Original Article

Implementation of a model-based design in a phase 1b study of combined targeted agents

Nolan A. Wages[1*], Craig A. Portell[2], Michael E. Williams[2], Mark R. Conaway[1], Gina R. Petroni[1]

[1]Division of Translational Research & Applied Statistics, Department of Public Health Sciences, University of Virginia, Charlottesville, VA, USA
[2]Division of Hematology / Oncology, University of Virginia, Charlottesville, VA, USA

Running title: Phase 1b design for combined targeted therapies

Key words: Targeted therapy, Clinical trials, Phase 1b, Combination, Lymphoma

Correspondence*:
Nolan A. Wages, Ph.D.
Division of Translational Research & Applied Statistics
University of Virginia
Public Health Sciences
P.O. Box 800717
Charlottesville, Virginia, USA
(434) 924-8222
nwages@virginia.edu

Disclosure: The authors declare no potential conflicts of interest.

Word count: 4419

Number of figures: 0

Number of tables: 5

1

Original Article

Abstract:

In recent years, investigators have recognized the rigidity of single agent, safety only, traditional designs, rendering them ineffective for conducting contemporary early-phase clinical trials, such as those involving combinations and /or biological agents. Novel approaches are required to address these research questions, such as those posed in trials involving targeted therapies. We describe the implementation of a model-based design for identifying an optimal treatment combination, defined by low toxicity and high efficacy, in an early-phase trial evaluating a combination of two oral targeted inhibitors in relapsed / refractory mantle cell lymphoma. Operating characteristics demonstrate the ability of the method to effectively recommend optimal combinations in a high percentage of trials with reasonable sample sizes. The proposed design is a practical, early-phase, adaptive method for use with combined targeted therapies. This design can be applied more broadly to early-phase combination studies, as it was used in an ongoing study of a melanoma helper peptide vaccine plus novel adjuvant combinations.

2

Original Article

INTRODUCTION

Historically, in single agent dose finding oncology studies, the main objective has been to identify the maximum tolerated dose (MTD) among a range of predefined dose levels. Each increasing dose level is associated with an assumed increasing probability of dose-limiting toxicity (DLT). The underlying assumption is that the MTD is the highest dose that satisfies some safety requirement, and therefore provides the most promising outlook for efficacy; thus, decisions about which dose to recommend for further research were based upon safety outcomes. In this framework, dose finding trials do not incorporate an efficacy endpoint in the decision process. Two recent articles in *Clinical Cancer Research* (1, 2), among others, have highlighted the fact that the current landscape of oncology drug development is challenging the traditional "MTD-based" approach to early-phase clinical trials. Sachs et al. (1) provides evidence that determination of Phase II recommended doses using MTD-based approaches have resulted in inappropriate dosing for some therapies, such as targeted agents. The authors contend that, in addition to a DLT endpoint, dose finding strategies should incorporate an "effect marker," with the goal of locating an effective dose. Examples of effect markers include an early measure of efficacy (e.g. clinical response), pharmacokinetic/pharmacodynamics, biological targets (e.g. immune response or binding/inhibition of therapeutic targets), and more. Figure 1 in Sachs et al. (1) illustrates that doses below the MTD are being approved, indicating that targeting the MTD may not be the appropriate primary trial objective. Nie et al. (2) argue that the traditional 3+3 design is inadequate for meeting the objectives of studies involving targeted therapies, and called for wider use of more innovative methods with the goal of answering more complex research questions. The statistical and medical literature is full of reviews, justification, and recommendations on the use of novel designs (2-4). Contemporary dose-finding problems have created the need to adapt early-phase trial design to include additional endpoints in the decision process, thereby conducting Phase I/II studies. For a more comprehensive study of the Phase I/II paradigm, we refer the reader to a recent textbook by Yuan, Nguyen, and Thall (5).

3

Original Article

Another challenge presented in contemporary early-phase trials is in the design and conduct of combination studies. It is reasonable to assume that toxicity increases with dose in a single-agent setting, but it may be difficult to characterize the toxicity relationship between some of the combinations being tested. One approach to this problem is to reduce the two-dimensions to one dimension by pre-specifying an escalation path in which the toxicity ordering is completely known between combinations, and applying the single agent traditional 3+3 method along this path. However, such an approach can potentially miss more promising dose combinations located outside the path. Increasing the complexity, several combinations may have an acceptable toxicity profile that meet the definition of the MTD combination, so an efficacy or activity endpoint would distinguish one combination as the optimal dose for that combination. Several recent Phase I/II methods allow for the assessment of both toxicity and efficacy in drug combination studies (6-10).

