Published in final edited form as:

Contemp Clin Trials. 2015 March ; 41: 172–179. doi:10.1016/j.cct.2015.01.016.

A Phase I/II adaptive design to determine the optimal treatment regimen from a set of combination immunotherapies in high-risk melanoma

Nolan A. Wages^{1,*}, Craig L. Slingluff², and Gina R. Petroni¹

¹Division of Translational Research & Applied Statistics, Department of Public Health Sciences University of Virginia, Charlottesville, Virginia 22908, U.S.A.

²Division of Surgical Oncology, Department of Surgery, University of Virginia, Charlottesville, Virginia 22904-4135, U.S.A.

Abstract

In oncology, vaccine-based immunotherapy often investigates regimens that demonstrate minimal toxicity overall and higher doses may not correlate with greater immune response. Rather than determining the maximum tolerated dose, the goal of the study becomes locating the optimal biological dose, which is defined as a safe dose demonstrating the greatest immunogenicity, based on some predefined measure of immune response. Incorporation of adjuvants, new or optimized peptide vaccines, and combining vaccines with immune modulators may enhance immune response, with the aim of improving clinical response. Innovative dose escalation strategies are needed to establish the safety and immunogenicity of new immunologic combinations. We describe the implementation of an adaptive design for identifying the optimal treatment strategy in a multi-site, FDA-approved, phase I/II trial of a novel vaccination approach using long-peptides plus TLR agonists for resected stage IIB-IV melanoma. Operating characteristics of the design are demonstrated under various possible true scenarios via simulation studies. Overall performance indicates that the design is a practical Phase I/II adaptive method for use with combined immunotherapy agents. The simulation results demonstrate the method's ability to effectively recommend optimal regimens in a high percentage of trials with manageable sample sizes. The numerical results presented in this work include the type of simulation information that aid review boards in understanding design performance, such as average sample size and frequency of early trial termination, which we hope will augment early-phase trial design in cancer immunotherapy.

Keywords

Phase I/II; Optimal dose; Oncology; Immunotherapy; Biological agent

© 2015 Published by Elsevier Inc.

*nwages@virginia.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1 Introduction

Historically, the primary objective of dose-finding clinical trials in oncology is to identify the maximum tolerated dose (MTD) of the agent being investigated from a discrete set of available doses. Numerous designs have been proposed for identifying the MTD from a set of I doses $\mathcal{D} = \{d_1, \dots, d_I\}$, in which dose-limiting toxicity (DLT) is measured as a binary outcome (DLT; yes/no), a thorough review of which is given in a recent manuscript by Braun [1]. In oncology trials of a single cytotoxic agent, identification of the MTD is often determined by considering DLT information only, with the assumption that the MTD is the highest dose that satisfies some safety requirement, and therefore provides the most promising outlook for efficacy. In a subsequent Phase II trial, the therapy is evaluated for efficacy at the recommended dose (estimated MTD or one dose level below). Phase I trials evaluating the toxicity of single agents are becoming less common, giving way to more complex studies involving combinations of agents, biological therapies, or multiple treatment schedules. These more complex problems have created the need to adapt early-phase trial design to the specific agents being studied and the corresponding endpoints. Frequently, Phase I and Phase II trials are performed independently, without formally sharing information across the separate phases. Recently, there has been growing interest in oncology to integrate Phase I and Phase II trials to accelerate the drug development process, while potentially reducing costs [2]. To this end, several recent Phase I/II methods have extended dose-finding methodology to allow for the exploration of both toxicity and efficacy, some within the context of the more complex problems mentioned above [3, 4, 5].

Immunotherapy has emerged as an effective approach for many cancers, with particular promise for checkpoint blockade agents (e.g. antibodies blocking CTLA-4 and PD-1). Cancer vaccines have also been studied for decades, but with disappointing results in many trials. On the other hand, recent studies have shown clinical benefit and promise with vaccines for lymphoma, melanoma, and prostate cancer patients [6, 7, 8], with a cancer vaccine now FDA-approved for hormone-refractory prostate cancer. Vaccine-based immunotherapy offers potential low-toxicity alternatives or additions to chemotherapy, targeted molecular therapies, or other immunotherapies.

2 Designing immunotherapy trials

Immunotherapy is posing a challenge to the accepted methods of early-phase trial design [9]. The Food and Drug Administration (FDA) recently published a “Guidance for Industry: Clinical Considerations for Therapeutic Cancer Vaccines,” which acknowledged the need for alternative dose-escalation methods in cancer vaccines [10, 11]. In contrast to chemotherapeutic agents, dose-finding in immunotherapy often investigates treatments that demonstrate minimal toxicity overall, where higher doses may not induce greater immunologic effect. The biologic activity of the treatment may increase at low doses, and then begin to level off, or plateau, at higher doses [12, 13, 14, 15, 16]. Rather than determining the maximum tolerated dose (MTD), the goal of dose-finding with immunologic agents is to locate the optimal biological dose, which we define as a safe dose which demonstrates the greatest immunogenicity, based on some predefined measure of immune response. Protocol-specific immunological endpoints give the investigator a

measure of biologic activity that serves as a driving factor in the trial design. Although the relationship between clinical outcome and immunologic activity is not clear at this point, it is generally assumed that the absence of biologic effect will accompany a lack of clinical efficacy. Before an immunologic agent can be taken into large-scale trials to test for clinical efficacy, early-phase trials are needed to establish that the therapy can produce an immunologic effect with the potential to translate to clinical benefit. Some published Phase I/II methods [5, 17, 18, 19, 20] can reasonably be applied in a single-agent immunotherapy setting by replacing “efficacy” with “immune response” because most account for the potential plateau of biological effects at higher doses.

