

Evaluating multiple treatment courses in clinical trials

Peter F. Thall^{1,*}, Randall E. Millikan² and Hsi-Guang Sung¹

¹ *Department of Biostatistics, Box 213, The University of Texas, M.D. Anderson Cancer Center, 1515 Holcombe Boulevard, Houston, Texas 77030, U.S.A.*

² *Department of Genitourinary Medical Oncology, Box 013, The University of Texas, M.D. Anderson Cancer Center, 1515 Holcombe Boulevard, Houston, Texas 77030, U.S.A.*

SUMMARY

In oncology, a patient's treatment often involves multiple courses of chemotherapy. The most common medical practice in choosing treatments for successive courses is to repeat a treatment that is successful in a given course and otherwise switch to a different treatment. Patient outcome thus consists of a sequence of dependent response variables and corresponding treatments. Despite the widespread use of such adaptive 'play-the-winner-and-drop-the-loser' algorithms in medical settings involving multiple treatment courses, most statistical methods for treatment evaluation characterize early patient outcome as a single response to a single treatment, resulting in a substantial loss of information. In this paper, we provide a statistical framework for multi-course clinical trials involving some variant of the play-the-winner-and-drop-the-loser strategy. The aim is to design and conduct the trial to more closely reflect actual clinical practice, and thus increase the amount of information per patient. The proposed design is similar to a multi-stage cross-over trial, with the essential difference that here all treatments after the first course are assigned adaptively. We illustrate the method by application to a randomized phase II trial for androgen independent prostate cancer. We consider the goals of selecting one best treatment, or selecting a best ordered pair of treatments with the second given if the first fails to achieve a patient success. A simulation study is reported, and extensions to trials involving toxicity or regimen-related death are discussed. Copyright © 2000 John Wiley & Sons, Ltd.

1. INTRODUCTION

Patients undergoing cancer chemotherapy often receive several successive courses of treatment. In typical medical practice, after the first course of chemotherapy the physician decides whether there has been sufficient evidence of response to justify another course of the same treatment. If not, an alternative treatment may be chosen for the second course. The physician continues in this manner, choosing a treatment for each course based on the sequence of treatments given and the outcomes observed in the previous courses, until a particular criterion for terminating the patient's therapy is met. This practice is quite common in settings where several active regimens are available for

* Correspondence to: Peter F. Thall, Department of Biostatistics, Box 213, The University of Texas, M.D. Anderson Cancer Center, 1515 Holcombe Boulevard, Houston, Texas, 77030, U.S.A.

a particular disease but the level of each regimen's activity is highly variable between patients, even after accounting for individual patient covariates related to treatment response.

An important example arises in therapy for solid tumours, where a single course of chemotherapy often succeeds in partially shrinking the patient's tumour but does not eradicate the disease, so that one or more additional courses are required. If little or no shrinkage is achieved in the first course, however, then a different regimen may be given subsequently. In such settings the objective criterion for treatment success on the first course, typically >50 per cent tumour shrinkage ('partial response'), may be quantitatively or qualitatively different from the definition of success on the second course of the same treatment, which may be that 100 per cent shrinkage ('complete response') has been achieved. Similarly, if a leukaemia patient does not achieve a complete remission after an initial series of one or more courses of a particular chemotherapy combination, then it is typical practice to declare the patient's cancer 'resistant' to that treatment and switch to a different treatment. Because the physician must rely on clinical experience and numerous variables to decide whether a patient's leukaemia is resistant, the number of courses of the initial regimen varies from patient to patient. A third example arises in treatment of life-threatening bacterial or fungal infections, two common adverse side-effects of chemotherapy, where the patient may be given several different agents in succession until either the infection is resolved or the patient dies.

The treatment selection algorithm common to these examples may be described as 'play-the-winner-and-drop-the-loser'. Despite the widespread use of this general approach by physicians to meet the practical demands of treating patients in successive courses when multiple active regimens are available, to our knowledge there have been no attempts to design clinical trials based on this paradigm. The most common statistical method for evaluating multiple treatment courses in clinical research is to define each patient's outcome in terms of the treatment given in the first course while ignoring any other treatments given to the patient subsequently if the initial course is not successful. Thus, a patient for whom treatment u is unsuccessful in the first course is scored as a failure with u , even if success is achieved subsequently with some other treatment t . Alternatively, the patient may be scored as a success with t while ignoring the initial failure with u . Evidently, both approaches ignore much of the available information.

An important phenomenon in the multi-course setting that is ignored by this sort of simplification is that of *cross-resistance* between u and t . This occurs when a failure with u on a given course renders the probability of success with t on the next course lower than if the patient had failed with some treatment other than u . An interesting aspect of the play-the-winner-and-drop-the-loser algorithm, illustrated in Section 3.2, is that a patient randomized into a trial of several treatments using this method may have a higher overall success probability than a patient treated initially with the single best regimen. This design property is especially important given concerns in the medical community regarding the ethics of randomized trials [1].

The purpose of this article is to provide a general statistical framework for a randomized multi-course clinical trial where some variant of the play-the-winner-and-drop-the-loser strategy is used. The proposed design is adaptive in that treatment decisions for an individual patient in all courses after the first are based on the patient's outcomes in previous courses. The inferential goals that we consider are selection of one best treatment, selection of a best ordered pair of treatments with the second given if the first fails to achieve a patient success, and estimation of cross-resistance between consecutive treatments. Our proposed methodology is a phase II design in that it is based on a patient outcome that, within each course of therapy, may be observed relatively soon, and the goal is to select a treatment or treatments for future study in a phase III trial having survival or disease-free survival time as the primary outcome.

The remainder of the paper is organized as follows. In Section 2, we describe the prostate cancer trial that initially motivated this research. We establish multinomial and regressive logistic probability models in Section 3. In Section 4 we summarize a simulation study of selection designs based on these models in the context of the prostate cancer trial. Generalizations of the method, including versions of the ‘play-the-winner-and-drop-the-loser’ strategy different from that used in the prostate cancer trial, are described in Section 5. We close with a discussion in Section 6.

2. THE PROSTATE CANCER TRIAL

About 200 000 men are diagnosed with prostate cancer each year in the United States. Despite the fact that it has an indolent course in comparison with most malignant diseases, about 40 000 deaths are due to prostate cancer in the U.S. annually, a mortality toll second only to lung cancer [2]. These deaths result from the failure of androgen deprivation therapy. Essentially, all patients treated by androgen deprivation will exhibit androgen-independent disease progression after a median time of about one year. There is no standard therapy in this setting, and no systemic therapy has been shown to provide a substantial improvement in survival.

Prostate-specific antigen (PSA) is an extremely useful marker for carcinoma of the prostate [3]. Declines in PSA following treatment are strongly associated with improved survival as well as other indicators of patient benefit, including improved haemoglobin concentration and decreased pain due to bone metastases. Although the use of PSA decline as a surrogate endpoint in clinical trials is still controversial, there is an emerging consensus that in the setting of androgen-independent disease an 80 per cent decline in PSA maintained for at least eight weeks is a legitimate marker of disease-altering ‘response’.

Recently, the use of chemotherapy regimens in which patients are exposed to treatment over several consecutive weeks has resulted in response rates in the range of about 35 per cent to 60 per cent [4]. With this development, there is now considerable interest in the oncology community in conducting a phase III trial to determine whether this recently recognized approach prolongs survival. However, before a phase III trial can be organized, the problem of selecting the most promising of the several contending regimens must be addressed. Thus, four candidate regimens were identified for evaluation in a randomized phase II selection trial aimed at providing a basis for deciding which to study in phase III. These are: 1, paclitaxel+estramustine+etoposide; 2, ketoconazole+doxorubicin alternating weekly with vinblastine+estramustine; 3, cyclophosphamide+vincristine+dexamethasone; 4, paclitaxel+estramustine+carboplatin. The trial, which is currently ongoing at M.D. Anderson Cancer Center (MDACC), is limited to prostate cancer patients who have developed androgen-independent disease. The goals are to select a best treatment, or a best two-treatment strategy, and to estimate cross-resistance between different treatments given in consecutive courses. As with any phase II trial, it is designed to serve as a precursor to a confirmatory phase III trial having survival as the primary outcome. Treatment evaluation is based on reduction of PSA level as a surrogate outcome for survival time.

