Biometrics , 1–??
2010
DOI: 000

Bayesian Enrichment Strategies for Randomized Discontinuation Trials

Lorenzo Trippa[1,*], Gary L. Rosner[2,], Peter Müller[3,***]**

[1] Harvard School of Public Health and

Department of Biostatistics, Dana–Farber Cancer Institute Boston, Massachusetts 02115

[2]Oncology Biostatistics/Bioinformatics, Sidney Kimmel Comprehensive Cancer Center at Johns Hopkins, Baltimore, Maryland 21205

[3]Department of Biostatistics, University of Texas M. D. Anderson Cancer Center, Houston, Texas 77030

email: lorenzo.trippa@jimmi.harvard.com

**email:* grosner@jhmi.edu

***email:* pmueller@mdanderson.org

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 putative cytostatic agents out of tumor growth considerations, specify a suitable utility function, and develop a computational procedure for selecting the optimal tuning parameters.

Key words: Clinical trials; Enrichment designs; Randomized discontinuation design; Tumor growth models.

1. Introduction

We propose a decision-theoretic procedure for choosing an optimal randomized discontinuation design (RDD). We use a tumor growth model to specify a prior on the 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 et al., 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 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 effect of cytostatic agents. Moreover, some agents may be active for a subgroup of the target population. Millar and Linch (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). But 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 pre-randomization 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 et al. (2009), Galsky et. al (2009), Millar and Linch (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. To this end, 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 $p(\phi)$ 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 $p(\phi)$ 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 as samples from the 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 framework or 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

Over the last 5 years, several trials have investigated the use of targeted agents to treat patients with advanced renal cell carcinoma (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 paper, we discuss designing a new study using a 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.

[Figure 1 about here.]

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 Gompertzian diffusion process (Ferrante et al., 2000).

The Gompertz model solves the differential equation

$$dX_t/dt = a \cdot X_t - b \cdot X_t \cdot \log(X_t) \qquad X_0 = x_0, \tag{1}$$

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 et al. (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.):

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

where $\{W_{it}\}_{i \geqslant 1}$ 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:

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

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).

If two therapeutic regimens are involved, the resulting process will solve the same s.d.e. with

a_i and b_i substituted by two piecewise-constant functions: $a_i(t) = a_i^0 I(t \in T_0) + a_i^1 I(t \in T_1)$ and $b_i(t) = b_i^0 I(t \in T_0) + b_i^1 I(t \in T_1)$, where T_0 and T_1 are contiguous time intervals during which the two therapies are administered. 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:

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

We note that the process has marginal lognormal distributions and that we can easily sample from the marginal distributions. Expression (3) provides a complete and synthetic definition of the process; for our purpose we will not need other theoretical results on this model.

The parametrization of the Markov process, when the patent is treated with two distinct therapies, can be slightly simplified. We can argue that $b_i^0 = b_i^1 = b_i$ 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 expectations

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

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 $p(X_{t+s} \leqslant x_{t+s} \mid X_t = x_t, a^1, b^1, \sigma) \geqslant p(X_{t+s} \leqslant x_{t+s} \mid X_t = x_t, a^0, b^0, \sigma)$ 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 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 ($E_i = 0$), $a_i^0 = a_i^1, b_i$ 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 the 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,

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

regulate the tumor growth process under $E_i = 0$. We observe that φ_{1i}, φ_{2i} and φ_{3i} have a clear interpretation which follows from the equalities:

$$E\Big(\log(X_{is}) \mid X_{i0}, a_i^0, b_i, \sigma_i, E_i = 0\Big) = e^{-b_i s}\log(X_{i0}) + \frac{a_i^0 - \sigma^2/2}{b_i}(1 - e^{-b_i s}) \quad \text{and}$$