Studies involving combinations of agents are another area of trial design that challenge accepted methods in dose-finding. Gao et al. [6] argue that the treatment benefits seen in patients who have received single-agent immunotherapies has reached a limitation and that trials combining multiple agents are necessary in order to advance cancer immunotherapy. Incorporation of adjuvants, such as Toll-like receptor (TLR) agonists, new or optimized peptide vaccines and combining vaccines with immune modulators such as checkpoint blockade agents (e.g. blocking antibodies to CTLA-4 or PD-1/PD-L1) may aid in improving immune response, with the goal of improving clinical outcome [6, 21]. Innovative dose escalation strategies are needed to establish the safety and immunogenicity of new immunologic combinations. Novel dose-finding designs for multiple immunologic agents could potentially have a significant impact on the growth of immunotherapy as a treatment for cancer. The purpose of this article is to present a Phase I/II design for identifying the optimal treatment regimen in a multi-center clinical trial (Mel 60) of a long-peptide vaccine (LPV7) plus TLR agonists for resected stage IIB-IV melanoma. This design may be applicable for other trials of combination immunotherapy.

3 Mel 60 trial

The trial design described in this article is a multi-site, FDA-approved, phase I/II trial of a novel melanoma vaccination approach using long-peptides plus TLR agonists. This study was designed at the University of Virginia (UVA) Cancer Center, and is currently open to accrual at UVA and MD Anderson Cancer Center. The primary objective of the trial is to evaluate the safety and immunogenicity of vaccination with a mixture of long peptides for patients with histologically or cytologically proven Stage IIB-IV melanoma rendered clinically free of disease by surgery, other therapy, or spontaneous remission. In line with the primary objective, the trial was designed to determine a range of optimal regimens, among the adjuvant preparations provided in Table 1, where an optimal regimen is one that is estimated to have an acceptable toxicity profile as measured by DLT's and a high immunologic response rate (IRR), as measured by peak immune response. Within the range of optimal regimens, a co-primary objective is, if more than one treatment strategy is contained within the range, to estimate the difference in IRR between regimens. For this study there are seven regimens, which we label $\mathcal{D} = \{d_1, \dots, d_7\}$, including TLR agonists and/or incomplete Freund's adjuvant IFA (Table 1).

3.1 Design considerations

This study incorporates an adaptive design with sequential initiation of accrual to study Zones, $Z = \{Z_1, Z_2, Z_3\}$, with increasing numbers of adjuvants; $Z_1 = \{d_1, d_2, d_3\}$, $Z_2 = \{d_4, d_5, d_6\}$, $Z_3 = \{d_7\}$, as displayed in Table 1. There is a 3 week interval between subjects 1 and 2 of each regimen of the study. The design allows, during that 3 week interval, randomization to regimens that remain open within a zone. We assume that both DLT, X , and immune response, Y , are defined as binary measures

$$X = \begin{cases} 0 & \text{if no DLT,} \\ 1 & \text{if DLT} \end{cases}, \quad Y = \begin{cases} 0 & \text{if no immune response,} \\ 1 & \text{if immune response.} \end{cases}$$

A detailed description of the trial can be found at <https://clinicaltrials.gov/ct2/show/NCT02126579?term=Mel+60&rank=1>.

Safety—In monitoring safety, adverse events are being assessed and acute toxicity graded using the National Cancer Institute Common Terminology Criteria for Adverse Events (CTCAE) Version 4.03 and a patient is classified as experiencing a DLT (yes/no) based on protocol-defined dose limiting adverse events. For escalation decisions, subjects must be observed for a minimum of 3 weeks after the initial vaccine; however any DLT observed through 12 weeks will be used for escalation decisions in future patients. Acceptability, with regards to safety, is defined by an estimated DLT probability less than or equal to a maximum toxicity tolerance of $\phi_T = 33\%$, which was chosen based on the expectedness of adverse events. Denoting the probability of DLT at regimen d_i as $\pi_T(d_i)$, the primary safety objective is to identify an acceptable set of regimens defined by

$\mathcal{A} = \{d_1 : \pi_T(d_i) \leq \phi_T; i=1, \dots, 7\}$. Safety assessments are based on the assumption that, as the number of adjuvants increases, the probability of DLT is non-decreasing. In other words, it is reasonable to assume that regimens in higher zones do not have lower probabilities of DLT than regimens in lower zones (i.e. $\pi_T(d_1) \leq \pi_T(d_4)$). This assumption is based on data from a series of previous Mel studies in which these adjuvant preparations were combined with other peptide vaccines. It is unknown whether regimens have higher or lower DLT probabilities than other regimens within the same zone. It could be that $\pi_T(d_2) \leq \pi_T(d_3)$ or that $\pi_T(d_3) \leq \pi_T(d_2)$. Our overall strategy is to formulate possible orderings of the DLT probabilities. There are six possible ways of ordering the three regimens within each of the first two zones, Z_1 and Z_2 , with non-decreasing DLT probabilities from Z_1 to Z_3 . The ordering within Z_1 is used in making assumptions about the ordering within Z_2 . For instance, within Z_1 , if it assumed that $\pi_T(d_1) \leq \pi_T(d_2) \leq \pi_T(d_3)$, then, within Z_2 , it assumed that $\pi_T(d_4) \leq \pi_T(d_5) \leq \pi_T(d_6)$, since d_4 is a combination of d_1 and d_2 , d_5 combines d_1 and d_3 , and so on. Using this information, we specify the following six orderings for the DLT probabilities, which are indexed by m ; $m = 1, \dots, 6$.

1. $m = 1 : \pi_T(d_1) \leq \pi_T(d_2) \leq \pi_T(d_3) \leq \pi_T(d_4) \leq \pi_T(d_5) \leq \pi_T(d_6) \leq \pi_T(d_7)$
2. $m = 2 : \pi_T(d_1) \leq \pi_T(d_3) \leq \pi_T(d_2) \leq \pi_T(d_5) \leq \pi_T(d_4) \leq \pi_T(d_6) \leq \pi_T(d_7)$
3. $m = 3 : \pi_T(d_2) \leq \pi_T(d_1) \leq \pi_T(d_3) \leq \pi_T(d_4) \leq \pi_T(d_6) \leq \pi_T(d_5) \leq \pi_T(d_7)$

4. $m = 4 : \pi_T(d_2) \quad \pi_T(d_3) \quad \pi_T(d_1) \quad \pi_T(d_6) \quad \pi_T(d_4) \quad \pi_T(d_5) \quad \pi_T(d_7)$
5. $m = 5 : \pi_T(d_3) \quad \pi_T(d_1) \quad \pi_T(d_2) \quad \pi_T(d_5) \quad \pi_T(d_6) \quad \pi_T(d_4) \quad \pi_T(d_7)$
6. $m = 6 : \pi_T(d_3) \quad \pi_T(d_2) \quad \pi_T(d_1) \quad \pi_T(d_6) \quad \pi_T(d_5) \quad \pi_T(d_4) \quad \pi_T(d_7)$

We allow the accumulating information to guide us as to which ordering is most represented by the data. Throughout the trial duration, we use a working model to continuously monitor safety and adaptively update a set $\mathcal{A}$ of acceptable regimens, with which we make allocation decisions based on immune response rates.

