Published in final edited form as:

Stat Med. 2016 October 30; 35(24): 4285–4305. doi:10.1002/sim.6989.

Utility-based designs for randomized comparative trials with categorical outcomes

Thomas A. Murray^{a,*}, Peter F. Thall^a, and Ying Yuan^a

Thomas A. Murray: TAMurray@MDAnderson.org; Peter F. Thall: Rex@MDAnderson.org; Ying Yuan: YYuan@MDAnderson.org

^aDepartment of Biostatistics, The University of Texas MD Anderson Cancer Center, Unit 1411, P.O. Box 301402, Houston, TX, 77030-1402, USA

Abstract

A general utility-based testing methodology for design and conduct of randomized comparative clinical trials with categorical outcomes is presented. Numerical utilities of all elementary events are elicited to quantify their desirabilities. These numerical values are used to map the categorical outcome probability vector of each treatment to a mean utility, which is used as a one-dimensional criterion for constructing comparative tests. Bayesian tests are presented, including fixed sample and group sequential procedures, assuming Dirichlet-multinomial models for the priors and likelihoods. Guidelines are provided for establishing priors, eliciting utilities, and specifying hypotheses. Efficient posterior computation is discussed, and algorithms are provided for jointly calibrating test cutoffs and sample size to control overall type I error and achieve specified power. Asymptotic approximations for the power curve are used to initialize the algorithms. The methodology is applied to re-design a completed trial that compared two chemotherapy regimens for chronic lymphocytic leukemia, in which an ordinal efficacy outcome was dichotomized and toxicity was ignored to construct the trial's design. The Bayesian tests also are illustrated by several types of categorical outcomes arising in common clinical settings. Freely available computer software for implementation is provided.

Keywords

Bayesian Methods; Dirichlet-multinomial; Multiple Outcomes; Oncology; Randomized Comparative Trials; Utility Elicitation

*Correspondence to: Department of Biostatistics, The University of Texas MD Anderson Cancer Center, Unit 1411, P.O. Box 301402, Houston, TX, 77030-1402, USA.

Supplementary Materials

The Web Appendices referenced in Sections 2, 4, 5, and 6 are available with this paper at the journal's website. An example spreadsheet mentioned in Section 2 for utility elicitation in the context of the CLL example, and a suite of R functions for implementing the computational algorithms in Sections 5.2 and 5.3, along with annotated example R programs for replicating each illustration are available at:

<https://biostatistics.mdanderson.org/SoftwareDownload/>

1. Introduction

Medical outcomes often are complex and multivariate. Physicians routinely select each patient's treatment based on consideration of risk-benefit tradeoffs between desirable and undesirable clinical outcomes. Conventional designs for randomized comparative trials (RCTs) seldom reflect this aspect of medical practice. Rather, most designs in clinical trial protocols are based on one outcome, identified as "primary," with all other outcomes given the nominal status of "secondary." This dichotomy often is codified in institutionally required protocol formats. For example, in cancer studies of chemotherapies for solid tumors, the primary outcome may be *objective response*, defined as 30% or greater tumor shrinkage compared to baseline evaluation, while regimen-related adverse events, called "toxicities," are listed as secondary outcomes [1]. This approach is convenient because it facilitates sample size and power computations in terms of the probabilities of a one-dimensional outcome in the treatment arms. It does not reflect the way that practicing physicians actually think and behave, however. Alternative design approaches include defining a composite outcome that treats efficacy and safety events equally [2, 3], using a test statistic that is a weighted average [4], or basing a test on a quadratic form, such as Hotelling's T-squared statistic, with weights estimated to reflect variability [5]. These approaches ignore the relative clinical importance of beneficial and adverse outcomes, however.

Safety is never a secondary concern in a clinical trial. In actual trial conduct, if interim data from a randomized clinical trial (RCT) show that one treatment has a much higher adverse event rate than the other, or that both arms are unacceptably toxic in a trial comparing two experimental agents, the physicians conducting the trial will terminate accrual whether the protocol's design includes a formal safety stopping rule or not. Such a decision shows that, due to their unwillingness to continue the trial, the physicians have decided that one treatment is inferior to the other in terms of safety. While stopping a trial due to an unacceptably high adverse event rate is an ethical decision, it also is part of the general consideration of how much risk of an adverse outcome is acceptable as a tradeoff for a given level of therapeutic benefit.

This paper is motivated by the consideration that, because clinical trial conduct must accommodate medical practice, a trial design should account formally for risk-benefit tradeoffs between all clinically relevant outcomes. That is, in actual trial design and conduct, scientific and ethical considerations should not be separated. We provide a practical framework for including such tradeoffs explicitly in the treatment comparison underlying the design of two-arm RCTs. We focus on settings where the clinically relevant events are categorical, and thus the outcome Y is a realization from a finite set of elementary patient outcomes. The clinically relevant events, and the resulting set of elementary outcomes, are determined in collaboration with the physician(s) planning the trial. The proposed framework accommodates most discrete outcome structures that occur in practice, including univariate ordinal, bivariate binary indicators of efficacy and safety, bivariate ordinal variables, and such bivariate variables with death as a separate event.

1.1. A Trial in Chronic Lymphocytic Leukemia

We illustrate the proposed methodology by applying it to re-design a RCT reported by [6] that compared two chemotherapy regimens for untreated chronic lymphocytic leukemia (CLL), FC = fludarabine plus cyclophosphamide versus F = fludarabine alone. Patients in this study were treated for up to six 28-day cycles. Following the recommended guidelines at the time of the trial [7], patients were monitored for *clinical response*, with categories CR = Complete response, PR = Partial response, SD = Stable disease, and PD = Progressive disease. Patients also were monitored for several adverse events (AEs), including infections, with severity grades {None, Minor, Major, Fatal}, hematological toxicities with severity grades 0–5, and non-hematological toxicities graded 0–5, according to the National Cancer Institute (NCI) Common Terminology Criteria for Adverse Events (CTCAE). Detailed definitions of the levels of clinical response and the AEs are given in [7].

In the CLL trial design, CR was designated as the primary outcome, with all other outcomes designated as secondary. Thus, the comparison of FC to F was based on the probabilities of CR in the two arms. For this comparison, since clinical response was not evaluable for patients that died during the observation period, these patients were counted as non-responders. This approach is sensible since it counts death during response evaluation as a treatment failure. In contrast, the non-fatal AEs were not included in the study design, despite the fact that the safety of FC was an important concern. Because the above approach to constructing the design for this trial is quite typical, it serves as a useful illustration of our proposed methodology.

To apply our methodology to design this trial would have required working with the physicians planning the trial to determine the clinically relevant outcomes and elicit their utilities. Thus, for the sake of illustration, we first assume that the physicians decided that the relevant outcomes were clinical response, specifically the ordinal variable with possible values {CR, PR, SD, PD}, and also the worst AE with levels {Minimal, Moderate, Severe, Fatal}. Here, “minimal” is defined as no AE requiring medical intervention, “moderate” as a non-life-threatening AE requiring medical intervention without hospitalization, “severe” as an imminently life-threatening AE requiring hospitalization, and “fatal” as an AE resulting in death. Using these definitions, a moderate AE includes grade 3 hematologic and non-hematologic toxicities and minor infections, and a severe AE includes grade 4 hematologic and non-hematologic toxicities and major infections. To define the values of Y , we denote the $12 = 4 \times 3$ non-fatal elementary patient outcomes by the pairs (r, s) , for $r = \{CR, PR, SD, PD\}$ and $s = \{Min, Mod, Sev\}$, with the 13^{th} elementary event D = a fatal AE. Thus, for example, (PR, Mod) is the elementary outcome that the patient had a partial response and a moderate worst AE level. Our design requires numerical utilities for the 13 elementary outcomes, which in practice would be elicited from the physicians. Since we cannot do this retrospectively, we specify numerical utilities (Table 1) for the CLL trial’s 13 outcomes that may be considered a reasonable representation of what would be obtained in practice. In Section 6, once our methodology has been established, we will compare our proposed design to a design that compares the two regimens based on the probabilities of CR . Because the numerical utilities are a key component our methodology, we also include an analysis of the sensitivity of the final inferences to alternative utilities (Table 6).

1.2. Mean Utilities

For the general development, we index the elementary outcomes by $k = 1, \dots, K$, and denote their numerical utilities by $U_k = U(Y = k)$, with $\mathbf{U} = (U_1, U_2, \dots, U_K)'$. These are elicited from the physician(s) planning the trial. For some specific examples, we will replace these integer indices with more descriptive indexing schemes. For convenience, we assign the most desirable outcome utility 100, the least desirable outcome utility 0, with all other outcomes assigned utilities between these two extremes. The domain $[0, 100]$ is chosen to facilitate communication with the physician(s), although in general any compact domain will work. In Section 2, we provide practical strategies for utility elicitation, and illustrate them for the case of bivariate-ordinal outcomes that include the possibility of death.

For treatments $j = A$ and B , we denote the patient response probabilities $\theta_{j,k} = \Pr(Y = k | \text{trt} = j)$, with $\boldsymbol{\theta}_j = (\theta_{j,1}, \theta_{j,2}, \dots, \theta_{j,K})'$, and $\boldsymbol{\theta} = (\boldsymbol{\theta}_A, \boldsymbol{\theta}_B)$. The mean utility of treatment j is

$$\bar{U}(\boldsymbol{\theta}_j) = \mathbf{U}' \boldsymbol{\theta}_j = \sum_{k=1}^K U_k \theta_{j,k}. \quad (1)$$

Our testing methodology relies on the mean utilities $\bar{U}(\boldsymbol{\theta}_A)$ and $\bar{U}(\boldsymbol{\theta}_B)$ as one-dimensional criteria to compare overall treatment effects, since $\bar{U}(\boldsymbol{\theta}_A) > \bar{U}(\boldsymbol{\theta}_B)$ corresponds to the mean clinical desirability of patient outcome being higher for A than B , and conversely. The Bayesian comparative test relies on the posterior of $\delta_{U,A-B}(\boldsymbol{\theta}) = \bar{U}(\boldsymbol{\theta}_A) - \bar{U}(\boldsymbol{\theta}_B)$.

As a first illustration, suppose that the clinically relevant outcome is trinary where a treatment may result in response (R), failure (F), or neither response nor failure (N) so, temporarily suppressing j , for a single treatment $\boldsymbol{\theta} = (\theta_R, \theta_N, \theta_F)$. In particular, R and F are not complementary events. Since $U_R = 100$ and $U_F = 0$ in this case, only $U_N \in (0, 100)$ need be elicited, and the mean utility is $\bar{U}(\boldsymbol{\theta}) = \theta_R \times 100 + \theta_N \times U_N$, which increases with U_N for any $\boldsymbol{\theta}$. If, for example, $U_N = 60$ and the true outcome probabilities are $(\theta_R, \theta_N, \theta_F)' = (0.30, 0.60, 0.10)'$ then the mean utility is $\bar{U}(\boldsymbol{\theta}) = \mathbf{U}' \boldsymbol{\theta} = 0.30 \times 100 + 0.60 \times 60 + 0.10 \times 0 = 66$. Next, consider a trial to compare two clot dissolving agents, A and B , for rapid treatment of stroke, with the outcome evaluated within 24 hours from the start of treatment. Response, R , is defined as the clot that caused the stroke being dissolved without a brain hemorrhage or death, failure, F , is defined as a brain hemorrhage or death, and N is the third event that no brain hemorrhage occurred, the patient did not die, but the clot was not dissolved. Suppose that the true outcome probabilities are $\boldsymbol{\theta}_A = (\theta_{A,R}, \theta_{A,N}, \theta_{A,F})' = (0.50, 0.30, 0.20)'$ and $\boldsymbol{\theta}_B = (\theta_{B,R}, \theta_{B,N}, \theta_{B,F})' = (0.60, 0.30, 0.10)'$. Since B has both a larger response probability and a smaller failure probability compared to A , it is clear that B is clinically superior to A . The mean utilities reflect this, since $\bar{U}(\boldsymbol{\theta}_B) = 60 + \theta_{B,N}U_N$ and $\bar{U}(\boldsymbol{\theta}_A) = 50 + \theta_{B,N}U_N$, so $\bar{U}(\boldsymbol{\theta}_B) - \bar{U}(\boldsymbol{\theta}_A) = 10$ for all $U_N \in (0, 100)$.

