



OPEN ACCESS

EDITED BY

Ken Chen,
University of Texas MD Anderson
Cancer Center, United States

REVIEWED BY

Wei Wei,
Yale University, United States
Thomas Murray,
University of Minnesota Twin Cities,
United States

*CORRESPONDENCE

Emily C. Zabor
zabore2@ccf.org

SPECIALTY SECTION

This article was submitted to
Cancer Immunity
and Immunotherapy,
a section of the journal
Frontiers in Oncology

RECEIVED 28 May 2022

ACCEPTED 15 November 2022

PUBLISHED 06 December 2022

CITATION

Zabor EC, Kaizer AM, Pennell NA and
Hobbs BP (2022) Optimal predictive
probability designs for randomized
biomarker-guided oncology trials.
Front. Oncol. 12:955056.
doi: 10.3389/fonc.2022.955056

COPYRIGHT

© 2022 Zabor, Kaizer, Pennell and
Hobbs. This is an open-access article
distributed under the terms of the
[Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is
permitted, provided the original author
(s) and the copyright owner(s) are
credited and that the original
publication in this journal is cited, in
accordance with accepted academic
practice. No use, distribution or
reproduction is permitted which does
not comply with these terms.

Optimal predictive probability designs for randomized biomarker-guided oncology trials

Emily C. Zabor^{1*}, Alexander M. Kaizer², Nathan A. Pennell³
and Brian P. Hobbs⁴

¹Lerner Research Institute & Taussig Cancer Institute, Department of Quantitative Health Sciences, Cleveland Clinic, Cleveland, OH, United States, ²Colorado School of Public Health, Department of Biostatistics and Informatics, University of Colorado, Aurora, CO, United States, ³Taussig Cancer Institute, Department of Hematology and Medical Oncology, Cleveland Clinic, Cleveland, OH, United States, ⁴Department of Population Health, University of Texas-Austin, Austin, TX, United States

Introduction: Efforts to develop biomarker-targeted anti-cancer therapies have progressed rapidly in recent years. With efforts to expedite regulatory reviews of promising therapies, several targeted cancer therapies have been granted accelerated approval on the basis of evidence acquired in single-arm phase II clinical trials. And yet, in the absence of randomization, patient prognosis for progression-free survival and overall survival may not have been studied under standard of care chemotherapies for emerging biomarker subpopulations prior to the submission of an accelerated approval application. Historical control rates used to design and evaluate emerging targeted therapies often arise as population averages, lacking specificity to the targeted genetic or immunophenotypic profile. Thus, historical trial results are inherently limited for inferring the potential “comparative efficacy” of novel targeted therapies. Consequently, randomization may be unavoidable in this setting. Innovations in design methodology are needed, however, to enable efficient implementation of randomized trials for agents that target biomarker subpopulations.

Methods: This article proposes three randomized designs for early phase biomarker-guided oncology clinical trials. Each design utilizes the optimal efficiency predictive probability method to monitor multiple biomarker subpopulations for futility. Only designs with type I error between 0.05 and 0.1 and power of at least 0.8 were considered when selecting an optimal efficiency design from among the candidate designs formed by different combinations of posterior and predictive threshold. A simulation study motivated by the results reported in a recent clinical trial studying atezolizumab treatment in patients with locally advanced or metastatic urothelial carcinoma is used to evaluate the operating characteristics of the various designs.

Results: Out of a maximum of 300 total patients, we find that the enrichment design has an average total sample size under the null of 101.0 and a total average sample size under the alternative of 218.0, as compared to 144.8 and 213.8 under the null and alternative, respectively, for the stratified control arm design. The pooled control arm design enrolled a total of 113.2 patients under the null and 159.6 under the alternative, out of a maximum of 200. These average sample sizes that are 23–48% smaller under the alternative and 47–64% smaller under the null, as compared to the realized sample size of 310 patients in the phase II study of atezolizumab.

Discussion: Our findings suggest that potentially smaller phase II trials to those used in practice can be designed using randomization and futility stopping to efficiently obtain more information about both the treatment and control groups prior to phase III study.

KEYWORDS

predictive probability, oncology, clinical trial, phase II, randomized, futility monitoring, two-sample, Bayesian

Introduction

Over the last decade the focus of drug discovery in oncology has shifted away from cytotoxic treatments and toward biomarker-targeted agents. For these types of drugs, such as small molecule inhibitors, antibody drug conjugates, immune checkpoint inhibitors, and monoclonal antibodies, the traditional approach to clinical trial design is not always appropriate. Traditionally, new chemotherapeutic treatments were evaluated in phase I dose-escalation trials to assess safety and identify the maximum tolerated dose. Next, the maximum tolerated dose would be tested for preliminary efficacy in single-arm phase II trials, with a historical control rate forming the basis of comparison. Finally, successful drugs would proceed to phase III, where randomized trials would be used to directly compare efficacy against a standard of care treatment. But the historical control rates used in single-arm phase II studies may not be valid in the context of biomarker-targeted agents. Historical control rates used to design and evaluate emerging targeted therapies often arise as population averages, lacking specificity to the targeted genetic or immunophenotypic profile of interest. Patient prognosis for objective response, progression-free survival, and overall survival may not have been studied under standard of care chemotherapies for emerging biomarker subpopulations prior to phase III. Other factors, such as patient population drift or stage shift, add heterogeneity and bias (1). Consequently, expectations for response and survival for the current biomarker delineated patient populations may differ meaningfully from population averages observed in prior studies of current standard of care therapies. Additionally, in

the specific context of biomarker-targeted agents, heterogeneity of response to standard of care treatments based on the biomarker of interest is also possible, so that the historical control rate may represent an averaging of effect across levels of the biomarker of interest. If the biomarker of interest is prognostic, then response to standard of care treatment in the biomarker-targeted subgroup will differ from the population-averaged response regardless of the treatment being given (2).