Immunologic response—Immune response is measured by the levels of peptide-reactive CD8⁺ cells in the peripheral blood reactive to short peptide sequences with the long peptides in the vaccines, using a direct ELISpot assay. For response, a subject is identified as an “immune responder” (yes/no) based on whether or not his/her immune response assay results are sufficiently large post-treatment compared to pre-treatment to be confident that the agent had an effect on the immune status of the patient. The presence or absence of an immune response is defined by induction of at least a 2-fold increase in IFN-gamma-secreting cells over background, and over pre-existing responses, and at least 30 IFN-gamma secreting spots per 100,000 CD8 T cells, as described in Slingsluff et al. [22]. This assessment will be performed on blood through week 12, although a minimum of 4 weeks of data will be used to guide decisions about the range of optimal dose combinations. IRR is defined as the proportion of patients where presence (as defined above) of an immune response has been observed. Other measures of immune response and durable immunogenicity will also be evaluated at week 26 by ELISpot assay against the defined nonamer peptides in the peripheral blood.

In contrast to toxicity, regimens in higher zones may not necessarily produce a higher immunologic response rate. Immune response may increase initially at lower zones, and then plateau at higher zones. The primary objective, both within and at the conclusion of the trial, is to select the regimen that is estimated to maximize $\pi_R(d_i)$ for $d_i \in \mathcal{A}$. In general, regimen allocation will occur in two stages. The initial stage will accrue patients until a patient experiences a DLT. The second stage will allocate patients based upon a continual reassessment method (CRM; [23]) for combinations of agents [24]. The design incorporates randomization, which will be based on equal allocation among allowable regimens, until a maximization scheme is triggered (See Section 5). Randomization will not be stratified by institution. Patients must be observed for a minimum of 3 weeks after the initial vaccine for initial escalation between zones, and between the 1st and 2nd patients within a regimen.

4 Estimation

4.1 Toxicity

The motivating trial described in Section 3 is testing $I = 7$ regimens $\mathcal{D} = \{d_1, \dots, d_7\}$. We specify a working model for the set of DLT probabilities corresponding to the $m = 1, \dots, 6$ possible orderings. For a particular ordering, m , we model $\pi_T(d_i)$, the true probability of DLT response at regimen d_i by

$$\pi_T(d_i) = \Pr[X=1|d_i] \approx F_m(d_i, \beta_m) = p_{mi}^{\exp(\beta_m)},$$

where the parameter $\beta_m \in (-\infty, \infty)$ is to be estimated from the data. The p_{mi} values are pre-specified constants, often termed ‘skeleton,’ with values between 0 and 1 which can be chosen based on clinical experience or using the algorithm of Lee and Cheung [25]. We allow the plausibility of each ordering to be described by a set of prior probabilities

$\{\xi(1), \dots, \xi(6)\}$, where $\xi(m) \geq 0$ and $\sum \xi(m) = 1$. Even though there is no prior information available on the possible orderings in the current trial, we formally proceed in the same way by specifying a discrete uniform so that $\xi(m) = 1/6$. At any point in the trial, the accumulated toxicity data can be represented by $\mathcal{X} = (x_1, x_2, \dots, x_7)$, the number of DLT's, and $\mathcal{N} = (n_1, n_2, \dots, n_7)$, the number of patients observed on each regimen. For each of the assumed orderings, $m = 1, \dots, 6$, the likelihood is given by

$$\mathcal{L}(\beta_m) = \prod_{i=1}^7 \{F_m(d_i, \beta_m)\}^{x_i} \{1 - F_m(d_i, \beta_m)\}^{n_i - x_i},$$

from which we can obtain the maximum likelihood estimate (MLE), $\hat{\beta}_m$, of the parameter β_m for each of the 6 orderings. We need a value of m so we weight each of the candidate orderings as we make progress and appeal to sequential model selection techniques to guide decision-making. A plausible choice is driven by the maximization of

$$\omega(m) = \frac{\mathcal{L}(\hat{\beta}_m) \xi(m)}{\sum_{m=1}^6 \mathcal{L}(\hat{\beta}_m) \xi(m)}; \quad m=1, \dots, 6,$$

where $\omega(m)$ is considered as the weight of evidence in favor of ordering m , and $\mathcal{L}(\hat{\beta}_m)$ is the value of the likelihood evaluated at its MLE. If there is a tie between the likelihood values of two or more orderings, then the selected order is randomly chosen from among the tied orderings. At the inclusion of each new patient, we choose a single model, m^* , that maximizes $\omega(m)$, and implement its working model in generating DLT probability estimates, $\hat{\pi}_T(d_i)$, at each regimen so that

$$\hat{\pi}_T(d_i) \approx F_{m^*}(d_i, \hat{\beta}_{m^*}), \quad \text{where } m^* = \arg \max_m \omega(m).$$

Based on these estimates, we obtain a set of acceptable (safe) regimens to guide allocation decisions as described below in the trial conduct.

4.2 Immunologic response

For immune response, we appeal to a simple, non-parametric approach by estimating response probabilities, $\pi_R(d_i) = \Pr[Y=1|d_i]$, using observed IRR's. The shape of the

regimen-response curve is not known at every regimen. It is possible for regimen-response to exhibit unimodal (increasing then decreasing) or plateau (increasing then leveling off) patterns. With the aim of practical and computational feasibility, estimation of immune response probabilities does not rely on a parametric dose-response model. At patient inclusion, the accumulated immune response data can be represented by

$\mathcal{Y} = (y_1, y_2, \dots, y_7)$, the number of immune responses, and $\mathcal{N} = (n_1, n_2, \dots, n_7)$, the number of patients observed on each regimen. Using this information, we can compute a set of estimated immune response probabilities

$$\hat{\pi}_R(d_i) = \frac{y_i}{n_i}; \quad i=1, \dots, 7.$$

At each patient inclusion, these immune response probability estimates are used to make allocation decisions as described below in the trial conduct.