In this article, we present a model-based, early-phase design for combining two targeted agents that accounts for both safety and efficacy in order to identify an optimal dose for the combination. The statistical modeling structure is outlined in Wages and Conaway (11). We describe the implementation of the method in an ongoing, multi-institution trial (NCT02419560) designed at the University of Virginia (UVA) Cancer Center studying venetoclax (ABT-199) in combination with ibrutinib for the treatment of relapsed or refractory mantle cell lymphoma. Recently published recommendations (12, 13) for conducting novel early-phase methods were adhered to in implementing the described design.

METHODS

The study described is a multi-institution phase 1b evaluation of the safety and efficacy of two dose levels of venetoclax (ABT-199) in combination with three dose levels of ibrutinib, as shown in Table 1. As venetoclax has been associated with significant tumor lysis when initiating therapy at maximum dose, careful dose escalation of venetoclax is required prior to reaching

4

Original Article

the assigned dose level (14). Thus, all participants start treatment with 1 week of venetoclax monotherapy before starting the allocated ibrutinib dose in combination. Venetoclax is then dose escalated to the assigned dose level. Treatment combinations are grouped into toxicity "zones" (1, 2, 3 or 4), based on the dose level of each agent in the treatment combination. The trial was designed to find the optimal dose combination, defined as a combination estimated to have acceptable toxicity and a good response profile. An adaptive design is being used to guide accrual decisions with toxicity and efficacy assessments characterizing the main decision measures. The decision endpoints are dose-limiting toxicities (DLT) and efficacy, as measured by the attainment of a partial or complete response 2 months post start of treatment.

In monitoring safety, adverse events (AE) are being assessed and acute toxicity graded using the National Cancer Institute (NCI) Common Terminology Criteria (CTCAE) Version 4.03. A participant is classified as experiencing a DLT (yes/no) based on protocol-specified criteria. In this study, a DLT is defined as any unexpected adverse event that is possibly, probably, or definitely related to treatment and meets the following criteria; (i) Any non-hematological toxicity Grade ≥ 3, except for alopecia, fatigue, and nausea uncontrolled by medical management (ii) Grade 4 neutropenia lasting more than 5 days, (iii) febrile neutropenia of any duration, (iv) Grade 4 thrombocytopenia, Grade 3 thrombocytopenia with bleeding, or any requirement for platelet transfusion, or (v) Grade 4 anemia, unexplained by underlying disease. Efficacy is defined by response to treatment, which is determined using criteria modified from Cheson et al. (15). An important consideration is that the result of the efficacy response measure must be available in a reasonable timeframe if it is to be useful in guiding trial enrollment. Thus, it is important to select an efficacy endpoint that occurs early enough to be meaningful and to design processes for collecting data rapidly so that it may guide participant enrollment in accord with the study design. If the efficacy endpoint occurs much later than the toxicity endpoint and/or accrual is rapid, the trial could be exposed to higher amounts of missing data. In this case, methods that account for delayed outcomes in the modeling procedure could be explored (16,

5

Original Article

17).   As data for each endpoint accumulate, each participant is classified as experiencing a

DLT (yes/no) and experiencing a response (yes/no). Based on the expectedness of adverse

events, the DLT tolerance level was chosen to be 25% (i.e., any optimal combination that we

are satisfied has an estimated DLT probability ≤25% to be considered "acceptable" in terms of

safety).

Model-based estimation

Model-based allocation is being based upon a continual reassessment method (CRM) (18) that

accounts for two binary endpoints (DLT, efficacy) in combinations of agents (11). Safety

assessments are based on the assumption that, as the dose of one agent remains fixed and the

dose of the other agent increases, the probability of DLT is increasing. In other words, toxicity

increases across rows and up columns of Table 1. It is reasonable to assume that combinations

in higher zones have higher probabilities of DLT than combinations in lower zones. It is

unknown whether combinations have higher or lower DLT probabilities than other combinations

within the same zone. It could be that B < C or C < B in terms of their respective DLT

probabilities. This uncertainty is expressed through specification of the multiple one-parameter

models in Table 2 that reflect different orderings of the DLT probabilities. Model selection

techniques are used to choose the model most consistent with the observed data. A common

model choice in the CRM is to raise a set of initial DLT probability estimates, also referred to as

the 'skeleton' of the model, to a power θ_m that is a parameter to be estimated by the data,

where m indexes the skeleton. For each possible DLT ordering, $m = 1, \dots, 5$ in Table 2, the DLT

probabilities are modeled via a one-parameter power model $\Pr(\text{DLT at combination i}) \approx p_{mi}^{\theta_m}$,

where the p_{mi} are the skeleton values for ordering m also given in Table 2. The skeleton values

for each model were generated using the algorithm of Lee and Cheung (19). Using the available

toxicity data, the CRM is fit for each DLT probability working model, and the parameter θ_m is

6

Original Article

estimated for each ordering by maximum likelihood estimation, where the likelihood is given by

$\prod_{i=1}^{6}\left(p_{mi}^{\theta_m}\right)^{y_i}\left(1-p_{mi}^{\theta_m}\right)^{n_i-y_i}$ where y_i = the number of DLTs and n_i = the number of participants

evaluated for DLT on combination i. The working model with the largest likelihood is chosen

and, using the selected model, DLT probability estimates are updated for each combination. If

there is a tie between the highest likelihood values of two or more models, then the selected

model is randomly chosen from among those with tied likelihood values.