Consider the recent clinical trial of atezolizumab for use in metastatic urothelial carcinoma (NCT01375842). Atezolizumab is a programmed death-ligand 1 (PD-L1) blocking monoclonal antibody that was given accelerated approval by the U.S. Food and Drug Administration in May 2016 for the treatment of patients with locally advanced or metastatic urothelial carcinoma who had disease progression following platinum-containing chemotherapy. The approval was based on the results of a single-arm phase II study in 310 patients (3). The phase II study used a hierarchical fixed-sequence testing procedure to test increasingly broad subgroups of patients based on PD-L1 status, and found overall response rates of 26% (95% CI: 18–36), 18% (95% CI: 13–24), and 15% (95% CI 11–19) in patients with $\geq 5\%$ PD-L1-positive immune cells (IC2/3 subgroup), in patients with $\geq 1\%$ PD-L1-positive immune cells (IC1/2/3 subgroup), and in all patients, respectively (3). All three rates exceeded the historical control rate of 10%. Then, in March 2021, the approval in this indication was voluntarily withdrawn by the sponsor following negative results from a randomized phase III study (NCT02302807) (4). In the phase III study, 931 patients were randomly assigned to receive atezolizumab or chemotherapy in a 1:1 ratio, and the same hierarchical fixed-

sequence testing procedure as in the phase II study was used. The phase III study found that overall survival did not differ significantly between the atezolizumab and chemotherapy groups of the IC2/3 subgroup (median survival 11.1 months [95% CI: 8.6-15.5] versus 10.6 months [95% CI: 8.4-12.2]), so no further testing was conducted for the primary endpoint (4). Further analyses revealed that while the response rates to atezolizumab were comparable to those seen in the phase II study, the response rates to chemotherapy were much higher than the historical control rate of 10%. The overall response rates to chemotherapy were 21.6% (95% CI: 14.5-30.2), 14.7% (95% CI: 10.9-19.2), and 13.4% (95% CI: 10.5-16.9) for the IC2/3 subgroup, IC1/2/3 subgroup, and all patients, respectively. The overall response rates to atezolizumab were 23% (95% CI: 15.6-31.9), 14.1% (95% CI: 10.4-18.5), and 13.4% (95% CI: 10.5-16.9) for the IC2/3 subgroup, IC1/2/3 subgroup, and all patients, respectively. These results indicate that PD-L1 status is a prognostic biomarker, with higher response rates to both the standard of care chemotherapies that comprised the control arm and to atezolizumab treatment in the biomarker-enriched subgroup (2).

The example of atezolizumab in metastatic urothelial carcinoma is one of many. Between 2015 and 2021, the U.S. Food and Drug Administration (FDA) approved six antibodies against PD-L1 or programmed death 1 (PD-1) for 75 cancer indications, and 35 of these approvals were accelerated based on early phase trial results (5). This extremely rapid pace of development within a single drug class was unprecedented, and led to ten such “dangling” accelerated approvals, which are approved indications for which the confirmatory trial showed no benefit, yet the drug remained on the market for that indication (5). Other voluntary withdrawals following confirmatory trial results include durvalumab treatment for metastatic urothelial carcinoma, and nivolumab and pembrolizumab treatments for metastatic small-cell lung cancer (5–9). These failed confirmatory phase III trials highlight both the need for rapid development of new treatments in patient populations with few therapeutic options, and the need for innovations that facilitate more rigorous designs of phase II trials for targeted therapies. To overcome many issues, including those associated with the use of historical control rates, randomization may be unavoidable in this setting. Arguments for the use of randomization in the phase II setting have been prominent for over a decade (10–13). In addition to addressing the inconvenient reality that historical control rates often used in single arm studies may have limited value for novel targets, randomized phase II trials can also overcome issues of selection bias and patient heterogeneity. Randomized designs that incorporate futility stopping can provide information on current control rates to the treatment under study while also stopping ineffective treatments early. The FDA’s Project Frontrunner (<https://www.fda.gov/about-fda/oncology-center-excellence/project-frontrunner>) encourages the use of

randomized controlled trials earlier in the drug development process, especially if accelerated approval is one of the trial goals.

This article proposes three different biomarker-guided randomized phase II trial designs with optimal efficiency predictive probability monitoring for futility. Using the trial of atezolizumab for metastatic urothelial carcinoma as a case study and motivating example, we compare designs based on their traditional statistical properties of type I error and power through simulation study. The designs are also evaluated based on the number of patients enrolled, the number of patients treated, the number of patients who undergo biomarker testing, and accurate estimation of the response rates of interest. Our findings suggest that potentially smaller phase II trials to those used in practice can be designed using randomization and futility stopping to efficiently obtain more information about both the treatment and control groups prior to phase III study.

Materials and methods

This paper focuses on the setting of a two-sample randomized trial with a binary outcome. We will refer to the binary outcome as “response” and use “response rate” to describe the probability of a response throughout the article, in line with the motivating example of the phase II study of atezolizumab in metastatic urothelial carcinoma, which estimated response rates among biomarker subpopulations and compared to the historical average in the primary analysis. Any hypothetical measure of efficacy, however, such as progression-free survival, could be used with the design methodology proposed. Each patient enrolled in the trial is denoted by i , and they either have a response such that $x_i=1$ or do not have a response such that $x_i=0$. Then $X = \sum_{i=1}^n x_i$ represents the total number of responses out of n currently observed patients, up to a maximum sample size of N total patients. The probability of response is denoted p , where p_0 represents the null response rate under the standard of care treatment and p_1 represents the alternative response rate under the experimental treatment. We wish to test the null hypothesis $H_0: p_1 \leq p_0$ versus the alternative hypothesis $H_1: p_1 > p_0$.