5 Trial conduct

5.1 Regimen allocation in stage 1

The initial stage is accruing patients on regimens in cohorts of one until a patient experiences a dose-limiting toxicity (DLT). The escalation plan for the first stage is based on grouping treatment strategies into zones. With this escalation scheme, patients can be accrued and assigned to other open regimens within a zone but escalation will not occur outside the zone until the minimum follow-up period is observed for the first patient accrued to a regimen. Allocation will begin in Z_1 . Initial allocation within a zone is based upon random allocation between the possible regimens. The first patient in a zone will be entered onto a regimen chosen at random from within the occupied zone. Subsequent patients in a zone will be randomized to one of the remaining regimens that have not yet been tried. Escalation to a higher zone occurs only when all regimens in the lower zone have accrued at least 1 patient and no patients have experienced a DLT. This allocation strategy is followed for accrual to increasing zones until a patient experiences a DLT. In the absence of DLT's, patients will continue to be randomized to regimens in $\mathcal{D}$ until a DLT occurs, upon which the modeling stage, Stage 2, begins. If at least 3 patients have been treated at every regimen and there is still no observed DLT, allocation will be based on which regimen maximizes $\hat{\pi}_R(d_i)$. If DLT is never observed on any regimen, allocation will continue until sufficient information has been obtained regarding the optimal regimen, according to the stopping rules described below.

5.2 Regimen allocation in stage 2

After Stage 1 described above, we sequentially allocate each new patient to the regimen that maximizes $\hat{\pi}_R(d_i)$, among those with acceptable toxicity. Upon obtaining DLT probabilities, $\hat{\pi}_T(d_i)$, for each regimen using the estimation procedure outlined in the previous section, we restrict our attention to those regimens with estimated DLT rates less than the maximum allowable DLT rate $\phi_T = 0.33$. By estimating an acceptable set, $\tilde{A} = \{\hat{\pi}_T(d_i) \leq \phi_T; i=1, \dots, 7\}$, we exclude overly toxic regimens. If, at any point in the

trial, $\tilde{A}$ is empty, the trial is terminated for safety. After each cohort inclusion the acceptable set is adaptively redefined based on the current DLT probability estimates, so it is possible, once more data have been observed, for $\tilde{A}$ to include regimens that were previously excluded when a limited amount of data existed. The allocation algorithm depends upon the amount of data that have been observed so far in the trial. In the presence of limited data, we rely on a randomization phase to allocate future patients to acceptable regimens.

Randomization phase—Early in the trial, there may not be enough immune response data to accurately assign patients to the regimen with maximum $\hat{\pi}_R(d_i)$ and acceptable toxicity. There may be regimens in $\tilde{A}$ that have never been tried and information on these can only be obtained through experimentation. It is possible for a sequential design relying on an algorithm that maximizes $\hat{\pi}_R(d_i)$ to become “greedy” and develop a tendency to repeatedly assign a suboptimal dose [26]. For example, the first assignment could, by chance, not result in an immunologic response, meaning $\hat{\pi}_R(d_i) = 0/1$ at that regimen. Implementation of a maximization algorithm at this point could result in this regimen never being tried again as a result of this one observation. A common practical solution is to randomly assign a small number of patients to suboptimal regimens. This added randomization allows for information to be obtained on competing regimens and prevents the method from getting “stuck” on a suboptimal regimen that has been tried early in the trial, thus allowing information to be obtained more broadly. The overall strategy is to randomly assign patients to regimens in $\tilde{A}$ until enough data have been obtained at competing regimens in $\tilde{A}$ in order to feel comfortable implementing a maximization algorithm. The current trial randomizes patients, with equal probability, to safe regimens in $\tilde{A}$ until a minimum of *three* patients have been treated at each regimen in $\tilde{A}$. If $n_i < 3$ for any $d_i \in \tilde{A}$, then the recommended regimen for the next patient will be chosen at random from acceptable regimens with less than 3 patients treated.

Maximization phase—After each regimen in $\tilde{A}$ has accrued the minimum number of patients, randomization will cease. In the latter portion of the trial, when a sufficient amount of data have been observed, we will utilize a maximization phase. The recommended regimen, d_i , is defined as the acceptable regimen with the highest observed IRR so that

$$\max \hat{\pi}_R(d_i).$$

After each patient, a new recommendation is obtained, and the next entered patient is allocated to the recommended regimen. The trial will stop once sufficient information about the optimal dose range has been obtained, according to the stopping rules outlined in Section 5.4.

5.3 Sample size/accrual

Maximum sample size is based upon acquiring sufficient information to assess the objective of selecting the regimen with the highest IRR, assuming at least one potential optimal regimen has been found. Data from our prior Mel 43 trial [27] indicated a high T cell response (5-fold increase in reactivity, at least 1200 cells per 100,000 CD8, and no overlap

of 2 SD) to the 12 peptide vaccine without GM-CSF in the peripheral blood in 30 of 60 patients (50%; 95% C.I. (37%, 63%)). Using the range as an approximate guideline, maximum sample size is established to ensure that if IRR among all seven regimens differ by 30%, with the smallest expected rate of 50%, then the regimen with the higher IRR would be selected with high probability (i.e. 90%). The initial maximum target accrual of 20 patients per regimen is based upon the ranking and selection procedure, BSH [28], which allows selection of the best treatment worthy of further investigation in single-factor Bernoulli response experiments. If each regimen reached the maximum of 20 patients, the total maximum accrual would be 140 patients. However, it is not anticipated that all regimens will fall in the range of optimal regimens, thus, sample size is estimated from the simulations and will be determined by the stopping rules below. In the simulation results, on average a total of between 7 and 46 patients were required to complete the study.

