Contemporary Clinical Trials 44 (2015) 56–63

Contents lists available at ScienceDirect

Contemporary Clinical Trials

journal homepage: www.elsevier.com/locate/conclintrial

Bayesian dose escalation in oncology with sharing of information between patient populations

Kristine R. Broglio [a,*], Larissa Sandalic [b], Tina Albertson [b], Scott M. Berry [a]

[a] Berry Consultants, LLC, 4301 Westbank Drive Suite 140 Bldg B, Austin, TX 78746, United States
[b] Seattle Genetics, Inc., 21823 30th Drive S.E., Bothell, WA 98021, United States

ARTICLE INFO

Article history:
Received 17 April 2015
Received in revised form 1 July 2015
Accepted 3 July 2015
Available online 11 July 2015

Keywords:
Maximum tolerated dose
Pediatrics
Continual reassessment method
Phase I
Bayesian
Adaptive clinical trial

ABSTRACT

We present a Phase I dose escalation trial design based on a modified continual reassessment method that allows for sharing of information between populations. We describe our approach in the context of a trial for patients with acute lymphoblastic leukemia (ALL) that is currently being conducted. The ALL trial enrolls both adult and pediatric patient populations. Dose escalation and the determination of the maximum tolerated dose (MTD) are performed separately for each population, but to increase efficiency, information about the dose–toxicity curve is shared. Dose escalation rules allow pediatric patients to skip dose levels provided safety has been shown in adults and the dose level is estimated to be safe for pediatric patients. Trial objectives are to efficiently determine the MTD for each population and to minimize the number of pediatric patients required for dose escalation.

1. Introduction

The maximum tolerated dose (MTD) determined in a Phase 1 study is partly dependent on the characteristics of the patients enrolled. The dose–toxicity relationship may vary by population, defined for example by age or prior extent of treatment, yet this variability is rarely explored in Phase 1 studies. Trials that enroll a heterogeneous patient population may estimate the dose–toxicity relationship for the overall population, but may not accurately estimate the MTD for particularly low- or high-risk subgroups [1,2]. Sample sizes tend to be small, particularly at dose levels below the MTD, and so post hoc analyses of toxicity by patient or disease characteristics are limited. On the other hand, Phase I trials in a targeted population do not capture the expected toxicity for other populations that may receive the drug in the future [3,4].

Conducting a series of Phase I trials across different targeted patient populations to determine an MTD for each is not efficient, as this approach ignores the fact that the dose–toxicity relationship in one population provides information about the dose–toxicity relationship in another population [5,6]. In addition, some diseases are rare and it is not feasible to enroll a separate Phase I trial. In this situation, data available from one population may dictate the MTD for a more rare population. A middle ground is to enroll multiple patient populations in the same trial and allow for sharing of information [5].

Including the pediatric patient population in Phase 1 studies during drug development is particularly challenging [7–12]. Historically, the schedule of administration and MTD determined for adults has been extrapolated to pediatric patients based on weight or body surface area [7,13,14]. Yet, the pediatric population differs also in pharmacokinetics and metabolism, generally resulting in greater tolerance than adults [7,13,15]. For pediatric diseases, the FDA and EMA have both encouraged earlier testing in the pediatric population [9,12,16], however, due to logistical challenges as well as the rarity of pediatric oncology patients, enrollment in independent dose escalation studies may not be feasible.

We present a Phase I dose escalation trial design for both adult and pediatric patients in oncology. The primary objective is to determine the MTD for each patient population. We jointly model the dose–toxicity relationship between adults and pediatrics in order to share information between the two populations. Dose escalation rules allow pediatric patients to skip dose levels when safety has been demonstrated in the adult population. This trial allows open enrollment such that patients may be enrolled as they become available rather than enrolling in dosing cohorts. Finally, the trial is continuously monitored for safety and for early success in characterizing the MTD.

This trial was designed to enroll adult and pediatric patients with acute lymphoblastic leukemia (ALL). Therefore, we present the design in this context. We describe a single simulated trial as an example of the behavior of the design and then summarize the trial's operating

* Corresponding author.
E-mail addresses: kristine@berryconsultants.com (K.R. Broglio), lsandalic@seagen.com (L. Sandalic), tina.albertson@gmail.com (T. Albertson), scott@berryconsultants.com (S.M. Berry).

http://dx.doi.org/10.1016/j.cct.2015.07.003
1551-7144/

characteristics. We compare our design to a CRM conducted in each population independently.

2. Methods

This trial includes 7 dose levels. The maximum sample size is 60 patients including 35 adult and 25 pediatric patients. Dose limiting toxicities (DLTs) are assessed throughout the first 3-week cycle of therapy. We assume that adults will be accrued at the rate of 3 patients per month and that pediatric patients will be accrued at the rate of 1 patient per month.