The working models for efficacy probabilities are formulated under three different

assumptions (i) the probabilities are increasing with increasing zone, (ii) the probabilities

increase initially and then plateau after a certain dose of ibrutinib, or (iii) the probabilities

increase initially and then plateau after a certain zone, as displayed in Tables 3 and 4. Like

toxicity, these possible shapes for the combination-efficacy curve are expressed through

multiple skeletons of CRM models. We again rely on a class of one-parameter power models,

indexed by k, and the algorithm of Lee and Cheung (19) to formulate working models for the

efficacy probabilities. For each efficacy working model, $k = 1, \dots, 10$, in Tables 3 and 4, the

efficacy probabilities are modeled via a one-parameter power model

P r(response at combination i) $\approx q_{ki}^{\beta_k}$, where the q_{ki} are the skeleton values for order k. Using

the accumulated efficacy data, the CRM is fit for each efficacy probability working model, and

the parameter β_k is estimated for each model by maximum likelihood estimation, where the

likelihood is given by $\prod_{i=1}^{6}\left(q_{ki}^{\beta_k}\right)^{z_i}\left(1-q_{ki}^{\beta_k}\right)^{n_i-z_i}$ where z_i = the number of responses and n_i =

the number of participants evaluated for efficacy on combination i. Again, the working model

with the largest likelihood is chosen and, using the selected model, efficacy probability

estimates are updated for each combination. We make allocation decisions based on the

probability estimates for both DLT and efficacy.

For combination B-F, a two-sided 80% confidence interval is calculated using the

estimated DLT probability for that combination, based on confidence interval estimation for CRM

Original Article

models (20). If the lower bound of this confidence interval exceeds the maximum toxicity tolerance of 25%, then this combination is deemed too toxic and excluded from the acceptable set of combinations. If combination A is excluded from the acceptable set then no combination is considered acceptable and the trial is stopped for safety. Therefore, for combination A the level of confidence is set at 90% instead of 80%.

In sequentially obtaining model-based dose recommendations, an estimated probability of efficacy will be calculated for each combination in the acceptable set. The recommended combination will be based upon how many participants have been entered into the study to that point. For the first third of the trial (i.e. 1/3 the maximum sample size), the combination recommendation is based on randomization using a weighted allocation scheme. Randomization prevents the design from getting "stuck" at a sub-optimal combination based on limited data (21). The recommended combination for the next entered participant is chosen at random from the "acceptable" combinations, with each acceptable combination weighted by its estimated efficacy probability. That is, acceptable combinations with higher estimated efficacy probabilities have a higher chance of being randomly chosen as the next recommended combination. For the latter two-thirds of the trial (i.e. final 2/3 of maximum sample size), the recommended combination for the next entered participant is defined as the "acceptable" combination with the highest estimated efficacy probability.

Stopping the trial

Accrual to the study will be halted and trigger a safety review by the study investigators and the Data and Safety Monitoring Committee to determine if the study should be modified, or permanently closed to further accrual according to the following: (i) accrual will be halted for safety if the first three entered participants in Zone 1 experience a DLT, (ii) if at any point in Stage 2, the set of acceptable combinations is empty and no combinations are considered safe, the trial will stop for safety, (iii) otherwise, accrual to the study will end if the recommendation is

8

Original Article

to assign the next participant to a combination that already has 10 participants treated at that combination or the pre-specified maximum sample size of 28 eligible patients has been reached.

Sample size and accrual

Maximum target sample size for the optimal combination is based upon acquiring sufficient information to assess the objective of estimating efficacy rates while satisfying safety conditions, assuming at least one optimal combination has been found. Based upon simulation results, 10 eligible participants treated at the optimal combination will provide adequate data to assess efficacy. The target of 10 participants at the optimal combination was chosen based on having sufficient information to determine if the optimal combination shows an improved 2-month PR rate compared to the single-agent ibrutinib rate which was estimated at 34% (95% CI: 25 to 44%) (22). For this study the optimal combination would be considered promising with an observed 2-month PR of at least 70% (7/10), a doubling, which results in the lower limit of a one-sided 90% CI exceeding the 44%. Total study sample size is estimated from the simulations, and is determined by the stopping rules in the 'Stopping the trial' section. The maximum total sample size was set at to 28 eligible participants; however, as indicated in the simulation results the maximum average trial size over all scenarios is closer to 20 participants.

