely followed before the next patient or group is assigned. In situations such as when evaluating late-onset effects of radiation or toxicities from chemopreventive agents, this may result in trials of impractically long duration. We propose a new method, called the time-to-event continual reassessment method (TITE-CRM), that allows patients to be entered in a staggered fashion. It is an extension of the continual reassessment method (CRM; O'Quigley, Pepe, and Fisher, 1990, *Biometrics* 46, 33–48). We also note that this time-to-toxicity approach can be applied to extend other designs for studies of short-term toxicities. We prove that the recommended dose given by the TITE-CRM converges to the correct level under certain conditions. A simulation study shows our method's accuracy and safety are comparable with CRM's while the former takes a much shorter trial duration: a trial that would take up to 12 years to complete by the CRM could be reduced to 2–4 years by our method.

KEY WORDS: Continual reassessment method; Dose limiting; Late-onset toxicities; Likelihood-based design; Phase I trial; Time-to-event.

1. Introduction

Phase I clinical trials are experiments in which a drug or radiation treatment is administered to humans to determine the maximum tolerated dose (MTD) as an upper safety limit for further clinical research on efficacy. While a strict quantitative definition of the MTD is seldom provided in clinical protocols, Storer and DeMets (1987) define it as some percentile of a tolerance distribution of the individual doses that result in an objectively defined clinical toxicity in a population of patients. Thereafter, several new designs for phase I clinical trials were proposed using this percentile interpretation of the MTD (e.g., Whitehead and Brunier, 1995; Durham, Flournoy, and Rosenberger, 1997; Babb, Rogatko, and Zacks, 1998).

Although not stated explicitly in the literature, the MTD is related to a given duration of time. To examine, e.g., if the treatment will cause late-onset toxicities, the MTD should be defined with respect to a long period of time or else the trial may underestimate the harm incurred. A phase I trial of perillyl alcohol was conducted at the the University of Wisconsin Medical School during the period from October 1995 to June 1996. In the study, a dose-limiting toxicity (DLT) was defined as any grade 3 or higher toxicity that occurred in the first 4 weeks of the patient being on study. Fifteen patients were evaluated and three experienced grade 3 toxicities. However, since they took place at times beyond the 4-week period, they were not noted as dose limiting according to protocol. Had the DLT been defined with respect to 3 months or longer, these toxic outcomes would have affected dose escalation in

the study. And they should affect dose escalation for the study to be applicable: perillyl alcohol is a cytostatic agent that is intended to block further tumor growth (as opposed to cytotoxic agents designed to quickly kill tumor cells). If the agent is shown successful, it is intended to be given to patients over a long term, perhaps years or even permanently. We use a 6-month observation period, which is a practical but useful compromise.

On the other hand, since current designs require complete follow-up of the current patient (or group) before admitting a new patient or group, timeliness may suffer as a result. This is indeed a major criticism of the continual reassessment method (CRM), a phase I design proposed in O'Quigley, Pepe, and Fisher (1990), which otherwise shows promise in yielding a good MTD estimate. The CRM has a more serious timeliness problem than other methods because it takes one patient at a time. As a remedy, Goodman, Zahurak, and Piantadosi (1995) propose a modified CRM that accrues a small group of patients at a time. This certainly cuts down trial duration but can take impractically long periods of time when the DLT corresponds to an extended observation period. In this paper, we propose a method that incorporates the time-to-event of each patient into the CRM (TITE-CRM). Using TITE-CRM, we allow patients to be entered in a staggered fashion and hence shorten the study duration. In addition, Goodman (1992) recommends that participants in chemoprevention studies should be enrolled in a staggered fashion, allowing weeks to pass before escalating to the next dose. While the author did not

provide any specific design suggestions, our proposal appears to achieve the goals he set for conducting a phase I trial.

We review the CRM and describe our method in the following section. In Section 3, an example will be worked out with respect to practical and ethical considerations. An asymptotic result of our method is stated in Section 4. A simulation study that compares the TITE-CRM to the CRM in finite sample situations is described in Section 5. The results are to be presented and discussed in Section 6.