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

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

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

Let $\text{Ga}(x;\ \nu, r)$ indicate a Gamma distributed random variable with mean ν/r. Let $\text{NIG}(\theta_1, \theta_2;\ m, v^2, \nu, r)$ 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 $(\theta_1 \mid \theta_2) \sim N(m, \theta_2 v^2)$. We complete the model with lognormal (LN) random effects distributions for $\varphi_{1i}, \varphi_{2i}, \varphi_{3i}$ and ψ_i,

$$\varphi_{ji} \mid \theta \sim \text{LN}(\theta_{j1}, \theta_{j2}), \text{ and } \psi_i \mid \theta, E_i \sim \begin{cases} 0 & \text{if } E_i = 0 \\ \text{LN}(\theta_{41}, \theta_{42}) & \text{if } E_i = 1 \end{cases}$$

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 $\theta = (\theta_{j1}, \theta_{j2};\ j = 1, \ldots, 4)$:

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

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 Jeffreys' prior for lognormal 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 for example Garthwaite el al. (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 judgement 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 effects ψ_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 fact that we use the interpolation leaves the proposed strategy very robust against possible deviations from the model assumptions. The architecture of the procedure is easy to be interpreted and implemented. We assume a probability model consistent with an extensive literature on solid tumors growth. The historical trajectories are imputed on the basis of discrete time observations. In the end, we have a prior distribution for tumor growth trajectories that represents a priori beliefs about the treatment effect and the historical data through imputation. 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:

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

We use $\tilde{X}$ to distinguish the historical data from future experimental observations X. The restriction to fixed measurement times $(t_1, \ldots, t_k)$ across patients is not important. We only use this assumption to simplify presentation. The proposed procedure can easily be adapted to historical data with heterogeneous follow up and data observed at different times across patients. 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 $\mathbf{GP}(a, b, \sigma|\triangle)$ be the conditional law of the process given the event $\{\triangle\}$. We assume

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

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 (Gutierretz et al., 2006).

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

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

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

$$\log(\tilde{X}^1_{jt_{\ell+1}}) = \gamma(\tilde{b}_j, t_\ell - t_{\ell+1}) \log(\tilde{X}^1_{jt_\ell}) + \beta(\tilde{a}_j - \tilde{\psi}_j \tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j, t_\ell - t_{\ell+1}) + \delta(\tilde{b}_j, \tilde{\sigma}_j, t_\ell - t_{\ell+1}) Z_{j\ell}. \quad (7)$$

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 is representative of the expectations on the cytostatic activity of the treatment.

The hypothetical observations $(\tilde{X}^1_{jt_0}, \ldots, \tilde{X}^1_{jt_k};\ j = 1 \ldots, M)$ are defined as the random solutions of equations (7) assuming that for every historical patient $\tilde{X}^0_{jt_0} = \tilde{X}^1_{jt_0}$. Finally, conditional on $\tilde{X}^1 = (\tilde{X}^1_{jt_0}, \ldots, \tilde{X}^1_{jt_k};\ j = 1 \ldots, M)$ and on the random vector $\tilde{\psi} = (\tilde{\psi}_1, \ldots, \tilde{\psi}_M)$,

future tumor growth trajectories under the treatment regimen are characterized as follows:

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

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 (5,8) are estimated through the standard maximum likelihood technique and the random quantities $\{\tilde{\psi}_j\}_{j=1}^M$ are modeled through the prior discussed in Section 3. We impute, for each historical patient, the hypothetical observations at $(t_1, \ldots, t_k)$ and T_1. For each patient j the distribution of $\tilde{X}^0_{jT_1}$ conditional on $\{\tilde{X}^0_{jt_0}, \tilde{X}^0_{jt_2}, \ldots, \tilde{X}^0_{jt_k}\}$ is computed. We then slightly modify the equations (6) and (7) by adding to the vector $(t_0, t_1, \ldots, t_k)$ the additional observation time T_1, and substitute $\beta(\tilde{a}_j - \psi_i\tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j, t_\ell - t_{\ell+1})$ in (7) with

$$\beta(\tilde{a}_j - \psi_i\tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j, t_\ell - t_{\ell+1})I(t_{\ell+1} \leqslant T_1) + \beta(\tilde{a}_j, \tilde{b}_j, \tilde{\sigma}_j, t_\ell - t_{\ell+1})I(t_{\ell+1} > T_1).$$

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 $(\tilde{X}^0_{1,T_1}, \ldots, \tilde{X}^0_{M,T_1})$ 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 simulation the predictive probability that the trial detects the effect of the agent for several assumed levels of efficacy of the new agent.

