Published in final edited form as:

Biometrics. 2012 March ; 68(1): 203–211. doi:10.1111/j.1541-0420.2011.01623.x.

Bayesian Enrichment Strategies for Randomized Discontinuation Trials

Lorenzo Trippa^{1,*}, Gary L. Rosner^{2,**}, and Peter Müller^{3,***}

¹Harvard School of Public Health and Department of Biostatistics, Dana–Farber Cancer Institute, Boston, Massachusetts 02115, U.S.A.

²Division of Biostatistics & Bioinformatics, Sidney Kimmel Comprehensive Cancer Center at Johns Hopkins, Baltimore, Maryland 21205, U.S.A.

³Department of Biostatistics, University of Texas M. D. Anderson Cancer Center, Houston, Texas 77030, U.S.A.

Summary

We propose optimal choice of the design parameters for random discontinuation designs (RDD) using a Bayesian decision-theoretic approach. We consider applications of RDDs to oncology phase II studies evaluating activity of cytostatic agents. The design consists of two stages. The preliminary open-label stage treats all patients with the new agent and identifies a possibly sensitive subpopulation. The subsequent second stage randomizes, treats, follows, and compares outcomes among patients in the identified subgroup, with randomization to either the new or a control treatment. Several tuning parameters characterize the design: the number of patients in the trial, the duration of the preliminary stage, and the duration of follow-up after randomization. We define a probability model for tumor growth, specify a suitable utility function, and develop a computational procedure for selecting the optimal tuning parameters.

Keywords

Clinical trials; Enrichment designs; Randomized discontinuation design; Tumor growth models

1. Introduction

We propose a decision-theoretic approach to choosing an optimal randomized discontinuation design (RDD). We use a tumor growth model to specify a prior probability model for experimental outcomes across alternative choices of design parameters and we introduce a novel prior elicitation process for the tumor growth model.

RDDs have found use in phase II clinical studies in oncology for assessing activity of new agents with potentially cytostatic properties. The RDD proceeds in two stages, an open-label stage and a randomization stage (Rosner, Stadler, and Ratain, 2002). All (N) patients enrolling in the study enter the first stage and receive the new agent. At the end of some predetermined treatment duration (T_1), each patient undergoes an evaluation of disease status. Only patients with stable disease after stage 1 participate in the randomized second

stage, with randomization either to the new agent or to a control, usually placebo. After some period of follow-up after randomization (T_2), a subsequent evaluation of disease status yields the primary outcome data for assessing the treatment's activity.

The structure of the RDD is motivated by the characteristics that differentiate potentially cytostatic agents from more commonly used cytotoxic anticancer therapies. Traditionally, clinical trials in oncology have evaluated anticancer activity based on tumor shrinkage. In contrast, a potentially cytostatic agent, such as many of the new targeted anticancer agents, could demonstrate activity by slowing tumor growth, leading to stable disease. The great heterogeneity of tumor growth rates across patients with certain cancers and similar prognostic characteristics can confound the observed effect of cytostatic agents. Moreover, some agents may be active only for a subgroup of the target population. Millar and Lynch (2003) emphasize that for these reasons, traditional clinical trial designs for cytotoxic therapies are inappropriate for cytostatic agents. See Stadler et al. (2005) and Ratain et al. (2006) for applications of RDDs in cancer treatment studies that highlight how these features dictate the design structure.

With a good choice of tuning parameters, the first stage selects a possibly sensitive population, and the second stage evaluates whether the experimental therapy is active in this subpopulation. For the design to be efficient, it is critical that one choose these tuning parameters well. Compared to a two-arm randomized trial, a properly designed RDD can significantly improve the probability of detecting an active agent (Fedorov and Liu, 2005). However, this is only true for a good choice of the RDD design parameters. Stadler (2007) points out the importance of the tuning parameters and illustrates that a small T_1 (i.e., too short a prerandomization phase) can reduce the RDD's ability to detect an active agent.

Most of the literature on the RDD for cancer studies has focused on comparisons of this design with upfront randomization. Extensive discussions are given in Capra (2004); Freidlin and Simon (2005); Fu, Dowlati, and Shluchter (2009); Galsky et al. (2009); Millar and Lynch (2003); Ratain et al. (2006); Stadler et al. (2005); and Stadler (2007, 2009).

We discuss a Bayesian decision-theoretic approach for determining the RDD tuning parameters. The Bayesian approach accounts for all relevant uncertainties and proposes a choice of tuning parameters to best achieve the study goals. It also provides a unified framework for representing the impact of alternative tuning parameters on the design's operating characteristics and for optimally choosing among possible parameterizations. We propose a procedure for predicting the tumor growth trajectories on the basis of historical data and a utility function that allows comparison of alternative RDD parameterizations. Importantly, the application of our proposed approach only requires the elicitation of the prior on a single scalar parameter that can be interpreted as the treatment effect.

Formalizing the choice of the tuning parameters as a Bayesian decision problem (DeGroot, 2004) means that the decision maker considers different hypothetical scenarios (ϕ). These scenarios could include various degrees of activity of the new agent. A probability measure $\rho(\phi)$ on the set of scenarios reflects a priori beliefs. The decision problem—in this case the choice of $d = (N, T_1, T_2)$ —includes a utility function $u(d, \phi, X)$ that represents the decision maker's relative preferences for each combination of decision (i.e., design) d , hypothetical scenario ϕ , and experimental outcome X . The optimal decision is defined by maximizing the utility function in expectation. One can base the prior $\rho(\phi)$ on many sources, such as earlier studies of the agent under study and historical data under standard of care. Korn et al. (2001) give a detailed account of the clinical and preclinical information one should consider when designing phase II trials for cytostatic agents. The utility function $u(d, \phi, X)$ synthesizes the most relevant aspects of the study; it balances the costs of the trial against possible benefits.