The Bayesian statistical paradigm is based on a mathematical approach to combine prior distributions, which reflect prior beliefs about parameters such as the true response rate, with observed data (e.g., the observed number of responses in a given trial) to obtain posterior distributions of the model parameters. Here we assume a beta-binomial model, based on the computational ease and its popularity for use in the context of sequential trial monitoring. The prior distribution of the response rate has a beta distribution $Beta(a_0, b_0)$. We specifically use a $Beta(0.5, 0.5)$ prior distribution, which reflects the effective information of a single patient’s observation. We also perform sensitivity analyses using $Beta(1, 1)$, $Beta(0, 0)$, $Beta(2, 2)$, $Beta(0.75, 0.25)$, and $Beta(0.25, 0.75)$

priors, to examine the variation in operating characteristics across a range of priors. Our data X follow a binomial distribution $\text{bin}(n, p)$. We combine the likelihood function for the observed data $L_x(p) \propto p^x(1-p)^{n-x}$ with the prior to obtain the posterior distribution of the response rate, which follows the beta distribution $p|x \sim \text{Beta}(a_0+x, b_0+n-x)$. Posterior probabilities represent the probability that the experimental response rate exceeds the null response rate based on the data accrued so far in the trial. Posterior decision can be obtained by applying a clinically relevant threshold, θ , to the posterior distribution. We would declare a treatment efficacious if the posterior probability exceeded the posterior threshold, i.e. $\text{Pr}(p > p_0 | X) > \theta$.

Bayesian predictive probability monitoring has been a popular approach for designing clinical trials with sequential futility monitoring (14–17). It is a natural fit for this type of trial, as it allows for flexibility in both the timing and the number of looks. In addition, predictive probability is an intuitive interim monitoring strategy because it tells the investigator what the chances are of declaring the treatment efficacious at the end of the trial if enrollment is continued to the maximum planned sample size. At any given interim look, the posterior predictive distribution of the number of future responses X^* in the remaining $n^* = N - n$ future patients follows a beta-binomial distribution $\text{Beta-binomial}(n^*, a_0+x, b_0+n-x)$. The posterior predictive probability (PPP) represents the probability that the experimental treatment will be declared efficacious at the end of the trial when full enrollment is reached, conditional on the currently observed data and the specified priors. The posterior predictive probability is calculated as $PP = \sum_{x^*=0}^{n^*} \text{Pr}(X^* = x^* | x) \times I(\text{Pr}(p > p_0 | X, X^* = x^*) > \theta)$. A second predictive threshold θ^* is defined, and we would stop the trial early for futility if the predictive probability dropped below the given threshold, i.e. $PPP < \theta^*$. Predictive thresholds closer to 0 lead to less frequent stopping for futility whereas predictive thresholds closer to 1 lead to frequent stopping in the absence of almost certain probability of success.

When designing a trial with sequential predictive probability monitoring for futility, it is essential to ensure the trial conforms to traditional standards for type I error control and power. To do so, we must examine the operating characteristics of a variety of designs based on combinations of the posterior threshold θ and the predictive threshold θ^* and select a single design for use in the trial. In earlier work, we proposed two optimization criteria to help select from among a variety of designs in the setting of a one-sample study (18). Here we focus on the optimal efficiency design, defined as the combination of posterior and predictive thresholds with minimal average sample size under the null and maximal average sample size under the alternative, subject to constraints on the type I error and power, and extend the approach to the setting of a two-sample study for targeted therapy.

The simulation study is based on the phase II trial of atezolizumab in metastatic urothelial carcinoma. There are three independent biomarker subgroups based on the

percentage of PD-L1-expressing immune cells: IC0 (<1%), IC1 ($\geq 1\%$ and <5%), and IC2/3 ($\geq 5\%$). The subgroups have equal prevalence of 33% in the study population. We consider a standard of care arm denoted “chemotherapy” and an experimental treatment arm denoted “atezolizumab”. The null response rate was based on the stated historical control rate of 10% (3). As no specific alternative was specified, we examine subtype-specific alternative rates of 10%, 20%, and 30% in the IC0, IC1, and IC2/3 subgroups, respectively, in line with what we expect for a predictive biomarker, for which the treatment effect differs according to levels of the biomarker of interest (2). We further investigate a simulation setting where the treatment effect is a homogeneous 30% across the three biomarker subgroups. Interim looks for futility are planned after every 10 patients. A random number of responses was generated for every 10 patients up to the maximum sample size, based on a binomial distribution with the setting-specific response rate. 1000 simulated datasets were generated under the null and 1000 simulated datasets were generated under the alternative. We considered posterior thresholds θ of 0.7, 0.74, 0.78, 0.82, 0.86, 0.9, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, and 0.99, and predictive thresholds θ^* of 0.05, 0.1, 0.15, and 0.2. For each combination of posterior and predictive threshold, the predictive probability that the experimental treatment arm response rate exceeds the standard of care arm response rate at the end of the trial is calculated at each interim look until it either fell below the given predictive threshold or the end of the trial was reached, whichever came first. If the end of the trial was reached, the trial was considered positive if the predictive probability was greater than the given posterior threshold and negative otherwise. If halted early for futility, the trial was considered negative. We propose and compare three strategies for conducting randomized two-sample biomarker-guided designs that use optimal efficiency predictive probability monitoring for futility: a pooled control arm design, a stratified control arm design, and an enrichment design.

The pooled control arm design is depicted in Figure 1A. In this design, patients are randomized to atezolizumab or chemotherapy in a 3:1 ratio. PD-L1 testing is performed only on patients randomized to receive atezolizumab. The atezolizumab arm is separated into three biomarker-specific treatment subgroups. This design has a maximum sample size of 200: $n=50$ patients for the pooled chemotherapy control arm and $n=50$ for each PD-L1 biomarker-specific atezolizumab arm.