The specific parameters that we use to characterize treatment efficacy and estimators of these parameters, which will form the basis for the selection design, will be developed in Sections 3 and 4. The trial design was based on the multi-course treatment algorithm described in the next paragraph. To avoid confusion, we will refer to a particular chemotherapy regimen given to a patient in a particular course as a *treatment*, whereas we will refer to the entire sequence of treatments given to a patient as that patient’s *therapy*. We also make a sharp distinction between

treatment success in a particular course and overall patient success, since the latter may require more than one successful course. In the prostate cancer trial, *patient success with treatment t* is defined as two consecutive successful courses with t , while *patient failure* is defined as a total of two unsuccessful courses regardless of treatment.

For the first course of therapy, patients are randomized fairly among the four treatments. Success in the first course is defined as a decrease of 40 per cent or more in PSA from its baseline level at diagnosis, with no evidence of disease progression at any site. For each patient, a treatment that is successful in a given course is given to that patient in the next course, while a treatment that is unsuccessful in a given course is dropped from the set of acceptable treatments for that patient in any subsequent course. A second consecutive success with a given treatment is defined as a decrease of 80 per cent or more in PSA from baseline, again with no evidence of progression. Following an unsuccessful course, the treatment for the next course of the patient's therapy is chosen fairly from that patient's new set of acceptable treatments, which excludes any previous treatment that was unsuccessful for that patient. For example, patients who fail initially with treatment 1 are subsequently randomized among treatments $\{2, 3, 4\}$. The baseline PSA used to evaluate the next treatment is updated to equal the PSA level at the end of the unsuccessful course.

Patients are first stratified into two prognostic strata, low and high volume disease. The latter is defined as involvement of long bones or viscera, or more than three spots on a bone scan. Within each prognostic stratum, patients are assigned to their first treatment in a random series of blocks of size four, each block containing a random permutation of the four treatments. Patients from either stratum who fail with their first treatment in the first or second course are then assigned to a second, different treatment for the next course, according to a random series of blocks of size three, each consisting of a permutation of the three treatments not received initially. Thus, the stratification factor for the first course is the prognostic factor disease volume, whereas patients whose initial treatment is unsuccessful are subsequently stratified adaptively to exclude their initial treatment.

The essential difference between the prostate cancer trial and a more typical randomized phase II trial [5–7] is that we account for multiple treatment courses. Moreover, treatment assignment in courses after the first is adaptive in that it depends on the individual patient's previous history in the trial. Our proposed design thus follows actual clinical practice. This leads to a much richer data structure, which is in accordance with Fisher's principle of designing an experiment to maximize the amount of information per sampling unit [8]. Given the fact that the sampling unit in a clinical trial is a human being, the additional complexity of our method compared to simpler conventional methods seems warranted.

3. PROBABILITY MODELS

In this section we establish a general framework for determining the possible patient outcomes, defining a probability model, and designing a multi-course clinical trial. Our main focus is selection trials, similar to the prostate cancer trial, where the goals may include choosing one best treatment, a set of treatments, or a pair of treatments (u, t) with u given initially and t given if u is unsuccessful. Our approach to choosing a design parameterization in a particular clinical setting is to first study the design's operating characteristics under a range of parameterizations and clinical scenarios. Because the multi-course structure is relatively complex compared to trials based on a

single binary patient outcome, we evaluate the designs by computer simulation. We first develop a probability model for a k -treatment selection trial having the treatment assignment algorithm and definitions of patient success and failure used in the prostate cancer trial. Subsequently, we will describe variants of the model motivated by alternative algorithms appropriate in other clinical settings, including trials that require a total of three unsuccessful courses for the patient's therapy to be declared a failure and, similarly, trials in which patient success is defined as a single successful course.

3.1. A general probability model

A total of $n = km$ patients are randomized fairly among the k treatments for the first course of therapy. For the i th patient, denote the total number of courses by c_i and let $T_{i,j}$ denote the treatment and $Y_{i,j}$ the indicator of success for the j th course of chemotherapy, where $j = 1, \dots, c_i$ and $i = 1, \dots, n$. The treatment a patient receives in any subsequent course depends on his/her history, with $T_{i,j}$ chosen on the basis of $\mathbf{Y}_{i,j-1}$ and $\mathbf{T}_{i,j-1}$ for each $j \geq 2$. The data for the i th patient through j courses thus consist of the treatment vector $\mathbf{T}_{i,j} = (T_{i,1}, \dots, T_{i,j})$ and outcome vector $\mathbf{Y}_{i,j} = (Y_{i,1}, \dots, Y_{i,j})$. In general c_i varies randomly since the criteria for deciding whether to terminate a patient's therapy after j courses, and to then declare the patient's therapy either a success or a failure, are functions of $\mathbf{Y}_{i,j}$ and $\mathbf{T}_{i,j}$.

For simplicity we will denote a particular sequence of treatments and outcomes by writing S_u for success and F_u for failure with treatment u on a given course. For example, suppressing the patient index i , a failure with u in the first course followed by two successes with t in the second and third courses is denoted by $F_u S_t S_t = [Y_1 = 0, Y_2 = Y_3 = 1, T_1 = u, T_2 = T_3 = t]$. Thus, a patient treated initially with u and subsequently with t if there is a failure with u , may achieve overall success in three different ways, $S_u S_u, F_u S_t S_t$ or $S_u F_u S_t S_t$. Similarly, this patient may have an overall failure in four possible ways, $F_u F_t, F_u S_t F_t, S_u F_u F_t$ or $S_u F_u S_t F_t$. The best possible outcome for the patient, denoted $S_u S_u$, may occur in $k = 4$ possible ways, one for each treatment u , while each of the other outcomes can occur in $k(k - 1) = 12$ ways, the number of ordered pairs (u, t) . Thus, there are $k + 6k(k - 1) = 76$ possible elementary outcomes. Each patient has a total of 2, 3 or 4 courses. In general, the outcome sequences that are possible in a given trial depend upon the treatment assignment algorithm and the definitions of patient success and patient failure employed by the physicians conducting the trial. Because the statistical considerations must follow these medical definitions and the particular medical structure of a multi-course therapy will differ from trial to trial, this structure plays an essential role in constructing a probability model.

Let $\mathcal{A}_{i,j}$ denote the set of acceptable treatments for patient i at course j . Since the entire set of treatments is acceptable for all patients in the first course, $\mathcal{A}_{i,1} = \{1, \dots, k\}$ for all i , and the probability of assignment to any treatment t is simply $\tau_{i,1}(t) = k^{-1}$ for all t and i . If $Y_{i,j} = 0$ with treatment $T_{i,j}$ at course j then $T_{i,j}$ is removed from all subsequent sets of acceptable treatments for patient i , formally $\mathcal{A}_{i,r} = \mathcal{A}_{i,1} - \bigcup_{j=1}^{r-1} \{T_{i,j} : Y_{i,j} = 0\}$. Thus, for $j \geq 2$, randomizing fairly among the acceptable treatments at each course, the probability that patient i receives treatment t at course j given his/her history through $j - 1$ courses is

$$\tau_{i,j}(t) = \Pr[T_{i,j} = t \mid \mathbf{Y}_{i,j-1}, \mathbf{T}_{i,j-1}] = |\mathcal{A}_{i,j}|^{-1}, \quad t \in \mathcal{A}_{i,j} \tag{1}$$

with $\tau_{i,j}(t) = 0$ if $t \notin \mathcal{A}_{i,j}$, where $|\mathcal{A}|$ denotes the cardinality of the set \mathcal{A} . Temporarily suppress the patient index i and denote the conditional probability of success in the j th course given the j th treatment and the history through $j - 1$ courses in general by $\theta_j = \theta_j(\mathbf{Y}_{j-1}; \mathbf{T}_j) = \Pr[Y_j = 1 \mid \mathbf{Y}_{j-1}; \mathbf{T}_j]$.