Dose escalation is conducted according to a modified CRM [17,18]. The CRM jointly models both adult and pediatric populations [1,2]. As each patient's DLT status becomes known (yes/no), the distributions of all parameters in the CRM, those for adult and pediatric populations, are updated. The next dose level is assigned for each population based on the posterior probability of DLT at each dose level. We further customize the CRM with enrollment and dose escalation rules.

2.1. Dose–toxicity model

Dose levels are the same for both populations. Pediatric patients are expected to have lower levels of toxicity at each dose level as compared to adults but the shape (i.e. slope) of the dose–toxicity curve is expected to be the same [10,19]. Therefore, the dose–toxicity model assumes a common slope between adults and pediatrics, but allows for a different intercept. For each dose $d = 1,\ldots,7$ the probability of a dose limiting toxicity (DLT) is π_d. We model the log odds of π_d,

$$\theta_d = \log\left(\frac{\pi_d}{1-\pi_d}\right)$$

in the adult population with a two parameter model,

$$\theta_d = \alpha + \beta d.$$

The pediatric population is modeled with an additive effect in the intercept as

$$\theta_d = (\alpha + \alpha_{ped}) + \beta d.$$

We place the following independent prior distributions on the parameters

$$\alpha \sim \mathrm{N}\left(-3,\ 1^2\right)$$

$$\alpha_{ped} \sim \mathrm{N}\left(0,\ 1^2\right)$$

$$\beta \sim \mathrm{N}\left(0.5,\ 0.2^2\right).$$

The prior for the pediatric intercept term is centered at 0, and so does not assume that the shift will be either positive or negative. In addition to the above priors, we use "pseudo-prior" weighting to impose the appropriate weight in the prior distribution. We assume ½ an observation on the highest and lowest dose levels, with 0 DLTs on the lowest dose level and ½ DLT on the highest dose level.

2.2. Interim monitoring

Each population is continuously monitored for safety and for success in characterizing the MTD. Monitoring is performed independently in each population. Thus, dose escalation in one population may be stopped while it continues in the other. If dose escalation in a population is not stopped for either safety or success in identifying the MTD it will continue to enroll to the maximum sample size.

A dose level is considered safe if there is at least a 50% probability that the DLT rate is less than 30%,

$$\Pr(\pi_d<30\%)>50\%.$$

Safety is monitored continuously and patients may never be assigned to a dose level that is unsafe by this definition. If no dose levels are safe, dose escalation for that population will stop and no MTD will be declared. Formally if

$$\Pr(\pi_d<30\%)<50\% \ \text{ for all } \ d = 1,\ldots,7.$$

We define the MTD as the highest safe dose, where safe is defined as above. Dose escalation may be stopped early when the MTD has been sufficiently characterized [18]. We define this with two conditions. First, a minimum number of patients must be treated at or around the MTD. In the adult population, at least 8 patients must be treated at the MTD and at least 20 patients treated within one dose level (lower or higher) of the MTD. In the pediatric population, at least 5 patients must be treated at the MTD and at least 10 patients treated within one dose level of the MTD.

The second condition is to ensure that the selection of the MTD is robust and that a higher dose level should not still be considered. For this, we assume three additional hypothetical patients treated at the current MTD with no DLT. If, with this additional information, an update of the CRM would not recommend escalation, the MTD is considered well characterized.

2.3. Rules governing dose escalation

We customize the dose escalation recommended by the underlying dose–toxicity model with additional rules to ensure patient safety. In either population, no patients may be assigned to dose levels that are considered unsafe. Additionally, a minimum amount of observed DLT information at each dose level is required to escalate to the next higher level. Specifically, at the lower dose levels (1 to 3) complete DLT information from at least 2 patients is required prior to escalation, and at the higher dose levels (4 to 6) complete DLT information is required from at least 3 patients.

The adult population may escalate as described above and no dose levels may be skipped. However, for the pediatric population, escalation may occur based on information from patients enrolled in either population. Pediatric dose escalation can proceed as described for adults above based on information observed only within the pediatric population. Alternatively, the pediatric population may escalate to a higher dose, possibly skipping dose levels. This is allowed when there is complete DLT information from at least two adults at that dose level and the dose level above, and both are considered safe for adults.

2.4. Rules governing enrollment

The trial must begin with enrollment of adult patients to dose level 1. The ALL trial is the first time the compound will be given to humans; therefore, to proceed cautiously, the first 6 adult patients are enrolled in cohorts of 2 such that complete DLT information must be available on each cohort before the next patient can be enrolled. After these first 6 patients, there is open enrollment into the study such that patients can be enrolled as they become available.

Open enrollment has advantages over cohort enrollment in that no pause is required while DLT information is collected on all patients in the cohort. This may shorten total trial duration [6,10]. However, if the speed of accrual is too rapid, lower dose levels could enroll a large number of patients and the maximum sample size could be reached prior to observing adequate DLT information to escalate to the MTD. Additionally, rapid enrollment could result in many patients being treated at a dose level that is later determined to be unsafe. Therefore, to balance