The model that we use to characterize tumor growth trajectories builds on a parametric stochastic process proposed in Ferrante et al. (2000). We formalize the ideas of (i) imputing each patient's tumor growth trajectories in an historical cohort on the basis of observations at discrete times, (ii) assuming that one can model future trajectories under the control therapy by resampling imputed historical trajectories, while (iii) modeling future trajectories under the experimental therapy as samples from the imputed historical trajectories after a simple geometric transformation involving an unknown parameter representing the treatment effect. The attractive features of the outlined approach are that the prior for future observations under the control regimen is highly representative of the historical studies and that, overall, the a priori distribution on the treatment-effect parameter is easily interpreted. The probability model can be used both in a Bayesian decisional framework and for simulations to assess an RDD's operating characteristics under hypothetical scenarios representative of the historical data and of prior expectations of the treatment's effect.

In the next section, we describe an RDD we were asked to design. We then review the various steps in finding the optimal design and conclude with a discussion.

2. A Clinical Trial for Renal Cell Carcinoma (RCC)

Over the last 5 years, several trials have investigated the use of targeted agents to treat patients with advanced RCC (Larkin et al., 2009). We have been asked to design several RDDs for evaluating anticancer activity of new targeted agents. The first is described in Stadler et al. (2005), a study carried out by the Cancer and Leukemia Group B (CALGB) to evaluate a putative antiangiogenic agent for treatment of metastatic RCC. In this article, we discuss designing a new study using an RDD. The study will investigate a possibly cytostatic agent for treating metastatic RCC, as in Stadler et al. (2005).

Evaluation of an RDD requires prior probabilities for the following events. We need the probability (p_e) that a patient attains stable disease at time T_1 and the probabilities of response for a patient randomized in stage 2 to control and treatment (p_0 and p_1 , respectively). These are design-specific probabilities; alternative choices of (T_1 , T_2) typically correspond to different values of p_e , p_0 , and p_1 . Optimal choice of the design parameters (N , T_1 , T_2) requires a prior probability model for (p_e , p_0 , p_1) that is coherent across possible values of T_1 and T_2 . The easiest way to specify such a prior is through an underlying tumor growth model.

We construct a realistic prior probability model for tumor growth trajectories based on tumor growth data for 61 patients from the study reported in Stadler et al. (2005). In Section 3 we describe a parametric tumor growth model. Then, in Section 4, we use these data to derive the desired prior probability model for tumor growth trajectories. The probability model allows us to efficiently simulate (p_e , p_0 , p_1) for arbitrarily chosen values of (T_1 , T_2). Figure 1, for example, shows the prior distribution of p_e under alternative values of T_1 . The boxplots show how increasing T_1 leads to a decrease in the expected proportion of patients eligible for randomization. An important feature of the adopted elicitation approach is that the parametric model is only used for interpolation between actually observed data points. This leaves the approach robust against alternative model choices.

3. Tumor Growth Model and Prior Specification in Absence of Historical Data

3.1 A Model for Tumor Growth

Several tumor growth models exist in the literature. In the following discussion, we model tumor growth with a Gompertz diffusion process (Ferrante et al., 2000).

The Gompertz model solves the differential equation

$$dX_t/dt = a \cdot X_t - b \cdot X_t \cdot \log(X_t) \quad X_0 = x_0, \quad (1)$$

where a and b are two constants and X_t represents the tumor volume at time t . The Gompertz model reflects initial exponential dynamics and a subsequent decreasing growth rate. It often approximates tumor growth data well (Ribba et al., 2006). Reviews of relevant applications in oncology are given in Clare, Nakhli, and Panetta (2000) and Waliszewski (2005).

One limitation of (1) is the deterministic nature of the implied tumor growth curves X_t . We therefore assume that tumor proliferation for each patient is characterized by the above equation randomly perturbed by subject-specific variations that are modeled through a Brownian motion. The resulting random process, for the i th patient, is the solution to the following stochastic differential equation (s.d.e.):

$$\begin{aligned} dX_{it} &= \{(a_i X_{it} - b_i X_{it} \log(X_{it})) dt + \sigma_i X_{it} dW_{it}, \\ X_{i0} &= x_{i0}, \quad t \in [0, T], \end{aligned}$$

where $\{W_{it}\}_{i=1}$ are independent Brownian motions. The stochastic model describes tumor growth if the individual patient receives the same therapy during the entire period $[0, T]$. The solution of the s.d.e. (Ferrante et al., 2000) is:

$$\exp \left\{ \left(\frac{a_i - \sigma_i^2/2}{b_i} + \left(\log(x_{i0}) - \frac{a_i - \sigma_i^2/2}{b_i} \right) e^{-b_i t} + \sigma_i \int_0^t e^{-b_i s} dW_{is} \right) \right\}. \quad (2)$$

An in-depth discussion of the probabilistic properties of the Gompertz diffusion model and its application in oncology is given in Albano and Giorno (2006, 2009).

For a patient in an RDD whose therapy switches at time T , the resulting process will solve the same s.d.e. Now, a_i and b_i are two piecewise-constant functions:

$a_i(t) = a_i^0 I(t < T) + a_i^1 I(t \geq T)$ and $b_i(t) = b_i^0 I(t < T) + b_i^1 I(t \geq T)$. The Markovian property of the process allows easy computation of the finite-marginal distributions. The transition probability density of the process (2), omitting patient index i , is:

$$\begin{aligned} p(X_{t+s}|X_t, a, b, \sigma) &\propto \frac{1}{X_{t+s}} \\ &\times \exp \left[-\frac{\{\log(X_{t+s}) - e^{-bs} \log(X_t) - (1 - e^{-bs})(a - \sigma^2/2)/b\}^2}{\sigma^2(1 - e^{-2bs})/b} \right]. \quad (3) \end{aligned}$$

We note that the process has marginal log-normal distributions and that we can easily sample from the marginal distributions. Expression (3) provides a complete definition of the process. In the following sections we will use only these log-normal conditional distributions.