If a third agent, C , has $\boldsymbol{\theta}_C = (0.60, 0.10, 0.30)'$, then C has a larger response probability than A but also a larger failure probability, so it is not obvious which of the treatments A or C is superior. If $U_N = 50$, then $\bar{U}(\boldsymbol{\theta}_A) = 65$ compared to $\bar{U}(\boldsymbol{\theta}_B) = 75$, so B is superior to A for this utility. The large difference $\delta_{U,B-A}(\boldsymbol{\theta}) = \bar{U}(\boldsymbol{\theta}_B) - \bar{U}(\boldsymbol{\theta}_A) = 75 - 65 = 10$ is due to the

fact that B increases $\theta_{A,R}$ by 0.10 and also decreases $\theta_{A,F}$ by 0.10. This might be described as a “win-win” scenario for B versus A . Comparing C to A , since $(\theta_C) = (\theta_A) = 65$, that is, A and C have identical mean utilities with $\delta_{U,C-A}(\theta) = 0$, they are equally desirable despite the fact that $\theta_A \neq \theta_C$. This is because the increases in both the response and failure probabilities with C compared to A , specifically $\theta_{C,R} - \theta_{A,R} = 0.60 - 0.50 = 0.10$ and $\theta_{C,F} - \theta_{A,F} = 0.30 - 0.20 = 0.10$, cancel each other out if $U_N = 50$. If $U_N = 20$ rather than 50, however, then $\delta_{U,C-A}(\theta) = 62 - 56 = 6$, so for this utility C is slightly superior to A since the increase in failure probability with C versus A is considered a favorable tradeoff for the increase in response probability.

1.3. Utility-Based Design Framework

Given this general categorical outcome and utility structure, since θ_A and θ_B are not known they must be estimated, and data for doing this must be obtained. The statistical problem thus is how to design and conduct a clinical trial to obtain the necessary data. This requires specification of decision rules, a trial design, and a practical method for establishing a consensus among the investigators for the numerical values in U , since the methodology requires one utility and one utility only. This provides a transparent, formal structure that reflects what physicians actually do in practice, rather than constructing a trial design that focuses on a single primary outcome and then, formally or informally, also monitors secondary outcomes. For the Bayesian version, we call the methodology categorical outcome Bayesian utility-based (CAT-BUB) tests. To implement the proposed design framework, in cooperation with the physician(s) planning the trial, one should take the following steps:

- a. Specify the clinically relevant outcomes and resulting set of elementary patient responses.
- b. Elicit numerical utilities.
- c. Specify design parameters, including targeted alternative treatment differences that will be identified with a specified power, type I error, timing of interim analyses for a group sequential test, and test cut-offs.
- d. Implement the design algorithm, developed below, to determine maximum and interim sample sizes and operating characteristics.
- e. Repeat steps (a–d) until a design with satisfactory operating characteristics is identified.

1.4. Outline

In Section 2 we provide practical guidelines for utility elicitation, illustrated for bivariate categorical outcomes. The Dirichlet-multinomial model is reviewed in Section 3. In Section 4 we present the Bayesian utility-based comparative testing procedure. For the Bayesian test, we provide a scaled-beta approximation for the posterior distribution of the mean utility to facilitate calculation of the test statistic and derive frequentist properties, including an approximate sample size calculation that we use to initialize our computational algorithms. In Section 5, we discuss designs for a single test or a group sequential procedure, and

provide guidelines for eliciting targeted alternatives, and computational algorithms to derive a CAT-BUB design having given overall type I error and power. In Section 6, we illustrate how to implement the CAT-BUB procedure in several settings and report simulation results, including comparison of the CAT-BUB design for the CLL trial to the design based on a binary indicator of CR . We conclude with a brief discussion in Section 7. The Web Supplement provides additional illustrations for several categorical outcome structures often encountered in practice. To facilitate application, freely available user-friendly software is provided (see Supplementary Materials).

2. Utility Elicitation

Since a utility function is required for implementing the proposed methods, we provide practical utility elicitation guidelines. In our experience, specifying U is an intuitive process for the physician(s) that they find to be quite natural. An extension of the previously discussed trinary outcome case with elementary events $\{R, N, F\}$ is an ordinal Y with four or more categories. For example, in oncology trials it is very common to characterize solid tumor response from the start of chemotherapy as an ordinal variable. Following the RECIST tumor evaluation guidelines [1], the outcome may be defined using tumor size relative to baseline, with a 100% decrease a complete response (CR), a 30% to 99% decrease a partial response (PR), a 19% increase to 19% decrease stable disease (SD), and a 20% or greater increase progressive disease (PD). In this and similar contexts, the statistician can simply provide each physician a spreadsheet with the outcomes ordered by desirability and, given $U(CR) = 100$ and $U(PD) = 0$, the physicians can specify numerical utilities for the intermediate outcomes. When there are multiple physicians planning the trial, one approach to establish a consensus utility is the “Delphi” method [8, 9], wherein one asks each physician independently to specify their numerical utilities, then shows the mean of all elicited utilities to all physicians and allows them to adjust their utilities if desired on that basis, and if needed iterates the process until a consensus is reached.

Another common categorical structure is a bivariate binary (efficacy, toxicity) outcome. An example from chemotherapy for acute myelogenous leukemia (AML) defines efficacy as complete remission, C , in terms of recovery of circulating white cells, platelets, and blastic (undifferentiated) cells to normal levels, and toxicity, T , as severe (NCI grade 3 or 4) non-hematologic toxicity, both scored within 42 days. Denoting the respective complementary events by $\bar{C}$ and $\bar{T}$, the statistician can again simply provide each physician with a spreadsheet that contains a 2×2 utility table with $U(C, \bar{T}) = 100$, $U(\bar{C}, T) = 0$, so only the two intermediate utilities $U(\bar{C}, \bar{T})$ and $U(C, T)$ must be specified.

A refinement of the bivariate binary (efficacy, toxicity) outcomes is to define these events for patients who are alive, and include death as a fifth event. This is appropriate for treatment of rapidly fatal diseases, such as AML, where death during therapy has a non-trivial probability. In the AML example, the four elementary events determined by C and T are defined only for patients alive at day 42, and the fifth event is $D = [\text{death within 42 days}]$. This structure may motivate the question of whether assigning a finite utility to death is ethically appropriate, since the utilities will be the basis for medical decision making. If the value $U_D = -\infty$ were assigned, however, the mean utility is $-\infty$ whenever the probability

of D is non-zero, so in practice a single death would terminate the trial. Thus, when death has a non-trivial probability, if one wishes to actually do utility-based decision making then death must be assigned a finite numerical utility having magnitude comparable to the numerical utilities of the other possible patient outcomes. We recommend that the physician(s) first specify $U(\bar{C}, T)$, i.e., the worst outcome for a patient who is alive, relative to $U(C, \bar{T}) = 100$, i.e., the best outcome, and $U(D) = 0$, and then specify $U(\bar{T}, \bar{C})$ and $U(C, T)$ relative to the $U(C, \bar{T}) = 100$ and the selected $U(\bar{C}, T)$. To implement this, the statistician may ask the physician(s) to fill in the following two tables sequentially,

$(C, \bar{T})$	$(\bar{C}, T)$	D
100		0

	C	$\bar{C}$
$\bar{T}$	100	
T		$U(\bar{C}, T)$

where $U(\bar{C}, T)$ in the right-hand table takes the specified value from the left-hand table. This sequence decomposes utility elicitation into two intuitive steps. It also provides a partial motivation for establishing utilities for our re-design of the CLL trial.

To establish or elicit utilities for the CLL trial outcomes, and in general for bivariate ordinal outcomes with death as a separate event, we propose the following two alternative strategies, one *direct* and the other *indirect*. The *direct* elicitation strategy simply requires the statistician to provide the physician(s) with a utility table and suggest a specification order. For the CLL outcomes, using the direct strategy one would provide the physician(s) the table below and tell them to fill the empty cells in alphabetical order.

	CR	PR	SD	PD	
Min	100	C	C	B	Death
Mod	C	D	D	C	0
Sev	B	C	C	A	

The basic idea is to first specify the utility of the worst non-fatal outcome, then the two most extreme (efficacy, toxicity) trade-off outcomes, then the intermediate outcomes where either the best efficacy or worst toxicity event occurs, and finally the remaining outcomes in the interior portion of the table.

In contrast, the *indirect* strategy decomposes elicitation into a series of intuitive, mutually independent steps that induce numerical utilities. For the CLL outcome, we would implement the indirect strategy by having the physician(s) specify the following sub-tables,

(CR,Min)	(PD,Sev)	D
100	$100 \times \nu$	0

	CR	PD
Min	100	$100 \times \zeta_1$
Sev	$100 \times \zeta_2$	0

(CR,Min)	(PR,Min)	(SD,Min)	(PD,Min)
100	$100 \times \phi_{1,1}$	$100 \times \phi_{1,2}$	0

(CR,Sev)	(PR,Sev)	(SD,Sev)	(PD,Sev)
100	$100 \times \phi_{2,1}$	$100 \times \phi_{2,2}$	0

(CR,Min)	(CR,Mod)	(CR,Sev)
100	$100 \times \xi_1$	0

(PD,Min)	(PD,Mod)	(PD,Sev)
100	$100 \times \xi_2$	0

In the above sub-tables, we denote the proportions that will be specified by the physician with Greek symbols, e.g. ν and ζ_1 , which we use to determine the induced numerical utilities later. When the statistician provides the sub-tables to the physician(s), these entries will be left blank for the physician(s) to fill in, with the instruction that, for example, ν is the proportion quantifying the desirability of (PD,Sev) relative to (CR,Min), and so on. The sub-tables are mutually independent, i.e., the values in a particular sub-table are not restricted by, or dependent on the values from any other sub-table. Therefore, the sub-tables can be specified in whatever order the physicians prefer, and each can be revisited and adjusted during the specification process until the physicians are satisfied.

Based on the previous sub-tables, the induced numerical utilities can be determined sequentially as follows,

$$\begin{aligned}
U(CR, Min) &= 100, & U(D) &= 0, & U(PD, Sev) &= 100\nu, \\
U(PD, Min) &= \zeta_1[U(CR, Min) - U(PD, Sev)] + U(PD, Sev), \\
U(CR, Sev) &= \zeta_2[U(CR, Min) - U(PD, Sev)] + U(PD, Sev), \\
U(PR, Min) &= \phi_{1,1}[U(CR, Min) - U(PD, Min)] + U(PD, Min), \\
U(SD, Min) &= \phi_{1,2}[U(CR, Min) - U(PD, Min)] + U(PD, Min), \\
U(PR, Sev) &= \phi_{2,1}[U(CR, Sev) - U(PD, Sev)] + U(PD, Sev), \\
U(SD, Sev) &= \phi_{2,2}[U(CR, Sev) - U(PD, Sev)] + U(PD, Sev), \\
U(CR, Mod) &= \xi_1[U(CR, Min) - U(CR, Sev)] + U(CR, Sev), \\
U(PD, Mod) &= \xi_2[U(PR, Min) - U(PR, Sev)] + U(PR, Sev), \\
U(PR, Mod) &= \left[\frac{\xi_2(\phi_{1,1} - \phi_{2,1}) + \phi_{2,1}}{1 - (\xi_1 - \xi_2)(\phi_{1,1} - \phi_{2,1})} \right] [U(CR, Mod) - U(PD, Mod)] + U(PD, Mod), \text{ and} \\
U(SD, Mod) &= \left[\frac{\xi_2(\phi_{1,2} - \phi_{2,2}) + \phi_{2,2}}{1 - (\xi_1 - \xi_2)(\phi_{1,2} - \phi_{2,2})} \right] [U(CR, Mod) - U(PD, Mod)] + U(PD, Mod).
\end{aligned}$$

To aid elicitation, we recommend that the statistician provide the physician(s) with a spreadsheet that contains the relevant sub-tables and a numerical utility table that automatically populates based on the physician's specified values. As an example, we provide such a spreadsheet for the CLL outcome (see Supplementary Materials). In the Web Appendix A, we provide a generalization and detailed derivation of the induced numerical utilities for the indirect elicitation strategy with a $K \times L$ bivariate ordinal outcome plus death.

The proposed indirect strategy facilitates utility elicitation in several important ways. First, when an individual physician is selecting numerical utilities, they can adjust the values in any sub-table and the resulting numerical utilities will repopulate automatically while preserving the partial ordering constraints. In contrast, for the direct strategy, adjusting a single numerical utility may require changing several other values, perhaps even the entire table, which may become impractical if the elementary patient outcome set is large. Second, when the physicians convene to obtain consensus utilities, each sub-table can be addressed independently in turn. Therefore, should a disagreement arise, the physicians can focus on a specific low-dimensional sub-table rather than the entire numerical utility table. Third, the indirect strategy requires the physician(s) to specify fewer values than the direct strategy, which can be a great practical advantage when K and L are both moderately large, say 4. An advantage of the indirect approach for the statistician is that the sub-tables provide low-dimensional bases for conducting a utility sensitivity assessment, which we discuss below in Section 6.