2. Methods

2.1 The CRM

In a typical phase I trial setting, binary responses \mathbf{y} are observed at doses d_1, \dots, d_K . The CRM assumes a parametric model $F(d, \beta)$ to describe the correct model that delineates the stochastic relationship between dose and toxicity. The dose-response model F should be monotone increasing in d and monotone in β ; also, given any p in $(0, 1)$ and d , there should exist β_0 such that $F(d, \beta_0) = p$. Other than these, choice of F is usually made on grounds of convenience subject to a few regularity conditions (see Shen and O'Quigley, 1996). With the first n observations, an efficient estimate $\hat{\beta}_n$ of the model parameter β is computed and the next dose level $[n+1]$ is chosen such that

$$|F(d_{[n+1]}, \hat{\beta}_n) - p| \leq |F(d_k, \hat{\beta}_n) - p| \quad \text{for } k = 1, \dots, K, \quad (1)$$

where p is the target probability of toxicity. Estimation of β can be based on the likelihood

$$L_n(\beta) = \prod_{i=1}^n F(d_{[i]}, \beta)^{y_i} \{1 - F(d_{[i]}, \beta)\}^{1-y_i},$$

where y_i is the indication of toxic response for the i th patient and $[i]$ is the dose level assigned to him. O'Quigley and colleagues (1990) used the Bayes estimator with respect to squared error loss when they first introduced the method. The maximum likelihood estimator was considered in a more recent paper by O'Quigley and Shen (1996).

2.2 The TITE-CRM

We extend the method by considering a weighted dose-response model $G(d, w, \beta)$ that is monotone increasing in w with marginal constraints $G(d, 0, \beta) = 0$ and $G(d, 1, \beta) = F(d, \beta)$ for all d, β . The weight w is a function of the time-to-event of a patient. Motivated by simplicity, we incorporate the weight linearly into F , i.e., $G(d, w, \beta) = wF(d, \beta)$ for $0 \leq w \leq 1$. Dose allocation is facilitated by the marginal model as in (1). Under this modeling framework, we can estimate β at any time based on the weighted likelihood

$$\tilde{L}_n(\beta) = \prod_{i=1}^n G(d_{[i]}, w_{i,n}, \beta)^{y_{i,n}} \{1 - G(d_{[i]}, w_{i,n}, \beta)\}^{1-y_{i,n}},$$

where $y_{i,n}$ and $w_{i,n}$ are, respectively, the indication of toxic response for the i th patient and the weight assigned to this observation just prior to the entry time of the $(n+1)$ th patient.

Using arguments similar to O'Quigley et al. (1990), it follows that the weighted likelihood \tilde{L}_n is a regular likelihood if G is the assumed model. In fact, the model G can be viewed as a failure time regression model. Suppose that each patient is followed for a fixed period of time T and U_i is the time-to-

toxicity of patient i . Then, for $u \leq T$,

$$\begin{aligned} P(U_i \leq u) &= P(U_i \leq u \mid U_i \leq T)P(U_i \leq T) \\ &\equiv w(u; T)F(d_{[i]}, \beta), \end{aligned}$$

i.e., the weight function w is identified with a truncated probability distribution and the dose-response curve F with the marginal model at time T . Consequently, \tilde{L}_n is a likelihood function based on conditionally independent current status data. In this paper, we will focus on the function

$$w(u; T) = \frac{u}{T}, \quad (2)$$

which might appear to be an oversimplified choice but will be shown adequate in many cases via simulation.

2.3 Other Weight Functions

From another point of view, a weight function, typically between zero and one, should reflect the amount of information available from a patient. If a toxic outcome is observed, we should claim full information from this observation by letting $w = 1$; otherwise, define, e.g., $w = w(u; T)$ as in (2). The weight function just defined and that defined in (2) are equivalent in the sense that any weighted likelihood-based estimation resulting from both of them coincide. To verify this, it suffices to realize that \tilde{L}_n is proportional to a function that does not depend on the weights assigned to the toxic outcomes when the two functions differ.