5.4 Stopping rules

The study will be stopped according to the following rules:

1. The study will stop early for safety if the first three entered patients experience DLT on regimens in Z_1 in Stage 1.
2. If at any point in Stage 2, the set of acceptable regimens is empty, $\tilde{A}=\emptyset$, the trial will stop for safety.
3. The study will stop if the recommendation is to assign the next patient to a regimen that already has 20 patients treated at that regimen.

6 Statistical properties

6.1 Illustration

In this section, we illustrate the behavior of the method described in this article under a set of true DLT and immune response probabilities, which will serve as Scenario 1 in our simulation studies in Table 4. The true (DLT, immune response) probabilities, (π_T, π_R) , are $\{(0.01, 0.35), (0.02, 0.35), (0.05, 0.40), (0.07, 0.55), (0.11, 0.60), (0.14, 0.65), (0.20, 0.80)\}$, indicating that regimen d_7 is optimal since it is safe and has the largest IRR. All regimens, $d_1, \dots, d_7$, are considered acceptable in terms of safety (DLT rates $\leq 33\%$), with nondecreasing probabilities from d_1 to d_7 . This would indicate that ordering $m = 1$ is most consistent with the true underlying DLT probabilities. The estimation of DLT probabilities embodies characteristics of the CRM, so we appeal to its features in specifying design parameters. For instance, the skeleton values, p_{mi} , were chosen according to the algorithm of Lee and Cheung [25], and are reflected in the working models given in Table 2. It has been demonstrated that CRM designs are robust and efficient with the implementation of “reasonable” skeletons, defined by adequate spacing between adjacent values [29]. The values in Table 2 were generated using the `getprior` function in R [30] package `dfcrm` [31]. We assumed that each model was equally likely at the beginning of the trial and set $\xi(m) = 1/6$.

The data from an entire simulated trial are provided in Table 3. The first 9 eligible patients are escalated in cohorts of one on each regimen in the absence of DLT. The first DLT occurs

in patient 10 on regimen d_7 , at which point the modeling stage begins. With this limited amount of data, $m^* = 2$ is estimated to be the true ordering, and $\hat{\beta}_{m^*} = 0.7764$. These values are used to calculate DLT probability estimates at each regimen via $\hat{\pi}_T(d_i) \approx p_{2i}^{\exp(0.7764)}$, yielding $\hat{\pi}_T(d_i) = \{0.00, 0.01, 0.00, 0.06, 0.03, 0.11, 0.18\}$. These estimates indicate that all regimens have acceptable toxicity. Note that all regimens will remain safe until a second DLT is observed, although DLT rates will continue to be estimated after each patient inclusion. Since three patients have not yet been observed on each acceptable regimen, the next patient is randomized to one of the 7 acceptable regimens. This randomization yields a recommendation of d_6 for patient 11. This patient does not experience a DLT but does have an immune response. Patient 12 is then randomized to d_2 and has neither a DLT nor a response. Patient 13 is administered d_4 and experiences a DLT. The toxicity data are then used to update the estimates, from which $m^* = 6$ is estimated to be the true ordering, and $\hat{\beta}_{m^*} = 0.4536$. The updated DLT estimates become $\hat{\pi}_T(d_i) = \{0.04, 0.01, 0.00, 0.20, 0.14, 0.08, 0.29\}$, again indicating that all regimens are acceptable, although this most recent DLT caused an increase in the DLT probability estimates, as expected. After patient 21, the minimum of three patients have been treated on each regimen, so the method begins to base allocation on which regimen has the maximum observed IRR, which, under this true scenario, is d_7 . All regimens remain acceptable until a string of consecutive DLT's on d_6 and d_7 for patients 27–30. At this point, the estimated acceptable set, $\tilde{A}$, is reduced to 6 regimens, which excludes d_7 . After a string of 6 consecutive non-DLT's on d_6 , all regimens are again estimated to be acceptable after patient 35, illustrating the adaptive nature of the design. Regimen d_7 is allowed to “re-enter” $\tilde{A}$ after additional non-toxic data is accumulated at “lower” regimens. After 45 patients, the trial settles on d_7 for the final 7 patients entered into the trial. Upon observation of 52 patients, 20 patients have been allocated to d_7 , triggering stopping criteria (3) and the trial is terminated. Overall, in this simulated trial, 52 patients are treated, yielding optimal regimen recommendation of d_7 .

6.2 Simulation studies

Simulation results were run in **R** [30] to display the performance of the design characteristics. For each scenario, 1000 simulated trials were run. The results report the true DLT rate at each regimen, the true IRR at each regimen, the proportion of trials in which each regimen was recommended as the optimal regimen, the average number (and quartiles) of patients treated in a simulated study, the percentage of times the trial was stopped for safety, and the overall DLT/immune response rate. The results displayed in Table 4 were based upon a maximum target accrual of 140 patients where accrual stopped when 20 patients had been treated on the recommended optimal regimen. With this type of design and stopping rules, the results indicate that on average the trial would achieve this goal with accrual well below 140 patients. The true scenarios were chosen to reflect a range of situations, with various locations of true optimal regimens, involving an assortment of safety profiles. The true optimal regimen or range of optimal regimens is indicated in bold type. In Scenario 1, all true DLT probabilities are safe (i.e. less toxic than 33%) and one regimen has the highest IRR. In Scenario 2, one regimen (d_7) has true DLT probability more toxic than 33% and the three regimens in Z_2 have the highest IRR among safe regimens. In Scenario 3,

regimens in Z_2 and Z_3 have true DLT probabilities more toxic than 33% and the two regimens in Zone 1 have the highest IRR among safe regimens. In Scenario 4, two of the three regimens in Z_2 are considered safe, and both are more immunogenic than regimens in Z_1 . In Scenario 5, two of the three regimens in Z_1 are considered safe, and all other regimens are unsafe. Finally, in Scenario 6, all true DLT probabilities are unsafe (i.e. more toxic than 33%).