For our re-design of the CLL trial, suppose that the physician(s) specified sub-table entries corresponding to the following parameters: $\nu = 0.10$, $\zeta_1 = 0.10$, $\zeta_2 = 0.20$, $\phi_{1,1} = \phi_{2,1} = 0.80$, $\phi_{1,2} = \phi_{2,2} = 0.20$, $\xi_1 = 0.90$, and $\xi_2 = 0.40$. The numerical utilities induced by these values are given in Table 1. Our choice to specify $\nu = 0.10$ in this illustration reflects the belief that (PD, Sev) is very undesirable relative to (CR, Min) . Specifying $\zeta_1 = 0.10$ and $\zeta_2 = 0.20$ reflects that (CR, Sev) is more desirable than (PD, Min) , yet both responses have desirabilities more similar to (PD, Sev) than (CR, Min) , i.e., both are undesirable outcomes with utilities < 30 . Specifying $\phi_{1,1} = \phi_{2,1} = 0.80$ and $\phi_{1,2} = \phi_{2,2} = 0.20$ reflects the belief that PR and SD have desirabilities similar to CR and PD , respectively, and moreover their desirabilities relative to CR and PD are invariant across the AE severity levels. In contrast,

specifying $\xi_1 = 0.90$ and $\xi_2 = 0.40$ reflects the belief that a moderate AE is more tolerable given an efficacious clinical response. Conditional on CR , a moderate AE has similar desirability compared to a minimal AE, whereas, conditional on PD , it has desirability more similar to a severe AE than to a minimal AE. In summary, these choices reflect the general belief that (CR, Min) , (CR, Mod) , (PR, Min) , and (PR, Mod) are all desirable patient outcomes with numerical utilities > 75 , whereas all other patient outcomes are relatively undesirable with numerical utilities < 35 .

3. Dirichlet-Multinomial Model

Let $\mathbf{X}_j = (X_{j,1} \ X_{j,2} \ \cdots \ X_{j,K})'$ denote the count vector of patient outcomes, with probabilities $\boldsymbol{\theta}_j = (\theta_{j,1} \ \cdots \ \theta_{j,K})'$ and $n_j = \sum_{k=1}^K X_{j,k}$ the number of observations for treatment $j = A, B$. For the Bayesian tests presented in Section 4, we will assume the Dirichlet-multinomial model

$$\begin{aligned} \mathbf{X}_j | \boldsymbol{\theta}_j &\sim \text{Mult}(n_j, \boldsymbol{\theta}_j), & (\text{Likelihood}) \\ \boldsymbol{\theta}_j &\sim \text{Dir}(n_j^* \boldsymbol{\theta}_j^*), & (\text{Prior}) \end{aligned} \quad (2)$$

where $\boldsymbol{\theta}_j^* = (\theta_{j,1}^* \ \cdots \ \theta_{j,K}^*)'$ is the prior mean of $\boldsymbol{\theta}_j$ and n_j^* is the effective sample size (ESS) of the prior (cf. Morita, et al., 2008, 2010). This is well known for the important special case $K = 2$, which is the beta distribution, where the ESS of $f(\boldsymbol{\theta}_j | n_j^*, \boldsymbol{\theta}_j^*)$ is $n_j^* = n_j^* (\theta_{j,1}^* + \theta_{j,2}^*)$. The model (2) has a simple conjugate structure, with each $\boldsymbol{\theta}_j | \mathbf{X}_j \sim \text{Dir}(\mathbf{X}_j + n_j^* \boldsymbol{\theta}_j^*)$, *a posteriori*, which greatly facilitates posterior computation. The Dirichlet-multinomial model is quite general, and accommodates any categorical outcome structure. The multinomial pdf is

$$f(\mathbf{X}_j | \boldsymbol{\theta}_j) = \Gamma(n_j + 1) \prod_{k=1}^K \frac{\theta_{j,k}^{X_{j,k}}}{\Gamma(X_{j,k} + 1)}, \quad (3)$$

and the Dirichlet pdf is

$$f(\boldsymbol{\theta}_j | n_j^* \boldsymbol{\theta}_j^*) = \Gamma(n_j^*) \prod_{k=1}^K \frac{\theta_{j,k}^{n_j^* \theta_{j,k}^* - 1}}{\Gamma(n_j^* \theta_{j,k}^*)}, \quad (4)$$

where $\Gamma(\cdot)$ is the gamma function. The Dirichlet has

$E(\boldsymbol{\theta}_j | n_j^*, \boldsymbol{\theta}_j^*) = \boldsymbol{\theta}_j^*$, $\text{Var}(\theta_{j,k} | n_j^*, \boldsymbol{\theta}_j^*) = \theta_{j,k}^* (1 - \theta_{j,k}^*) / (n_j^* + 1)$, and

$\text{Cov}(\theta_{j,k}, \theta_{j,\ell} | n_j^*, \boldsymbol{\theta}_j^*) = -\theta_{j,k}^* \theta_{j,\ell}^* / (n_j^* + 1)$, for $k \neq \ell = 1, \dots, K$. The posterior has a conjugate form with pdf

$$p(\boldsymbol{\theta}_j | \mathbf{X}_j, n_j^*, \boldsymbol{\theta}_j^*) = \Gamma(n_j + n_j^*) \prod_{k=1}^K \frac{\theta_{j,k}^{X_{j,k} + n_j^* \theta_{j,k}^* - 1}}{\Gamma(X_{j,k} + n_j^* \theta_{j,k}^*)}, \quad (5)$$

and posterior mean $E(\boldsymbol{\theta}_j | \mathbf{X}_j, n_j^*, \boldsymbol{\theta}_j^*) = (\mathbf{X}_j + n_j^* \boldsymbol{\theta}_j^*) / (n_j + n_j^*)$.

For prior specification, the two priors for $\boldsymbol{\theta}_j$ should match, i.e. $n_A^* = n_B^*$ and $\boldsymbol{\theta}_A^* = \boldsymbol{\theta}_B^*$, so that any statistical comparisons are unbiased, and the priors should not include an inappropriate amount of information, which is quantified by ESS. For this reason, we drop the treatment subscript on these hyperparameters in the sequel. As a default prior, i.e., in the absence of prior information, we will assume $n^* = 1$ and $\theta_k^* = K^{-1}$ so that each ESS = 1 and all elementary events are equally likely *a priori*. This default choice allows the accruing data to quickly overwhelm the prior while shrinking response probabilities away from 0 and 1 in small samples. When historical information is available, n^* and $\boldsymbol{\theta}^*$ can be specified to reflect that experience and its relevance to the investigation. Alternatively, for a more robust use of historical information, power priors [10] or commensurate prior methods [11] could be applied. Because the use of historical data to construct informative priors for Bayesian models underlying RCTs is a complex and controversial issue, however, we will not use such priors here, and assume $n^* = 1$ in the sequel.

4. Comparative Tests

Treatment differences are characterized by the mean utility difference, and we test the hypotheses

$$H_0: \delta_{U,A-B}(\boldsymbol{\theta}) = 0 \quad \text{versus} \quad H_1: \delta_{U,A-B}(\boldsymbol{\theta}) \neq 0. \quad (6)$$

If desired, a one-sided version of (6) may be appropriate. For example, to test whether A is superior to B , the hypotheses would be $H_0: \delta_{U,A-B}(\boldsymbol{\theta}) \leq 0$ versus $H_1: \delta_{U,A-B}(\boldsymbol{\theta}) > 0$. In what follows, we will focus on two-sided hypotheses, since the one-sided case is a straightforward modification.

Let $\mathbf{X} = (\mathbf{X}_A, \mathbf{X}_B)$ denote the observed elementary outcome count data. We conduct a CAT-BUB comparative test using the following symmetric decision criteria. If

$$T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) = Pr\{\delta_{U,A-B}(\boldsymbol{\theta}) > 0 | \mathbf{X}, n^*, \boldsymbol{\theta}^*\} > p_{cut} \quad (7)$$

then conclude superiority of A over B , denoted by $A > B$. If

$$T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) = Pr\{\delta_{U,A-B}(\boldsymbol{\theta}) < 0 | \mathbf{X}, n^*, \boldsymbol{\theta}^*\} > p_{cut}, \quad (8)$$

then conclude superiority of B over A , denoted by $B > A$. We select the probability cutoff, p_{cut} to ensure an approximate level α test for all $\boldsymbol{\theta} = (\boldsymbol{\theta}_A, \boldsymbol{\theta}_B)$ with $\delta_{U,A-B}(\boldsymbol{\theta}) = 0$. We discuss technical details for doing this below.

4.1. Efficient Posterior Computation

While the posterior distributions of $(\boldsymbol{\theta}_A)$ and $(\boldsymbol{\theta}_B)$ are not analytically tractable, because mean utilities are linear combinations of Dirichlet random vectors there are several feasible numerical approximations. With a Monte Carlo (MC) approach, one would generate M samples from the posterior mean utility (PMU) distribution for treatment j , i.e. $p(\boldsymbol{\theta}_j | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*)$, by drawing $\boldsymbol{\theta}_j^{(m)} \sim Dir(\mathbf{X}_j + n^* \boldsymbol{\theta}^*)$, since $p(\boldsymbol{\theta}_j | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*) \equiv Dir(\mathbf{X}_j + n^* \boldsymbol{\theta}^*)$ (see (5)), and defining $\bar{U}(\boldsymbol{\theta}_j^{(m)}) = \mathbf{U}' \boldsymbol{\theta}_j^{(m)}$, for $m = 1, \dots, M$ and $j = A, B$ (see [12], Chapter 3.3). These samples provide estimates of $T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$ and $T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$ in (7) and (8). For data analysis, any desired level of accuracy can be obtained by increasing M , since it only needs to be conducted once. In contrast, the MC approach is computationally expensive for constructing a clinical trial design, since it requires iterative simulations to assess frequentist operating characteristics in a variety of scenarios, and thus a very large number of MC calculations.

For the CAT-BUB design, a more computationally efficient method for estimating the posterior quantities in (7) and (8) is a parametric approximation to the PMU distribution based on a scaled-beta distribution. To implement this approach, we exploit the following well known forms of the posterior moments of a Dirichlet:

$$\begin{aligned} E[\boldsymbol{\theta}_j | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*] &= \tilde{\boldsymbol{\theta}}_j = \frac{\mathbf{X}_j + n^* \boldsymbol{\theta}^*}{(n_j + n^*)}, \\ Var[\theta_{j,k} | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*] &= \frac{\tilde{\theta}_{j,k} (1 - \tilde{\theta}_{j,k})}{(n_j + n^* + 1)} = \frac{(X_{j,k} + n^* \theta_k^*)[(n_j + n^*) - (X_{j,k} + n^* \theta_k^*)]}{(n_j + n^*)^2 (n_j + n^* + 1)}, \text{ and} \\ Cov[\theta_{j,k}, \theta_{j,\ell} | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*] &= \frac{-\tilde{\theta}_{j,k} \tilde{\theta}_{j,\ell}}{(n_j + n^* + 1)} = \frac{-(X_{j,k} + n^* \theta_k^*)(X_{j,\ell} + n^* \theta_\ell^*)}{(n_j + n^*)^2 (n_j + n^* + 1)}, \text{ for } k \neq \ell. \end{aligned} \quad (9)$$

It follows that

$$\tilde{\mu}_j = E[\bar{U}(\boldsymbol{\theta}_j) | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*] = \mathbf{U}' \tilde{\boldsymbol{\theta}}_j, \text{ and } \tilde{\sigma}_j^2 = Var[\bar{U}(\boldsymbol{\theta}_j) | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*] = \mathbf{U}' \sum_j \tilde{\Sigma}_j \mathbf{U}, \quad (10)$$

where $\tilde{\Sigma}_j = Var[\boldsymbol{\theta}_j | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*]$ with entries defined in (9). Using (10), we match the support, mean, and variance of each PMU distribution with those of a scaled-beta distribution. Let $Beta(\lambda, \gamma)$ denote a beta distribution with mean $\mu = \lambda/(\lambda + \gamma)$ and variance $\sigma^2 = \mu(1 - \mu)/(\lambda + \gamma + 1)$. We approximate $p(\boldsymbol{\theta}_j | \mathbf{X}_j, n^*, \boldsymbol{\theta}^*)$ with $100 \times Beta(\tilde{\lambda}_j, \tilde{\gamma}_j)$, where

$$\tilde{\lambda}_j = \tilde{\mu}_j \left[\frac{\tilde{\mu}_j(1-\tilde{\mu}_j)}{\tilde{\sigma}_j^2} - 1 \right], \quad \tilde{\gamma}_j = (1-\tilde{\mu}_j) \left[\frac{\tilde{\mu}_j(1-\tilde{\mu}_j)}{\tilde{\sigma}_j^2} - 1 \right], \quad (11)$$

and the mean $\tilde{\mu}_j$ and variance $\tilde{\sigma}_j^2$ are defined in (10). When $K=2$ the PMU distribution is precisely this scaled-beta distribution. We provide the derivation for (11) in Web Appendix B.

Using this approximation, the posterior decision criterion is

$$T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) \approx \int_0^1 [1 - B(x|\tilde{\lambda}_A, \tilde{\gamma}_A)] b(x|\tilde{\lambda}_B, \tilde{\gamma}_B) dx, \quad (12)$$

where $B(x|\lambda, \gamma)$ and $b(x|\lambda, \gamma)$ denote the cdf and pdf of a $Beta(\lambda, \gamma)$ distribution. The approximation for $T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$ follows by symmetry. We use adaptive quadrature via the `integrate()` function in R to evaluate (12) efficiently [13]. In Web Appendix B, we confirm the validity of (12) by comparing it to the usual MC approach using simulation under a variety of settings. The scaled-beta approximation is 1,000 times faster than the usual MC approximation with $M=100,000$, and it works well even with very small sample sizes, such as $n_A = n_B = 10$. We will use the scaled-beta approximation for the remainder of the paper, and recommend its use in practice.