The stratified control arm design is depicted in Figure 1B. In this design, PD-L1 testing is conducted on all patients. Then, within each subgroup, patients are randomized to atezolizumab or chemotherapy in a 1:1 ratio. This design has a maximum sample size of 300: $n=50$ for each PD-L1 biomarker-specific chemotherapy and atezolizumab arm.

The enrichment design is depicted in Figure 1C. This design is equivalent to the pooled design at stage 1. If all subgroups stop for futility in stage 1, then the trial is stopped. Otherwise, at the

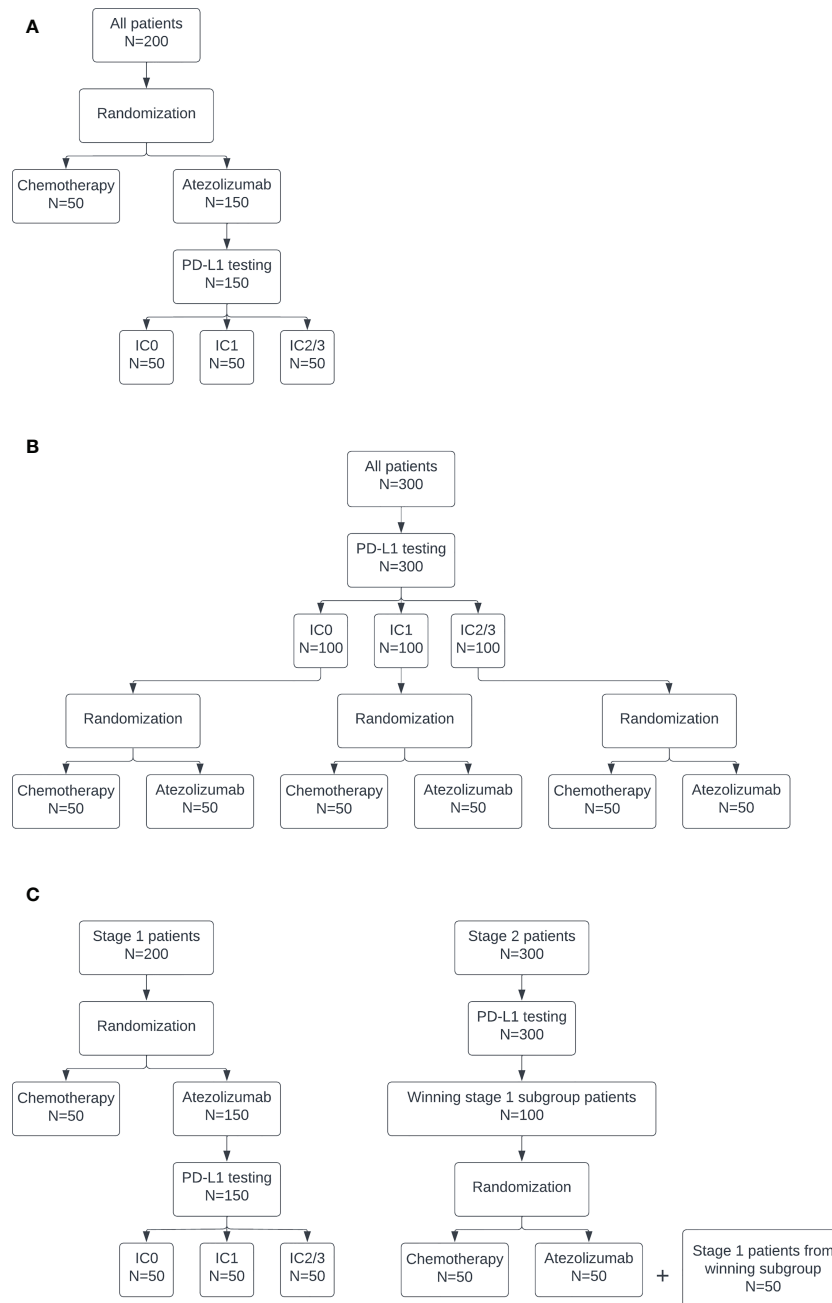


FIGURE 1
Diagrams of the (A) pooled control arm, (B) stratified control arm, and (C) enrichment designs.

end of stage 1, the subgroup with the highest posterior predictive probability, subject to some lower bound, continues to stage 2. The lower bound was selected as the 80th percentile of maximum posterior predictive probability across the three subgroups at stage 1 under the null. This percentile was used to target a 20% rate of moving a subgroup forward when all of the subgroups are truly null. This higher rate of stage 1 type I error is consistent with the phase objective, which emphasizes

acquiring more data on safety and efficacy for promising treatments in early phase trial designs of this type. The actual type I error at stage 1 was calculated as the proportion of simulated trials under the global null, i.e. if all three biomarker-specific subgroups had a true response rate of 10%, in which the subgroup with maximum posterior predictive probability exceeded the lower bound and did not stop early for futility, so was selected to continue to stage 2. The power at

stage 1 was calculated as the proportion of simulated trials under the alternative in which the IC2/3 subgroup was selected as having the maximum posterior predictive probability, subject to the lower bound, and did not stop early for futility. In stage 2, PD-L1 testing is conducted on all patients. Only those patients belonging to the subgroup selected in stage 1 are enrolled on the trial and randomized 1:1 to atezolizumab or chemotherapy. The stage 1 treatment group results for the selected subgroup, if any, are carried forward into stage 2. An additional $n=100$ biomarker-specific patients are enrolled at stage 2, for a total maximum sample size of 300.