Denoting $\theta_{i,j} = \theta_j(\mathbf{Y}_{i,j-1}; \mathbf{T}_{i,j})$ for the i th patient, a simple conditioning argument shows that the probability of a particular sequence of outcomes, $\mathbf{y}_{i,r}$, and treatments, $\mathbf{t}_{i,r}$, through r courses of treatment, $r \leq c_i$, may be expressed as

$$\mathcal{L}_{i,r} = \Pr[\mathbf{Y}_{i,r} = \mathbf{y}_{i,r}, \mathbf{T}_{i,r} = \mathbf{t}_{i,r}] = \prod_{j=1}^r \theta_{i,j}^{y_{i,j}} (1 - \theta_{i,j})^{1-y_{i,j}} \tau_{i,j}(t_{i,j}) \quad (2)$$

Given a parameterization of $\theta_{i,j}$, the full likelihood is the usual product $\mathcal{L} = \prod_{i=1}^n \mathcal{L}_{i,c_i}$.

3.2. Multinomial models

The probabilities of compound events of interest may be computed by referring to the appropriate multinomial distribution based on the possible elementary outcomes. Patient success with treatment t in the prostate cancer trial is the event $S_t^* = [S_t S_t] \cup [\bigcup_{u \neq t} F_u S_t S_t] \cup [\bigcup_{u \neq t} S_u F_u S_t S_t]$, corresponding to patient success with t in two, three or four courses of chemotherapy. The probability of this event is

$$\pi_t = \Pr(S_t^*) = \Pr[S_t S_t] + \sum_{u, u \neq t} \{\Pr[F_u S_t S_t] + \Pr[S_u F_u S_t S_t]\} \quad (3)$$

We will use estimates of the π_t 's as the basis for selecting a single best treatment. An important property of π_t is that it is the sum of the probabilities of many different elementary patient outcomes that together involve all of the treatments in the trial. Thus, the value of π_1 in a trial of treatments $\{1, 2, 3, 4\}$ may be quite different from π_1 in a trial of treatments $\{1, 5, 6, 7\}$. This is very different from the common statistical formulation wherein each patient's outcome is characterized as a single binary success/failure variable corresponding to a single treatment and the data consist of k independent binomial samples each of size m . As noted in the Introduction, this commonly used formulation is typically based on simplifying assumptions that essentially ignore or collapse most of the multi-course structure. In that case, each $\pi_t = \Pr[S_t]$ as defined depends on t alone, regardless of the other treatments in the trial. While this distinction may appear to be a drawback of the multi-course setting, we will show that the information in (\mathbf{Y}, \mathbf{T}) provides a basis for treatment selection not available from the naive approach of reducing patient outcome to a single binary variable. In particular, accounting for the multi-course structure provides a basis for constructing desirable treatment combinations. The power of this approach, as illustrated by our simulation results, is a probabilistic validation of the intuitive process by which physicians actually practice.

An alternative goal is to select an ordered pair of treatments (u, t) . This denotes the treatment strategy in which u is given initially and, if u is unsuccessful, the patient is then given t . The probability of patient success with (u, t) is $\zeta(u, t) = \zeta_u + (1 - \zeta_u) \zeta_{t|u}$, where ζ_u denotes the probability of a patient success in the first two courses given that the patient was treated initially with u , and $\zeta_{t|u}$ denotes the *salvage probability* that the patient has two successful courses with t given that the strategy (u, t) was used and the initial treatment with u was unsuccessful in either course 1 or 2. In the prostate cancer trial, $\zeta_u = \Pr[S_u S_u | (u, t)]$ and $\zeta_{t|u} = \Pr[F_u S_t S_t \text{ or } S_u F_u S_t S_t | (u, t) \text{ and } F_u \text{ or } S_u F_u]$.

MLEs of these probabilities may be obtained as follows. Denote the number of patients treated initially with u by n_u and let $X_u = |S_u S_u|$ be the number in this group who succeed in their first two courses. Similarly, denote the number of patients treated initially with u who fail either the first or second course and are then treated with t by $n_{t|u}$, and let $X_{t|u} = |F_u S_t S_t \cup S_u F_u S_t S_t|$ be the number of

patient successes with t in this subgroup. Since the patients who have an initial failure with u are randomized among the remaining treatments, $\sum_{t,t \neq u} n_{t|u} = n_u - X_u$. This yields the MLEs $\hat{\xi}_u = X_u/n_u$ and $\hat{\xi}_{t|u} = X_{t|u}/n_{t|u}$, and thus $\hat{\xi}(u, t)$. These estimators may be used as the basis for selecting a best two-treatment strategy among the $k(k - 1)$ ordered pairs $\{(u, t), 1 \leq u, t \leq k, u \neq t\}$. Recall that the π_t 's are defined unconditionally, in that they include the randomization probabilities in the trial. Specifically, (3) may be expressed equivalently as

$$\pi_t = \frac{1}{4}\xi_t + \frac{1}{12} \sum_{u, u \neq t} (1 - \xi_u)\xi_{t|u} \tag{4}$$

and thus the above MLEs also provide estimates of the π_t 's.

The clinical utility of treating patients in a randomized trial of this sort is a very important issue. Given the choice between either randomizing a patient into the trial or treating the patient with the single best regimen, some physicians might choose the latter option. In general this may be a mistake, since the patient would miss the chance of being salvaged if the initial treatment were not successful. For example, consider a trial of three treatments $\{1, 2, 3\}$ having initial two-course success probabilities $\xi_1 = 0.50, \xi_2 = 0.40, \xi_3 = 0.30$ and salvage probabilities $\xi_{2|1} = \xi_{2|3} = \xi_{1|2} = 0.20, \xi_{3|1} = \xi_{3|2} = \xi_{1|3} = 0.40$. The overall patient success probability is $\sum_{(u,t)} [\xi_u + (1 - \xi_u)\xi_{t|u}]/6 = 0.58$, which is larger than the probability 0.50 of initial success with treatment 1 alone.

3.3. Regressive logistic probability models

In addition to relying on the above multinomial models, it will also be useful to use parametric probability models defined in terms of the conditional probability of success at each course given the patient's history and current treatment, since this reflects the way a physician regards the patient's prognosis at each course of therapy. In this section, we define regressive logistic probability models [9, 10] for the possible elementary outcomes described earlier. A regressive logistic model for a vector of dependent binary variables (Y_1, \dots, Y_k) decomposes their joint probability into $\prod_{j=1}^k \Pr[Y_j | Y_1 \dots Y_{j-1}]$, with each conditional probability expressed as a logistic regression function of the conditioning variables and possibly additional covariates. The likelihood (2) is a slightly extended version of this decomposition in that it includes the treatment selection probabilities as well as the patient outcome probabilities, although only the latter will involve treatment effect parameters.

We develop regressive logistic models in terms of the following three probabilities characterizing patient outcome through the first two courses of therapy:

$$\theta_1(t) = \Pr[Y_1 = 1 | T_1 = t] \tag{5}$$

$$\theta_2(1; (t, t)) = \Pr[Y_2 = 1 | Y_1 = 1, T_1 = T_2 = t] \tag{6}$$

$$\theta_2(0; (u, t)) = \Pr[Y_2 = 1 | Y_1 = 0, T_1 = u, T_2 = t] \tag{7}$$

Expression (5) is the probability of success in the first course, while (6) and (7) are the conditional probabilities of success in the second course given either success or failure in the first course. Thus, for example, $\xi_t = \theta_1(t)\theta_2(1; (t, t))$.