4.2. Type I Error, Power, and Sample Size

We derive an expression for the approximate power function of the CAT-BUB procedure based on (7) and (8), and use this result to show control of type I error and to obtain a sample size formula. We first apply the Bayesian central limit theorem, and use the resulting posterior asymptotic normality to obtain tractable expressions for $T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$ and $T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$. We will show that the resulting approximate test statistics are tractable functions of the data, $\mathbf{X}$. We then take the frequentist perspective, treating $\boldsymbol{\theta} = (\boldsymbol{\theta}_A, \boldsymbol{\theta}_B)$ as a fixed quantity, and apply the classical central limit theorem to derive the asymptotic sampling distributions of $T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$ and $T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*)$, and an approximate power function.

Since $\mathbf{X}_j$ is multinomial with parameter $\boldsymbol{\theta}_j$, the MLE is $\hat{\boldsymbol{\theta}}_j = \mathbf{X}_j/n_j$ and the estimated Fischer

information is $n_j \hat{\Sigma}_j^{-1}$, where $\hat{\Sigma}_j$ has k -th diagonal entry $\hat{\theta}_{j,k}(1 - \hat{\theta}_{j,k})$ and (k, ℓ) -th off-diagonal entry $-\hat{\theta}_{j,k} \hat{\theta}_{j,\ell}$, $k, \ell = 1, \dots, K, j = A, B$. Applying the Bayesian central limit theorem (see [14], Chapter 4)

$$\boldsymbol{\theta}_j | \mathbf{X}_j, n^*, \boldsymbol{\theta}^* \sim \mathcal{N}_K(\hat{\boldsymbol{\theta}}_j, n_j^{-1} \hat{\Sigma}_j), \quad j = A, B.$$

Since $\mathbf{X}_A$ and $\mathbf{X}_B$ are independent, $\mathbf{U}'(\boldsymbol{\theta}_A - \boldsymbol{\theta}_B)|\mathbf{X}, n^*, \boldsymbol{\theta}^* \sim \mathcal{N}(\hat{\delta}_{U,A-B}, \hat{\sigma}_{+,n}^2)$, where $\hat{\delta}_{U,A-B} = \mathbf{U}'(\hat{\boldsymbol{\theta}}_A - \hat{\boldsymbol{\theta}}_B)$ and $\hat{\sigma}_{+,n}^2 = \mathbf{U}'\left(\sum_A /n_A + \sum_B /n_B\right) \mathbf{U}$. It follows that

$$T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) \approx \Phi\left(\frac{\hat{\delta}_{U,A-B}}{\hat{\sigma}_{+,n}}\right) \quad \text{and} \quad T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) \approx \Phi\left(-\frac{\hat{\delta}_{U,A-B}}{\hat{\sigma}_{+,n}}\right), \quad (13)$$

where $\Phi(\cdot)$ denotes the standard normal cdf. We use the notation “ $\approx$ ” to mean that an approximation can be made arbitrarily accurate for sufficiently large sample size.

To derive an approximate power function, we treat the posterior quantities in (13) as functions of the data $\mathbf{X}$ given a fixed $\boldsymbol{\theta}$, and derive asymptotic approximations for their sampling distributions. First, the exact power function is the probability of rejecting the null for a fixed $\boldsymbol{\theta}$, i.e.

$$\psi(\boldsymbol{\theta}) = Pr\{T_{A>B}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) > p_{cut} | \boldsymbol{\theta}\} + Pr\{T_{B>A}(\mathbf{X}; n^*, \boldsymbol{\theta}^*) > p_{cut} | \boldsymbol{\theta}\}. \quad (14)$$

Applying the classical central limit theorem, $(\hat{\delta}_{U,A-B} - \delta_{U,A-B}(\boldsymbol{\theta})) / \hat{\sigma}_{+,n} \sim \mathcal{N}(0, 1)$, so plugging (13) into (14) gives the approximate power function

$$\psi(\boldsymbol{\theta})^{approx} = \Phi\left[\left(\frac{\delta_{U,A-B}(\boldsymbol{\theta})}{\sigma_{+,n}(\boldsymbol{\theta})}\right) - \Phi^{-1}(p_{cut})\right] + \Phi\left[-\left(\frac{\delta_{U,A-B}(\boldsymbol{\theta})}{\sigma_{+,n}(\boldsymbol{\theta})}\right) - \Phi^{-1}(p_{cut})\right], \quad (15)$$

where $\sigma_{+,n}(\boldsymbol{\theta})^2 = \mathbf{U}'(\boldsymbol{\Sigma}_A(\boldsymbol{\theta})/n_A + \boldsymbol{\Sigma}_B(\boldsymbol{\theta})/n_B)\mathbf{U}$. is a function of $\boldsymbol{\theta}$.

The type I error is $\sup\{\psi(\boldsymbol{\theta}) : \delta_{U,A-B}(\boldsymbol{\theta}) = 0\}$, and since $\psi(\boldsymbol{\theta})^{approx} = 2(1 - p_{cut})$ for all $\boldsymbol{\theta}$ with $\delta_{U,A-B}(\boldsymbol{\theta}) = 0$, using $p_{cut} = 1 - \alpha/2$ provides an asymptotic level α test. To derive an approximate sample size formula, we set $p_{cut} = 1 - \alpha/2$ and define $n = n_A = n_B$. If desired, one could instead define $n = n_A$ and $n_B = \eta \times n_A$, where η controls the randomization ratio. For a given fixed target alternative $\boldsymbol{\theta}^{Alt}$, e.g., the hypothesized outcome probabilities, we equate $\psi(\boldsymbol{\theta}^{Alt})^{approx} = 1 - \beta$ and solve for n , which gives approximate sample size

$$n_f(\boldsymbol{\theta}^{(Alt)}, \alpha, \beta) = \frac{[\Phi^{-1}(1-\beta) + \Phi^{-1}(1-\alpha/2)]^2 \sigma_{+,n}^2(\boldsymbol{\theta}^{(Alt)})}{\delta_{U,A-B}^2(\boldsymbol{\theta}^{(Alt)})}. \quad (16)$$

We discuss elicitation of $\boldsymbol{\theta}^{Alt}$ in Section 5.

5. Designing a CAT-BUB Trial

In this section, we derive design parameters that control overall type I error at level α and provide $1-\beta$ power for targeted alternatives, i.e., the set of treatment effects that we want to identify with the specified power. For this computation, we distinguish between fixed sample designs with one comparative test at the end of the trial, and group sequential designs with up to S comparative analyses over the course of the trial, allowing early termination with rejection of the null at each interim analysis. We first present guidelines for eliciting targeted alternatives, then discuss fixed sample CAT-BUB designs, followed by group sequential CAT-BUB designs. For each design setting, we provide a computational algorithm for deriving the probability cut-offs and sample size, given α , β and the targeted alternatives.

5.1. Eliciting Targeted Alternatives

Consider a fixed sample CAT-BUB test with type I error α for all θ for which $\delta_{U,A-B}(\theta) = 0$, and power $1 - \beta$ for a set of fixed targeted alternatives with $|\delta_{U,A-B}(\theta)| > 0$. The approximate power function in (15) shows that selecting p_{cut} to control type I error for one fixed null response probability vector, say $\theta^{(Null)} = (\theta_A^{(Null)}, \theta_B^{(Null)})$ with $\delta_{U,A-B}^{(Null)} = 0$, will control type I error for all fixed θ with $\delta_{U,A-B}(\theta) = 0$. In contrast, the power varies with both the targeted utility difference and the fixed θ from which this difference arises, via $\sigma_{+,n}(\theta)$. Therefore, targeted alternatives must be elicited in the θ domain. We denote a fixed target by $\theta^{(Alt)} = (\theta_A^{(Alt)}, \theta_B^{(Alt)})$ and its utility difference by $|\delta_{U,A-B}^{(Alt)}| > 0$.

Since it may not be intuitively obvious how to specify $\theta^{(Alt)}$, we provide the following guidelines, which require a discussion between the statistician and the physicians. For simplicity, we will treat A as the null or standard treatment, although the algorithm works if A and B are both experimental and considered to be symmetric. The statistician begins by eliciting an expected probability vector corresponding to historical experience with standard therapy, say $\theta_A^{(Alt)}$, which may be based on the physician(s)' experience or analysis of historical data. Given $\theta_A^{(Alt)}$, the statistician asks the physician(s) to specify one or more alternative probability vectors, $\theta_B^{(Alt,1)}, \dots, \theta_B^{(Alt,m)}$, that are considered equally desirable improvements over $\theta_A^{(Alt)}$. In practice, m should be reasonably small, in the range $1 \leq m \leq K$. Each elicited alternative $\theta_B^{(Alt,r)}$ gives standardized utility difference $s^{(Alt,r)} = \delta_{U,B-A}^{(Alt,r)} / \sigma_+^{(Alt,r)}$, where $\delta_{U,B-A}^{(Alt,r)}$ and $\sigma_+^{(Alt,r)}$ are evaluated at $\theta^{(Alt,r)} = (\theta_A^{(Alt)}, \theta_B^{(Alt,r)})$ for $r = 1, \dots, m$. For the sample size calculation in (16), one then selects the targeted alternative $\theta_B^{(Alt)}$ giving smallest $s^{(Alt,r)}$, formally

$$\theta^{(Alt)} = \left\{ \left(\theta_A^{(Alt)}, \theta_B^{(Alt,r^*)} \right) : s^{(Alt,r^*)} = \min \left\{ s^{(Alt,1)}, \dots, s^{(Alt,m)} \right\} \right\}. \quad (17)$$

This choice is conservative since it ensures the test will achieve the desired power for all elicited $\theta^{(Alt,r)}$.

In practice, if this computation gives a sample size that is not feasible, then the physician(s) should be asked to re-consider their set of specified alternatives. This is not unlikely, since it may not be intuitively obvious, when specifying one or more fixed target probability vectors, how they translate into a required sample size. To help guide the physician(s) in this process, one should show them the numerical values of $\delta_{U,B-A}^{(Alt,r)}$, $s^{(Alt,r)}$, and $\theta_B^{(Alt,r)}$, for $r = 1, \dots, m$, possibly as a table with m rows and three columns to facilitate comparison and interpretation. Since smaller values of $s^{(Alt,r)}$ and $\delta_{U,B-A}^{(Alt,r)}$ require a larger sample size to detect the corresponding $\theta_B^{(Alt,r)}$, this provides a quantitative index of the relative difficulty of detecting each target, and it also identifies the targeted alternative $\theta_B^{(Alt,r^*)}$ having the smallest $s^{(Alt,r)}$ that produced the sample size. If a modified set of targets is specified, the sample size may be recomputed, with this process iterated if desired. This may be considered a multidimensional analog of a conventional power and sample size computation in terms of a one-dimensional parameter. If desired, the CAT-BUB test's power function computed over a set of (θ_A, θ_B) values also may be examined.

Recall the trinary outcome example where $U = (U_R, U_N, U_F)' = (100, 60, 0)'$. Given standard vector $\theta_A^{(Alt)} = (0.30, 0.50, 0.20)'$, suppose that the three equally desirable targets $\theta_B^{(Alt,1)} = (0.40, 0.50, 0.10)'$, $\theta_B^{(Alt,2)} = (0.50, 0.35, 0.15)'$, and $\theta_B^{(Alt,3)} = (0.35, 0.60, 0.05)'$ are elicited. Then $s^{(Alt,1)} = 10/45.8 = 0.218$, $s^{(Alt,2)} = 11/49.2 = 0.224$, and $s^{(Alt,3)} = 11/42.6 = 0.258$, so we would take $\theta^{(Alt)} = (\theta_A^{(Alt)}, \theta_B^{(Alt,1)})$. If the standardized utility differences, $s^{(Alt,r)}$, differ substantially, then the physician(s) may instead select the *a priori* most likely alternative, perhaps sacrificing power for some alternatives as a trade-off for a smaller sample size. If the utility differences, $\delta_{U,B-A}^{(Alt,r)}$, differ substantially, then the physician(s) may wish to reconsider the choices of equally desirable targets, or possibly may decide to modify some entries of the numerical utility vector U .

5.2. Computational Algorithm for Fixed Sample CAT-BUB Design

Given α , β and the targeted alternative $\theta^{(Alt)}$ defined in (17), we jointly select a sample size and cutoff p_{cut} for a fixed sample CAT-BUB design using the following algorithm:

Step 0. Set $\hat{n} = n_f(\theta^{(Alt)}, \alpha, \beta)$, where $n_f(\cdot)$ is defined by (16).

Step 1. Generate G_0 null datasets as follows. For $g_0 = 1, \dots, G_0$,

- i. generate $\mathbf{X}_j^{(g_0)} \sim Mult(\hat{n}, \theta_A^{(Alt)})$ for $j = A, B$.
- ii. store $\mathbf{X}^{(Null, g_0)} = (\mathbf{X}_A^{(g_0)}, \mathbf{X}_B^{(g_0)})$.
- iii. calculate and store

$$T^{(Null, g_0)} = \max \left\{ T_{A>B}(\mathbf{X}^{(Null, g_0)}; n^*, \theta^*), T_{B>A}(\mathbf{X}^{(Null, g_0)}; n^*, \theta^*) \right\}.$$

Step 2. Set $\hat{p}_{cut}$ to the empirical $(1 - \alpha)\%$ -tile of $\{T^{(Null, 1)}, \dots, T^{(Null, G_0)}\}$.