It is clear from examining the results in Table 4 that the proposed design is performing well in terms of recommending optimal treatment regimens, as well as allocating patients to these regimens. In Scenario 1, the design selects, as the optimal treatment regimen, the target regimen in 62% of simulated trials, while assigning 33% of 45 patients to this regimen. In Scenario 2, recommendation of target regimens as the optimal treatment regimen occurs in approximately 67% of simulated trials, while more than half (53%) of the patients enrolled are treated at these regimens. In Scenario 3, the design identifies target regimens as the optimal treatment regimen in approximately 72% of simulated trials while allocating 83% of patients to one of these three regimens. In Scenario 4, the optimal regimens in Z_2 are selected in 56% of trials, while the overly toxic d_5 is only selected 9% of the time. Similar findings are obtained from Scenario 5, in which the method is able to locate the two optimal regimens in Z_1 in 66% of simulated trials. Finally, in Scenario 6, all regimens are overly toxic. The method correctly terminates the trial in 96% of simulated trials, and treats 95% of the 7 accrued patients to regimens in Z_1 . Overall the simulation results indicate that the design outlined in this article is a practical Phase I/II adaptive method for use with combined immunotherapy agents.

7 Discussion

In this article, we have outlined a novel Phase I/II adaptive design, implemented in an FDA/IRB-approved trial of a novel regimen of immunologic agents for patients with high-risk melanoma. The simulation results demonstrated the method's ability to effectively recommend optimal regimens, defined by acceptable toxicity and high immune response rates, in a high percentage of trials with manageable sample sizes. The method we outline in this work can be viewed as an extension of the CRM, utilizing possible orderings of DLT probabilities, increasing the ability of CRM designs to handle more complex dose-finding problems [32]. The development of novel methods in early-phase dose-finding has been rapid in recent years, yet, the use of innovative designs remains infrequent [33, 34]. This can be attributed to several causes, not least of which includes (1) clinician skepticism, and (2) difficulty or assumed difficulty in obtaining approval of entities such as IRB's, pharmaceuticals and the FDA. These complications are likely to be enhanced in the coming years as the recent paradigm of oncology drug development involves a shift to more complex dose-finding problems. The numerical results presented in this work include the type of simulation information that aid review boards in understanding design performance, such as average sample size, frequency of early trial termination, etc., which we hope will augment early-phase trial design in oncology.

Patient allocation decisions presented in this manuscript depend on definition (and measure) of an immune response as a yes/no (binary) endpoint. In particular, the endpoint must be

defined prospectively and ideally is one that has biologic relevance and/or is predictive of clinical response. Many immunologic endpoints depend on continuous variables, and it is important to define prospectively a criterion for defining a “positive” or “negative” result along a continuum of data values. In the present study, we have defined criteria based on prior reports [35]. Another important consideration is that the result of the immune response measure must be available promptly if it is to be useful in guiding trial enrollment. Thus, it is important to select an immunologic endpoint that occurs early enough to be meaningful and to design processes for collecting samples and assaying them rapidly so that the data may guide patient enrollment in accord with the study design.

Acknowledgements

Dr. Wages is partially supported by Grant Number K25CA181638-01A1 from the National Cancer Institute. Dr. Petroni is partially supported by Grant Number 1R01CA142859 from the National Cancer Institute. Drs. Slingsluff and Petroni are supported by a Team Science Award from the Melanoma Research Alliance.

References

1. Braun T. The current design of oncology phase I clinical trials: progressing from algorithms to statistical models. *Chinese Clinical Oncology*. 2014; 3:2.
2. Yin, G. *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*. John Wiley & Sons; Hoboken, New Jersey: 2012.
3. Thall PF, Nguyen HQ, Braun T, Qazilbash MH. Using Joint Utilities of the Times to Response and Toxicity to Adaptively Optimize Schedule-Dose Regimes. *Biometrics*. 2013; 69:673–682. [PubMed: 23957592]
4. Wages NA, Conaway MR. Phase I/II adaptive design for drug combination oncology trials. *Statistics in Medicine*. 2014; 33:1990–2003. [PubMed: 24470329]
5. Zang Y, Lee JJ, Yuan Y. Adaptive designs for identifying optimal biological dose for molecularly targeted agents. *Clinical Trials*. 2014; 11:319–327. [PubMed: 24844841]
6. Gao J, Bernatchez C, Sharma P, Radvanyi LG, Hwu P. Advances in the development of cancer immunotherapies. *Trends in Immunology*. 2012; 34:90–98. [PubMed: 23031830]
7. Marshall NA, Galvin KC, Corcoran AM, Boon L, Millis KH. Immunotherapy with PI3K inhibitor and Toll-like receptor agonist induces IFN-gamma+IL-17+ polyfunctional T cells that mediate rejection of murine tumors. *Cancer Research*. 2012; 72:581–591. [PubMed: 22158905]
8. Sikora AG, et al. IFN-alpha enhances peptide vaccine-induced CD8+ T cell numbers, effector function, and antitumor activity. *Journal of Immunology*. 2009; 182:7398–7407.
9. Petroni, G. Design issues for early-stage clinical trials.. In: Disis, M., editor. *Immunotherapy of cancer*. Humana Press; Totowa, New Jersey: 2006. p. 479-85.
10. FDA USDoHaHS, Administration FaD, Research CfBEa. [October 2011] Guidance for Industry Clinical Considerations for Therapeutic Cancer Vaccines. 2011. Available from: <http://www.fda.gov/BiologicsBloodVaccines/GuidanceComplianceRegulatoryInformation/Guidances/default.htm>
11. Rahma OE, Gammoh E, Simon RM, Khleif SM. Is the “3+3” dose-escalation phase I clinical trial design suitable for therapeutic cancer vaccine development? A recommendation for alternative design. *Clinical Cancer Research*. 2014; 20:4758–4767. [PubMed: 25037736]
12. Simon R, Steinberg S, Hamilton M. Clinical trial designs for the early clinical development of therapeutic cancer vaccines. *Journal of Clinical Oncology*. 2001; 19:1848–54. [PubMed: 11251017]
13. Korn E. Nontoxicity endpoints in phase I trial designs for targeted, non-cytotoxic agents. *Journal National Cancer Institute*. 2004; 96:977–978.
14. Parulekar WR, Eisenhauer EA. Phase I trial design for solid tumors studies of targeted, non-cytotoxic agents: theory and practice. *Journal National Cancer Institute*. 2004; 96:970–977.