RESULTS

Accrual to combinations occurs in two stages. Within the modeling framework described in this article, an initial stage of escalation is needed to get the trial underway. Model-based estimates for outcome probabilities do not exist until some heterogeneity in the data for each endpoint has been observed (23). That is, we need at least one DLT and one non-DLT to estimate the DLT probabilities, and we need at least one response and one non-response to estimate efficacy probabilities at each combination. The initial stage accrued eligible participants

9

Original Article

in cohorts of one on each combination, until a participant experienced a DLT. The role of stage 1 is to have a "dose-escalation" beginning to the study in order to test the safety of "lower" combinations that are assumed to be less toxic. Also, the safety data are available slightly faster than are the efficacy data, so Stage 1 allows allocation to combinations based only on a toxicity endpoint, while waiting for the first few efficacy response observations to be observed. The second stage is allocating eligible participants in cohorts of one according to the procedure described in the 'Model-based estimation' section.

Allocation in completed stage 1

The escalation plan for the first stage was based on the zones. With this design, participants could be accrued and assigned to other open combinations within a zone but escalation would not occur outside the zone until the minimum required time frame has elapsed for the first participant accrued to combination A. The minimum follow-up period for determination of escalation between Zones was 6 weeks from the start of cycle 1. If the minimum follow-up period is not satisfied at the time a new participant is ready to be put on-study, then the participant may be accrued to any arm within the highest zone being assessed, by random allocation, with the intent of minimizing halts to accrual and trial duration. Initial allocation within a zone was based upon random allocation (1:1) between the possible combinations. Escalation to a higher zone occurred only when all combinations in the lower zone had been tried, and no DLT had been observed. Participant allocation to subsequent combinations within the new zone followed the same accrual strategy. This allocation strategy was followed for accrual to increasing zones until a participant experienced a DLT or a stopping rule was triggered. Accrual began at UVA in September 2015 and was slow to initiate. The 1st participant was given combination A, and he/she did not experience a DLT. The 2nd participant was randomized in January 2016 to receive a combination in Zone 2 (B or C). The randomization resulted in this participant receiving combination B, and this patient did not have a DLT. The 3rd participant filled

10

Original Article

the remaining combination in Zone 2 beginning in February 2016, and did not have a DLT on combination C. The study then was opened at other sites, and the next participant was accrued to the study in May 2016, which allowed for the 6-week DLT and 2-month efficacy window to be assessed for the first three participants. At this time, the randomization strategy continued into Zone 3, with the 4th participant receiving combination D, and no DLT being observed. The 5th participant was accrued to the study in early June 2016, and he/she experienced the first DLT on Combination E, at which point the second stage allocation strategy using multidimensional CRM modeling began (24). While awaiting outcomes for the minimum required follow-up period for Zone 3, two additional participants (6th and 7th) were enrolled and randomly assigned to combination D and E, respectively, in July 2016.

Allocation in ongoing stage 2

Stage 2 is allocating eligible participants based upon the multidimensional CRM modeling approach described in the 'Model-based estimation' section. Model-based estimation of DLT probabilities began for the accrual of the 8th participant to the study. After each new accrual in Stage 2, the estimated DLT probabilities are being updated and used to define a set of "acceptable" combinations in terms of safety. If the minimum follow-up period for participants already on-study is not satisfied at the time a new participant is ready to be put on-study, then the participant may be accrued to any combination, by random allocation, which has accrued at least one participant and is in the acceptable set.

Model-based estimation of efficacy probabilities began at the beginning of Stage 2, since at least one efficacy response and one non-efficacy response was observed in Stage 1. All participants on arms A, B, and C achieved at least a partial response, and one non-response occurred on arm E (24). If no responses had been observed in Stage 1, patients would have been randomized to acceptable arms until a response was observed. At the time of combination allocation for the next participant, model-based estimates are calculated for both DLT and