This equivalence property will hold in general if the weight function does not depend on the unknown parameter β . By contrast, this property does not hold when the weight function corresponds to an accelerated failure time model,

$$w(u; T, d, \beta) = \frac{F_o\{\log(u/T)^c - a + \beta d\}}{F_o\{\beta(d - a/\beta)\}}, \quad \beta > 0 \text{ and some } c > 0,$$

for some distribution F_o and constant a . This weight function prescribes a fuller survival data analysis approach and has the advantage of assigning weights adaptively. Alternatively, we may employ some adaptive weighting scheme that does not depend on β , such as

$$w(u; T) = \frac{\kappa}{z+1} + \frac{1}{z+1} \left(\frac{u - t_{(\kappa)}}{t_{(\kappa+1)} - t_{(\kappa)}} \right), \quad (3)$$

where z is the total number of toxic observations, $0 \equiv t_{(0)} < t_{(1)} \leq \dots \leq t_{(z)} < t_{(z+1)} \equiv T$ are the ordered failure times, and $\kappa = \max_{0 \leq j \leq z} \{j : u \geq t_{(j)}\}$. When a toxic outcome is reported at a time later than half of the complete follow-up period, weight function (3) will assign lesser weight for each fixed u than function (2) does. If most toxic responses occur near the end of the follow-up period, much less weight will be given to patients who have entered the trial for a short period of time.

Using any of the above weight functions, the TITE-CRM is an extension of the original CRM that takes place when all patients are censored at time T .

3. Practical Considerations

3.1 Patient Admission

Suppose each patient in a trial remains in the study until the end of follow-up T or a toxic response is reported. The CRM will typically admit one patient every T time units since




Figure 1. Dose allocation of a simulated trial using the TITE-CRM.

not many toxic responses would be expected. By contrast, TITE-CRM allows patients to be entered at considerably shorter intervals. Figure 1 summarizes the observations from a simulated trial with 25 patients. There were six dose levels on trial, and the first patient was entered at level 3 according to prior information. The MTD was defined to be the 20th percentile with respect to a 6-month follow-up period. One patient was admitted every .5 month. Thus, the last patient was accrued at 1 year since the study began bringing the whole trial to an end at the 18th month. More details about the design parameters are given in Section 5. In brief, the trial operated under scenario A1 (see Table 1).

The first escalation occurred at the second month, by which four patients had shown no toxicity. Another four patients were entered at the next higher dose subsequently in the following 2 months. No dose-limiting toxicity was reported from any of the eight patients by the fourth month of the study. But since no patient had yet been completely followed, not much information was provided to lower the dose-response curve enough to recommend higher doses. At the end of the trial, dose-limiting toxicity was reported from four patients, two at level 3 and the other two at level 4. This example illustrates the TITE-CRM as a reasonably cautious dose escalation scheme, even if it compresses the trial duration considerably. Certainly, if the admission schedule is too tight, not much information is available for efficient dose allocation at the initial stage and it might lead to undue assignment at high doses. However, as the simulation results (Section 6) suggest, the TITE-CRM and the CRM cause similar numbers of toxicities on average under certain conditions, even with a substantial compression of trial duration.

3.2 Initial Design

TITE-CRM coupled with maximum likelihood estimation will not be operative until an estimate exists; in a Bayesian paradigm, where the first patient is entered at the lowest dose instead of the *a priori* best dose, the Bayes estimate that always exists may result in aggressive escalation at an early stage of the trial. These two situations demand an initial design that should consist of a dose escalating scheme and a transition rule. A simple and reasonable way is to include

Table 1

The 15 toxicity scenarios in the simulation study. The A scenarios correspond to the uniform model, B to the log-logistic model, and C to the Weibull model.

Scenario	Toxicity probability
A1, B1, C1	.05, .10, .20 , .30, .50, .70
A2, B2, C2	.30 , .40, .52, .61, .76, .87
A3, B3, C3	.05, .06, .08, .11, .19 , .34
A4, B4, C4	.06, .08, .12, .18 , .40, .71
A5, B5, C5	.00, .00, .03, .05, .11, .22

three patients at a time and start at the lowest level, escalate to the next dose only if no toxicity is observed among the previously accrued patients, and switch to TITE-CRM when first toxicity occurs. This initial design was also considered in O'Quigley and Shen (1996).