Step 3. Generate G_1 alternative datasets as follows. For $g_1 = 1, \dots, G_1$,

- i. generate $\mathbf{X}_j^{(g_1)} \sim \text{Mult}(\hat{n}, \boldsymbol{\theta}_j^{(Alt)})$ for $j = A, B$.
- ii. store $\mathbf{X}^{(Alt, g_1)} = (\mathbf{X}_A^{(g_1)}, \mathbf{X}_B^{(g_1)})$.
- iii. If $\delta_{U, A-B}^{(Alt)} > 0$, calculate and store $T^{(Alt, g_1)} = T_{A>B}(\mathbf{X}^{(Alt, g_1)}; n^*, \boldsymbol{\theta}^*)$.
Otherwise, calculate and store $T^{(Alt, g_1)} = T_{B>A}(\mathbf{X}^{(Alt, g_1)}; n^*, \boldsymbol{\theta}^*)$.

Step 4. Set $\hat{\beta} = G_1^{-1} \sum_{g_1=1}^{G_1} [T^{(Alt, g_1)} \leq \hat{p}_{cut}]$, where $[E] = 1$ if E is true, and 0 otherwise.

Step 5. If $\hat{\beta} \in [\beta - \varepsilon, \beta + \varepsilon]$, stop and select $n = \hat{n}$ and $p_{cut} = \hat{p}_{cut}$.

Otherwise, update $\hat{n} = \hat{n} \left(\frac{\Phi^{-1}(1-\beta) + \Phi^{-1}(\hat{p}_{cut})}{\Phi^{-1}(1-\beta) + \Phi^{-1}(\hat{p}_{cut})} \right)^2$ and return to Step 1.

In practice, $\hat{n}$ is rounded to its nearest integer value and $\hat{p}_{cut}$ is rounded to its nearest larger thousandth. We use default values $\varepsilon = 0.005$, $G_0 = 50,000$ and $G_1 = 25,000$. Since choosing G_0 and G_1 is non-intuitive, detailed guidelines are given in Web Appendix C. Briefly, these default values allow us to estimate p_{cut} accurately to three digits, and be certain that the power for $\boldsymbol{\theta}^{(Alt)}$ at the selected n is within $2 \times \varepsilon = 0.01$ of $1 - \beta$. The sample size adjustment in step 5 is motivated by (16), and allows $\hat{n}$ to be increased or decreased by a magnitude proportional to the current discrepancy between the estimated and desired power.

5.3. Computational Algorithm for Group Sequential CAT-BUB Design

In typical practice, RCTs require group sequential tests [15]. Here, we discuss implementation of the CAT-BUB test in this context, denoting the sample sizes where an analysis occurs by n_s , $s = 1, \dots, S$. We take use the α -spending approach proposed by [16] and extended by [17], with an α -spending function $f(n_s; \alpha, \rho, n_S) = \alpha(n_s/n_S)^\rho$ suggested by [18]. The design parameter $\rho \geq 0$ controls the α -spending rate, with larger values spending less α at early looks. This approach is appealing in practice because the actual analysis schedule need not follow the planned schedule. At the first interim look with n_1 of the planned n_S samples, we calibrate the probability threshold, $p_{cut,1}$, to spend $f(n_1; \alpha, \rho, n_S)$ of the overall type I error. Similarly, at s -th interim look, we calibrate $p_{cut,s}$ to spend $f(n_s; \alpha, \rho, n_S) - f(n_{s-1}; \alpha, \rho, n_S)$ of the overall type I error. So if the trial reaches a final analysis at n_S samples, the overall type I error is exactly α .

To determine a maximum sample size, n_S , for a group sequential CAT-BUB design with up to S tests, power $1 - \beta$ for the elicited alternative $\boldsymbol{\theta}^{(Alt)}$, we specify a complete analysis schedule using the proportions of n_S , denoted by t_s , $s = 1, \dots, S$. We determine n_S using the following algorithm:

Step 0. Set $\hat{n}_S = n_f(\boldsymbol{\theta}^{(Alt)}, \alpha, \beta)$, where $n_f(\cdot)$ is defined by (16), and $\hat{n}_s = t_s \times \hat{n}_S$, for $s = 1, \dots, S-1$.

Step 1. Generate G_0 null sequential datasets as follows. For $g_0 = 1, \dots, G_0$,

- i. generate $\mathbf{X}_{j,s}^{(g_0)} \sim \text{Mult}(\hat{n}_s, \boldsymbol{\theta}_A^{(Alt)})$ for $j = A, B$ and $s = 1, \dots, S$.
- ii. store $\mathbf{X}_{s,+}^{(Null,g_0)} = (\mathbf{X}_{A,s,+}^{(g_0)}, \mathbf{X}_{B,s,+}^{(g_0)})$, where $\mathbf{X}_{j,s,+}^{(g_0)} = \sum_{m=1}^s \mathbf{X}_{j,m}^{(g_0)}$ for $j = A, B$ and $s = 1, \dots, S$.
- iii. calculate and store, for $s = 1, \dots, S$,

$$T_s^{(Null,g_0)} = \max \left\{ T_{A>B}(\mathbf{X}_{s,+}^{(Null,g_0)}; n^*, \boldsymbol{\theta}^*), T_{B>A}(\mathbf{X}_{s,+}^{(Null,g_0)}; n^*, \boldsymbol{\theta}^*) \right\}.$$

Step 2. Calculate $\hat{p}_{cut,1}, \dots, \hat{p}_{cut,S}$ as follows.

- i. Set $\hat{p}_{cut,1}$ to the empirical $\{1 - f(n_1; \alpha, \rho, n_S)\}$ %-tile of $\{T_1^{(Null,1)}, \dots, T_1^{(Null,G_0)}\}$.
- ii. Set $\hat{p}_{cut,s}$ to the empirical $[\{1 - f(n_s; \alpha, \rho, n_S)\} / \{1 - f(n_{s-1}; \alpha, \rho, n_S)\}]$ %-tile of $\{T_s^{(Null,g_0)} : T_1^{(Null,g_0)} \leq \hat{p}_{cut,1}, \dots, T_{s-1}^{(Null,g_0)} \leq \hat{p}_{cut,s-1}, g_0 = 1, \dots, G_0\}$ for $s = 2, \dots, S$.

Step 3. Generate G_1 alternative sequential datasets as follows. For $g_1 = 1, \dots, G_1$,

- i. generate $\mathbf{X}_{j,s}^{(g_1)} \sim \text{Mult}(\hat{n}_s, \boldsymbol{\theta}_j^{(Alt)})$ for $j = A, B$ and $s = 1, \dots, S$.
- ii. store $\mathbf{X}_{s,+}^{(Alt,g_1)} = (\mathbf{X}_{A,s,+}^{(g_1)}, \mathbf{X}_{B,s,+}^{(g_1)})$, where $\mathbf{X}_{j,s,+}^{(g_1)} = \sum_{m=1}^s \mathbf{X}_{j,m}^{(g_1)}$ for $j = A, B$ and $s = 1, \dots, S$.
- iii. If $\delta_{U,A-B}^{(Alt)} > 0$, calculate and store $T_s^{(Alt,g_1)} = T_{A>B}(\mathbf{X}_{s,+}^{(Alt,g_1)}; n^*, \boldsymbol{\theta}^*)$, for $s = 1, \dots, S$.
Otherwise, calculate and store $T_s^{(Alt,g_1)} = T_{B>A}(\mathbf{X}_{s,+}^{(Alt,g_1)}; n^*, \boldsymbol{\theta}^*)$, for $s = 1, \dots, S$.

Step 4. Set $\hat{\beta} = G_1^{-1} \sum_{g_1=1}^{G_1} [T_1^{(Alt,g_1)} \leq \hat{p}_{cut,1}, \dots, T_S^{(Alt,g_1)} \leq \hat{p}_{cut,S}]$, where $[E] = 1$, if E is true, and 0, otherwise.

Step 5. If $\hat{\beta} \in [\beta - \varepsilon, \beta + \varepsilon]$, stop and select $n_S = \hat{n}_S$.

Otherwise, update $\hat{n}_S = \hat{n}_S \left(\frac{\Phi^{-1}(1-\beta) + \Phi^{-1}(\hat{p}_{cut,S})}{\Phi^{-1}(1-\hat{\beta}) + \Phi^{-1}(\hat{p}_{cut,S})} \right)^2$, $\hat{n}_S = t_S \times \hat{n}_S$ for $s = 1, \dots, S - 1$, and return to Step 1.

We use the same default values as the fixed sample algorithm, that is $\epsilon = 0.005$, $G_0 = 50,000$ and $G_1 = 25,000$. Using the planned analysis schedule, we can assess the operating characteristics at a variety of alternatives. The actual analysis schedule may differ from the planned schedule, so the realized power may differ from $1 - \beta$; however, [15] show that the realized power is quite robust to deviations from the planned analysis schedule. During an actual trial, we can follow steps 1–2 to re-estimate $p_{cut,s}$ for the actual n_s being used, given the previous interim analysis sample sizes $n_1, \dots, n_{s-1}$ and their corresponding $p_{cut,1}, \dots, p_{cut,s-1}$ values.

6. Illustrations

In this section, we illustrate CAT-BUB tests and report results of various simulation studies comparing both fixed sample and group sequential CAT-BUB designs with beta-binomial designs. We investigate the proposed procedure in the contexts of a trinary outcome, a bivariate-binary outcome, and the CLL trial, which actually had a bivariate ordinal outcome including death. We also report the results of utility sensitivity analyses.

6.1. Trinary Outcomes

6.1.1. Fixed Sample Tests—Returning to the example involving clot dissolving agents for rapid treatment of stroke with a trinary outcome and utility $U = (100, 50, 0)'$, we investigate the frequentist operating characteristics of the proposed CAT-BUB approach for a variety of fixed response probability vectors. We consider a CAT-BUB test with type I error $\alpha = 0.05$, and power $1 - \beta = 0.80$ for targeted alternative

$\theta^{(Alt)} = (\theta_A^{(Alt)}, \theta_B^{(Alt)}) = ((0.50, 0.30, 0.20)', (0.60, 0.30, 0.10)')$ with $\delta_{U,B-A}^{(Alt)} = 10$. In this context, the fixed sample CAT-BUB design algorithm, given in Section 5.2, gives $p_{cut} = 0.976$ and $n = 208$.

We compare the CAT-BUB approach for trinary outcomes $\{R, N, F\}$ with a Bayesian design that follows the more common approach of combining the events N and F so that outcome may be considered binary, specifically $R = \text{“success,”}$ versus $N \cup F = \text{“failure,”}$ and compares therapies in terms of the probabilities $\pi_j = \Pr(Y = R | j)$ for $j = A, B$. For this design, we assume a Bayesian beta-binomial model with common beta priors $\pi_j | q_j \sim \text{Beta}(q_{j,1} = 0.50, q_{j,2} = 0.50)$ for $j = A, B$, which has ESS = 1. The posterior is $\text{Beta}(S_j + 0.50, n_j - S_j + 0.50)$, where S_j is the number of successes out of n_j in arm j . Denoting $W = (S_A, n_A - S_A, S_B, n_B - S_B)$ and $q = (q_{A,1}, q_{A,2}, q_{B,1}, q_{B,2})$, we use the test statistic $S_{A>B}(W; q) = \Pr(\pi_A > \pi_B | W, q)$, which we calculate similarly to (12). This is the special case of the Dirichlet-multinomial model and CAT-BUB test with $K = 2$, since the mean utility for treatment j is $100 \times \pi_j$ so the utility is superfluous. To ensure comparability, for the binary test we also use $n = 208$, and set $p_{cut} = 0.975$ to obtain a 0.05-level test when $\pi_A = \pi_B$.