[Figure 2 about here.]

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

$$U(d) = \int \int u(d, \phi, X)\, p(\mathrm{d}X \mid \phi, d)\, p(\mathrm{d}\phi). \tag{9}$$

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).

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

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

We assume that the recommendation for phase III is 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. When the agent is recommended for phase III, we assume that the benefits are proportional to expected efficacy of the therapy in a future $(N+1)$-st patient. We use

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

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, while 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

$$\delta_1(X_{i0}, T_1) \geqslant X_{iT_1} \geqslant \delta_2(X_{i0}, T_1).$$

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. 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 before, 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.

In order to efficiently obtain approximations of $\tilde{U}(d)$ for every design d the algorithm carry out one (or few) simulation for H designs d_h, $h = 1, \ldots, H$, and fit 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 borrowing strength from results associated with neighboring designs. In summary,

Step 1. Select H designs $\{d_h \in D, h = 1, \ldots, H\}$, 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), $h = 1, \ldots, H$. 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 polyonomials 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 Sections 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 $\delta_2(X_0, T_1) = 0.7X_0$. This corresponds to the RECIST response criteria (Therasse et al., 2000) used in many oncology trials, such as Stadler et al. (2005). The primary end point 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 $D = \{(N, T_1, T_2) \in (0, 1, \ldots, 400) \times (0, 1, \ldots, 300)^2\}$. In particular, the alternative of not performing the trial ($N = 0$), as well as a traditional two-arm randomized trial ($T_1 = 0$) are options.

[Figure 3 about here.]

The optimal strategies corresponding to alternative utility functions are determined using Monte Carlo samples of $H = 10^7$ simulated trials with tuning parameters $(d_h; h = 1, \ldots, H)$ 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. The efficiency of the outlined computational procedure (Müller and Parmigiani, 1995) is particularly helpful for comparing the optimal designs under different utility functions and prior distributions. Such comparisons provide the investigator with a clear representation of the sensitivity of the decision framework.

[Figure 4 about here.]

In the example illustrated in Figure 4, a rational decision maker favors the RDD under

all three considered utility functions. The optimal discontinuation strategy $d^\star$ guarantees a superior expected utility compared to a standard two-arm trial designs. 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 $d^* = (N^* = 216, T_1^* = 66, T_2^* = 102)$ for hypothetical values of p_e, p_0 and p_1.

[Table 1 about here.]

We can then assess the operating characteristics of the optimal design d^* by comparisons with two arm randomized trials. Table 2 shows the power of alternative upfront randomized trials for alternative 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 whole population of patients, as with upfront randomization. A feature of the RDD is that the first phase of the design selects a subpopulation of patients likely of benefitting from the novel agent. This selection mechanism results in substantial differences between the response probabilities (p_0, p_1) of a RDD evaluating a cytostatic agent compared to those of an upfront randomized trial; several examples are given in Fedorov and Liu (2005). The upfront randomization, in most cases, corresponds to narrower differences between p_0 and p_1.

[Table 2 about here.]

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 renal cell carcinoma (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 ψ_i is parameterized such that, conditionally on $\{E = 1\}$, the mean 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 binary. 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 $D = \{d = (N, T_1, T_2) \in (0, 1, \ldots, 300)^3\}$. 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 operative 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 in evaluating the operative characteristics for possible parameterization of the probability model. We computed, for example, the design 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 power decreases as the susceptible subpopulation proportion π decreases; power values

above the 75% are observed for $\pi \geqslant 50\%$. A synthetic table reporting these computations is provided in the online 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. Several recent studies have shown potential benefits of RDD protocols over alternative up-front randomization, but also cautioned that 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 a randomized discontinuation design (RDD). We assumed that the final result of the clinical trial—typically the decision to perform a phase III clinical trial or to consider the novel treatment lacks activity—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, viz., 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 new 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, at least 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 judgement. Various simplifications in the proposed probability model mitigate some of these difficulties.

We finally 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.