We summarize the patient's failure history in terms of the following constructed variable. Let $W_{i,j} = \sum_{r=1}^j (1 - Y_{i,r})/(j + \frac{1}{2})$ denote the mean number of failures for patient i through j courses.

We define the patient's sequence of failure history variables $\{Z_{i,1}, \dots, Z_{i,c_j}\}$ recursively as

$$\begin{aligned} Z_{i,j} &= W_{i,j} & \text{if } Y_{i,j} &= 0 \\ &= W_{i,j-1} & \text{if } Y_{i,j} &= 1 \end{aligned} \quad (8)$$

with $W_{i,0} = 0$. Thus, $Z_{i,j}$ is the mean number of failures through the most recent course up to the j th that was a failure. For the prostate cancer trial, this must be either the j th or $(j-1)$ th course, since patient success is defined as two consecutive successful courses. That is, if $Y_{i,j} = 0$ then $Z_{i,j} = W_{i,j}$, while $Y_{i,j} = 1$ implies that $Y_{i,j-1} = 0$ and hence $Z_{i,j} = W_{i,j-1}$; otherwise, the patient's therapy would have been completed successfully in the j th course. We use the divisor $j + \frac{1}{2}$ rather than j in the definition of $W_{i,j}$ as a device to obtain a slightly more refined numerical value of $Z_{i,j}$. In particular, for trials with a 'three-and-out' definition of patient failure this yields $Z_{i,j} = \frac{2}{3}$ given the history F_1S_2 and $Z_{i,j} = \frac{4}{5}$ given the history $F_1F_2S_2$, whereas these would both equal 1 if the divisor were j rather than $j + \frac{1}{2}$.

We define a regressive logistic probability model in terms of the linear component

$$\eta_{i,j} = \text{logit}(\theta_{i,j}) = \mu_{t_{i,j}} + \alpha_{t_{i,j}} Y_{i,j-1} + \beta_{t_{i,j}} Z_{i,j-1} \quad (9)$$

characterizing the conditional probability of success on the j th course of patient i , where $\text{logit}(\theta) = \log[\theta/(1-\theta)]$. For a given $t = t_{i,j}$, the parameter $\mu_t = \text{logit}[\theta_1(t)]$ characterizes the probability of success on the initial course with t . Recall that a second consecutive success with t given a success with t on the previous course typically is harder to achieve than an initial success with t because the clinical definition for a second successful course is usually more demanding. Unfortunately, the most common circumstance in oncology is that the second success is less likely than the first, as in the treatment of prostate cancer. Formally, $\theta_2(1; (u, t)) < \theta_1(t)$, which holds if and only if $\alpha_t < 0$. From a clinical perspective, the larger the value of α_t in the negative direction, the more difficult it is to achieve a second success with t following an initial success with t . The closer α_t is to 0, the higher the quality of the initial success with t . There are some settings where $\theta_2(1; (u, t)) \geq \theta_1(t)$, however, such as treatment of hairy cell leukaemia with 2-chloro-deoxy adenosine, where about 90 per cent of patients who respond initially are cured. The parameter β_t accounts for the cost of the patient's previous failures, as summarized in terms of $Z_{i,j-1}$. The regressive logistic model is obtained by combining (2) and (9). For example, denoting $g(\eta) = \text{logit}^{-1}(\eta) = e^\eta/(1+e^\eta)$

$$\Pr[F_u S_t S_t] = \frac{1}{4} [1 - g(\mu_u)] \times \frac{1}{3} g\left(\mu_t + \frac{2}{3}\beta_t\right) \times 1g\left(\mu_t + \alpha_t + \frac{2}{3}\beta_t\right) \quad (10)$$

An important point is that this model only makes sense medically if $\beta_t < 0$ for all t , since previous failures reduce the probability of success on any given course. This will affect the design of the simulations, described below, for evaluating the selection methods using this model. The treatment selection probabilities in (10) are $\tau_1(s) = \frac{1}{4}$ for each $s = 1, 2, 3, 4$ in the first course, then $\tau_2(u) = 0$ and $\tau_2(s) = \frac{1}{3}$ for each $s \neq u$ since there is a failure with u in the first course, and $\tau_3(t) = 1$ since the success with t in the second course implies that t is certain to be given in the third course. The one-to-one correspondence between $\{\mu_t, \alpha_t, \beta_t\}$ and the three outcome probabilities $\{\theta_1(t), \theta_2(1; (t, t)), \theta_2(0; (u, t))\}$ for each treatment t implies that this $3k$ -parameter model characterizes all outcome probabilities in terms of the probabilities of the outcomes in the first two courses. As with all regressive logistic models, a practical advantage enjoyed by this parametric model is that it may be fit using any statistical software that fits standard logistic regression models, since the patient's history enters the model as variables in the linear component.

The definition of $\theta_j(\mathbf{Y}_{j-1}; \mathbf{T}_j)$ given by expression (9) ignores the particular treatments given in the first $j - 1$ courses and thus does not account for cross-resistance between each distinct treatment pair. A more refined version of the regressive logistic model that characterizes cross-resistance more specifically than (9) is given as follows. Extend the definition of $Z_{i,j}$ so that, for each treatment u , the history variable after the first course is

$$\begin{aligned} Z_{i,j}(u) &= W_{i,j} \quad \text{if } Y_{i,j} = 0 \text{ and } T_{i,j} = u \\ &= W_{i,j-1} \quad \text{otherwise} \end{aligned} \tag{11}$$

again with $W_{i,0} = 0$. The refined model is given by

$$\text{logit}(\theta_{i,j}) = \mu_{t_{i,j}} + \alpha_{t_{i,j}} Y_{i,j-1} + \sum_{u \neq t_{i,j}} \beta_{u,t_{i,j}} Z_{i,j-1}(u) \tag{12}$$

Under this model, $\beta_{u,t}$ characterizes the cross-resistance between the current treatment t and the treatment u given in the most recent course that was a failure. As with the simpler model (9), we require all $\beta_{u,t} < 0$. The probability of success in the second course with t following a failure with u in the first course is $\theta_2(0; (u, t)) = \text{logit}^{-1}(\mu_t + \frac{2}{3}\beta_{u,t})$ for each $u \neq t$. This model has $2k + k(k - 1)$ parameters, which is 20 for the prostate cancer trial. Thus, the multinomial model and second regressive logistic model (12) have the same number of parameters. Because they describe the phenomenon in different ways, however, they often yield different fits and thus different selection probabilities for a particular data set, as will be shown in the simulation study below. We will refer to (9) and (12) as the first and second regressive logistic models, RLM1 and RLM2.

4. SIMULATION STUDY

4.1. Selecting one best treatment

In this section we summarize results from a simulation study of the prostate cancer trial. The purpose of the simulation was to study the design's behaviour under a range of clinical scenarios and use this as a basis for calibrating design parameters, including sample size. Our aim here is to illustrate both the prostate cancer trial application and how the proposed methodologies may be applied in similar clinical settings involving multiple treatment courses. Let the subscript $[j]$ denote the j th smallest of a set of values indexed by $\{1, \dots, k\}$, so that in particular $[k]$ indexes the maximum of the k values. We selected a single treatment t as best if it had the largest estimated patient success probability $\hat{\pi}_t = \hat{\pi}_{[k]}$. We selected (u, t) as the best two-treatment strategy if $\hat{\xi}(u, t)$ was largest among all $k(k - 1)$ estimates. Ties were broken by fair randomization. Using maximum likelihood throughout, we computed parameter estimates under one or both of the regressive logistic models and under the appropriate multinomial model (MM). We determined the sample size empirically to ensure that, under clinical Scenario 1 described below in Section 4.2, the probability of correctly selecting the treatment t having true $\pi_t = \pi_{[4]}$ as best was at least a specified value PCS*. Each case was simulated 4000 times, with a run time of about two hours per case. All computations were done on a DEC AlphaServer 4100 5/400 running Digital UNIX 4.0D in S-plus using a customized program for the simulations and MM fits and the S-plus program glim to fit the regressive logistic models. The program is available from the third author on request.