Operating characteristics of the fixed sample CAT-BUB and beta-binomial tests are given in Table 2. Scenario 1.0 is the null case used to calibrate p_{cut} for each design, so the type I error for both designs is 0.05, with equal probabilities for concluding $A > B$ or $B > A$. Scenario 2.4 is the alternative used to select a sample size that provides power 0.80, so the estimated power is in the interval $[0.80 - \epsilon, 0.80 + \epsilon]$. For Scenarios 2.1–2.5, π_B is fixed at 0.60 versus

$\pi_A = 0.50$, so the beta-binomial design always has power 0.55, despite obvious differences between these four scenarios. For example, in Scenarios 2.1 and 2.2, the beta-binomial design fails by concluding $B > A$ 55% of the time even though A is clinically superior or equal to B in terms of $\delta_{U,B-A}(\theta)$. In contrast, the CAT-BUB test distinguishes between these scenarios, correctly concluding $A > B$ 21% of the time in Scenario 2.1 and controlling type I error at 0.05 in Scenario 2.2. Scenarios 2.3–2.5 exhibit various tradeoffs that favor B over A in an increasing manner in terms of $\delta_{U,B-A}(\theta) = 5, 10, 15$ and the CAT-BUB test reflects this with increasing power figures 0.246, 0.798, 0.997. In particular, the CAT-BUB test has substantially more power than the beta-binomial test for the “win-win” Scenarios 2.4 and 2.5. Scenarios 3.1–3.4 and 4.1–4.4 respectively fix the probability of response at 0.65 or 0.70, for which the beta-binomial test has 0.88 and 0.99 power figures. In contrast, the power of the CAT-BUB test increases as the true utility difference $\delta_{U,B-A}(\theta)$ increases, and equals or exceeds that of the beta-binomial design in “win-win” scenarios where the probability of failure is also reduced (Scenarios 3.3–3.4 and 4.3–4.4). The failure of the beta-binomial design is due to B providing an unfavorable trade-off between the probability of response and failure versus A . Such tradeoffs cannot be identified by the naive binary outcome design, which is used very commonly. The price of the CAT-BUB approach is potentially less power for “tradeoff” scenarios when the treatment redistributes probability away from N to R and/or F (Scenarios 2.3, 3.2, 4.1 and 4.2). However, we feel that this price is well worth being able to distinguish between, for example, Scenarios 2.1–2.5 in practice. Lastly, the CAT-BUB test has varying power over the set of θ with the same utility difference. For example, Scenarios 2.3 and 4.1 have utility difference 5, yet power figures 0.25 and 0.22, respectively. The CAT-BUB design’s power varies more substantially with $\delta_{U,A-B}(\theta)$ than over the set of θ with the same utility difference.

6.1.2. Sensitivity to Elicited Utilities—The power of the CAT-BUB design at the targeted alternatives, and thus the required sample size, depend on the particular elicited utilities. The sensitivity of the CAT-BUB test’s power to the elicited utilities can be assessed by fixing the sample size and calculating the power for targeted alternatives using other numerical utilities. Continuing with the example involving clot dissolving agents, we fix $n = 208$, $p_{cut} = 0.976$ and $\theta_A = (0.50, 0.30, 0.20)'$, and calculate the CAT-BUB test’s power for alternative Scenarios 2.1, 2.2, 2.4 and 2.5 in Table 2 over the entire domain $U_N \in [0, 100]$. We calculate power using (15), which is quite accurate when $n = 208$.

Figure 1 plots the *overall* power as a function of U_N at each alternative; that is, we do not explicitly distinguish between the decisions $A > B$ and $B > A$. Scenarios 2.1 and 2.2 are trade-off scenarios, wherein B relative to A has a higher probability of R and F , and lower probability of N , so $\delta_{U,A-B}(\theta)$ varies substantially with U_N and the power is thus quite sensitive. This sensitivity is a desirable property, because the numerical value of U_N determines whether a particular trade-off favors $B > A$ or $A < B$. For Scenarios 2.1 and 2.2, although not explicitly depicted, the CAT-BUB test has power primarily for $A > B$ ($B > A$) to the left (right) of the numerical utility with minimum power. In Scenario 2.4, the probability of N is equal for both A and B , so $\delta_{U,B-A} = 10$ for all U_N , and the sensitivity merely reflects the relationship between U_N and δ_+ . In win-win Scenario 2.5, power increases with U_N because $\delta_{U,B-A}(\theta)$ increases with U_N . Lastly, each scenario in Figure 1

fixes $\theta_{R,B} = 0.60$ and $\theta_{R,A} = 0.50$, and for $U_N = 0$ the CAT-BUB test is based exclusively on $100 \times \theta_R$, so it is identical to the usual beta-binomial test, providing 54% power. Therefore, even in settings where selecting a particular U_N may be challenging, the proposed CAT-BUB test obviously is more sensible than the usual beta-binomial test, which implicitly sets $U_N = 0$.

It is useful to consider sensitivity of inferences to the elicited utilities. We illustrate how this may be done for bivariate binary outcomes. Figure 2 depicts the posterior probability for $B > A$, defined by (8), for the AML bivariate binary example from Section 1, given three different realizations of $\mathbf{X}_B$ while fixing $\mathbf{X}_A = (X_{A,[C,\bar{T}]}, X_{A,[C,T]}, X_{A,[\bar{C},\bar{T}]}, X_{A,[\bar{C},T]})' = (15, 20, 25, 40)'$ and enumerating over $(U_{C,T}, U_{\bar{C},\bar{T}}) \in [0, 100]^2$, i.e. all possible intermediate utilities. In Scenario 1, $\mathbf{X}_B = (10, 40, 20, 30)'$ and $\mathbf{X}_A = (15, 20, 25, 40)'$. For these data, inference is more sensitive to $U_{C,T}$ than $U_{\bar{C},\bar{T}}$, because the two treatments appear to differ greatly for the probability of $[C, T]$ (40 versus 20 observations) and little for the probability of $[\bar{C}, \bar{T}]$ (20 versus 25 observations). The data in Scenario 2 reflect a similar yet smaller treatment difference, and inference is less sensitive to the utilities. In Scenario 3, the data suggest that B is a win-win relative to A in that B has both higher marginal probability of C and lower marginal probability of T , and the CAT-BUB design's inference always supports the conclusion $B > A$. For these data, posterior evidence supporting $B > A$ becomes stronger as $U_{C,T}$ is increased.

6.1.3. Group Sequential Tests—To assess the operating characteristics of the group sequential tests, we continue with the trinary versus binary example. We assume the analysis schedule has $S = 3$ equally spaced looks at $t_1 = 0.33$, $t_2 = 0.66$ and $t_3 = 1$. We use the same targeted alternative as the fixed sample design for calibration, and compare the operating characteristics of the group sequential versions of the CAT-BUB design and beta-binomial design for Scenarios 1.0 and 2.1–2.5 used for the fixed sample simulation. We applied the group sequential CAT-BUB design algorithm, given in Section 5.3, to maintain $\alpha = 0.05$ with $\rho = 3$. This gave $n_S = 213$, $p_{cut,1} = 0.999$, $p_{cut,2} = 0.993$ and $p_{cut,3} = 0.978$. Scenarios 1.0 and 2.4 were used to jointly calibrate the planned sample size and probability thresholds to provide type I error of 0.05 and overall power of 0.80, respectively. For the beta-binomial design, to maintain size 0.05 we used $p_{cut,1} = 0.999$, $p_{cut,2} = 0.992$ and $p_{cut,3} = 0.979$. We used $n_S = 213$ for both designs to ensure comparability.

The results of the group sequential simulations are reported in Table 3. For the null Scenario 1.0, the operating characteristics of the CAT-BUB and beta-binomial designs are practically identical. Both designs have an average sample size of 212 and overall type I error of 0.05. In contrast, for Scenarios 2.1–2.5, the operating characteristics of the two designs differ dramatically. The beta-binomial design does not distinguish between these scenarios because $\pi_{B,R} = 0.60$ and $\pi_{A,R} = 0.50$ for all 5 scenarios, whereas the CAT-BUB test distinguishes between them quite well. In Scenario 2.1, A is preferred over B due to an unfavorable tradeoff between R and F . The beta-binomial design incorrectly selects B over A 54% of the time with an average sample size of 193, whereas the CAT-BUB design correctly selects A over B 21% of the time with an average sample size of 208. In Scenario 2.2, B and A are equivalent due to the increase in response probability being canceled out by the increase in failure probability. Here, the CAT-BUB design controls type I error at level 0.05. Scenarios

2.3–5 have increasing magnitudes of the benefit for B over A , and the CAT-BUB design has increasing power for concluding $B > A$. As the true benefit of B over A increases, the average sample size of the CAT-BUB design decreases because the probability of early termination increases. In the most favorable Scenario 2.5, the CAT-BUB design has power 0.998 and terminates early nearly 95% of the time, with average sample size 124 that is 42% smaller than the planned maximum sample size. In contrast, the beta-binomial design has 54% power and average sample size 193 in this case, as in all Scenarios 2.1 – 2.5, essentially because it ignores the distinction between N and F .

6.2. Redesigning the CLL Trial

Returning to the CLL trial, we illustrate how to implement the CAT-BUB design in this context. We assume that the elicited numerical utilities are those in Table 1. Recall that, since we cannot elicit utilities for this trial retrospectively, as explained in Section 1 the utilities in Table 1 are specified to be a reasonable representation of what one actually would elicit in practice. We compare the CAT-BUB design with a beta-binomial design based on an efficacy test, which we denote by BB-EO. Like the actual trial, the BB-EO design defines efficacy using a binary indicator for CR , where the comparative test was based on targeted alternative CR probability $\pi_{CR,FC} = 0.45$ versus null $\pi_{CR,F} = 0.25$ [6]. We also compare the CAT-BUB design to an alternative approach that is based on a hierarchical testing procedure. This alternative design first compares the probabilities of efficacy (here, CR) as the primary endpoint and, if this test fails to reject the null, then the procedure compares the probabilities of toxicity (here, severe or fatal AE) in a second test. This design, which we denote by BB-ET, assumes independent beta-binomial models for the two outcomes. Based on this hierarchical testing procedure, the BB-ET design recommends a treatment if it is found to have either better efficacy or toxicity compared to the other treatment.

Because the actual CLL trial outcome is bivariate ordinal plus death, to implement the CAT-BUB design, a practical approach for eliciting the targeted alternative(s) is as follows. First, ask the physicians to hypothesize the marginal probabilities of the AE levels, {Min, Mod, Sev, Fatal}, in each treatment group. Denote these probabilities by

$$\theta_{j,T} = (\theta_{j,Min}, \theta_{j,Mod}, \theta_{j,Sev}, \theta_{j,Fatal}), \text{ where } \theta_{j,Min} + \theta_{j,Mod} + \theta_{j,Sev} + \theta_{j,Fatal} = 1, j = F, FC.$$

Next, ask the physicians to hypothesize probabilities of the clinical response events, {CR, PR, SD, PD}, conditional on being alive. Denote these conditional probabilities by

$$\theta_{j,E} = (\theta_{j,CR}, \theta_{j,PR}, \theta_{j,SD}, \theta_{j,PD}), \text{ where } \theta_{j,CR} + \theta_{j,PR} + \theta_{j,SD} + \theta_{j,PD} = 1, j = F, FC.$$

Assuming independence for simplicity, set $\theta_{j,Fatal}^{(Alt)} = \theta_{j,Fatal}$ and $\theta_{j,k,\ell}^{(Alt)} = \theta_{j,k} \theta_{j,\ell}$ for $j = F, FC$, $k = \{Min, Mod, Sev\}$, and $\ell = \{CR, PR, SD, PD\}$. We assume that the targeted alternative arises from $\theta_{F,T} = \theta_{FC,T} = (0.67, 0.25, 0.05, 0.03)$, i.e., FC and F have equivalent toxicity, and $\theta_{F,E} = (0.25, 0.35, 0.20, 0.20)$ versus $\theta_{FC,E} = (0.45, 0.35, 0.10, 0.10)$, i.e., FC compared to F has higher efficacy. This alternative maintains similar marginal CR probabilities $\pi_{CR,FC}$

$= 0.4365$ versus $\pi_{CR,F} = 0.2425$ specified for the actual trial design, and it results in a large mean utility difference $\delta_{U,FC-F}^{(Alt)} = 13.5$ for the utilities given in Table 1. Specifying $n^* = 1$ and $\theta^* = \theta_F$ for the Dirichlet priors, a fixed sample CAT-BUB test requires slightly more patients than a beta-binomial test to achieve 90% power, $n_F = n_{FC} = 127$ versus 120. To ensure comparability, we determine the power figures for all three designs using the larger sample size 128. We compare the designs for 12 scenarios covering a wide range of different possibilities. The response probabilities for F and FC in each scenario are reported in Table 4. These probabilities for F are fixed at the targeted alternative values throughout, whereas the probabilities for FC vary with the combinations of the ordinal efficacy and toxicity outcomes. For each outcome, we characterize these numerical probability vector pairs nominally as being “equivalent” ($=$), or having “moderate” ($>$), “large” ($>>$), or “very large” ($>>>$) differences.