15. Le Tourneau C, Lee JJ, Siu LL. Dose escalation methods in phase I cancer clinical trials. *Journal National Cancer Institute*. 2009; 101:708–720.
16. LoRusso P, Boerner S, Seymour L. An overview of the optimal planning, design, and conduct of phase I studies of new therapeutics. *Clinical Cancer Research*. 2010; 16:1710–1718. [PubMed: 20215546]
17. Hunsberger S, Rubinstein L, Dancey J, Korn E. Dose escalation trial designs based on a molecularly targeted endpoint. *Statistics in Medicine*. 2005; 24:2171–2181. [PubMed: 15909289]
18. Zhang W, Sargent D, Mandrekar S. An adaptive dose-finding design incorporating both toxicity and efficacy. *Statistics in Medicine*. 2006; 25:2365–2383. [PubMed: 16220478]
19. Polley MY, Cheung YK. Two-stage designs for dose-finding trials with a biologic endpoint using stepwise tests. *Biometrics*. 2008; 64:232–241. [PubMed: 17573866]
20. Wages NA, Tait C. Seamless phase I/II adaptive design for oncology trials of molecularly targeted agents. *Journal of Biopharmaceutical Statistics*. Epub ahead of print 6 June 2014. DOI: 10.1080/10543406.2014.920873.
21. Cannistra S. Challenges and pitfalls of combining targeted agents in phase I studies. *Journal of Clinical Oncology*. 2008; 26:3665–3667. [PubMed: 18669449]
22. Slingluff CL Jr, Petroni GR, Chianese-Bullock KA, et al. A randomized multicenter trial of the effects of melanoma-associated helper peptides and cyclophosphamide on the immunogenicity of a multi-peptide melanoma vaccine. *Journal of Clinical Oncology*. 2011; 29:2924–2932. [PubMed: 21690475]
23. O'Quigley J, Pepe M, Fisher L. Continual reassessment method: a practical design for phase I clinical trials in cancer. *Biometrics*. 1990; 46:33–48. [PubMed: 2350571]
24. Wages NA, Conaway MR, O'Quigley J. Dose finding design for multi-drug combinations. *Clinical Trials*. 2011; 8:380–389. [PubMed: 21652689]
25. Lee SM, Cheung YK. Model calibration in the continual reassessment method. *Clinical Trials*. 2009; 6:227–238. [PubMed: 19528132]
26. Thall PF, Nguyen HQ. Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics*. 2012; 22:785–801. [PubMed: 22651115]
27. Slingluff CL Jr, Petroni GR, Olson WC, et al. Effect of GM-CSF on circulating CD8+ and CD4+ T cell responses to a multi-peptide melanoma vaccine: Outcome of a multicenter randomized trial. *Clinical Cancer Research*. 2009; 15:7036–7044. [PubMed: 19903780]
28. Bechhofer RESTJaGDM. Design and analysis of experiments for statistical selection, screening, and multiple comparisons. John Wiley & Sons; New York: 1995.
29. O'Quigley J, Zohar S. Retrospective robustness of the continual reassessment method. *Journal of Biopharmaceutical Statistics*. 2010; 5:1013–1025. [PubMed: 20721788]
30. R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria: 2013. ISBN 3-900051-07-0, URL <http://www.R-project.org>
31. Cheung, K. dfrm: Dose-finding by the continual reassessment method.. R package version 0.2-2. 2013. <http://CRAN.R-project.org/package=dfrm>
32. O'Quigley J, Conaway MR. Extended model-based designs for more complex phase I clinical trials. *Statistics in Medicine*. 2011; 30:2062–69. [PubMed: 21351287]
33. Jaki T, Clive S, Weir CJ. Principles of dose finding studies in cancer: a comparison of trial designs. *Cancer Chemotherapy and Pharmacology*. 2013; 71:1107–1114. [PubMed: 23299793]
34. Iasonos A, O'Quigley J. Adaptive dose-finding studies: a review of model-guided phase I clinical trials. *Journal of Clinical Oncology*. 2014; 32:2505–2511. [PubMed: 24982451]
35. Slingluff CL, Petroni GR, Chianese-Bullock KA, et al. Immunologic and Clinical Outcomes of a Randomized Phase II Trial of Two Multi-peptide Vaccines for Melanoma in the Adjuvant Setting. *Clinical Cancer Research*. 2007; 13:6386–6395. [PubMed: 17975151]

Table 1

Available treatment regimens for a study of a long-peptide vaccine (LPV7) plus TLR agonists for resected stage IIB–IV melanoma.

Zone	Regimen	Peptide Vaccine	Adjuvant Preparation	Dose of TLR agonists
Z ₁	d ₁	LPV7 + tet	IFA	—
Z ₁	d ₂	LPV7 + tet	PolyICLC	1 mg
Z ₁	d ₃	LPV7 + tet	Resquimod	112.5 mcg
Z ₂	d ₄	LPV7 + tet	IFA + PolyICLC	1 mg / 112.5 mcg
Z ₂	d ₅	LPV7 + tet	IFA + Resquimod	1 mg
Z ₂	d ₆	LPV7 + tet	PolyICLC + Resquimod	112.5 mcg
Z ₃	d ₇	LPV7 + tet	IFA + PolyICLC + Resquimod	1 mg / 112.5 mcg

Table 2

Working model of DLT probabilities under each possible ordering, chosen according to the algorithm of Lee and Cheung [25]

Ordering	d_1	d_2	d_3	d_4	d_5	d_6	d_7
$m = 1$	0.01	0.05	0.12	0.20	0.28	0.36	0.45
$m = 2$	0.01	0.12	0.05	0.28	0.20	0.36	0.45
$m = 3$	0.05	0.01	0.12	0.20	0.36	0.28	0.45
$m = 4$	0.12	0.01	0.05	0.28	0.36	0.20	0.45
$m = 5$	0.05	0.12	0.01	0.36	0.20	0.28	0.45
$m = 6$	0.12	0.05	0.01	0.36	0.28	0.20	0.45