For the pooled and stratified designs, the type I error was calibrated in the IC0 subgroup null setting as the proportion of simulated trials in which the IC0 subgroup was declared positive as compared to the control group. The power was calibrated in the IC2/3 subgroup alternative setting as the proportion of simulated trials in which the IC2/3 subgroup was declared positive as compared to the control group. The IC1 subgroup was considered an intermediate setting and no results were calibrated based on this subgroup. For the enrichment design, the type I error was calibrated based on the stage 2 results as the proportion of simulated trials under the null in which the selected subgroup, if any, was declared positive. The power was calibrated based on the stage 2 results as the proportion of simulated trials under the alternative in which the selected subgroup, if any, was declared positive. The stage 2 calibration occurs in fewer than 1000 simulated trials, as only the specific trials in which a subgroup was selected to continue to stage 2 were used. The resulting design options were limited to those that resulted in a type I error rate between 0.05 and 0.1 and a power of at least 0.8. Then, the efficiency distance metric was calculated as described in Zabor et al. (18) using total trial-level sample sizes. The design with the minimal efficiency distance metric was identified as the optimal design.

All results were generated using R software version 4.2.0 (19) along with the 'ppseq' R package (20).

Results

The accuracy and efficiency results for all 56 possible posterior and predictive threshold combinations are plotted in Figure 2. Each point represents the combination of one posterior threshold and one predictive threshold, and the orange diamond on each plot indicates the design that was identified to have optimal efficiency while maintaining type I error between 0.05 and 0.1 and with power of at least 0.8. For the enrichment design, only threshold combinations that ever proceeded to stage 2 are plotted so there are only 36 points, as 20 threshold combinations never resulted in designs that continued to stage 2. The optimal efficiency pooled control arm design had posterior threshold 0.9 and predictive threshold 0.1, the optimal efficiency stratified control arm design had posterior

threshold 0.9 and predictive threshold 0.2, and the optimal efficiency enrichment design had posterior threshold 0.96 and predictive threshold 0.15. We see that the different threshold combinations result in a wide range of results, some with low power < 0.5 or high type I error > 0.2 , and some with too low average sample size under the alternative or too high average sample size under the null. By applying the optimal efficiency criteria, we are able to identify a design that seeks to maximize sample size under the alternative and minimize sample size under the null, within the pre-specified range of type I error and minimum power.

The type I error and power for each biomarker-specific subgroup under the pooled control arm and stratified control arm designs, and the overall type I error and power for the enrichment design, are presented in Table 1. We see that both the pooled control arm and stratified control arm designs result in reasonable power to detect an effect for the IC2/3 subgroup, with slightly higher power of 0.82 in the stratified control arm design as compared to 0.8 in the pooled control arm design. The pooled control arm design and stratified control arm design both have type I error of 0.07 for the IC2/3 subgroup. Both the pooled control arm and stratified control arm designs have very low power < 0.5 to detect the IC1 subgroup and < 0.1 to detect the IC0 subgroup. Only overall results are available for the enrichment design, which results in a type I error and power of 0.09 and 0.73, respectively, for stage 1. The type 1 error rate of 0.09 for stage 1 means that under the null 91% of simulated trials did not proceed to stage 2; however, 4.2% proceeded to stage 2 with the IC2/3 subgroup, 2.9% proceeded to stage 2 with the IC1 subgroup, and 1.9% proceeded to stage 2 with the IC0 subgroup. The power of 0.73 for stage 1 means that under the alternative 73% of simulated trials proceeded to stage 2, and all of them did so with the IC2/3 subgroup; the remaining 27% of simulated trials did not proceed to stage 2. The overall type I error rate for stage 2 of the enrichment design was 0.09 and the overall power for stage 2 of the enrichment design was 0.86. Since the IC2/3 subgroup was exclusively carried forward to stage 2 of the enrichment design under the alternative, this could also be considered the power for the IC2/3 subgroup, and it exceeds the power of 0.82 of the stratified control arm design and the power of 0.8 of the pooled control arm design.

The resulting average sample sizes under the null and alternative for each selected optimal efficiency design are presented in Table 2. We can directly compare the total sample sizes between the stratified control arm and enrichment designs, which can each enroll a maximum of 300 patients. We find that the enrichment design has a lower average total sample size under the null of 101.0 and a higher total average sample size under the alternative of 218.0, as compared to 144.8 and 213.8 under the null and alternative, respectively, for the stratified control arm design. This occurs because the use of the pooled control arm design at stage 1 combined with the low stage 1 type I error rate and high stage 2 power means that