The results of our simulation reported in Table 5 show that, in general, the CAT-BUB design is sensitive to efficacy-toxicity tradeoffs characterized by the utilities, while the beta-binomial design with an efficacy test (BB-EO) is not. In contrast, if there is low power for detecting an efficacy difference between the two treatments, then the hierarchical beta-binomial design (BB-ET) is sensitive to toxicity, otherwise it is not. Because the BB-ET design is based on two tests, rather than one test like the BB-EO design, it requires a more stringent cut-off to control the type I error, and thus has lower power compared to the BB-EO design for selecting the treatment with superior efficacy, which is FC in all the scenarios we considered. Scenario 1.0 is the null, i.e., $\theta_{FC} = (\theta_{FC,E}, \theta_{FC,T}) \equiv (\theta_{F,E}, \theta_{F,T}) = \theta_F$ and the cut-off for all three designs was calibrated to control type I error at the $\alpha = 0.05$ -level, where F and FC are selected with the same 0.025 probabilities. In Scenarios 1.1 and 1.2, FC has equivalent efficacy with moderate and high toxicity, respectively, and both the CAT-BUB design and the BB-ET design are increasingly likely to select F , whereas the BB-EO design is unable to distinguish between these clinically very different scenarios since it ignores toxicity. The BB-ET design is more likely to correctly select F than the CAT-BUB design, 0.39 versus 0.22 and 0.98 versus 0.78, respectively. In Scenario 2.0, FC has a moderate efficacy advantage with equivalent toxicity, and the BB-EO, CAT-BUB, and BB-ET designs select FC with probabilities 0.40, 0.35, and 0.29, respectively. In Scenario 2.1, FC has a moderate efficacy advantage and toxicity disadvantage, where this tradeoff that slightly favors FC for the assumed utilities. The CAT-BUB design is unlikely to select either treatment, whereas the BB-EO design selects FC with probability 0.33, and the BB-ET design selects FC and F with probabilities 0.23 and 0.31, respectively. In Scenario 2.2, because the toxicity disadvantage for FC increases, the tradeoff moderately favors F for the assumed utilities. The CAT-BUB design selects F with higher probability 0.31 and does not select FC , whereas the BB-EO design selects FC with probability 0.28 and does not select F , and the BB-ET design selects F with probabilities 0.82 and FC with probability 0.17. Scenario 3.0 is the targeted alternative, which is a case where FC has higher efficacy and equivalent toxicity compared to F . In this ideal case, the CAT-BUB design has 90% power compared to 91% power for the BB-EO design and 85% for the BB-ET design. Scenarios 3.1, 3.2, 4.1, and 4.2 are tradeoff settings where FC has an a large or very large efficacy advantage, and either a moderate or large toxicity disadvantage. In these cases, the CAT-BUB design selects treatments with probabilities that are sensitive to the assumed utilities,

whereas the beta-binomial designs consistently select FC with high probability, regardless of the toxicity burden of FC .

Scenarios 1.2, 2.2, and 3.2 show very undesirable potential consequences of using the BB-EO design, which completely ignores toxicity. The BB-EO design based on the probability of CR treats Scenario 1.2 like a null case since $\pi_{CR,FC} = 0.2275$ versus $\pi_{CR,F} = 0.2425$, while in fact the two pairs of toxicity probability vectors θ_{FC} and θ_F are very different, with $\theta_{FC,T} = (0.26, 0.45, 0.20, 0.09)$ versus $\theta_{F,T} = (0.67, 0.25, 0.05, 0.03)$, so that FC has a much lower minor toxicity probability but much higher moderate, severe, and fatal AE probabilities compared to F . The CAT-BUB design recognizes this, concluding that $F > FC$ with power 0.78 compared to 0.05 for the BB-EO design. Scenario 2.2 is an intermediate case, since FC has moderate efficacy with $\theta_{FC,E} = (0.35, 0.35, 0.15, 0.15)$ versus $\theta_{F,E} = (0.25, 0.35, 0.20, 0.20)$ but also high toxicity, with $\theta_{FC,T} = (0.26, 0.45, 0.20, 0.09)$ versus $\theta_{F,T} = (0.67, 0.25, 0.05, 0.03)$. The BB-EO design detects the $0.3185 - 0.2425 = 0.076$ difference in CR probabilities in favor of FC with power 0.28, but since the probability of severe toxicity or death is 0.29 with FC versus 0.08 with F , the CAT-BUB design concludes F is superior to FC with power 0.314 and never concludes that FC is superior to F . In Scenario 3.2, FC has a large efficacy advantage but also high toxicity burden, with $\theta_{FC,E} = (0.45, 0.35, 0.10, 0.10)$ versus $\theta_{F,E} = (0.25, 0.35, 0.20, 0.20)$ but, as in Scenario 2.2, $\theta_{FC,T} = (0.26, 0.45, 0.20, 0.09)$ versus $\theta_{F,T} = (0.67, 0.25, 0.05, 0.03)$. The BB-EO design has power 0.82 of concluding that FC is superior to F , whereas the CAT-BUB design recognizes both the much better efficacy and much worse toxicity with FC compared to F , and based on the assumed utilities does not recommend either treatment over the other with probability $1 - (0.041 + 0.015) = 0.94$. Scenario 4.2 shows that the BB-TE design can have a similar undesirable behavior as the BB-EO design. If the efficacy advantage of FC is very large, because the efficacy test will detect a difference with high probability, the BB-TE design effectively ignores toxicity, since the toxicity test is unlikely to be applied. In Scenarios 2.1 and 3.2, which have less extreme tradeoffs than Scenario 4.2, the BB-TE design is likely to recommend a particular treatment, despite that neither treatment may be strongly preferred under the assumed utilities. For example, in Scenario 3.2, the BB-TE design recommends FC and F with probabilities 0.72 and 0.28, respectively, and thus recommends either treatment with probability 0.99. Lastly, the BB-TE design has lower power than the CAT-BUB design for the targeted alternative in the CLL trial, i.e., Scenario 3.0.

There are several key points in these comparisons. First, basing a test on the probability of CR is equivalent to using a degenerate utility-based test that assigns utilities 100 to CR and 0 to its complement, while completely ignoring toxicity. An elaboration of this that accounts for the ordinal categories of efficacy is the two-sample test of [19], although this test still suffers from the fact that it ignores toxicity. Considering Scenarios 3.0, 3.1, and 3.2 together shows how the CAT-BUB design adjusts its conclusions depending on the varying $\theta_{FC,T}$ vectors, essentially agreeing with the BB-EO design when toxicity is equivalent but very likely reaching the opposite conclusion when FC has much higher toxicity than F . The same pattern can be seen when considering Scenarios 4.0, 4.1, and 4.2 together. This also illustrates the benefits from considering the ordinal level of each outcome rather than dichotomizing it, since the probabilities of concluding that FC is superior to F vary from 1 to

0.20 as the probabilities of the levels of each outcome change across scenarios. In practice, if a conventional design based on efficacy alone is used one might hope that, in such cases, at some point during actual trial conduct someone would notice an excessively higher toxicity rate in one arm compared to the other, and ask the Principal Investigator or Institutional Review Board to halt accrual to the trial. Hope is not a strategy, however. Moreover, if in fact a trial designed based on efficacy alone will be stopped due to such a toxicity difference, then the nominal size and power of the design are incorrect, and in fact they are conditional on the assumption that there will be no difference in toxicity sufficiently large that it would cause the trial to be stopped early. All of these concerns are taken care of automatically by the group sequential CAT-BUB test's structure. For the group sequential design of the CLL trial (see Web Supplement), its interim decision rules will stop the trial early with high probability when there is a large difference in terms of either efficacy or toxicity, as quantified by the joint utilities of the elementary (efficacy, toxicity) outcomes.

The utilities in Table 1 used for the CAT-BUB re-design of the CLL trial emphasize both avoiding severe toxicity or death and achieving good clinical response by specifying the relative utility parameters (Section 2) to be $\zeta_1 = 0.1$ and $\zeta_2 = 0.2$ near zero, respectively. To assess the CAT-BUB design's sensitivity to the utilities, we also considered the two alternative utilities given in Table 6. One alternative places greater importance on achieving better efficacy, and the other places greater importance on achieving lower toxicity. We obtained the first (second) set of alternative utilities by changing ζ_2 from 0.20 to 0.80 (ζ_1 from 0.10 to 0.60), while retaining all other indirect elicitation parameter values as detailed in Section 2. For the first alternative, the utilities of moderate or severe toxicity for CR or PR levels of efficacy are substantially increased from the original utility, for example, $U(CR, Sev)$ is increased from 28 to 82. For the second alternative, the utilities of minimal or moderate toxicity for SD or PD levels of efficacy are substantially increased from the original utility, for example, $U(SD, Min)$ is increased from 35 to 71. The required sample sizes for the two alternative utilities are $n_F = n_{FC} = 110$ and 271, respectively, which is in contrast to $n_F = n_{FC} = 127$ for the elicited numerical utilities. The CAT-BUB design based on the utilities that place greater importance on higher efficacy thus requires a smaller sample size to achieve 90% power for detecting the targeted alternative compared with the beta-binomial design that was used for the actual trial. The numerical utilities that place greater importance on higher efficacy (lower toxicity) provide more (less) power for detecting the targeted alternative, which has a large treatment difference for the marginal probabilities of clinical response and zero difference for AE probabilities.

Table 7 reports the power figures for the same scenarios as in Table 5 for the CAT-BUB design based on the two alternative utilities with the above sample sizes. These power figures illustrate that the numerical utilities affect the power of detecting specific treatment differences and determine which therapy is preferred for tradeoff scenarios. Scenario 4.2 is an extreme tradeoff example, wherein FC has very high efficacy and high toxicity compared to F . As shown by the mean utility difference and power figures, this tradeoff slightly favors FC under the original utility function, whereas FC is strongly favored under the alternative utilities that place greater importance on efficacy, and conversely, under the alternative utilities that place greater importance on lower toxicity. The dependence of the proposed

CAT-BUB design on the numerical utilities underscores the importance of eliciting values that actually reflect the clinical desirabilities of each patient response.

6.3. Additional Illustrations

In Section 1, we introduced several categorical outcome structures, including trinary, bivariate binary, bivariate binary plus death, ordinal, and bivariate ordinal. The computational algorithms given previously readily accommodate all of these cases. In Web Appendix D, we provide detailed illustrations of both fixed sample and group sequential CAT-BUB designs for a bivariate binary outcome, and for a bivariate ordinal outcome having $16 = 4 \times 4$ elementary events. We also include a group sequential CAT-BUB re-design of the CLL trial.

7. Discussion

Because clinical trial conduct must accommodate medical practice, a trial design should account formally for risk-benefit tradeoffs between all clinically relevant outcomes. The utility-based tests that we have proposed provide a practical approach for comparing treatments based on categorical outcomes in a RCT. We have provided both fixed sample and group sequential procedures, computational algorithms to derive design parameters, and freely-available user-friendly software. The CAT-BUB test directly addresses the problem of comparing treatments for all *clinically relevant* differences. The method deals with the problem of deciding whether one therapy is clinically superior to another based on its outcome probability vector by exploiting the elicited utilities of the elementary outcomes to reduce the multidimensional outcome to a one dimensional mean utility. This is used to construct comparative tests. The elicited utilities provide a rigorous framework for treatment comparison that makes explicit any subjective tradeoffs between outcomes.

We have demonstrated that designs based on a single binary outcome for efficacy often are unsafe, and do not reflect medical practice. Because safety is never a secondary concern in any clinical trial, conventional designs only based on an efficacy test presumably rely on informal stopping criteria for safety, which makes it difficult to determine operating characteristics. We also have demonstrated that designs based on a hierarchical testing procedure may be unsafe in scenarios where one treatment is more efficacious but also is more toxic than the other drug. These designs also do not reflect medical practice. If an efficacy difference is detected, it is naive to believe that the toxicity profiles of the available treatment options will not be considered by physicians when deciding which treatment is actually clinically preferable. In the proposed method, we have provided a practical tool to explicitly account for the tradeoffs between disparate outcomes that physicians routinely assess, and thus for designing RCTs that better reflect medical practice.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Contract/grant sponsor: NIH/NCI grant 5-R01-CA083932