The parametrization of the Markov process when the patient is treated with two distinct therapies can be slightly simplified. We can argue that $b_i^0 = b_i^1 = b_i = b^1_i = b_i$ is necessary to allow for the possibility of one treatment being superior to the other. The hypothesis that the new treatment, labeled 1, is superior to the control, labeled 0, can be formalized as the inequality

$$E(X_{t+s}|X_t = x_t, a^0, b^0, \sigma) > E(X_{t+s}|X_t = x_t, a^1, b^1, \sigma) \quad (4)$$

for every x_t and s . In words, treatment 1 permanently improves the patient's condition. Considering the conditional expectation

$$E(X_{t+s}|X_t=x_t, a, b, \sigma) = \exp \left\{ e^{-bs} \log(x_t) + \frac{a - \sigma^2/2}{b} (1 - e^{-bs}) + \frac{\sigma^2}{4b} (1 - e^{-2bs}) \right\}$$

and its derivative with respect to s , we observe that expression (4) holds if, and only if, $a_i^1 < a_i^0$ and $b_i^1 < b_i^0$. Moreover if $a^1 < a^0$ and $b^1 = b^0$, then, for every couple $(x_t, x_{t+s}) \in (0, \infty)^2$, the inequality $p(X_{t+s} > x_{t+s} | X_t = x_t, a^1, b^1, \sigma) > p(X_{t+s} > x_{t+s} | X_t = x_t, a^0, b^0, \sigma)$ holds. This fact can be straightforwardly verified using Theorem 4 in Levy (1973).

3.2 A Hierarchical Prior

We introduce a prior distribution for the subject-specific random effects $a_i^0, a_i^1, b_i, \sigma_i$ and σ_i . It includes a mixture over two cases: inhibition and no effect. We use latent binary variables $E_i \in \{0, 1\}$ to indicate whether the treatment inhibits tumor growth for patient i . In the case of no growth inhibition (i.e., $E_i = 0$), σ_i, b_i and σ_i completely characterize the i th patient's tumor growth.

We assume that the random effects vector $(a_i^0, a_i^1, b_i, \sigma_i)$ varies across a heterogeneous population according to a parametric distribution. The random effects distribution F_θ is defined for a reparametrization of $(a_i^0, a_i^1, b_i, \sigma_i)$ to $(\varphi_{1i}, \varphi_{2i}, \varphi_{3i}, \psi_i)$. The first three components,

$$\varphi_{1i} = \frac{a_i^0 - \sigma_i^2/2}{b_i}, \quad \varphi_{2i} = \frac{\sigma_i^2}{2b_i} \quad \text{and} \quad \varphi_{3i} = b_i$$

regulate the tumor growth process under $E_i = 0$. We observe that $\varphi_{1i}, \varphi_{2i}$ and φ_{3i} have a clear interpretation in terms of the conditional moments

$$E(\log(X_{is}) | X_{i0}, a_i^0, b_i, \sigma_i, E_i=0) = e^{-b_i s} \log(X_{i0}) + \frac{a_i^0 - \sigma_i^2/2}{b_i} (1 - e^{-b_i s})$$

and

$$\text{Var}(\log(X_{is}) | X_{i0}, a_i^0, b_i, \sigma_i, E_i=0) = \frac{\sigma_i^2}{2b_i} (1 - e^{-2b_i s}).$$

The fourth random effect (ψ_i) reports the treatment effect for patient i

$$\psi_i = \left(\frac{a_i^0 - \sigma_i^2/2}{b_i} - \frac{a_i^1 - \sigma_i^2/2}{b_i} \right).$$

Let $\text{Ga}(x; \nu, r)$ indicate a Gamma distributed random variable with mean ν/r . Let $\text{NIG}(\theta_1, \theta_2; m, \nu^2, \nu, r)$ indicate a normal-inverse-gamma distribution for the random vector (θ_1, θ_2)

defined as $\frac{1}{\theta_2} \sim \text{Ga}(\nu, r)$ and $(\theta_1 | \theta) \sim \mathcal{N}(m, \theta_2 \nu^2)$. We complete the model with log-normal (LN) random effects distributions for ϖ_{1j} , ϖ_{2j} , ϖ_{3j} and ψ_j

$$\varphi_{ji} | \theta \sim \text{LN}(\theta_{j1}, \theta_{j2}), \quad \text{and} \quad \psi_i | \theta,$$

$$E_i \sim \begin{cases} 0 & \text{if } E_i=0 \\ \text{LN}(\theta_{41}, \theta_{42}) & \text{if } E_i=1, \end{cases}$$

where $j \in \{1, 2, 3\}$, independent Bernoulli priors on the latent indicators for growth inhibition

$$p(E_i=1) = \pi,$$

and conjugate priors for the hyperparameters $\theta = (\theta_{j1}, \theta_{j2})$

$$(\theta_{j1}, \theta_{j2}) \sim \text{NIG}(m_j, v_j^2, v_j, r_j).$$

The normal-gamma hyperprior allows us to formulate an initial guess about the random effects distribution for the growth process parameters across the population. Frequently in Bayesian statistics, elicitation of informative priors postulates a priori the existence of an imaginary data set representing initially available knowledge. In our case, when Jeffrey's prior for log-normal parameters is updated (with hypothetical prior equivalent data points), a normal-gamma distribution is obtained (Padgett and Wei, 1977).

Prior elicitation is a critical part of a decision-theoretic approach. Although considerable research has been conducted to improve the elicitation process (see, e.g., Garthwaite, Kadane, and O'Hagan, 2005, and references therein), many difficulties remain. This is especially true when a considerable number of uncertain quantities is involved. Such difficulties can be mitigated in the application to clinical trial design by using frequentist properties, also known as operating characteristics, to critique a proposed trial design. For example, the frequentist properties of RDDs, for a fixed sample size N , are determined uniquely by three unknown parameters: the probability of a patient being eligible for randomization p_e and, conditionally on eligibility, the response probabilities for the treatment and control regimens (p_0, p_1). Validation of the optimal decision is particularly meaningful when the elicitation of expert opinion on the tumor growth process across the heterogeneous population is difficult, but, at the same time, expert judgment about the three pivotal probabilities is available.

4. Prior Construction Based on Historical Data

The study that we described in Section 2 is one of many similar studies of targeted agents for RCC. We have access to data from a similar earlier study carried out by CALGB and reported in Stadler et al. (2005). The prior construction should exploit such available historical data to minimize the need for prior elicitation from experts. In the following paragraphs, we describe a practical strategy for specifying an informative prior in such cases. The approach assumes that tumor growth trajectories under the control regimen in the planned trial are likely to be similar to those in the historical trial. The proposed strategy restricts the need for subjective prior elicitation to only the prior parameters related to the treatment effect φ_i .