11

Original Article

efficacy probabilities using the available observed data from all participants accrued to the study at that time. For instance, for the first participant accrued in Stage 2 (8[th] participant), estimates for the DLT and efficacy probabilities were based on available data from the first 5 patients accrued to the study, since participants 6 and 7 were allocated to arm E and D, respectively, while waiting for the minimum follow-up period in Zone 3. Accrual of the 8[th] participant occurred at the end of July 2016, at which time complete DLT and response data were available for each of the first 5 participants. For arms {A, B, C, D, E, F}, the DLT data at this point was {0/1, 0/1, 0/1, 0/1, 1/1, 0/0}, and the efficacy data was {1/1, 1/1, 1/1, 1/1, 0/1}. Using the procedure described in 'Model-based estimation', the estimated DLT probabilities for arm {A, B, C, D, E, F} were {0.05, 0.15, 0.09, 0.22, 0.31, 0.39}, from which arms A-E were deemed to be acceptable based on confidence interval estimation. The estimated efficacy probabilities were {0.8, 0.8, 0.8, 0.8, 0.8, 0.8}, and, based on these estimates, the 8[th] participant was randomly assigned arm C. The recommendation of the 9[th] participant used updated DLT and response data from the 6[th] and 7[th] participants to calculate estimates and make a decision. This adaptive decision process will continue until sufficient information about the optimal dose combination has been obtained, according to the stopping rules described in the 'Stopping the trial' section. It is important to note that in this design approach some model-based decisions may be made using slightly less efficacy data than DLT data, due to the longer minimum observation window for efficacy. For example, the recommendation for the 11[th] participant was based on DLT data from the first 9 participants, and efficacy data from the first 8 participants, since the 2-month response data for participant 9 had yet to be fully observed when participant 11 was accrued to the study. Currently, the trial has accrued 15 participants from September 2015 through February 2017, with model-based allocation (Stage 2) having been utilized after the first 7 participants. Allocation based on weighted randomization occurred for the first 2 participants in Stage 2, which is approximately one-third the maximum sample size when combined with the data from the Stage 1 participants. To date in Stage 2, the number of participants accrued to arms {A, B,

12

Original Article

C, D, E, F} are {1, 2, 4, 1, 1, 0}. The only DLT has occurred on arm E, and the only non-responses have occurred on arms E and D.

Statistical properties

Simulation results were run to display the performance of the study design, see Table 5. In order to evaluate operating characteristics, six scenarios of assumed DLT and efficacy probabilities were chosen to reflect the following four "themes," (1) all assumed DLT probabilities are acceptable in terms of safety (i.e. ≤ 25%) and the highest Zone has the combination with the highest assumed efficacy rate (scenarios 1 and 3), (2) three combinations have assumed DLT probabilities more toxic than 25% and the assumed efficacy probabilities begin to plateau at dose level two of ibrutinib (420 mg/day) (scenarios 2 and 4), (3) 5 combinations (B-F) have assumed DLT probabilities more toxic than 25% making combination A the only acceptable combination in terms of safety (scenario 5), and (4) when all combinations are too toxic (i.e. much more toxic than 25%) (scenario 6). For each scenario, 1000 simulated trials were run, with the optimal combination(s) indicated in bold type in Table 5. Displayed in the table within each scenario for each arm is the assumed true DLT probability, the assumed true efficacy response rate, the percentage of trials in which the combination was recommended as the optimal dose combination, and the number of participants treated on each combination. Displayed in the last four columns is the average and selected percentiles for the trial size at study closure, the percentage of times in the simulations that the trial closed due to safety concerns, the percentage of simulated participants that had a DLT, and the percentage of simulated participants that had an efficacy response. The results displayed in Table 5 were based upon a maximum target accrual of 28 participants where accrual stopped when 10 participants had been treated on the recommended 'optimal' dose combination or the maximum accrual had been reached. With this type of design and stopping rules the results indicated that on average the trial would achieve this goal with accrual of approximately 20 participants.

13

Original Article

It is clear from examining the results in Table 5 that the proposed design is performing well in terms of recommending optimal dose combinations, as well as allocating participants to these combinations. In Scenario 1, the design selects, as the optimal dose combination, the target combination in 67% of simulated trials, while assigning 7.55 of 19.1 participants (40%) on average to this combination. Similar findings are obtained from Scenario 3. In Scenario 2, recommendation of target combinations as the optimal treatment combination occurs in approximately 38% of simulated trials based on an average trial size of 18.1 participants, while allocating 5.47 participants on average to the optimal dose combination. It is important to note that when the target combination is not selected as the optimal dose combination, treatments with assumed DLT probabilities marginally outside the range of acceptable safety are selected in another 23% of simulated trials. Similar findings are obtained from Scenario 4. In Scenario 5, the design identifies the target combination as the optimal dose combination in approximately 26% of simulated trials while allocating 4.41 of 18.4 participants on average to this combination. When combination A is not selected, the method tends to either choose combination B with an assumed DLT rate just outside the window of acceptable safety (25% of the time), or stop the trial for safety (11.8% of the time). Finally, in Scenario 6, where all combinations are overly toxic, the method correctly terminates the study in 100% of simulated trials based on an average trial size of 4.3 participants, and treats 1.98 accrued participants on average to Zone 1. Overall the simulation results indicate that the design outlined in this article is a practical early-phase adaptive method for use with combined targeted therapies.