Table I. The first two model-based scenarios for the simulation.

	Probability	Treatment (t)			
		1	2	3	4
Scenario 1	$\theta_1(t)$	0.40	0.40	0.40	0.55
	$\theta_2(1; (t, t))$	0.37	0.37	0.37	0.52
	$\theta_2(0; (u, t))$	0.15	0.15	0.15	0.25
	π_t	0.043	0.043	0.043	0.087
Scenario 2	$\theta_1(t)$	0.40	0.55	0.55	0.55
	$\theta_2(1; (t, t))$	0.37	0.37	0.52	0.52
	$\theta_2(0; (u, t))$	0.15	0.42	0.15	0.42
	π_t	0.043	0.074	0.079	0.107

4.2. Model-based scenarios

We evaluated the design under a broad range of possible clinical scenarios. For purposes of illustration we summarize four of these. The first two are defined in terms of the regressive logistic model (7), in Table I. Under scenario 1, treatments $t=1, 2$ and 3 have the identical parameterization determined by $\theta_1(t)=0.40$, $\theta_2(1; (t, t))=0.37$, and $\theta_2(0; (u, t))=0.15$ for each $u \neq t$, corresponding to a typical prostate cancer treatment. In particular, given that one of $t=1, 2$ or 3 is chosen as the first treatment, the probability of success in the first two cycles with t is $\theta_1(t)\theta_2(1; (t, t))=0.40 \times 0.37=0.148$, while the unconditional probability is $\Pr[S_t S_t] = \frac{1}{4}\theta_1(t)\theta_2(1; (t, t))=0.037$. The probability of patient success with t in three courses is

$$\Pr[F_u S_t S_t] = \frac{1}{12} \sum_{u, u \neq t} \{[1 - \theta_1(u)]\theta_2(0; (u, t))\theta_3((0, 1); (u, t, t))\} = 0.0031$$

and the probability of patient success with t in four courses is

$$\begin{aligned} \Pr[S_u F_u S_t S_t] &= \frac{1}{12} \sum_{u, u \neq t} \{\theta_1(u)[1 - \theta_2(1; (u, u))]\theta_3((1, 0); (u, u, t))\theta_4((1, 0, 1); (u, u, t, t))\} \\ &= 0.0028 \end{aligned}$$

so that the overall probability of patient success with each of treatments $t=1, 2$ or 3 is $\pi_t = 0.037+0.0031+0.0028=0.0429$. Treatment 4 differs only in that μ_4 is larger than the common value of $\mu_1 = \mu_2 = \mu_3$. This is reflected in the larger values of all three success probabilities for the first two courses given in Table I, as well as the larger overall patient success probability $\pi_4 = 0.087$. Scenario 2 is qualitatively different from scenario 1 in that each treatment provides an overall improvement over the last in terms of π_t , with the largest jumps from $\pi_1 = 0.043$ to $\pi_2 = 0.074$ and $\pi_3 = 0.079$ to $\pi_4 = 0.107$. Thus, selecting the best treatment is statistically easier under scenario 1, since $\pi_{[4]} - \pi_{[3]} = 0.044$ while the corresponding difference under scenario 2 is 0.028.

An important probability from the patient's point of view is the overall probability of success $\sum_{t=1}^4 \pi_t$, regardless of the treatment with which it occurs. This is 0.216 under scenario 1, but is 0.30 under scenario 2. Thus, scenario 2 would be more desirable from the viewpoint of a patient entering the trial.

4.3. Determining sample size

We evaluated the estimators of the π_t 's and single treatment selection probabilities under the RLM1, the MM, and the naive approach of only recording the patient's response to the initial treatment. Specifically, the naive approach scores a patient given t in the first course as a success if $Y_1 = Y_2 = 1$ and a failure otherwise, and hence ignores any outcomes of subsequent courses with treatments other than t if either the first or second course is unsuccessful. We considered overall sample sizes $n = 92, 124, 156$, which were determined to ensure that the probabilities of correctly selecting treatment 4 as best under scenario 1 are, respectively, $\geq 0.80, 0.85$, and 0.90 . It is worthwhile to compare these sample sizes to what might be obtained using a method that relies on the approximate normality of the $\hat{\pi}_t$'s, bearing in mind that these arise from a multinomial setting and hence are negatively correlated. The procedure \mathcal{N}_B , proposed by Bechhofer [11] and described by Bechhofer *et al.* (Reference [12], Section 2.2.1), is based on independent samples from k normal populations with means μ_1, \dots, μ_k and common variance σ^2 . The treatment having the largest sample mean is selected as best, and the per-treatment sample size m ensuring that the treatment with largest true mean $\mu_{[k]}$ will be selected with probability at least PCS^* if $\mu_{[k]} - \mu_{[k-1]} = \delta^*$ is given by $m = 2[\sigma Z_{k-1, 1/2}^{(1-\text{PCS}^*)} / \delta^*]^2$, where $Z_{k-1, 1/2}^{(1-\text{PCS}^*)}$ is the upper $1 - \text{PCS}^*$ cut-off of the maximum of k independent standard normals with common correlation $1/2$. Some care must be taken in applying this formula here since, as shown in expression (4), we have defined patient success with t so that it includes the treatment selection probabilities. If the π_t 's were used in place of the μ_t 's to compute sample size, this would be analogous to using $\delta^* = (\pi_2 - \pi_1)/2$ rather than $\pi_2 - \pi_1$ in the usual two-sample setting, which would incorrectly increase the computed sample size. We thus base the computation on $4\pi_t$, which equals ξ_t plus the equally weighted average of $\xi_{t|u}(1 - \xi_u)$ over u such that $u \neq t$. Applying the above method in this way, and using the approximation $\sigma = [\sum_{t=1}^4 4\pi_t(1 - 4\pi_t)/4]^{1/2} = 0.4043$, since $Z_{3, 1/2}^{0.10} = 1.734$ and $4\pi_4 - 4\pi_3 = 4(0.0873 - 0.0429) = 0.0444$, this yields a per-treatment sample size $m = 32$, for a total trial size of 128 patients. The large difference between this value and our empirically determined total 156 may be explained by the heteroscedasticity, negative correlation and deviation from exact normality of the $\hat{\pi}_t$'s, since each of these conditions violates the assumptions underlying the above computation. Thus, we strongly suggest that trials using our proposed methods be sized empirically.

4.4. Simulation results under the model-based scenarios

The simulation results under scenarios 1 and 2 are summarized in Tables II and III. Since the $\hat{\pi}_t$'s obtained under RLM1, RLM2 and the MM and the actual π_t 's were identical to three decimal places for all scenarios, these appear in Table II as a single value labelled 'multi-course'. The smaller numerical values of the estimates under the naive approach are due to the fact that it is estimating $\Pr[S_t S_t] = \xi_t/4$ rather than π_t . The fact that the RLM1-based and MM-based estimates of the π_t 's are the same is reflected in the single treatment selection probabilities given in Table III, which are also identical aside from sample variation. The much smaller probabilities of selecting treatment 4 under the naive approach show what is lost by using this method for treatment selection rather than taking full advantage of the multi-course data. This difference is most pronounced under scenario 2, where the superiority of treatment 4 is due to its higher salvage probability $\theta_2(0; (u, 4))$, which is completely ignored by the naive approach. Since the naive method wastes data and performs poorly even for the simplest goal of selecting one best treatment, we do not consider it further.