4. Consistency of the Design

In this section, we state a consistency result for the maximum likelihood TITE-CRM. First, define $H_l = \{\beta \in \Theta : |F(d_l, \beta) - p| < |F(d_k, \beta) - p| \text{ for } k \neq l\}$, where p is the target probability of toxicity and l is the correct level. Note that, if the parameter estimate $\hat{\beta}_n$ resides in H_l , the TITE-CRM will recommend the correct dose. Further, let μ_k be the true (unknown) toxicity probability for dose level k such that $\mu_1 < \dots < \mu_K$ and define ϕ_k such that $\mu_k = F(d_k, \phi_k)$ for $k = 1, \dots, K$.

THEOREM 1. *Assume that the number of incomplete observations is of order $o(n)$, where n is the sample size. If the weight function does not depend on the unknown parameter β and F satisfies certain regularity conditions, then, under condition C1 that all ϕ_k are in H_l , the TITE-CRM will recommend the correct level l eventually with probability one.*

See Cheung and Chappell (1999) for the proof.

5. A Simulation Study

Three versions of the CRM are compared with their TITE-CRM counterparts. Two starting doses, namely the *a priori* best dose and the lowest dose, are considered for the Bayesian CRM (denoted respectively as CRM and B-CRM). The maximum likelihood CRM starts a trial at the lowest level (denoted as B-CRML). The initial design described in Section 3.2 is incorporated into B-CRM and B-CRML. The respective TITE-CRM schemes are denoted as TITE, B-TITE, and B-TITEL.

In each trial, patients were followed for 6 months. Trials conducted by the TITE-CRM admitted a patient every .5 month except that groups of three were admitted at 6-month intervals in the initial design in B-TITE and B-TITEL. All CRM schemes entered a new patient (or group) every 6 months. Sample sizes considered were 25 and 48. There were six dose levels on trial and the target probability of toxicity was .20.

The dose-response model and the weighted model in the simulation were, respectively, $F(d, \beta) = d^\beta$ and $G(d, w, \beta) = w F(d, \beta)$, with $w = w(u; T) = u/T$. The prior distribution on the model parameter β is $\Lambda(\beta) = 1 - \exp(-\beta)$ for $\beta > 0$. The prior belief on the toxicity probabilities were

Table 2

Proportion of correct recommendations and expected trial duration (in years) based on 1000 simulated trials. A correct dose is defined to be one with toxicity probability closest to .2.

<i>n</i>	Schemes	Failure time	Toxicity configuration					Expected duration
			1	2	3	4	5	
25	TITE	Uniform	.50	.92	.51	.61	.29	1.5
	TITE	Log logistic	.47	.92	.49	.60	.28	1.5
	TITE	Weibull	.48	.90	.58	.60	.28	1.5
	CRM	—	.48	.92	.59	.63	.31	12.5
	B-TITE	Uniform	.48	.93	.35	.50	.32	2.1–3.6
	B-TITE	Log logistic	.48	.93	.35	.48	.31	2.0–3.6
	B-TITE	Weibull	.47	.92	.41	.54	.30	2.1–3.6
	B-CRM	—	.48	.92	.42	.54	.29	6.8–11.0
	B-TITEL	Uniform	.47	.95	.33	.47	.53	2.0–3.6
	B-TITEL	Log logistic	.47	.94	.30	.44	.47	2.0–3.6
	B-TITEL	Weibull	.45	.94	.35	.50	.54	2.1–3.6
	B-CRML	—	.45	.94	.40	.49	.52	6.8–11.0
48	TITE	Uniform	.63	.98	.61	.74	.33	2.5
	TITE	Log logistic	.63	.99	.62	.72	.31	2.5
	TITE	Weibull	.63	.97	.70	.77	.36	2.5
	CRM	—	.62	.98	.68	.75	.39	24
	B-TITE	Uniform	.64	.98	.50	.67	.39	3.0–4.6
	B-TITE	Log logistic	.62	.98	.51	.67	.37	3.0–4.6
	B-TITE	Weibull	.62	.99	.51	.68	.38	3.0–4.6
	B-CRM	—	.62	.98	.51	.67	.36	18.3–22.5
	B-TITEL	Uniform	.62	.99	.46	.64	.57	3.0–4.6
	B-TITEL	Log logistic	.61	.99	.49	.62	.58	3.0–4.6
	B-TITEL	Weibull	.60	.99	.49	.64	.59	3.0–4.6
	B-CRML	—	.61	.99	.49	.66	.56	18.3–22.5