CONCLUSIONS

The development of new methods in early-phase dose-finding has been rapid in the last decade, yet, the use of innovative designs remains infrequent. In this article, we have outlined a novel early-phase adaptive design, implemented in an ongoing trial of six treatment dose combinations of two targeted agents for participants with relapsed or refractory mantle cell

14

Original Article

lymphoma. The method presented in this paper describes an innovative and appropriate approach for investigating combinations of targeted therapies, which are being called for by the FDA and by others (2-4). Simulation studies were performed to evaluate the performance of the design characteristics and are reported in Table 5. The results demonstrate the method's ability to effectively recommend the optimal dose combination, defined by acceptable toxicity and high efficacy rates, in a high percentage of trials with manageable sample sizes. Software in the form of **R** (25) code for both simulation and implementation of the method is available upon request of the first author. The method we outline in this work can be viewed as an extension of the CRM, utilizing multiple skeletons for DLT and efficacy probabilities, increasing the ability of CRM designs to handle more complex dose-finding problems. The numerical results presented include the type of simulation information that aid review entities in understanding design performance, such as average sample size, frequency of early trial termination, etc., which we hope will augment early-phase trial design for targeted therapy combinations in cancer. This design can be applied more broadly in early-phase combination studies that need to consider an 'effect marker' in addition to toxicity (26), as it was used in a recently completed study of a melanoma helper peptide plus novel adjuvant combinations ((27); NCT02425306). The design would work well with any well-defined, binary "activity" endpoint.

REFERENCES

1. Sachs JR, Mayawala K, Gadamsetty S, Kang SP, de Alwis DP. Optimal dosing for targeted therapies in oncology: drug development cases leading by example. *Clin Cancer Res* 2016; 22(6): 1318-1324.

2. Nie L, Rubin EH, Mehrotra N et al. Rendering the 3+3 design to rest: more efficient approaches to oncology dose-finding trials in the era of targeted therapy. *Clin Cancer Res* 2016; 22(11): 2623-2629.

3. Iasonos A, O'Quigley J. Adaptive dose-finding studies: a review of model-guided phase I clinical trials. *J Clin Onc* 2014; 32(23): 2505-2511.

15

Original Article

4. Paoletti X, Ezzalfani M, Le Tourneau C. Statistical controversies in clinical research: requiem for the 3 + 3 design for phase I trials. *Ann Oncol* 2015; 26(9): 1808-1812.

5. Yuan Y, Nguyen HQ, Thall PF. *Bayesian designs for phase I-II clinical trials.* Chapman and Hall/CRC Press: New York; 2016.

6. Yuan Y, Yin G. Bayesian phase I/II adaptively randomized oncology trials with combined drugs. *Ann Appl Stat* 2011; 5(2A): 924-942.

7. Cai C, Yuan Y, Ji Y. A Bayesian dose-finding design for oncology clinical trials of combinational biological agents. *J R Soc Ser C Appl Stat* 2014; 63(1): 159-173.

8. Riviere MK, Yuan Y, Dubois F, Zohar S. A Bayesian dose finding design for clinical trials combining a cytotoxic agent with a molecularly targeted agent. *J R Soc Ser C Appl Stat* 2015; 64(1): 215-229.

9. Guo B, Li Y. Bayesian dose-finding designs for combination of molecularly targeted agents assuming partial stochastic ordering. *Stat Med* 2015; 34(5): 859-875.

10. Thall PF, Nguyen HQ, Zinner RG. Parametric dose standardization for optimizing two-agent combinations in a phase I-II trial with ordinal outcomes. *J R Soc Ser C Appl Stat* 2017; 66(1): 201-224.

11. Wages NA, Conaway MR. Phase I/II adaptive design for drug combination oncology trials. *Stat Med* 2014; 33(12): 1990-2003.

12. Iasonos A, Gönen M, Bosl GJ. Scientific Review of Phase I Protocols With Novel Dose-Escalation Designs: How Much Information Is Needed? *J Clin Onc* 2015; 33(19): 2221-2225.

13. Petroni GR, Wages NA, Paux G, Dubois F. Implementation of adaptive methods in early-phase clinical trials. *Stat Med* 2017; 36(2): 215-224.

14. Davids MS, Roberts AW, Seymour JF, Pagel JM, Kahl BS, Weirda WG, et al. Phase I first-in-human study of venetoclax in patients with relapsed or refractory non-hodgkin lymphoma. *J Clin Onc* 2017; 35(8): 826-833.

15. Cheson BD, Pfistner B, Juweid ME, Gascoyne RD, Specht L, Horning SJ, et al. Revised response criteria for malignant lymphoma. *J Clin Onc* 2007; 25(5): 579-86.

16. Yuan Y, Yin G. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *J R Soc Ser C Appl Stat* 2009; 58(5): 719-736.

17. Guo B, Yuan Y. A Bayesian dose-finding design for phase I/II clinical trials with nonignorable dropouts. *Stat Med* 2015; 34(10): 1721-1732.

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

16

Original Article

19. Lee SM, Cheung YK. Model calibration in the continual reassessment method. *Clin Trials* 2009; 6(3): 227-238.

20. Natarajan L, O'Quigley J. Interval estimates of the probability of toxicity at the maximum tolerated dose for small samples. *Stat Med* 2003; 22(11): 1829-1836.