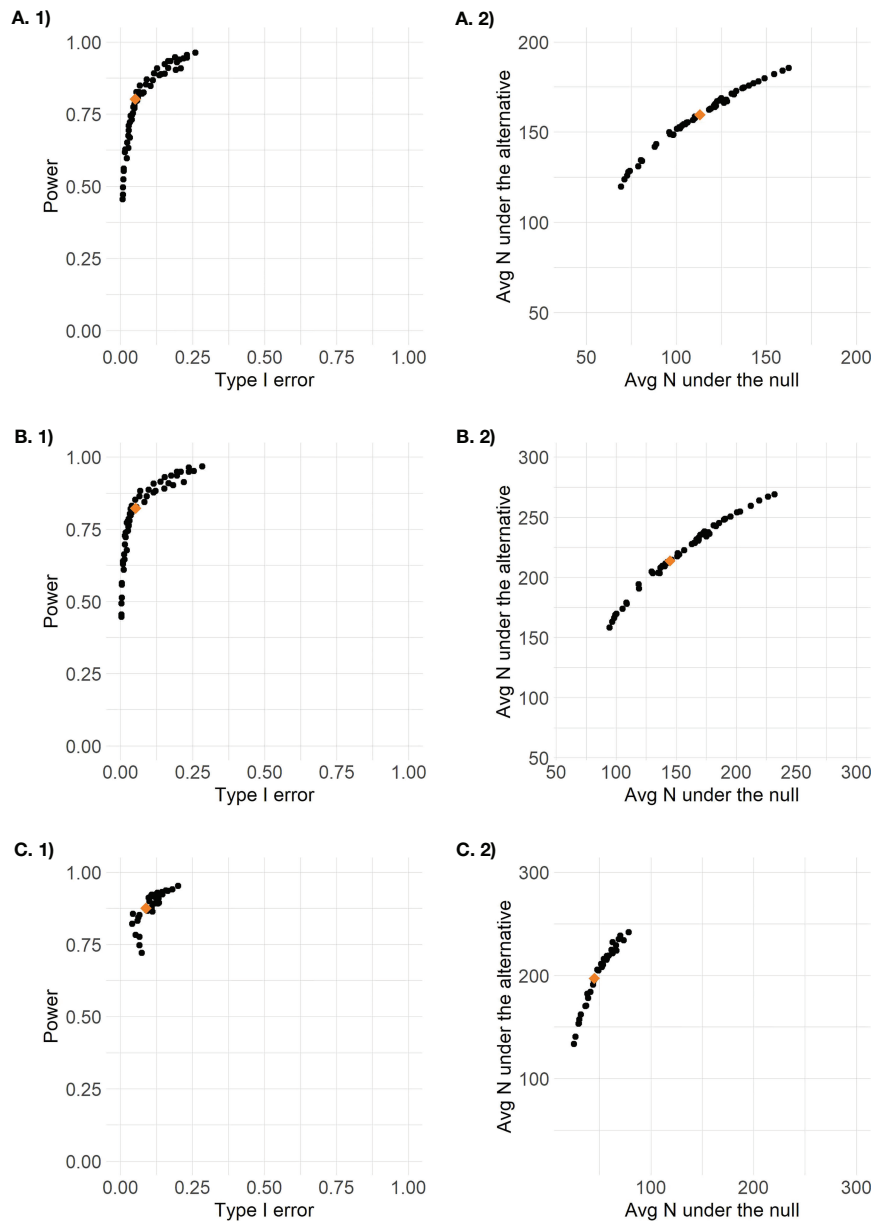


FIGURE 2

Plots of design options for (A) the pooled control design, (B) the stratified control design, and (C) the enrichment design (stage 2 results only) based on 1) accuracy defined as type I error by power and 2) efficiency defined as average total sample size under the null versus average total sample size under the alternative.

the enrichment design frequently stops after stage 1 under the null and frequently continues to full enrollment in stage 2 under the alternative. The pooled control arm design enrolled a total of 113.2 patients under the null and 159.6 patients under the alternative, out of a maximum of 200. The sample sizes of patients treated with atezolizumab are directly comparable for the pooled control arm and stratified control arm designs, which can each treat a maximum of 150 patients with atezolizumab, 50 per biomarker subgroup. The pooled control arm design treated

more patients with atezolizumab on average as compared to the stratified control arm design, 78.2 and 111.7 under the null and alternative, respectively, as compared to 72.4 and 106.9. The enrichment design can treat a maximum of 200 patients with atezolizumab across the two stages, and treats an average of 68.0 patients under the null and 137.0 patients under the alternative. We find that there is minimal variation in operating characteristics across the six prior distributions examined in sensitivity analysis (Supplemental Table 1).

TABLE 1 Type I error and power for each biomarker-specific subgroup under the pooled control arm and stratified control arm designs, and the overall type I error and power for the enrichment design.

	Pooled		Stratified		Enrichment			
	Type I error	Power	Type I error	Power	Stage 1		Stage 2	
					Type I error	Power	Type I error	Power
IC0	0.05	0.05	0.05	0.08	–	–	–	–
IC1	0.05	0.40	0.08	0.45	–	–	–	–
IC2/3	0.07	0.80	0.07	0.82	–	–	–	–
Overall	–	–	–	–	0.09	0.73	0.09	0.86

For the simulation setting where the treatment effect was homogeneous across the biomarker subgroups, we find that both the pooled control arm and stratified control arm designs perform well, with consistent type I error, power, and sample size across the three subgroups (Supplemental Tables 2 and 3). The enrichment design, however, has low stage 1 power (Supplemental Table 2). This is because there is no distinction between the three biomarker subgroups with respect to which should be the winning subgroup, so it is easy for the method to make the “wrong” choice, as we defined power in this setting based on a predictive biomarker as the proportion of simulated trials under the alternative in which the IC2/3 subgroup was selected as having the maximum posterior predictive probability, subject to the lower bound, and did not stop early for futility. However, the enrichment design performs well at stage 2, since any subgroup selected to continue would have equal chance of success in the second stage of the trial, given the homogeneous true response rates (Supplemental Table 2).

The enrichment design requires testing the largest number of patients, with 450 patients requiring PD-L1 testing if the design proceeds to stage 2, whereas the pooled design only tests 150 patients and the stratified design tests 300 patients. The pooled control arm and enrichment designs cannot address the question of whether the biomarker is predictive of response to the standard of care treatment, since they do not estimate response rates separately within each biomarker subgroup,

though the enrichment design can fully characterize the response rate to standard of care treatment within the selected stage 2 biomarker subgroup. Only the stratified control arm design fully characterizes the response rates to standard of care treatment within each biomarker subgroup, and can therefore address the question of whether the biomarker is predictive of response for both standard of care and targeted therapies.

Discussion

This article presented three different optimal efficiency predictive probability designs for randomized biomarker-guided oncology clinical trials. A simulation study was conducted to demonstrate that posterior and predictive thresholds can be selected to maintain appropriate levels of type I error between 5% and 10% and power of at least 80% in all three designs. This work was motivated by the case study of atezolizumab for the treatment of patients with locally advanced or metastatic urothelial carcinoma who had disease progression following platinum-containing chemotherapy. In the phase II trial that led to accelerated approval, 310 patients were enrolled and treated. The trial did not include any futility stopping rules and planned for a total sample size of 300 patients, expecting about 100 patients in each biomarker subgroup. So in practice the realized and expected sample sizes are broadly equivalent. By

TABLE 2 Average sample size under the null (“Avg N Null”) and average sample size under the alternative (“Avg N Alt”) by design and treatment subgroup.