(.05, .10, .20, .30, .50, .70). After substitution, the six doses were scaled accordingly as $d_1 = .05$, $d_2 = .10$, $d_3 = .20$, $d_4 = .30$, $d_5 = .50$, $d_6 = .70$.

Five toxicity probability configurations used in O'Quigley et al. (1990) were used in our simulation (see Table 1). For the TITE-CRM, the patients' failure times were generated under three models: (A) a conditionally uniform model, (B) a log-logistic model, and (C) a Weibull model. According to the conditional model, we first determined if a patient had a toxic response; if so, we would generate a number on the interval $(0, 6)$ uniformly as the failure time. As for the log-logistic model, the scale parameter was set to be one and the location parameters were chosen correspondingly so that the cumulative distribution function at $T = 6$ would be the toxicity probability. The scale parameters of the Weibull model (with a fixed shape parameter 4) were chosen in a similar manner. Scenarios A1 and A2 represent situations where the power model is correct (with true β being 1 and .4, respectively). Each scenario has 1000 simulation replicates.

6. Results and Discussion

The proportions of correct recommendation are shown in Table 2. Each TITE scheme produces results comparable to its CRM counterpart (i.e., CRM versus TITE, B-CRM versus B-TITE, B-CRML versus B-TITEL), except in configuration 3,

where TITE-CRM's accuracy is slightly worse than CRM's. Also, the proportions of correct recommendations are similar among trials generated under various failure time distributions. This is not surprising because the estimation of β at the end of each trial is based on the marginal model F rather than the weighted model G . The final recommendation, therefore, depends on the failure time distribution only through its effect on in-trial allocation. Simulation was also performed using the adaptive weight function (3). Similar results were obtained and hence we do not report them here.

Our method is robust to the model assumptions. For example, the accuracy in recommending the correct dose in scenario A4 is comparable with that in A1. On the other hand, with sample size 25, the Bayesian schemes are correct only for about 30% in configuration 5 while the maximum likelihood approaches yield about 50% correct recommendations. This is due to the inflexibility of the prior distribution—to recommend dose level 6, β needs to be greater than 3.35 while $P(\beta > 3.35) < .04$ under an $\exp(1)$ prior. This suggests we should consider a vague prior in order to ensure flexibility. Another potential problem of the power model used in the simulation can be posited using an asymptotic argument found in Shen and O'Quigley (1996). In brief, under toxicity configuration 5, this model violates condition C1 in Theorem 1 in Section 4.




Figure 2. Expected proportion of toxicity by dosage range.

Table 2 also shows the expected duration of trials by different schemes. Trials by CRM and TITE take a fixed amount of time to complete, while trials by the other four have random lengths depending on the toxicity configuration. However, the worst case of the TITE-CRM takes significantly less time than the best one of the CRM. If the trials were conducted by a group accrual CRM (group size 3) as discussed in Goodman et al. (1995), the trial duration would have been about 4 and 8 years, respectively, for $n = 25, 48$. It certainly saves a considerable amount of time, but Table 2 shows that the TITE-CRM is more impressive in duration reduction.

Figure 2 presents the expected proportion of toxicity by dosage range. The TITE schemes are generally more toxic than their CRM counterparts, but the degree of conservativeness is reasonable for B-TITE and B-TITEL. TITE appears to be aggressive, especially with Weibull failure time. The problem, as shown in the concerns of the associate editor, is that, since most toxicities occur near the end of the 6-month follow-up period under the Weibull model, erroneous escalation may result before any toxicities are seen. In cases like this, we may consider some other weight function (e.g., equation (3)) or adopt conservative modifications (cf., B-TITE and B-TITEL). More often than not, acute toxicities as well as late-onset tox-

icities are possible outcomes in phase I trials. This situation is roughly reflected by the log-logistic failure time model, under which the TITE-CRM is almost as safe as the CRM. It might appear otherwise under configuration 2; however, in the presence of preclinical information and the clinician's intuition and generally cautious inclination, configuration 2 does not reflect practical situations in general.