Table II. Estimates of patient success probabilities for trials of 92, 124 or 156 patients. Each entry is the mean from 4000 simulated trials. Standard errors, given in parentheses, are based on the 92-patient trial.

Treatment		Scenario	
		1	2
1	Multi-course	0.044 (0.021)	0.044 (0.020)
	Naive	0.038 (0.019)	0.037 (0.018)
2	Multi-course	0.044 (0.020)	0.074 (0.027)
	Naive	0.037 (0.019)	0.051 (0.022)
3	Multi-course	0.044 (0.020)	0.080 (0.025)
	Naive	0.037 (0.018)	0.071 (0.023)
4	Multi-course	0.088 (0.028)	0.107 (0.030)
	Naive	0.072 (0.024)	0.071 (0.023)
Total	Multi-course	0.219	0.305
	Naive	0.184	0.231

Table III. Probabilities of selection as best in terms of $\hat{\pi}_i$. Decisions were based on the first regressive logistic model (RLM1) or multinomial model (MM) for a multi-course trial, or on the naive method using only the first treatment (naive). Probabilities of correctly selecting treatment 4 are shown in bold type. Each entry is the mean from 4000 simulated trials.

n	Treatment	Scenario 1			Scenario 2		
		RLM1	MM	Naive	RLM1	MM	Naive
92	1	0.067	0.069	0.087	0.011	0.017	0.033
	2	0.065	0.066	0.084	0.163	0.148	0.119
	3	0.070	0.071	0.090	0.181	0.244	0.422
	4	0.800	0.794	0.739	0.646	0.591	0.425
124	1	0.048	0.055	0.068	0.004	0.007	0.020
	2	0.048	0.059	0.068	0.125	0.115	0.097
	3	0.050	0.057	0.066	0.160	0.231	0.437
	4	0.855	0.828	0.798	0.712	0.647	0.446
156	1	0.034	0.040	0.049	0.002	0.004	0.013
	2	0.035	0.041	0.051	0.108	0.098	0.080
	3	0.030	0.037	0.048	0.162	0.228	0.465
	4	0.900	0.882	0.851	0.729	0.670	0.441

4.5. Selecting a best treatment pair

We now discuss the somewhat different goal of selecting the best two-treatment strategy (u, t) having the largest patient success probability $\xi(u, t)$, as described in Section 3.2. The four treatments in the prostate cancer trial yield 12 such strategies. We examined the relative merits of using MLEs of the $\xi(u, t)$'s for achieving this goal based on the MM, RLM1 and RLM2. We simulated these three methods under scenario 3, which is based on RLM2, and under scenario 4, which is a non-model-based scenario that specifies the conditional probability of success in each course for each possible history. These two scenarios are summarized in Table IV. Scenario 3 is obtained from scenario 1 by changing the values of $\xi(4, t)$ from the common value 0.3086 to $\xi(4, 1) = 0.29$,

Table IV. Scenarios with varying cross-resistance probabilities. Scenario 3 is based on regressive logistic model 2. Scenario 4 is non-model-based.

		Treatment			
		$t = 1$	$t = 2$	$t = 3$	$t = 4$
Scenario 3	μ_t	-0.4055	-0.4055	-0.4055	0.2067
	α_t	-0.1268	-0.1268	-0.1268	-0.1268
	$\beta_{1,t}$	-	-1.9937	-1.9937	-1.9937
	$\beta_{2,t}$	-1.9937	-	-1.9937	-1.9937
	$\beta_{3,t}$	-1.9937	-1.9937	-	-1.9937
	$\beta_{4,t}$	-4.300	-0.9120	-0.0320	-
	ζ_t	0.148	0.148	0.148	0.287
	$\zeta(1,t)$	-	0.1723	0.1723	0.2103
	$\zeta(2,t)$	0.1723	-	0.1723	0.2103
	$\zeta(3,t)$	0.1723	0.1723	-	0.2103
	$\zeta(4,t)$	0.2900	0.3400	0.3900	-
Scenario 4	ζ_t	0.402	0.560	0.375	0.315
	$\zeta(1,t)$	-	0.5085	0.4379	0.4459
	$\zeta(2,t)$	0.5736	-	0.5755	0.5767
	$\zeta(3,t)$	0.4187	0.4907	-	0.4238
	$\zeta(4,t)$	0.3625	0.4406	0.3583	-

$\zeta(4,2)=0.34$, and $\zeta(4,3)=0.39$, so that the strategy (4,3) is optimal, (4,2) is second best, (4,1) is third best, and the remaining strategies have patient success probability 0.1723 or 0.2103. We chose these numerical values of the $\zeta(4,t)$'s because each is bounded below by $\zeta_4 = 0.2868$, where $\beta_{4,t} = -\infty$, and above by 0.3923, where $\beta_{4,t} = 0$, with the upper bound being due to the fact that the regressive logistic model only makes sense medically if $\beta_{u,t} < 0$ for all (u,t) . Under scenario 4, for practical purposes any of the strategies (2,1), (2,3), and (2,4) are optimal since each has patient success probability between 0.57 and 0.58, while the remaining nine strategies are clearly inferior. Scenario 4 is useful for evaluating the robustness of the selection procedures based on RLM1 or RLM2.

The treatment pair selection probabilities under scenario 3 are summarized in Table V. These probabilities should be evaluated relative to the probability $1/12 = 0.083$ of correct selection by simply guessing in the absence of any data. The RLM2-based estimators give the largest PCS for the best strategy (4,3), with the PCS dropping about 0.11 to 0.13 with the RLM1-based estimators and an additional 0.02 to 0.04 under the MM-based estimators. All three methods improve with increasing sample size. To address the question of how the three methods perform when the two-treatment strategies differ only in terms of their cross-resistance, we altered scenario 3 by setting $\mu_4 = -0.4055$ so that the treatments have identical μ_t 's and α_t 's and differ only in the values of $\beta_{4,1}, \beta_{4,2}$ and $\beta_{4,3}$, as given in Table IV. This yields patient success probabilities $\zeta(4,1) = 0.1513$, $\zeta(4,2) = 0.2098$ and $\zeta(4,3) = 0.2712$. The differences between these values are about the same as in scenario 3, but now they are due entirely to the varying cross-resistance parameters. In this more difficult case, for $n = 156$ the probabilities of correctly selecting strategy (4,3) as best for (MM, RLM1, RLM2) are (0.280, 0.173, 0.468). As might be expected, there is a drop in PCS from the corresponding values (0.449, 0.486, 0.607) under scenario 3, although the RLM2-based method still performs best. It thus appears that, when the $\beta_{u,t}$'s vary with both u and t , the regressive logistic model that accounts for this explicitly does the best job of detecting the best two-treatment

Table V. Selection probabilities for the 12 two-treatment strategies under scenario 3. Each ordered triple consists of the selection probabilities using the (MM, RLM1, RLM2)-based estimators. Probabilities of selecting the best treatment pair (4,3) are shown in bold type.