Table 3

Simulated sequential trial of 52 patients illustrating the proposed approach

Pt#	Regimen	x	y	m*	β_{m^*}	Size of A	Pt#	Regimen	x	y	m*	β_{m^*}	Size of A
1	d_1	0	0	-	-	7	28	d_7	1	1	6	0.3997	7
2	d_2	0	0	-	-	7	29	d_7	1	1	6	0.3167	6
3	d_3	0	0	-	-	7	30	d_6	1	1	4	0.2039	6
4	d_4	0	0	-	-	7	31	d_6	0	1	4	0.2305	6
5	d_5	0	1	-	-	7	32	d_6	0	1	4	0.2554	6
6	d_6	0	0	-	-	7	33	d_6	0	0	4	0.2789	6
7	d_7	0	0	-	-	7	34	d_6	0	1	4	0.3011	6
8	d_3	0	1	-	-	7	35	d_6	0	1	4	0.3220	6
9	d_5	0	1	-	-	7	36	d_6	0	1	4	0.3419	7
10	d_7	1	1	2	0.7764	7	37	d_7	0	1	4	0.3704	7
11	d_6	0	1	5	0.8005	7	38	d_7	0	0	4	0.3976	7
12	d_2	0	0	6	0.8024	7	39	d_6	0	1	4	0.4142	7
13	d_4	1	0	6	0.4536	7	40	d_6	0	1	6	0.3690	7
14	d_1	0	0	6	0.4544	7	41	d_6	0	1	6	0.3809	7
15	d_1	0	0	6	0.4731	7	42	d_6	0	1	6	0.3924	7
16	d_3	0	1	6	0.4901	7	43	d_6	0	0	6	0.4034	7
17	d_6	0	1	6	0.5182	7	44	d_6	0	1	6	0.4140	7
18	d_2	0	1	6	0.5226	7	45	d_6	0	0	6	0.4242	7
19	d_4	0	1	6	0.5704	7	46	d_7	0	1	6	0.4461	7
20	d_7	0	1	6	0.5707	7	47	d_7	1	1	6	0.3952	7
21	d_5	0	1	6	0.6028	7	48	d_7	0	1	6	0.4163	7
22	d_7	0	1	4	0.6312	7	49	d_7	0	0	6	0.4366	7
23	d_7	1	1	4	0.4825	7	50	d_7	0	1	6	0.4562	7
24	d_7	0	1	4	0.5296	7	51	d_7	0	1	6	0.4094	7
25	d_7	0	1	4	0.5731	7	52	d_7	0	1	6	0.4284	7
26	d_7	0	1	4	0.6132	7	53	d_7	-	-	-	-	-

Pt#	Regimen	x	y	m*	β_{m*}	Size of A	Pt#	Regimen	x	y	m*	β_{m*}	Size of A
27	d _L	1	0	4	0.4964	7							

Optimal regimen recommendation is d_L after N = 52 patients

Table 4

Results based on 1000 simulated trials. The optimal regimen in each scenario is indicated in bold-type. For each scenario, the table reports the proportion of optimal regimen recommendation (%rec), the proportion allocation to each regimen (%exp), the average (percentiles) sample size N, overall DLT (%DLT) and immune response (%IR) rates, and the percentage of trials stopped early for safety (%Stopped).

Scenario	Regimen	Zone 1			Zone 2			Zone 3			Avg N / Percentiles	%DLT/%IR	%stopped
		d_1	d_2	d_3	d_4	d_5	d_6	d_7					
1	π_T	0.01	0.02	0.05	0.07	0.11	0.14	0.20	45			11.8	0.4
	π_R	0.35	0.35	0.40	0.55	0.60	0.65	0.80	25th=40			60.9	
	%rec	0.00	0.01	0.01	0.07	0.10	0.18	0.62	50th=43				
	%exp	0.08	0.08	0.09	0.12	0.14	0.17	0.33	75th=48				
2	π_T	0.05	0.06	0.08	0.22	0.23	0.25	0.45	45			18.3	2.2
	π_R	0.45	0.50	0.50	0.65	0.65	0.65	0.80	25th=39			60.0	
	%rec	0.06	0.10	0.10	0.25	0.22	0.20	0.05	50th=44				
	%exp	0.12	0.14	0.13	0.18	0.18	0.17	0.08	75th=52				
3	π_T	0.18	0.20	0.25	0.45	0.48	0.54	0.64	29			25.4	25.2
	π_R	0.65	0.70	0.70	0.80	0.80	0.80	0.90	25th=21			70.6	
	%rec	0.19	0.28	0.25	0.01	0.01	0.00	0.00	50th=33				
	%exp	0.25	0.29	0.29	0.06	0.06	0.05	0.01	75th=39				
4	π_T	0.05	0.08	0.10	0.18	0.40	0.20	0.50	44.7			18.9	3.6
	π_R	0.45	0.50	0.50	0.65	0.65	0.65	0.80	25th=39			59.0	
	%rec	0.07	0.10	0.12	0.32	0.09	0.24	0.02	50th=44				
	%exp	0.12	0.13	0.14	0.21	0.14	0.20	0.06	75th=52				
5	π_T	0.05	0.40	0.18	0.50	0.55	0.60	0.65	32.2			24.4	16.2
	π_R	0.55	0.70	0.60	0.65	0.65	0.70	0.80	25th=29			62.5	
	%rec	0.27	0.18	0.39	0.00	0.00	0.00	0.00	50th=35				
	%exp	0.28	0.25	0.34	0.05	0.03	0.04	0.00	75th=41				
6	π_T	0.50	0.55	0.56	0.65	0.66	0.74	0.80	7			45.6	96.0

Scenario	Regimen	Zone 1			Zone 2			Zone 3	Avg N / Percentiles	%DLT/%IR	%stopped
		d_1	d_2	d_3	d_4	d_5	d_6	d_7			
	π_R	0.70	0.70	0.70	0.80	0.80	0.80	0.90	25th=2	70.3	
	%rec	0.01	0.01	0.03	0.00	0.00	0.00	0.00	50th=3		
	%exp	0.32	0.28	0.35	0.02	0.02	0.01	0.00	75th=6		