The main features of the proposed construction are as follows. We use model-based interpolation of the historical tumor growth data to generate realizations of tumor growth curves for future patients. The underlying model is the Gompertz diffusion model defined earlier. The use of interpolation leaves the proposed strategy very robust against possible deviations from the model assumptions. The construction is easily interpreted and implemented. We assume a probability model consistent with an extensive literature on solid tumor growth. The historical trajectories are imputed on the basis of discrete time observations. The details of the procedure formalize the idea of defining a prior, for the control regimen, identical to the empirical distribution of the historical trajectories and, similarly, a prior, for the treatment regimen, identical to a suitable geometric transformation of the same empirical distribution.

Tumor growth under the control regimen. Assume we have historical data on M patients under the control regimen:

$$\left\{ \tilde{X}_{1t_0}^0, \tilde{X}_{1t_1}^0, \dots, \tilde{X}_{1t_k}^0 \right\}, \\ \left\{ \tilde{X}_{2t_0}^0, \dots, \tilde{X}_{2t_k}^0 \right\}, \dots, \left\{ \tilde{X}_{Mt_0}^0, \dots, \tilde{X}_{Mt_k}^0 \right\}.$$

We use $\tilde{X}$ to distinguish the historical data from future experimental observations X . The restriction to fixed measurement times ($t_1, \dots, t_k$) across patients is not important. We only use this assumption to simplify notation. We specify a prior distribution for tumor growth curves (for future patients) under the control regimens as a mixture of diffusion processes. Each component of the mixture corresponds to one of the M historical patients. Let $\mathbf{GP}(a, b, \sigma)$ denote a Gompertz diffusion process as defined in Subsection 3.1 and let $\mathbf{GP}(a, b, \sigma|\Delta)$ be the conditional law of the process given the event $\{\Delta\}$. We assume

$$\left\{ X_{it}^0 \right\}_{t \geq 0} \sim \frac{1}{M} \sum_{j=1}^M \mathbf{GP}(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j | X_{it_0}^0 = \tilde{X}_{jt_0}^0, \dots, X_{it_k}^0 = \tilde{X}_{jt_k}^0). \quad (5)$$

That is, the law of the process is a mixture of M components and each component reflects one of the historical trajectories. The parameters $(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j)$ in (5) are estimated with the standard maximum likelihood technique (Gutierrez et al., 2006). We note that simulating from (5) only requires the log-normal conditionals of (3).

Alternatively, the scheme could be slightly modified. The investigator could first define clusters of historical patients (Schnatter and Kaufmann, 2008) and then assume a common set of parameters $(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j)$ for all patients in the same cluster and proceed again as in (5).

Tumor growth under the treatment regimen. We impute for each historical patient a hypothetical trajectory under the treatment regimen. We exploit the linear relationship

$$\log(\tilde{X}_{jt_{l+1}}^0) = \gamma(\tilde{b}_j, t_l - t_{l+1}) \log(\tilde{X}_{jt_l}^0) \\ + \beta(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1}) + \delta(\tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1}) Z_{jl}, \quad (6)$$

of the Gompertz model. The random components $(Z_{jl}; j = 1, \dots, M, l = 0, \dots, k-1)$ are independent standard Gaussian variables and $\gamma(\cdot)$, $\beta(\cdot)$ and $\delta(\cdot)$ are known functions. Solving (6) for Z_{jl} we find realizations of the Gaussian random variables. Next we substitute these values for Z_{jl} in

$$\log \left(\tilde{X}_{jt_{l+1}}^1 \right) = \gamma \left(\tilde{b}_j, t_l - t_{l+1} \right) \log \left(\tilde{X}_{jt_l}^1 \right) + \beta \left(\tilde{a}_j - \tilde{\psi}_j \tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1} \right) + \delta \left(\tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1} \right) Z_{jl}. \quad (7)$$

These are the analog of (6) for the treatment regimen. The only unknown random components in (7) are the treatment effects $\{\tilde{\psi}_j\}_{j=1}^M$; the distribution of the treatment effects representative of the expectations on the cytostatic activity of the treatment.

The hypothetical observations $(\tilde{X}_{jt_0}^1, \dots, \tilde{X}_{jt_k}^1; j=1, \dots, M)$ are defined as the random solutions of equations (7) assuming that for every historical patient $\tilde{X}_{jt_0}^0 = \tilde{X}_{jt_0}^1$. Finally, conditional on $\tilde{X}^1 = (\tilde{X}_{jt_0}^1, \dots, \tilde{X}_{jt_k}^1; j=1, \dots, M)$ and on the random vector $\tilde{\psi} = (\tilde{\psi}_1, \dots, \tilde{\psi}_M)$, future tumor growth trajectories under the treatment regimen are characterized as follows:

$$\left\{ X_{it}^1 \right\}_{t \geq 0} | \tilde{X}^1, \tilde{\psi} \sim \frac{1}{M} \sum_j \mathbf{GP} \left(\tilde{a}_j - \tilde{\psi}_j \tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j | X_{it_0}^1 = \tilde{X}_{jt_0}^1, \dots, X_{it_k}^1 = \tilde{X}_{jt_k}^1 \right). \quad (8)$$

Tumor growth under randomization to treatment and control. The outlined plug-in strategy can easily be extended to characterize the distribution of the tumor growth process for a patient receiving the new treatment until T_1 and who is subsequently randomized to control. Also in this case, the process law can be represented as a mixture of M components. Each component consists of the conditional law of a Gompertz diffusion process whose transition densities switch parametrization at time T_1 . The parameters $(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j)$ as in (5) are estimated through the standard maximum likelihood technique and the random quantities $\{\tilde{\psi}_j\}_{j=1}^M$ are through modeled the prior discussed in Section 3. We impute, for each historical patient, the hypothetical observations at $(t_1, \dots, t_k)$ and T_1 . For each patient j the distribution of $\tilde{X}_{T_1}^0$ conditional on $\{\tilde{X}_{jt_0}^0, \tilde{X}_{jt_2}^0, \dots, \tilde{X}_{jt_k}^0\}$ is computed. We then slightly modify the equations (6) and (7) by adding to the vector $(t_0, t_1, \dots, t_k)$ the additional observation time T_1 , and substitute $\beta(\tilde{a}_j - \psi_i \tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1})$ in (7) with

$$\beta \left(\tilde{a}_j - \psi_i \tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1} \right) I(t_{l+1} \leq T_1) + \beta \left(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j, t_l - t_{l+1} \right) I(t_{l+1} > T_1).$$

The hypothetical observations are then defined as the solutions of the slightly modified random equations. In this case, the equations' unknown random quantities consist of the independently distributed random vectors $(\tilde{X}_{1,T_1}^0, \dots, \tilde{X}_{M,T_1}^0)$ and $\{\tilde{\psi}_j\}_{j=1}^M$.