	First treatment	Second treatment			
		1	2	3	4
<i>n</i> = 92					
	1	–	(0.032, 0.020, 0.022)	(0.024, 0.029, 0.013)	(0.057, 0.064, 0.062)
	2	(0.024, 0.007, 0.017)	–	(0.026, 0.031, 0.014)	(0.057, 0.061, 0.056)
	3	(0.027, 0.006, 0.021)	(0.024, 0.024, 0.017)	–	(0.059, 0.062, 0.060)
	4	(0.058, 0.094, 0.018)	(0.221, 0.238, 0.228)	(0.347, 0.365, 0.472)	–
<i>n</i> = 124					
	1	–	(0.019, 0.010, 0.008)	(0.016, 0.025, 0.009)	(0.050, 0.049, 0.048)
	2	(0.016, 0.003, 0.010)	–	(0.019, 0.025, 0.010)	(0.052, 0.058, 0.052)
	3	(0.019, 0.002, 0.009)	(0.017, 0.014, 0.008)	–	(0.055, 0.057, 0.049)
	4	(0.058, 0.083, 0.010)	(0.239, 0.248, 0.230)	(0.408, 0.427, 0.558)	–
<i>n</i> = 156					
	1	–	(0.016, 0.009, 0.007)	(0.017, 0.017, 0.008)	(0.038, 0.042, 0.034)
	2	(0.012, 0.002, 0.005)	–	(0.013, 0.019, 0.006)	(0.045, 0.044, 0.034)
	3	(0.015, 0.002, 0.007)	(0.011, 0.009, 0.006)	–	(0.042, 0.041, 0.031)
	4	(0.042, 0.071, 0.005)	(0.276, 0.258, 0.248)	(0.449, 0.486, 0.607)	–

strategy. If the difference between the true values of the best and second best strategies is larger, specifically if scenario 3 is modified so that $\zeta_{4,1} = \zeta_{4,2} = 0.29$ while $\zeta_{4,3} = 0.39$, then the PCS values for (MM, RLM1, RLM2) with $n = 156$ increase to (0.476, 0.582, 0.789). Thus, each method behaves consistently in that it is more likely to detect a best strategy that has a larger advantage over the others.

To assess each method's robustness, we also evaluated them under scenarios not arising from the parametric models that are the basis for RLM1 and RLM2. Scenario 4, summarized in Table IV, was constructed based on clinical judgement to obtain a reasonable set of success probabilities given each possible patient history. As was the case under the model-based scenarios, under scenario 4 all three methods give $\hat{\pi}_t$'s that agree with the actual values $(\pi_1, \pi_2, \pi_3, \pi_4) = (0.11, 0.17, 0.10, 0.09)$ to two decimal places. For the goal of selecting a single best treatment in terms of the $\hat{\pi}_t$'s, for $n = 156$ the probability of correctly selecting treatment 2 under scenario 4 is 0.95 for all three methods. This higher correct selection probability, compared to the corresponding values under scenarios 1 and 2 given in Table III, may be attributed to the larger true difference $\pi_{[4]} - \pi_{[3]} = 0.06$ under scenario 4. Table VI shows that the probabilities of selecting one of the three best strategies (2, 1), (2, 3) or (2, 4) with $n = 156$ under scenario 4 for (MM, RLM1, RLM2) are (0.477, 0.766, 0.579). Thus, when there is no clear cross-resistance advantage for a specific pair of treatments, the simpler regressive logistic model given by equation (9) that has a single β_t for each t appears to perform best. This also indicates that the RLM1-based procedure is reasonably robust, since scenario 4 is not based on either model (9) or (12).

To take advantage of these results in practice, one must decide between the two regressive logistic models. This is straightforward, since model (9) is nested within model (12); hence, one may first fit both models to the data, decide between them on the basis of a likelihood ratio (LR) test, and then use the model giving a better fit as the basis for treatment selection. Since the regressive logistic models may be fit using standard statistical software, this is straightforward.

Table VI. Selection probabilities for the two-treatment strategies under scenario 4. Each ordered triple consists of the selection probabilities using the (MM, RLM1, RLM2)-based estimators. Probabilities of correctly selecting an optimum strategy (2,1), (2,3) or (2,4) are shown in bold type.

	First treatment	Second treatment			
		1	2	3	4
<i>n</i> = 92					
	1	–	(0.160, 0.131, 0.147)	(0.038, 0.012, 0.037)	(0.043, 0.010, 0.040)
	2	(0.119, 0.286, 0.156)	–	(0.117, 0.190, 0.142)	(0.131, 0.201, 0.164)
	3	(0.037, 0.010, 0.043)	(0.137, 0.106, 0.122)	–	(0.036, 0.007, 0.035)
	4	(0.017, 0.004, 0.023)	(0.086, 0.042, 0.077)	(0.013, 0.002, 0.014)	–
<i>n</i> = 124					
	1	–	(0.146, 0.126, 0.140)	(0.030, 0.006, 0.034)	(0.042, 0.006, 0.034)
	2	(0.132, 0.308, 0.178)	–	(0.146, 0.211, 0.175)	(0.148, 0.203, 0.173)
	3	(0.030, 0.007, 0.030)	(0.137, 0.096, 0.116)	–	(0.032, 0.005, 0.032)
	4	(0.010, 0.002, 0.015)	(0.072, 0.028, 0.062)	(0.008, 0.002, 0.009)	–
<i>n</i> = 156					
	1	–	(0.155, 0.116, 0.140)	(0.025, 0.001, 0.022)	(0.035, 0.004, 0.024)
	2	(0.149, 0.344, 0.199)	–	(0.162, 0.210, 0.178)	(0.166, 0.212, 0.202)
	3	(0.020, 0.005, 0.029)	(0.118, 0.086, 0.112)	–	(0.025, 0.002, 0.024)
	4	(0.008, 0.001, 0.010)	(0.067, 0.018, 0.057)	(0.006, 0.000, 0.004)	–

For example, a single data set simulated under scenario 3 gives a log LR statistic, equivalently a difference in residual deviances, of $461.30 - 443.67 = 17.63$ on $20 - 12 = 8$ d.f., $p = 0.024$, indicating that the more complex model (12) is more appropriate. Similarly, a data set simulated under scenario 4 gives $\log LR = 469.02 - 460.59 = 8.43$ on 8 d.f., $p = 0.39$, indicating that the simpler model (9) is more appropriate.

5. GENERALIZATIONS

To illustrate how this approach may be applied more generally, we first briefly consider two other versions of the play-the-winner-and-drop-the-loser algorithm. In a bladder cancer trial currently being planned, patient failure is defined as three rather than two unsuccessful courses, with patient success still defined as two consecutive successful courses. Thus, a patient who first fails with treatment u , then succeeds and fails with w , then has two successful courses with t , has elementary outcome $F_u S_w F_w S_t S_t$. As before, there are k possible $S_u S_u$ outcomes, $k(k - 1)$ possible outcomes of the form $F_u S_w S_w$, and now $k(k - 1)(k - 2)$ outcomes of the form $F_u F_w S_t S_t$. Denoting the probability of patient success with t in j courses following a treatment sequence \mathbf{u} by $\pi_{t,j}(\mathbf{u})$, the probability of overall patient success with t is

$$\pi_{t,2} + \sum_{u \neq t} [\pi_{t,3}(u) + \pi_{t,4}(u, u)] + \sum_{u \neq t, w \neq t, u \neq w} [\pi_{t,4}(u, w) + \pi_{t,5}(u, u, w) + \pi_{t,5}(u, w, w) + \pi_{t,6}(u, u, w, w)]$$

MLEs may be defined by partitioning the patient outcome space into $k + 1$ sets as in Section 3.2. For evaluating multi-treatment strategies, now there are three-treatment combinations (u, w, t) , with

u given first, w given if u is unsuccessful, and t then given if w is unsuccessful. The probability of patient success with this strategy is

$$\zeta(u, w, t) = \zeta_u + (1 - \zeta_u)\{\zeta_{w|u} + (1 - \zeta_{w|u})\zeta_{t|u,w}\}$$

where ζ_u and $\zeta_{w|u}$ are as defined earlier and $\zeta_{t|u,w}$ is the probability of two consecutive successful courses with t given two earlier failures, first with u and then with w . Since the number of such combinations is now rather large, a practical approach is to focus on consecutive treatment pairs using the probability $\zeta(u, t)$ of patient success when an unsuccessful course with u is followed by t . The regressive logistic models now may draw on a richer patient history, however. For example, with $k = 4$ treatments, model (9) gives

$$\Pr[F_u F_w S_t S_t] = \frac{1}{4} [1 - g(\mu_u)]^{\frac{1}{3}} [1 - g(\mu_w + \frac{2}{3}\beta_w)]^{\frac{1}{2}} g(\mu_t + \frac{4}{5}\beta_t) 1g(\mu_t + \alpha_t + \frac{4}{5}\beta_t) \tag{13}$$

A much simpler application is a trial where a single successful course achieves patient success, but patient failure still consists of two unsuccessful courses. Here the possible patient outcomes are the k events S_u and the $k(k - 1)$ combinations of each of $F_u S_t$ and $F_u F_t$. Evaluation of single-treatment and two-treatment success probabilities may proceed as before in terms of π_t and $\zeta(u, t)$.