21. Thall PF, Nguyen HQ. Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *J Biopharm Stat* 2012; 22(4): 785-801.

22. Wang ML, Rule S, Martin P, Goy A, Auer R, Kahl B, *et al.* Targeting BTK with ibrutinib in relapsed or refractory mantle-cell lymphoma. *N Engl J Med* 2013; 369: 507-516.

23. O'Quigley J, Shen LZ. Continual reassessment method: a likelihood approach. *Biometrics* 1996; 52(2): 673-684.

24. Portell CA, Chen RW, Wages N, Cohen J, Weber MJ, Petroni GR, et al. Initial report of a multi-institutional phase I/Ib study of ibrutinib with venetoclax in relapsed or refractory mantle cell lymphoma. *Blood* 2016; 128: 2958.

25. R Development Core Team. R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing 2008.

26. Wages NA, Conaway MR, Slingluff CL, Jr et al. Recent developments in the implementation of novel designs for early-phase combination studies. *Ann Oncol* 2015; 26(5): 1036-1037.

27. Wages NA, Slingluff CL, Petroni GR. Statistical controversies in clinical research: early-phase adaptive design for combination immunotherapies. *Ann Oncol* 2017; 28(4): 697-701.

Table 1: Zone and Arm Designation by Combination				
Venetoclax (mg per day)	400	Zone 2 / Arm C	Zone 3 / Arm E	Zone 4 / Arm F
	200	Zone 1 / Arm A	Zone 2 / Arm B	Zone 3 / Arm D
		280	420	560
		ibrutinib (week 1+) mg per day		

17

Original Article

Table 2: Possible orderings / working models (skeletons) of DLT probabilities				
Ordering	**Working model**	**Skeleton for DLT probabilities**		
A-B-C-D-E-F	m = 1	0.25^{θ_1}	0.42^{θ_1}	0.50^{θ_1}
		0.11^{θ_1}	0.17^{θ_1}	0.33^{θ_1}
A-C-B-E-D-F	m = 2	0.17^{θ_2}	0.33^{θ_2}	0.50^{θ_2}
		0.11^{θ_2}	0.25^{θ_2}	0.42^{θ_2}
A-B-C-E-D-F	m = 3	0.25^{θ_3}	0.33^{θ_3}	0.50^{θ_3}
		0.11^{θ_3}	0.17^{θ_3}	0.42^{θ_3}
A-C-B-D-E-F	m = 4	0.17^{θ_4}	0.42^{θ_4}	0.50^{θ_4}
		0.11^{θ_4}	0.25^{θ_4}	0.33^{θ_4}
A-B-D-C-E-F	m = 5	0.33^{θ_5}	0.42^{θ_5}	0.50^{θ_5}
		0.11^{θ_5}	0.17^{θ_5}	0.25^{θ_5}

Original Article

Table 3: Possible orderings / working models (skeletons) under assumption of increasing efficacy probabilities				
Ordering	**Working model**	**Skeleton for efficacy probabilities**		
A-B-C-D-E-F	k = 1	0.35^{β_1}	0.63^{β_1}	0.74^{β_1}
		0.10^{β_1}	0.21^{β_1}	0.50^{β_1}
A-C-B-E-D-F	k = 2	0.21^{β_2}	0.50^{β_2}	0.74^{β_2}
		0.10^{β_2}	0.35^{β_2}	0.63^{β_2}
A-B-C-E-D-F	k = 3	0.35^{β_3}	0.50^{β_3}	0.74^{β_3}
		0.10^{β_3}	0.21^{β_3}	0.63^{β_3}
A-C-B-D-E-F	k = 4	0.21^{β_4}	0.63^{β_4}	0.74^{β_4}
		0.10^{β_4}	0.35^{β_4}	0.50^{β_4}
A-B-D-C-E-F	k = 5	0.50^{β_5}	0.63^{β_5}	0.74^{β_5}
		0.10^{β_5}	0.21^{β_5}	0.35^{β_5}

19

Original Article

Table 4: Possible orderings / working models (skeletons) under assumption of plateau efficacy probabilities				
Ordering	**Working model**	**Skeleton for efficacy probabilities**		
A-B-C-D-E-F	k = 6	0.50^{β_6}	0.50^{β_6}	0.50^{β_6}
		0.35^{β_6}	0.50^{β_6}	0.50^{β_6}
A-B-C-E-D-F	k = 7	0.35^{β_7}	0.50^{β_7}	0.50^{β_7}
		0.10^{β_7}	0.21^{β_7}	0.50^{β_7}
A-C-B-E-D-F	k = 8	0.35^{β_8}	0.50^{β_8}	0.50^{β_8}
		0.21^{β_8}	0.50^{β_8}	0.50^{β_8}
A-C-B-E-D-F	k = 9	0.21^{β_9}	0.50^{β_9}	0.63^{β_9}
		0.10^{β_9}	0.35^{β_9}	0.63^{β_9}
A-B-C-D-E-F	k = 10	$0.50^{\beta_{10}}$	$0.50^{\beta_{10}}$	$0.50^{\beta_{10}}$
		$0.50^{\beta_{10}}$	$0.50^{\beta_{10}}$	$0.50^{\beta_{10}}$