References

1. Eisenhauer E, Therasse P, Bogaerts J, Schwartz L, Sargent D, Ford R, Dancey J, Arbuck S, Gwyther S, Mooney M, et al. New response evaluation criteria in solid tumours: Revised {RECIST} guideline (version 1. 1). *European Journal of Cancer*. 2009; 45(2):228–247. <http://dx.doi.org/10.1016/j.ejca.2008.10.026>. URL <http://www.sciencedirect.com/science/article/pii/S0959804908008733>, response assessment in solid tumours (RECIST): Version 1.1 and supporting papers. [PubMed: 19097774]
2. Sankoh AJ, D'Agostino RB, Huque MF. Efficacy endpoint selection and multiplicity adjustment methods in clinical trials with inherent multiple endpoint issues. *Statistics in Medicine*. 2003; 22(20):3133–3150. <http://dx.doi.org/10.1002/sim.1557>. DOI: 10.1002/sim.1557 [PubMed: 14518019]
3. Pocock SJ. Clinical trials with multiple outcomes: A statistical perspective on their design, analysis, and interpretation. *Controlled Clinical Trials*. 1997; 18(6):530–545. [http://dx.doi.org/10.1016/S0197-2456\(97\)00008-1](http://dx.doi.org/10.1016/S0197-2456(97)00008-1). URL <http://www.sciencedirect.com/science/article/pii/S0197245697000081>, Eighth International Symposium on Long-Term Clinical Trials. [PubMed: 9408716]
4. Freedman L, Anderson G, Kipnis V, Prentice R, Wang C, Rossouw J, Wittes J, DeMets D. Approaches to monitoring the results of long-term disease prevention trials: Examples from the Women's Health Initiative. *Controlled Clinical Trials*. 1996; 17(6):509–525. [http://dx.doi.org/10.1016/S0197-2456\(96\)00016-5](http://dx.doi.org/10.1016/S0197-2456(96)00016-5). URL <http://www.sciencedirect.com/science/article/pii/S0197245696000165>. [PubMed: 8974210]
5. Hotelling H. The economics of exhaustible resources. *Journal of Political Economy*. 1931; 39(2): 137–175. <http://www.jstor.org/stable/1822328>.
6. Flinn IW, Neuberg DS, Grever MR, Dewald GW, Bennett JM, Paietta EM, Hussein MA, Appelbaum FR, Larson RA, Moore DF, et al. Phase III trial of fludarabine plus cyclophosphamide compared with fludarabine for patients with previously untreated chronic lymphocytic leukemia: US Intergroup Trial E2997. *Journal of Clinical Oncology*. 2007; 25(7):793–798. <http://jco.ascopubs.org/content/25/7/793.abstract>. DOI: 10.1200/JCO.2006.08.0762 [PubMed: 17283364]
7. Cheson B, Bennett J, Grever M, Kay N, Keating M, O'Brien S, Rai K. National Cancer Institute-Sponsored Working Group guidelines for chronic lymphocytic leukemia: revised guidelines for diagnosis and treatment. *Blood*. 1996; 87(12):4990–4997. [PubMed: 8652811]
8. Dalkey, NC. The Delphi Method: An Experimental Study of Group Opinion. Santa Monica, CA: RAND Corporation; 1969. http://www.rand.org/pubs/research_memoranda/RM5888
9. Brook RH, Chassin MR, Fink A, Solomon DH, Koseoff J, Park RE. A method for the detailed assessment of the appropriateness of medical technologies. *International Journal of Technology Assessment in Health Care*. 1986; 2:53–63. http://journals.cambridge.org/article_S0266462300002774. DOI: 10.1017/S0266462300002774 [PubMed: 10300718]
10. Ibrahim JG, Chen MH. Power prior distributions for regression models. *Statist Sci*. 2000; 15(1): 46–60. <http://dx.doi.org/10.1214/ss/1009212673>. DOI: 10.1214/ss/1009212673
11. Murray T, Hobbs B, Carlin B. Combining nonexchangeable functional or survival data sources in oncology using generalized mixture commensurate priors. *Annals of Applied Statistics*. 2015 In Press.
12. Carlin, BP.; Louis, TA. *Bayesian Methods for Data Analysis*, 3rd edition. Boca-Raton, FL: Chapman & Hall/CRC Press; 2009.
13. Piessens, R.; Doncker-Kapenga, ED.; Überhuber, CW. QUADPACK: a subroutine package for automatic integration. Springer; 1983.
14. Gelman, A.; Carlin, JB.; Stern, HS.; Dunson, DB.; Vehtari, A.; Rubin, DB. *Bayesian Data Analysis*. 3. Boca-Raton, FL: Chapman & Hall/CRC Press; 2014.
15. Jennison, C.; Turnbull, BW. *Group Sequential Methods Applications to Clinical Trials*. Boca-Raton, FL: Chapman & Hall/CRC Press; 2000.
16. Slud E, Wei LJ. Two-sample repeated significance tests based on the modified wilcoxon statistic. *Journal of the American Statistical Association*. 1982; 77(380):862–868. <http://www.jstor.org/stable/2287319>.

17. Lan KKG, DeMets DL. Discrete sequential boundaries for clinical trials. *Biometrika*. 1983; 70(3): 659–663. <http://www.jstor.org/stable/2336502>.
18. Kim K, DeMets DL. Design and analysis of group sequential tests based on the type I error spending rate function. *Biometrika*. 1987; 74(1):149–154.
19. Whitehead J. Sample size calculations for ordered categorical data. *Statistics in Medicine*. 1993; 12(24):2257–2271. <http://dx.doi.org/10.1002/sim.4780122404>. DOI: 10.1002/sim.4780122404 [PubMed: 8134732]

Figure 1.

Sensitivity to U_N of the fixed sample CAT-BUB design's power for various alternative θ_B s, fixing $n = 208$, $p_{cut} = 0.976$, and $\theta_A = (0.50, 0.30, 0.20)'$. The thick dot denotes power for the elicited utilities, i.e. $U_N = 50$.

Figure 2.

Posterior probability of $B > A$ while varying $(U_{C,T}, U_{\bar{C},\bar{T}})$ for three different realizations of $\mathbf{X}_B$ and $\mathbf{X}_A = (\mathbf{X}_{A,[C,\bar{T}]}, \mathbf{X}_{A,[C,T]}, \mathbf{X}_{A,[\bar{C},\bar{T}]}, \mathbf{X}_{A,[\bar{C},T]})' = (15, 20, 25, 40)'$. The thick dot denotes our inferential result at the elicited utilities, i.e. $(U_{C,T}=80, U_{\bar{C},\bar{T}}=40)$.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 1

Numerical utilities for the CLL trial’s 13 elementary outcomes.

Level of Worst Adverse Event	Clinical Response			
	CR	PR	SD	PD
<i>Minimal</i>	100	84	35	19
<i>Moderate</i>	93	77	29	14
<i>Severe</i>	28	24	14	10
				Death
				0

Table 2

Power of a fixed sample CAT-BUB design for trinary outcome $\{R, N, F\}$ versus a beta-binomial design based on “success” probability $\pi_j = \theta_{j,R}$ for $j = A, B$. In all scenarios, $\theta_A = (0.50, 0.30, 0.20)'$ and $n_A, n_B = 208$. Results in the first row are based on 50,000 simulations (std.err. ≈ 0.001), whereas all other results are based on 25,000 simulations (std.err. 0.0032).

Scenario θ_B	CAT-BUB Design			Beta-Bin Design		
	$\delta_{U,B-A}(\theta)$	$B > A$	$A > B$	$B > A$	$A > B$	
1.0: (0.50, 0.30, 0.20)	0	0.025	0.025	0.025	0.025	0.025
2.1: (0.60, 0.00, 0.40)	-5	0.001	0.206	0.552	0.000	
2.2: (0.60, 0.10, 0.30)	0	0.024	0.025			
2.3: (0.60, 0.20, 0.20)	5	0.246	0.001			
2.4: (0.60, 0.30, 0.10)	10	0.798	0.000			
2.5: (0.60, 0.40, 0.00)	15	0.997	0.000			
3.1: (0.65, 0.05, 0.30)	2.5	0.088	0.006	0.877	0.000	
3.2: (0.65, 0.15, 0.20)	7.5	0.485	0.000			
3.3: (0.65, 0.25, 0.10)	0.936	0.000				
3.4: (0.65, 0.35, 0.00)	1.000	0.000				
4.1: (0.70, 0.00, 0.30)	5	0.217	0.001	0.989	0.000	
4.2: (0.70, 0.10, 0.20)	10	0.720	0.000			
4.3: (0.70, 0.20, 0.10)	15	0.987	0.000			
4.4: (0.70, 0.30, 0.00)	20	1.000	0.000			

Power figures of a group sequential CAT-BUB design for a trinary outcome $\{R, N, F\}$ versus a beta-binomial design based on “success” probabilities $\pi_j = \theta_{j,R}$ for $j = A, B$. In each scenario, $\theta_A = (0.50, 0.30, 0.20)'$, $n_1 = 71$, $n_2 = 142$, $n_3 = 213$, and $\rho = 3$.

Table 3

Scenario Specification		CAT-BUB Design			Beta-Binomial Design		
θ_B	$\delta_{U,B-A}(\theta)$	Ave SS	$B > A$	$A > B$	Ave SS	$B > A$	$A > B$
1.0: (0.50, 0.30, 0.20)	0	211.9	0.025	0.025	211.8	0.026	0.024
2.1: (0.60, 0.00, 0.40)	-5	207.7	0.001	0.214	192.8	0.541	0.000
2.2: (0.60, 0.10, 0.30)	0	211.8	0.026	0.025			
2.3: (0.60, 0.20, 0.20)	5	206.6	0.250	0.001			
2.4: (0.60, 0.30, 0.10)	10	177.8	0.800	0.000			
2.5: (0.60, 0.40, 0.00)	15	123.8	0.998	0.000			

Table 4

Response probabilities for the scenarios considered in our CLL trial simulation study. Toxicity probabilities correspond to {Min, Mod, Sev, Fatal}, and efficacy probabilities correspond to {CR, PR, SD, PD}, given that the patient is alive.

Scenarios	Abbreviation	Response Probabilities
All	NA	$\theta_{FT} = (0.67, 0.25, 0.05, 0.03)$
1.0, 2.0, 3.0, 4.0	=	$\theta_{FC,T} = (0.67, 0.25, 0.05, 0.03)$
1.1, 2.1, 3.1, 4.1	>	$\theta_{FC,T} = (0.44, 0.40, 0.10, 0.06)$
1.2, 2.2, 3.2, 4.2	>>	$\theta_{FC,T} = (0.26, 0.45, 0.20, 0.09)$
All	NA	$\theta_{FE} = (0.25, 0.35, 0.20, 0.20)$
1.0, 1.1, 1.2	=	$\theta_{FC,E} = (0.25, 0.35, 0.20, 0.20)$
2.0, 2.1, 2.2	>	$\theta_{FC,E} = (0.35, 0.35, 0.15, 0.15)$
3.0, 3.1, 3.2	>	$\theta_{FC,E} = (0.45, 0.35, 0.10, 0.10)$
4.0, 4.1, 4.2	>>>	$\theta_{FC,E} = (0.60, 0.30, 0.05, 0.05)$

Table 5

Power figures for the CLL trial based on the CAT-BUB design, the beta-binomial design with an efficacy test only (BB-EO), and the hierarchical beta-binomial design with an efficacy test followed by a toxicity test (BB-ET). Comparisons of $\theta_{FC,E}$ vs θ_{FE} and $\theta_{FC,T}$ vs θ_{FT} are characterized as being “equivalent” (=), or having “moderate” (>), “large” (>>), or “very large” (>>>) differences.

Scenario		Probability of Final Conclusion							
Efficacy FC vs F	Toxicity FC vs F	$\theta_{U,FC-F}$	CAT-BUB Design		Efficacy Only		Beta-Binomial Designs		
			$FC > F$	$F > FC$	$FC > F$	$F > FC$	$FC > F$	$F > FC$	
1.0:	=	=	0.0	0.025	0.025	0.026	0.024	0.024	0.025
1.1:	=	>	-5.2	0.001	0.222	0.019	0.035	0.007	0.388
1.2:	=	>>	-11.9	0.000	0.782	0.012	0.047	0.005	0.982
2.0:	>	=	6.8	0.352	0.000	0.402	0.000	0.289	0.010
2.1:	>	>	1.1	0.041	0.012	0.331	0.000	0.226	0.308
2.2:	>	>>	-6.5	0.000	0.314	0.278	0.000	0.173	0.818
3.0:	>>	=	13.5	0.903	0.000	0.910	0.000	0.846	0.002
3.1:	>>	>	7.3	0.397	0.000	0.873	0.000	0.778	0.097
3.2:	>>	>>	-1.1	0.041	0.015	0.816	0.000	0.716	0.281
4.0:	>>>	=	21.0	1.000	0.000	1.000	0.000	1.000	0.000
4.1:	>>>	>	14.2	0.917	0.000	1.000	0.000	0.999	0.001
4.2:	>>>	>>	5.0	0.201	0.001	0.999	0.000	0.997	0.003

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6

Three alternative utilities for the CLL trial's outcomes.

Level of Worst Adverse Event	Clinical Response			Death
Original Utilities from Table 1				
	CR	PR	SD	PD
Minimal	100	84	35	19
Moderate	93	77	29	14
Severe	28	24	14	10
Utilities Giving Better Efficacy Higher Value				
Minimal	100	84	35	19
Moderate	98	81	31	14
Severe	82	68	24	10
Utilities Giving Lower Toxicity Higher Value				
Minimal	100	93	71	64
Moderate	93	81	44	32
Severe	28	24	14	10

Table 7

Power figures for the CLL trial based on a fixed sample CAT-BUB design using two alternative sets of numerical utilities that place greater importance on either improving efficacy or reducing toxicity, compared to the original utilities.

Scenario		Alternative Utility Giving Efficacy Higher Value		Alternative Utility Giving Lower Toxicity Higher Value	
Efficacy	Toxicity	$\theta_{L,FC-F}$	$FC > F$	$\theta_{L,FC-F}$	$FC > F$
1.0:	=	0.0	0.025	0.0	0.024
1.1:	>	-3.2	0.003	-8.4	0.000
1.2:	>>	-6.7	0.000	-18.3	0.000
2.0:	=	7.1	0.349	3.6	0.365
2.1:	>	3.7	0.122	-4.6	0.000
2.2:	>>	-0.1	0.026	-14.6	0.000
3.0:	>>	14.2	0.897	7.3	0.902
3.1:	>	10.6	0.647	-0.8	0.011
3.2:	>>	6.5	0.288	-11.0	0.000
4.0:	>>>	22.1	0.999	11.3	0.999
4.1:	>>>	18.2	0.985	3.4	0.290
4.2:	>>>	13.8	0.859	-7.0	0.000
					0.748