The use of historical data greatly simplifies the decision problem. The only prior information that needs to be directly elicited from expert knowledge is related to the treatment effects ψ_i . Prior elicitation for these remaining parameters can be guided by plotting simulated tumor growth trajectories under alternative prior assumptions. Figure 2 shows an example. Moreover, the procedure allows the investigator to report operating characteristics of alternative designs d following a simple predictive approach. It is straightforward to compute via Monte Carlo simulations the predictive probability that the trial detects the effect of the agent for several assumed levels of efficacy of the new agent.

5. Decision Problem

The utility of the trial $u(d, \phi, X)$ is a function of the adopted design d , the parameters that characterize the tumor growth processes (ϕ), and of the observed data X . We generically use $p(\cdot)$ to denote a probability distribution. The expected utility is

$$U(d) = \int \int u(d, \phi, X) p(dX|\phi, d) p(d\phi).$$

The decision problem of selecting a design is solved by maximizing the expected utility, $U(d^*) = \max_{d \in D} U(d)$. The action space D includes alternative designs d that differ in the overall number of patients in the trial (N) and in the durations of the two phases (T_1, T_2).

We consider the utility associated with each combination (d, ϕ, X) to balance the costs of the trial ($C(d, \phi, X)$) and the benefits that derive from recommending a phase III trial for an effective agent ($B(d, \phi, X)$):

$$u(d, \phi, X) = B(d, \phi, X) - C(d, \phi, X). \quad (9)$$

We assume that the recommendation to move to phase III will be determined by a hypothesis test after completion of the trial. The null hypothesis is that the novel treatment is not superior to the control. Let $S_d(X)$ denote the test statistic, and let R_d denote the rejection region. If the agent should be recommended for phase III, we assume the benefits would be proportional to the therapy's expected efficacy in a future $(N+1)$ th patient. We use

$$B(d, \phi, X) = I\{S_d(X) \in R_d\} E\{\log(1 + \psi_{N+1}) | \phi\}.$$

The benefit is balanced with the cost of enrolling patients in the trial. The cost depends on the number of patients in the two trial phases (N and n) and on the durations T_1 and T_2 :

$$C(d, \phi, X) = c_1 N + c_2 n + c_3 (T_1 N + T_2 n).$$

Here c_1 and c_2 are the costs of enrolling patients in the first phase and treating patients in the second phase of the trial, whereas c_3 is the variable cost per unit of time of retaining a patient in the trial.

We assume that the eligibility criteria for a patient to enter the second phase of the trial are fixed. Upon completion of the first phase, the i th patient continues to phase 2 if

$$\delta_1(X_{i0}, T_1) \geq X_{i1} \geq \delta_2(X_{i0}, T_1).$$

The two bounds δ_1 and δ_2 depend on the tumor volume at the beginning of the trial X_{i0} and on the length of the first stage T_1 . The lower bound is motivated by the need to continue the treatment for a patient who is appreciably benefitting from the treatment (i.e., no randomization if below $\delta_2(\cdot)$). The upper bound implements the desired enrichment strategy of the design intended to select a homogenous group of patients who may be benefitting from the cytostatic activity of the new agent.

The optimal strategy d^* is computed by a Monte Carlo optimization algorithm. We approximate the expected utility $U(d)$ for any point d in D by iteratively simulating population parameters ϕ and hypothetical experimental outcomes X from the prior model.

At each iteration, the utility $u(d, \phi, X)$ is computed. Simulation of X is carried out using the prior construction described in Section 4. In similar applications, when no historical data were available one would proceed by simulating from the hierarchical model described in Section 3.2 to generate future data X . We will use the hierarchical prior for the simulation study reported later, in Section 6.

To obtain approximations of $U(d)$ efficiently for every design d , the algorithm carries out one or a few simulations for H designs d_h , $h = 1, \dots, H$, and fits a smooth surface $\tilde{U}(d)$ to the Monte Carlo samples (d_h, u_h) . That is, we do not carry out several simulations for each design d . The approach has been considered in detail in Müller and Parmigiani (1995). Note that the algorithm approximates the expected utility $U(d)$ using not only results of simulations associated with d but also by borrowing strength from results associated with neighboring designs. In summary,

Step 1. Select H designs $\{d_h \in D, h = 1, \dots, H\}$, and generate random quantities u_h equal in distribution to $u(d_h, \phi, X)$.

Step 2. Fit a least squares multivariate polynomial regression to (d_h, u_h) , $h = 1, \dots, H$. The fitted surface $\tilde{U}(d)$ approximates the expected utility surface $U(d)$.

Step 3. Maximize $\tilde{U}(d)$ to find the (approximate) optimal design d^* .

The algorithm reduces the evaluation of the expected utility surface to a standard statistical regression problem. The maximization in step 3 is fully deterministic. The multivariate polynomial regression in step 2 involves model selection; it is necessary to select from a large number of possible models, ranging from first-order polynomials to the saturated model. Many parsimonious model selection procedures have appeared in the literature (Bursham and Anderson, 1998). Most commonly used selection criteria guarantee the

almost-sure convergence of the estimates $\{\tilde{U}(d)\}_{d \in D}$ to the expected utility surface $U(d)$ with increasing Monte Carlo sample size H .