Original Article

Table 5: Simulation results for two agents, multiple endpoints (DLT, Response).
Optimal combinations are defined by an acceptable toxicity profile (≤25% DLT rate)
and the highest clinical response rate. Stop when an 11th participant is
recommended to a dose combination that already has 10 participants.
Maximum total accrual set at 28, 1000 simulations

Scenario	Ven (mg per day)	(true %DLT, true %Response) Proportion combination recommended Avg # pts treated on combination			Avg Size, %-tiles	% stop	% DLT	% Rsp
		Ibrutinib (mg per day)						
		280	**420**	**560**				
1 "Ideal"	**400**	(0.07, 0.55) 0.12 1.72	(0.15, 0.70) 0.12 2.30	**(0.25,0.95)** **0.67** **7.46**	19.1, 25th = 17 50th = 18 75th = 20 90th = 24 95th = 27	0.2	16.4	74.2
	200	(0.05, 0.40) 0.07 2.30	(0.07, 0.55) 0.08 2.30	(0.15,0.80) 0.14 3.06				
2 "Worst"	**400**	(0.25, 0.25) 0.06 2.54	(0.30, 0.40) 0.08 2.18	(0.40,0.50) 0.06 1.63	18.1, 25th = 16 50th = 18 75th = 21 90th = 25 95th = 27	7.0	25.6	38.3
	200	(0.20, 0.25) 0.19 3.81	**(0.25, 0.40)** **0.36** **5.08**	(0.30,0.50) 0.17 3.08				
3 Best Guess	**400**	(0.12, 0.55) 0.02 1.97	(0.18, 0.80) 0.10 2.77	**(0.25,0.95)** **0.58** **6.72**	19.8, 25th = 17 50th = 18 75th = 22 90th = 26 95th = 28	0.6	18.4	75.0
	200	(0.10, 0.40) 0.07 2.57	(0.15, 0.55) 0.08 2.77	(0.20,0.80) 0.16 3.16				
4 "Plateau"	**400**	(0.25, 0.35) 0.06 2.55	(0.30, 0.50) 0.08 2.19	(0.40,0.50) 0.03 1.46	18.2, 25th = 17 50th = 18 75th = 21 90th = 25 95th = 28	9.0	25.7	44.8
	200	(0.20, 0.35) 0.20 3.83	**(0.25, 0.50)** **0.38** **5.47**	(0.30,0.50) 0.15 2.92				
5 "Toxic"	**400**	(0.30, 0.40) 0.09 2.94	(0.40, 0.50) 0.09 2.20	(0.50,0.75) 0.03 1.10	18.4, 25th = 16 50th = 19 75th = 23 90th = 26 95th = 28	11.8	32.9	52.9
	200	**(0.25, 0.40)** **0.26** **4.41**	(0.30, 0.50) 0.25 4.59	(0.40,0.75) 0.17 3.12				
6 "Very Toxic"	**400**	(0.75, 0.40) 0.00 0.95	(0.85, 0.50) 0.00 0.09	(0.95,0.75) 0.00 0.00	4.3, 25th = 2 50th = 2 75th = 6 90th = 9 95th = 11	100.00	74.0	42.5
	200	(0.70, 0.40) 0.00 1.98	(0.75, 0.50) 0.00 1.07	(0.85,0.75) 0.00 0.17				

Notation:
Average (avg); percentiles (%-tiles) ; Response (Rsp) ; Average trial sample size (avg size);
Percent of trials stopped early for safety (% stop); Percent of DLTs (% DLT); Percent of responses (% Rsp)

21

Clinical Cancer Research

Implementation of a model-based design in a phase 1b study of combined targeted agents

Nolan A. Wages, Craig A Portell, Michael E Williams, et al.

Updated version	Access the most recent version of this article at: doi:10.1158/1078-0432.CCR-17-1069
Author Manuscript	Author manuscripts have been peer reviewed and accepted for publication but have not yet been edited.

E-mail alerts	Sign up to receive free email-alerts related to this article or journal.
Reprints and Subscriptions	To order reprints of this article or to subscribe to the journal, contact the AACR Publications Department at pubs@aacr.org.
Permissions	To request permission to re-use all or part of this article, use this link http://clincancerres.aacrjournals.org/content/early/2017/07/21/1078-0432.CCR-17-1069. Click on "Request Permissions" which will take you to the Copyright Clearance Center's (CCC) Rightslink site.