Our formulation also may be extended to accommodate trials where there is a non-trivial probability that the patient may die during therapy. Let $D_{i,j}$ be the indicator that patient i dies during the j th course and $\zeta_j = \Pr[D_{i,j} = 1]$. If we now define $\theta_{i,j} = \Pr[Y_{i,j} = 1 \mid \mathbf{Y}_{i,j-1}; \mathbf{T}_{i,j}; D_{i,j} = 0]$, so that each outcome probability is as previously defined but now includes the event that the patient survives the j th course in the conditioning event, then the likelihood (2) is extended as follows:

$$\mathcal{L}_{i,r} = \prod_{j=1}^r \{\theta_{i,j}^{y_{i,j}} (1 - \theta_{i,j})^{1-y_{i,j}} \tau_{i,j}(t_{i,j})(1 - \zeta_j)\}^{1-D_{i,j}} \zeta_j^{D_{i,j}} \tag{14}$$

Thus, the MLE $\hat{\zeta}_j$ is simply the mean number of deaths in the j th course, and the other parameter estimates are as before but based on the patients who do not die. We did not use this extended model in designing the prostate cancer trial since, among 185 patients treated with similar regimens at MDACC during 1993–1999, there was only one death during therapy and this was due to a heart attack unrelated to treatment.

If the probability of death may be related either to treatment or to the individual patient’s history, beyond the number of previous courses, then a model extension more detailed than (14) may be required. In such settings, $Y_{i,j}$ may be replaced by a multinomial variable that accounts for the more complex outcome. For example, a trinary variable may be used to record whether the course is successful, the patient dies during that course, or the course is unsuccessful but the patient is alive. Dealing with this case would then require extending the probability model, constructing interim decisions rules for discontinuing treatments having unacceptably high death rates, and formulating selection rules in terms of the probabilities of both response and death. We are currently pursuing this line of research in the context of a multi-course chemotherapy trial in acute leukaemia.

Patient covariates may be included by adding them to the linear term of either regressive logistic model. A practical advantage of this, aside from the usual statistical benefits of covariate adjustment, is that a larger patient group may be included in a given trial.

6. DISCUSSION

We have provided a statistical framework for clinical trials in which therapy may consist of multiple courses, and each patient's treatment for any course after the first is chosen adaptively based on that patient's treatments and outcomes in previous courses. Because this sort of adaptive treatment assignment is quite common in medical settings where two or more treatments are available, our model and design more closely reflect actual clinical practice than more commonly used statistical methods. While any statistical formulation necessarily involves data reduction, such as characterizing the outcome in each course as a binary variable in our application, we feel strongly that accounting for the multi-course structure provides a much more informative basis for treatment evaluation. In particular, the notion of selecting treatment sequences rather than individual treatments follows naturally from consideration of the multi-course structure.

Our simulation study has shown that designs with attractive operating characteristics under a variety of clinical scenarios can be obtained with moderate sample sizes. For example, the sample of 156 patients required to obtain a 90 per cent correct selection probability in terms of the π_i 's under scenario 1 (Table I) also provides good two-course treatment strategy selection probabilities (Tables V and VI). Two important provisos are that: (i) until formal methods are derived, the sample size for a given trial should be determined empirically, rather than by attempting to adapt available methods intended for simpler settings; and (ii) the model upon which the probability estimates and thus the treatment selection is based should itself be chosen based on the data from the trial once it is completed. The fact that the different models give different fits and selection probabilities may be attributed to the fact that the regressive logistic models borrow strength across the outcomes. Even when the MM and RLM1 models have the same number of parameters, however, they describe the phenomenon in qualitatively different ways, and so they should be expected to perform differently for a given data set.

Our design is similar to a multi-stage cross-over trial where, in its simplest form, patients are randomized between two treatments for a first stage and each patient is then 'crossed over' to the other treatment in a second stage. The essential difference is that our design chooses all treatments after the first adaptively based on the patient's outcomes in earlier courses, whereas in a cross-over trial of treatments u and t , all patients given u in stage 1 are given t in stage 2, and vice versa, regardless of their stage 1 outcome. There is also a very different objective. In a cross-over trial, patients are crossed over to reduce variability in estimating individual treatment effects, and carry-over effects between courses are viewed as a nuisance. In contrast, the conditional distribution of Y_j given (Y_1, \dots, Y_{j-1}) is an essential feature of our model and treatment evaluation. Indeed, what we call cross-resistance would be called a carry-over effect in a cross-over trial. A discussion of what does and does not constitute a cross-over trial is given by Senn [13].

The two-stage randomization described in Section 2 is formally equivalent to simply randomizing patients equally among the 12 two-treatment strategies. Under either approach, a given patient receives any u in the first course with probability 1/4 and receives (u, t) over either three or four courses with probability $(1 - \xi_u)/12$. Since the number of patients $n_u - X_u$ who fail with u in course 1 or 2 is random, however, the two-stage randomization in which the second treatment assignment is balanced within this smaller subgroup for each u provides a slightly better balance. In our conduct of the prostate cancer trial, the second randomization has not presented any practical difficulties.

ACKNOWLEDGEMENTS

We thank two referees for their thoughtful and constructive comments on an earlier draft of this manuscript.

REFERENCES

1. Emanuel EJ, Patterson WB. Ethics of randomized clinical trials. *Journal of Clinical Oncology* 1998; **16**:365–371.
2. Landis SH, Murray T, Bolden S, Wingo PA. Cancer statistics. *CA Cancer J. Clin.* 1998; **48**:6–29.
3. Catalona WJ, Smith DS, Ratliff TL, Dodds KM, Doplén DE, Yaun J, Petros JA, Andriole GL. Measurement of prostate-specific antigen in serum as a screening test for prostate cancer. *New England Journal of Medicine* 1991; **324**:1156.
4. Storlie JA, Buckner JC, Wiseman GA *et al.* Prostate specific antigen levels and clinical response to low dose dexamethasone for hormone-refractory metastatic prostate cancer. *Cancer* 1995; **76**:96–100.
5. Simon R, Wittes RE, Ellenberg SS. Randomized phase II clinical trials. *Cancer Treatment Reports* 1985; **69**:1375–1381.
6. Palmer CR. A comparative phase II clinical trials procedure for choosing the best of three treatments. *Statistics in Medicine* 1991; **10**:1327–1340.
7. Thall PF, Sung HG. Some extensions and applications of a Bayesian strategy for monitoring multiple outcomes in clinical trials. *Statistics in Medicine* 1998; **17**:1563–1580.
8. Fisher RA. *The Design of Experiments*. Hafner: London, 1935.
9. Bonney GE. Regressive logistic models for familial disease. *Biometrics* 1986; **42**:611–625.
10. Bonney GE. Logistic regression for dependent binary observations. *Biometrics* 1987; **43**:951–973.
11. Bechhofer RE. A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Annals of Mathematical Statistics* 1954; **25**:16–39.
12. Bechhofer RE, Santner TJ, Goldsman DM. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. Wiley: New York, 1995.
13. Senn S. *Crossover Trials in Clinical Research*. Wiley: Chichester, 1993.