	Pooled		Stratified		Enrichment	
	Avg N Null	Avg N Alt	Avg N Null	Avg N Alt	Avg N Null	Avg N Alt
Control	35.0	48.0	72.4	106.9	33.1	81.3
IC0	25.9	26.7	23.6	23.8	–	–
IC1	26.3	38.8	24.4	37.5	–	–
IC2/3	26.1	46.2	24.3	45.6	–	–
Total Atezolizumab	78.2	111.7	72.4	106.9	68.0	137.0
Total Enrolled	113.2	159.6	144.8	213.8	101.0	218.0

comparison, the three proposed designs, which incorporate futility stopping and randomization, result in average phase II sample sizes that are 23–48% smaller under the alternative and 47–64% smaller under the null, and therefore represent a more efficient use of both human and financial resources.

At the same time, the three proposed designs provide additional information about response rates to standard of care treatment in the control arms, thus potentially avoiding the pitfall of the atezolizumab trial, in which the historical control rate used to show efficacy in the phase II trial proved to be far below the actual response rate to standard of care treatment in the biomarker-targeted subgroup of patients. The stratified control arm design results in the most information, allowing one to determine if the biomarker of interest is predictive of response to either the standard of care treatment or the experimental treatment or both. The enrichment design only characterizes the response rate to the standard of care treatment in the biomarker subgroup that is selected to continue to stage 2. But both the pooled control arm design and stage 1 of the enrichment design are superior to use of a historical control rate, since the patient population of the control group is identical to that of the treatment group in both timing and characteristics as a result of randomization. So these designs not only have lower average expected sample sizes than the 310 used in the atezolizumab phase II trial, but also have properties such as control groups and sequential futility monitoring that facilitate valid inference of comparative effectiveness and improve decision-making for continuation to phase III. The phase III trial of atezolizumab for this patient population randomized 931 patients who could have been available to enroll in trials of more promising treatments, or could have avoided the rigors of a clinical trial altogether in favor of the established standard of care treatment.

The decision of which design to select will depend on a number of factors. One is the costs of biomarker testing, including invasiveness of the testing procedure, turnaround time, and actual financial cost. The enrichment design tests the most patients whereas the pooled control arm design tests the fewest patients. So in the case of extremely invasive or expensive tests, the pooled control arm design may be preferred. Another consideration is the prevalence of the biomarker in the population. The enrichment design in stage 2 requires testing all patients in order to identify and enroll only patients with the biomarker of interest, which could be prohibitively expensive or time consuming in the setting of a rare biomarker. In that case, the pooled control arm design may be preferable since all patients are enrolled and the control group will more easily reach full enrollment by containing a mix of patients regardless of biomarker status. But any of the proposed designs could result in a more efficient use of resources in the setting of a rare biomarker, considering both the ability to stop the trial early for futility, and the potential to avoid embarking

on a confirmatory trial without adequate information about the population under study. A third consideration is clinical evidence for the biomarker being prognostic in nature, leading to differential response rates across biomarker subgroups on even standard of care therapies. If there is preliminary evidence or biological plausibility that such an effect might exist, the stratified control arm design may be preferable since it fully characterizes the response rates of the control groups within each biomarker subgroup. While the trial would “fail” in this setting, since it would find no difference between the control and treatment arms within the biomarker subgroups, the information from the stratified control arm design would be useful for planning future studies. And a final consideration is clinical evidence for the biomarker being predictive of experimental treatment response. If there is a strong belief that only biomarker positive patients will benefit from the treatment under study, then the enrichment design may be best as it enrolls more patients in only the selected subgroup at stage 2. Additional sample size savings could be achieved by eliminating the control group at stage 1 of the enrichment design, though the properties of such a design were not investigated in detail here.

The main limitation to the use of these designs is the computational intensity required to perform calibration across a variety of posterior and predictive thresholds for the setting of interest in order to identify a design with the desired operating characteristics of type I error and power. While we have developed open-source R software for the design of single-arm and two-arm optimal sequential predictive probability designs, specialized programming using the functions from the ‘ppseq’ R package would be required to design a pooled control arm, stratified control arm, or enrichment design of the type presented here. Moreover, a large memory server is needed to complete the computations in any reasonable time span. However, once the thresholds have been selected, decision rule tables for early stopping can be generated so that no mid-trial computations would be necessary.

As rapid development of biomarker-targeted agents in oncology continues, new implementations of existing statistical methods such as those presented here will represent the most nimble way for the statistical design of trials to keep up with the changing context of cancer treatment. Randomization is an old statistical tool that has not traditionally been employed in early phase oncology clinical trials due to the sample size requirements. But in the context of increasingly large early phase clinical trials that can include hundreds of patients across multiple cancer types or multiple biomarker levels or both, randomization is no longer the constraint that it once was. This kind of efficient design also stresses the importance of mandating that all patients enrolled to biomarker-targeted trials have the biomarker of interest tested at enrollment (as opposed to only a subset of those with tissue available) so that the most

accurate information about efficacy within biomarker groups can be obtained. Here we have demonstrated that it is possible to conduct randomized phase II trials with smaller sample sizes than those being used in practice for single-arm trials. Moreover, Bayesian sequential design with predictive probability yields more efficient and informative early phase clinical trial results than the standard frequentist approaches commonly implemented in practice.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: https://github.com/zabore/article-code-repository/tree/master/Zabor_Randomized-Biomarker-Guided-Designs/data.

Author contributions

EZ, AK, and BP contributed to conception and design of the study. EZ performed the statistical analysis and wrote the first draft of the manuscript. All authors contributed to the article and approved the submitted version.

Funding

AK is supported by NHLBI K01HL151754.