6. Simulation Example

We apply the proposed optimization procedure to a hypothetical decision problem. We evaluate the operating characteristics for a new study without historical data by using the hierarchical model from Section 3. The prior probability that the novel agent has cytostatic activity is $P(\pi > 0) = 0.5$ and $P(\pi | \pi > 0) = \text{Unif}(0, 1)$. Figure 3 summarizes the first 250 days of the random trajectories after the tumors reach the minimal threshold $\varepsilon = 0.02$. The threshold ε defines the minimum detectable tumor size. The randomization in the second phase is balanced. Half the patients receive the control treatment. Patients enroll into the trial after a time varying from 4 to 6 months after the tumor has reached the threshold ε .

The eligibility criteria that determine participation in the second (randomized and blinded) phase are defined by two thresholds of the relative variation of the tumor mass during the first (open) phase: $\delta_1(X_0, T_1) = 1.2X_0$ and $\delta_2(X_0, T_1) = 0.7X_0$. The thresholds correspond to the RECIST response criteria (Therasse et al., 2000) used in many oncology trials, such as Stadler et al. (2005). The primary endpoint for the second phase of the trial is dichotomous as in many oncology phase II trials. The disease is considered stable if the tumor mass increase is less than 20%. The decision for recommending the agent for phase III is based on a hypothesis test. We use a Fisher's exact test with significance level $\alpha = 0.05$.

The action space is $D = \{(N, T_1, T_2) \in (0, 1, \dots, 400) \times (0, 1, \dots, 300)^2\}$. In particular, the alternative alter of not performing the trial ($N = 0$), as well as a traditional two-arm randomized trial ($T_1 = 0$) are options.

The optimal strategies corresponding to alternative utility functions are determined using Monte Carlo samples of $H = 10^7$ simulated trials with tuning parameters $(d_h; h = 1, \dots, H)$ randomly spread across the action space. Alternatively one could use a sparse regular grid. Figure 4 shows orthogonal sections of the fitted surfaces $\tilde{U}(d)$ intersecting the optimal designs $d^* = (N^*, T_1^*, T_2^*)$.

Recall that the utility function includes three trade-off parameters, c_1 , c_2 , and c_3 . Comparing the optimal choices under the alternative utility functions demonstrates the sensitivity of the decision problem with respect to c_1 , c_2 , c_3 .

In the example illustrated in Figure 4, a rational decision maker favors the RDD under all three considered utility functions. The optimal discontinuation strategy d^* guarantees a superior expected utility compared to a standard two-arm trial design. The optimal choices (T_1^*, T_2^*) are appreciably stable while, as expected, increasing the costs associated with enrolling a patient in the trial, leads to a clear decrease of the optimal sample size N^* .

As mentioned earlier, the operating characteristics of an RDD with dichotomous final end point are determined by three probabilities. These are the eligibility probability p_e and the response probabilities p_0 and p_1 . Table 1 illustrates the power of the optimal design

$d^* = (N^* = 2.16, T_1^* = 66, T_2^* = 102)$ for hypothetical values of p_e , p_0 , and p_1 .

We can then assess the operating characteristics of the optimal design d^* and compare it to two-arm randomized trials. Table 2 shows the power of alternative upfront randomized trials for different sample sizes N and hypothetical probabilities of stable disease under the control (p_0) and treatment (p_1) regimens. In this case, the follow-up period could differ from T_2^* . Such comparison has to consider that the a priori expectations of (p_0, p_1) for the two alternative designs are usually considerably different. In the RDD, p_0 and p_1 relate to patients with stable disease at T_1 and not to the overall population of patients, as with up-front randomization. A feature of the RDD is that the first phase of the design selects a subpopulation of patients who are likely to benefit from the novel agent. Therefore, the response probabilities (p_0, p_1) for the selected subpopulation during the second stage of an RDD is substantially different from the response probabilities (p_0, p_1) under up-front randomization of a traditional two-arm design. Several examples are given in Fedorov and Liu (2005). The up-front randomization, in most cases, corresponds to much narrower differences between p_0 and p_1 .

7. An RDD for a Cytostatic Agent Targeting Renal Cell Carcinoma

We apply the proposed approach to implement an RDD for the trial described in Section 2. In contrast to Section 6, we now use the historical data for specifying the prior, as described in Section 4. We use historical tumor growth data from a subgroup of 61 patients enrolled in the recent multicenter trial of a multikinase inhibitor for RCC (Stadler et al., 2005). We adopt a utility function similar to the earlier example, using $c_1 = c_2 = 0.001$ and $c_3 = 3.3 \times 10^{-5}$. The prior probability for cytostatic activity of the novel agent is fixed at $\pi = 0.7$. The prior for the treatment effects ψ_i is parameterized such that, conditionally on $\{E_i = 1\}$, the expected growth of the tumor mass after 4 months from the baseline measurement is 12% under the treatment regimen versus 24% under the control regimen. The primary end point for the randomized phase is still assumed to be a binary indicator for stable disease. The disease is considered stable if the tumor mass increase, from the time of enrollment, is less than 20%. The eligibility criteria for the randomized second stage are:

$$\delta_1(X_0; T_1) = 1.2X_0 \text{ and } \delta_2(X_0; T_1) = 0.8X_0.$$

The action space is $D = \{d = (N, T, T_2) \in (0, 1, \dots, 300)^3\}$. Again we use the proposed Monte Carlo procedure to find the optimal rule. We find the optimal design at $d^* = (N^* = 221, T_1^* = 72, T_2^* = 145)$, with time in days.

The validation of the optimal design includes the computation of operating characteristics under hypothetical scenarios; these can be evaluated for a grid of possible values of (p_e, p_0, p_1) as in Table 1. A complementary approach consists of evaluating the operating characteristics for possible parameterization of the probability model. We computed, for example, the design's power for several values of π , assuming the treatment reduces the average tumor growth of susceptible patients, after 4 months from baseline measurement, from 24% to 12%. The design's power decreases as the susceptible subpopulation proportion π decreases; power values above the 75% are observed for $\pi \geq 50\%$. A synthetic table reporting these computations is provided in the Web Appendix.