In conclusion, our method shortens trial duration significantly without shortchanging estimation accuracy. With trial safety taken into account, two-stage designs may be preferred. Another advantage is the flexibility our method allows in patient admission. In cases where delayed treatment to patients who have consented is unethical, the TITE-CRM can assign the best dose to the next patient at any time given all information available. Last, the TITE approach can be attached to any phase I method that involves likelihood-based estimation (e.g., escalation with overdose control; see Babb et al., 1998) and provides an extension to a design as investigators desire.

ACKNOWLEDGEMENTS

We would like to thank the associate editor and a referee, whose comments and suggestions led to an improved paper.

RÉSUMÉ

Les schémas traditionnels dans les essais cliniques de phase I exigent que chaque patient (ou un petit groupe de patients) soit complètement suivi avant que le patient (ou le groupe de patients) suivant puisse être inclus. Dans des situations telles que l'évaluation d'effets tardifs des radiations, ou des toxicités retardées de traitement de chimio-prévention, cette règle peut entraîner des durées d'essais tellement longues qu'il est impossible de les réaliser. Nous proposons une nouvelle méthode, appelée méthode continue de réévaluation du délai jusqu'à apparition de l'événement (TITE-CRM), qui permet l'inclusion des patients de façon échelonnée. Cette méthode est une extension de la méthode CRM [O'Quigley, Pepe, and Fisher, 1990, *Biometrics* **46**, 33–48]. Cette approche utilisant le "délai jusqu'à toxicité" peut être de plus étendue à d'autres schémas d'études de toxicité à court terme. Nous démontrons que la dose recommandée établie par la méthode TITE-CRM converge vers la bonne valeur sous certaines conditions. Une étude par simulations montre que la précision et la sécurité de cette méthode sont comparables à la méthode CRM, alors même que cette première permet une importante réduction de la durée de l'essai. Un essai qui pourrait durer jusqu'à 12 ans, pour être mené à son terme par la méthode CRM, pourrait être réduit à une durée de 2 à 4 ans avec notre méthode.

REFERENCES

- Babb, J., Rogatko, A., and Zacks, S. (1998). Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.
- Cheung, Y. K. and Chappell, R. (1999). Sequential designs for phase I clinical trials with late-onset toxicities. Technical Report 134, Department of Biostatistics and Medical Informatics, University of Wisconsin, Madison.
- Durham, S. D., Flournoy, N., and Rosenberger, W. F. (1997). A random walk rule for phase I clinical trials. *Biometrics* **53**, 745–760.
- Goodman, G. E. (1992). The clinical evaluation of cancer chemoprevention agents: Defining and contrasting phase I, II and III objectives. *Cancer Research* **52**(Suppl.), 2752s–2757s.
- Goodman, S. N., Zahurak, M. L., and Piantadosi, S. (1995). Some practical improvements in the continual reassessment method for phase I studies. *Statistics in Medicine* **14**, 1149–1161.
- O'Quigley, J. and Shen, L. Z. (1996). Continual reassessment method: A likelihood approach. *Biometrics* **52**, 673–684.
- O'Quigley, J., Pepe, M., and Fisher, L. (1990). Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33–48.
- Shen, L. Z. and O'Quigley, J. (1996). Consistency of continual reassessment method under model misspecification. *Biometrika* **83**, 395–405.
- Storer, B. and DeMets, D. (1987). Current phase I/II designs: Are they adequate? *Journal of Clinical Research and Drug Development* **1**, 121–130.
- Whitehead, J. and Brunier, H. (1995). Bayesian decision procedures for dose determining experiments. *Statistics in Medicine* **14**, 885–893.

Received July 1999. Revised May 2000.

Accepted June 2000.