References

- Cannistra SA. Phase II trials in journal of clinical oncology. *J Clin Oncol Off J Am Soc Clin Oncol* (2009) 27(19):3073–6. doi: 10.1200/JCO.2009.23.1811
- Ballman KV. Biomarker: Predictive or prognostic? *J Clin Oncol* (2015) 33(33):3968–71. doi: 10.1200/JCO.2015.63.3651
- Rosenberg JE, Hoffman-Censits J, Powles T, van der Heijden MS, Balar AV, Necchi A, et al. Atezolizumab in patients with locally advanced and metastatic urothelial carcinoma who have progressed following treatment with platinum-based chemotherapy: a single-arm, multicentre, phase 2 trial. *Lancet* (2016) 387(10031):1909–20. doi: 10.1016/S0140-6736(16)00561-4
- Powles T, Durán I, van der Heijden MS, Loriot Y, Vogelzang NJ, De Giorgi U, et al. Atezolizumab versus chemotherapy in patients with platinum-treated locally advanced or metastatic urothelial carcinoma (IMvigor211): a multicentre, open-label, phase 3 randomised controlled trial. *Lancet* (2018) 391(10122):748–57. doi: 10.1016/S0140-6736(17)33297-X
- Beaver JA, Pazdur R. "Dangling" accelerated approvals in oncology. *New Engl J Med* (2021) 384(18):e68. doi: 10.1056/NEJMp2104846
- Powles T, van der Heijden MS, Castellano D, Galsky MD, Loriot Y, Petrylak DP, et al. Durvalumab alone and durvalumab plus tremelimumab versus chemotherapy in previously untreated patients with unresectable, locally advanced or metastatic urothelial carcinoma (DANUBE): a randomised, open-label, multicentre, phase 3 trial. *Lancet Oncol* (2020) 21(12):1574–88. doi: 10.1016/S1470-2045(20)30541-6
- Owonikoko TK, Park K, Govindan R, Ready N, Reck M, Peters S, et al. Nivolumab and ipilimumab as maintenance therapy in extensive-disease small-cell lung cancer: CheckMate 451. *J Clin Oncol Off J Am Soc Clin Oncol* (2021) 39(12):1349–59. doi: 10.1200/JCO.20.02212
- Spigel DR, Vicente D, Ciuleanu TE, Gettinger S, Peters S, Horn L, et al. Second-line nivolumab in relapsed small-cell lung cancer: CheckMate 331(☆). *Ann Oncol Off J Eur Soc Med Oncol* (2021) 32(5):631–41. doi: 10.1016/j.annonc.2021.01.071
- Rudin CM, Awad MM, Navarro A, Gottfried M, Peters S, Csösz T, et al. Pembrolizumab or placebo plus etoposide and platinum as first-line therapy for extensive-stage small-cell lung cancer: Randomized, double-blind, phase III KEYNOTE-604 study. *J Clin Oncol Off J Am Soc Clin Oncol* (2020) 38(21):2369–79. doi: 10.1200/JCO.20.00793
- Gan HK, Grothey A, Pond GR, Moore MJ, Siu LL, Sargent D. Randomized phase II trials: inevitable or inadvisable? *J Clin Oncol Off J Am Soc Clin Oncol* (2010) 28(15):2641–7. doi: 10.1200/JCO.2009.26.3343
- Korn EL, Arbuck SG, Pluda JM, Simon R, Kaplan RS, Christian MC. Clinical trial designs for cytostatic agents: are new approaches needed? *J Clin Oncol Off J Am Soc Clin Oncol* (2001) 19(1):265–72. doi: 10.1200/JCO.2001.19.1.265
- Ratain MJ, Sargent DJ. Optimising the design of phase II oncology trials: the importance of randomisation. *Eur J Cancer* (2009) 45(2):275–80. doi: 10.1016/j.ejca.2008.10.029
- Rubinstein LV, Korn EL, Freidlin B, Hunsberger S, Ivy SP, Smith MA. Design issues of randomized phase II trials and a proposal for phase II screening trials. *J Clin Oncol Off J Am Soc Clin Oncol* (2005) 23(28):7199–206. doi: 10.1200/JCO.2005.01.149
- Dmitrienko A, Wang MD. Bayesian Predictive approach to interim monitoring in clinical trials. *Stat Med* (2006) 25(13):2178–95. doi: 10.1002/sim.2204

Conflict of interest

BH declares advising/consulting for Amgen, Bayer HealthCare Pharmaceuticals Inc., Telperian, and STCube Pharmaceuticals. NP declares advising/consulting for AstraZeneca, Merck, Pfizer, Eli Lilly/LOXO, Genentech, BMS, Amgen, Mirati, Inivata, G1 Therapeutics, Viosera, Xencor, Janssen, Boehringer Ingelheim, and Sanofi-Genzyme.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fonc.2022.955056/full#supplementary-material>

15. Lee JJ, Liu DD. A predictive probability design for phase II cancer clinical trials. *Clin trials (London England)* (2008) 5(2):93–106. doi: 10.1177/1740774508089279
16. Saville BR, Connor JT, Ayers GD, Alvarez J. The utility of Bayesian predictive probabilities for interim monitoring of clinical trials. *Clin trials (London England)* (2014) 11(4):485–93. doi: 10.1177/1740774514531352
17. Hobbs BP, Chen N, Lee JJ. Controlled multi-arm platform design using predictive probability. *Stat Methods Med Res* (2018) 27(1):65–78. doi: 10.1177/0962280215620696
18. Zabor EC, Kaizer AM, Garrett-Mayer E, Hobbs BP. Optimal sequential predictive probability designs for early-phase oncology expansion cohorts. *JCO Precis Oncol* (2022) 6:e2100390. doi: 10.1200/PO.21.00390
19. R Core Team. *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing (2022).
20. Zabor EC, Hobbs BP, Kane MJ. *Ppseq: Design clinical trials using sequential predictive probability monitoring*. R package (2022). Available at: <https://CRAN.R-project.org/package=ppseq>