8. Discussion

Recent applications of RDDs in cancer have shown that it can be an interesting alternative to a standard two-arm randomized design for assessing the cytostatic activity of novel agents. The main features of the RDD are the highly increased acceptance of the design by potential patients (Stadler, 2007) and the ability of the design to focus the investigation on a relatively homogeneous cohort of patients who may be benefitting from the new treatment. However, the benefits are easily lost with a bad choice of the design parameters (Fedorov and Liu 2005; Stadler, 2009).

We have proposed a decision-theoretic approach for choosing the tuning parameters for an RDD. We assumed that the final result of the clinical trial—typically a recommendation whether to proceed with a following phase III trial—has to be based exclusively on the experimental data and does not depend on the prior elicitation. On the other hand, we recognized the need to use initial knowledge about the treatment and control regimens to design the trial appropriately.

A decision-theoretic approach can be considered as a formalization of another widely recommended step in clinical trial design, namely, simulation. Simulations of possible realizations of a clinical trial are widely recommended (Nestorov et al., 2001) and are usually conducted to compare competing designs. A simulation study has two major components: a set of alternative designs (d_j) and a set of probability models or scenarios (ϕ_j^o) . The purpose is to identify the design with the best overall performance. This approach implies an a priori guess about likely or potentially harmful scenarios ϕ_j^o and the ability to evaluate each combination of design d , scenario ϕ , and experimental outcomes X . A Bayesian decision-theoretic approach formalizes the evaluation of combinations of designs, scenarios, and outcomes by means of the utility function $u(d, \phi, X)$.

An explicit utility function is particularly useful when evaluating different experimental schemes. As an example, consider a comparison of an RDD and a two-arm randomized trial. If different costs are associated with the randomization of the patients at the beginning of the trial and with the randomization of a subgroup after the first stage, then one should not base the evaluation exclusively on power or average sample size. Recent RDD trials have shown exceptionally high accrual rates, reflecting strong patient preference for receiving the experimental drug initially, followed by possibly later randomization. A good utility

function will reflect the different nature of the up-front randomization in the two-arm randomized trial versus the later randomization in the RDD.

A limitation of the proposed decision process is the need for extensive prior elicitation when no historical data are available. A prior probability model and prior elicitation need to capture complicated aspects of the relationship of tumor mass and growth rate, the degree of predictability of a future tumor growth trajectory conditional on observations for an initial time interval, the degree of heterogeneity in the relevant patient population, and the level of certainty about any such judgments.

Finally, we indicate a possible enhancement of the proposed approach that would allow one to incorporate information from the first stage in decision making. Rosner et al. (2002) applied a Beta prior to the proportion of patients going on the randomization at T_1 , that is, p_e . If the observed proportion of patients going on to randomization is too small, then one would consider stopping the study. One could develop the tumor growth model further to allow for extrapolation from the early stage of the trial to the subsequent outcomes. The extrapolation within the context of a predictive distribution of the future outcomes could allow for a coherent early stopping decision. This is a matter for further investigation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This research was supported in part by grant CA075981 from the U.S. National Cancer Institute. We thank the CALGB and Ms Susan Sutherland for the historical tumor growth data.

References

- Albano G, Giorno V. A stochastic model in tumor growth. *Journal of Theoretical Biology*. 2006; 241:329–336. [PubMed: 16620871]
- Albano G, Giorno V. On the first exit time problem for a Gompertz-type tumor growth. *Lecture Notes in Computer Science*. 2009; 5717:329–336.
- Bursham, KP.; Anderson, DR. *Model Selection and Inference*. Springer-Verlag; New York: 1998.
- Capra W. Comparing the power of the discontinuation design to that of the classic randomized design on time-to-event endpoints. *Controlled Clinical Trials*. 2004; 25:168–177. [PubMed: 15020035]
- Clare S, Nakhli F, Panetta J. Molecular biology of breast metastasis: The use of mathematical models to determine relapse and to predict response to chemotherapy in breast cancer. *Breast Cancer Research*. 2000; 2:430–435. [PubMed: 11250737]
- DeGroot, M. *Optimal Statistical Decisions*. Wiley Interscience; New York: 2004.
- Fedorov, F.; Liu, T. Technical Report. Collegeville. Glaxo Smith Kline; Pennsylvania: 2005. Randomized discontinuation trials: Design and efficiency.
- Ferrante L, Bompadre S, Possati L, Leone L. Parameter estimation in a Gompertzian stochastic model for tumor growth. *Biometrics*. 2000; 56:1076–1081. [PubMed: 11129463]
- Freidlin B, Simon R. Evaluation of randomized discontinuation design. *Journal of Clinical Oncology*. 2005; 23:5094–5098. [PubMed: 15983399]
- Fu P, Dowlati A, Shluchter M. Comparison of power between randomized discontinuation design and upfront randomization design on progression-free survival. *Journal of Clinical Oncology*. 2009; 27:4135–4141. [PubMed: 19636018]
- Galsky MD, Zaks T, Hassani H, Vocila L, Sonpavde G, Hutson TE, Von Hoff DD. Target specific randomized discontinuation trial design: A novel approach in molecular therapeutics. *Investigational New Drugs*. 2009; 28:194–198. [PubMed: 19262991]

- Garthwaite PH, Kadane JB, O'Hagan A. Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*. 2005; 100:680–701.
- Gutiérrez R, Gutiérrez-Sánchez R, Nafidía A, Ramosa E. A new stochastic Gompertz diffusion process with threshold parameter: Computational aspects and applications. *Applied Mathematics and Computation*. 2006; 183:738–747.
- Korn EL, Arbuck SG, Pluda JM, Simon R, Kaplan RS, Christian MC. Clinical trial designs for cytostatic agents: Are new approaches needed? *Journal of Clinical Oncology*. 2001; 19:265–272. [PubMed: 11134222]
- Larkin JMG, Kipps ELS, Powell CJ, Swanton C. Systemic therapy for advanced renal cell carcinoma. *Therapeutic Advances in Medical Oncology*. 2009; 1:15–27. [PubMed: 21789110]
- Levy H. Stochastic dominance among log-normal prospects. *International Economic Review*. 1973; 14:601–614.
- Millar AW, Lynch KP. Rethinking clinical trials for cytostatic drugs. *Nature Reviews Cancer*. 2003; 3:540–545.
- Müller P, Parmigiani G. Optimal design via curve fitting of Monte Carlo experiments. *Journal of the American Statistical Association*. 1995; 90:1322–1330.
- Nestorov I, Graham G, Duffull S, Aarons L, Fuseau E, Coates P. Modeling and simulation for clinical trial design involving a categorical response: A phase II case study with Naratriptan. *Pharmaceutical Research*. 2001; 18:1210–1219. [PubMed: 11587494]
- Padgett W, Wei LJ. Bayes estimation of reliability for the two-parameter lognormal distribution. *Communications in Statistics*. 1977; 6:443–457.
- Ratain MJ, Eisen T, Stadler WM, Flaherty KT, Kaye SB, Rosner GL, Gore M, Desai AA, Patnaik A, Xiong HQ, Rowinsky E, Abbruzzese JL, Xia C, Simantov R, Schwartz B, O'Dwyer PJ. Phase II placebo-controlled randomized discontinuation trial of sorafenib in patients with metastatic renal cell carcinoma. *Journal of Clinical Oncology*. 2006; 24:2505–2512. [PubMed: 16636341]
- Ribba B, Saut O, Colin T, Bresch D, Grenier E, Boissel JP. A multiscale mathematical model of avascular tumor growth to investigate the therapeutic benefit of anti-invasive agents. *Journal of Theoretical Biology*. 2006; 243:532–541. [PubMed: 16930628]
- Rosner GL, Stadler W, Ratain MJ. Randomized discontinuation design: Application to cytostatic antineoplastic agents. *Journal of Clinical Oncology*. 2002; 20:4478–4484. [PubMed: 12431972]
- Schnatter S, Kaufmann S. Model-based clustering of multiple time series. *Journal of Business and Economic Statistics*. 2008; 26:78–89.
- Stadler WM. The randomized discontinuation trial: A phase II design to assess growth-inhibitory agents. *Molecular Cancer Therapeutics*. 2007; 6:1180–1185. [PubMed: 17431101]
- Stadler WM. Other paradigms: Randomized discontinuation trial design. *The Cancer Journal*. 2009; 15:431–434. [PubMed: 19826364]
- Stadler WM, Rosner G, Small E, Hollis D, Rini B, Zaentz SD, Mahoney J, Ratain MJ. Successful implementation of the randomized discontinuation trial design: An application to the study of the putative antiangiogenic agent carboxyaminoimidazole in renal cell carcinoma—CALGB 69901. *Journal of Clinical Oncology*. 2005; 23:3726–3732. [PubMed: 15923569]
- Therasse P, Arbuck SG, Eisenhauer EA, Wanders J, Kaplan RS, Rubinstein L, Verweij J, Van Glabbeke M, van Oosterom AT, Christian MC, Gwyther SG. New guidelines to evaluate the response to treatment in solid tumors. *Journal of the National Cancer Institute*. 2000; 92:205–216. [PubMed: 10655437]
- Waliszewski P. A principle of fractal-stochastic dualism and Gompertzian dynamics of growth and self-organization. *Biosystems*. 2005; 82:61–73. [PubMed: 16024163]

Figure 1.

Boxplots of the a priori distribution of the proportion p_e of patients eligible for the second stage of the trial across discontinuation trials having alternative durations of the first stage. The a priori distributions are computed through a simple data driven procedure; the adopted procedure for specifying the prior is discussed in Section 4. The lines in boxplots indicate the 5th, 25th, 50th, 75th, and 95th percentiles.

Figure 2.

Panel (i): The upper trajectory shows observed historical tumor growth data (■) under the control regimen. The lower trajectory shows imputed measurements and trajectory under the treatment regimen. Panels (ii) and (iii): Median trajectories, 80% confidence bands and 50% confidence bands of $X_{i,t}^1 \sim \mathbf{GP}(\tilde{a}_j - \tilde{\psi}_j \tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j | X_{i,t_0}^1 = \tilde{X}_{j,t_0}^1, \dots, X_{i,t_k} = \tilde{X}_{j,t_k}^1)$ under two alternative prior distributions for $\{\tilde{\psi}\}_{j=1}^M$.

Figure 3. Summary curves showing the distribution of projected tumor growth during the first 250 days after the tumor reaches the minimal threshold $\varepsilon = 0.02$. Solid lines: the median function and the 80% confidence band of the growth process X_t . Dashed lines: the 80% confidence band of the conditional expectations $E(X_t|\theta)$.

Figure 4. Optimal RDDs. Orthogonal sections of expected utility surfaces corresponding to alternative utility functions.

Table 1

Probabilities of rejecting the null hypothesis of no efficacy for the experimental therapy under hypothetical values of p_e , p_0 , and p_1 at the optimal design*
 $= (N^* = 216, T_1^* = 66, T_2^* = 102)$

		$p_e = 0.2$				$p_e = 0.4$				$p_e = 0.6$				$p_e = 0.8$			
		p_0				p_0				p_0				p_0			
p_1		0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8	0.2	0.4	0.6	0.8
0.4	0.30					0.57				0.75				0.86			
0.6	0.79	0.27				0.98	0.51			1.00	0.68			1.00	0.81		
0.8	0.99	0.79	0.30			1.00	0.98	0.57		1.00	1.00	0.75		1.00	1.00	0.87	

Table 2

Probabilities of rejecting the null hypothesis of no efficacy for the experimental therapy under hypothetical values of (p_0, p_1) for alternative up-front two-arm randomized trials. Sample size N varies between 100 and 250.

		$N = 100$			$N = 150$			$N = 200$			$N = 250$		
		p_0			p_0			p_0			p_0		
p_1		0.2	0.4	0.6	0.2	0.4	0.6	0.2	0.4	0.6	0.2	0.4	0.6
0.4	0.64				0.80			0.91			0.95		
0.6	0.99	0.54			0.99	0.72		1.00	0.85		1.00	0.92	
0.8	1.00	0.99	0.64		1.00	0.99	0.80	1.00	1.00	0.90	1.00	1.00	0.95