Acknowledgements

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

References

Albano,G., Giorno, V. (2006) A stochastic model in tumor growth. *Journal of theoretical biology.* Vol. 241, 329-336

Albano,G., Giorno, V. (2009) On the first exit time problem for a Gompertz-type tumor growth. *Lecture notes in computer science* Vol. 5717, 329-336

Bursham, K.P., Anderson, D.R. (1998) *Model Selection and Inference.* Springer-Verlag, New York.

Capra, W. (2004) Comparing the power of the discontinuation design to that of the classic randomized design on time-to-event endpoints. *Controlled clinical trials*, Vol. 25, 168-177

Clare, S., Nakhlis, F., Panetta, J. (2000) Molecular biology of breast metastasis: The use of mathematical models to determine relapse and to predict response to chemotherapy in breast cancer. *Brest Cancer Research*, Vol. 2, 430-435

DeGroot, M. (2004) *Optimal Statistical Decisions.* New York:Wiley-Interscience.

Fedorov, F., Liu, T. (2005) Randomized Discontinuation Trials: Design and Efficiency. *Glaxo Smith Kline Technical Report* .

Ferrante, L., Bompadre, S., Possati, L., Leone, L. (2000) Parameter estimation in a gompertzian stochastic model for tumor growth. *Biometrics* 56, 1076-1081

Freidlin, B., Simon, R. (2005) Evaluation of randomized discontinuation design. *Journal of clinica oncology*, Vol. 23, 5094-5098

Fu, P., Dowlati, A., Shluchter, M (2009) Comparison of power between randomized discon-

tinuation design and upfront randomization design on progression-free survival. *Journal of clinica oncology*, Vol. 27, 4135-4141

Garthwaite, P.H., Kadane, J.B., O'Hagan, A. (2005) Statistical Methods for Eliciting Probability Distributions. *Journal of the American Statistical Association*, Vol.100, No 470, 680-701

Galsky, M.D. et. al (2009) Target specific randomized discontinuation trial design: a novel approach in molecular therapeutics. *Investigational new drugs*, in press

Gutirrez, R., Gutirrez-Sncheza, R., Nafidia, A., Ramosa, E.(2006) A new stochastic Gompertz diffusion process with threshold parameter: Computational aspects and applications. *Applied Mathematics and Computation*, 183, 738-747

Korn, E.L. et al. (2001) Clinical Trial Designs for Cytostatic Agents: Are New Approaches Needed? *Journal of Clinical Oncology* Vol. 19, No. 1, 265-272

Larkin, . et al. (2009) Sisthemic therapy for advanced renal cell carcinoma. *Therapeutic Advances in Medical Oncology* Vol. 1, 15-27

Levy, H. (1973) Stochastic dominance among Log-Normal prospects. *International Economic Review*, Vol. 14, No. 3, 601-614

Millar, A.W., Lynch, K.P. (2003) Rethinking clinical trials for cytostatic drugs. *Nature Reviews Cancer*, Vol. 3; No. 7, 540-545

Müller, P., Parmigiani, G. (1995) Optimal Design Via Curve Fitting of Monte Carlo Experiments. *Journal of the American Statistical Association*, Vol. 90, No 432, 1322-1330

Nestorov, I. et al.(2001) Modeling and Simulation for Clinical Trial Design Involving a Categorical Response: A Phase II Case Study with Naratriptan. *Pharmaceutical Research* Vol.18, No.8, 1210-1219

Padgett, W., Wei, L.J. (1977) Bayes estimation of reliability for the two-parameter lognormal

distribution. *Communications in Statistics*, Vol. 6, No. 5, 443 - 457

Ratain, M. et al. (2006) Phase II Placebo-Controlled Randomized Discontinuation Trial of Sorafenib in Patients With Metastatic Renal Cell Carcinoma. *Journal of Clinical Oncology*, Vol 24, No. 16, 2505-2512.

Ribba, B., Sautb, O., Colinb ,T., Breschc, D., Grenierd, E., Boissela, J.P. (2006) A multiscale mathematical model of avascular tumor growth to investigate the therapeutic benefit of anti-invasive agents. *Journal of Theoretical Biology*, Vol. 243, No.4, 532-541

Rosner G.L., Stadler W., Ratain M.J. (2002) Randomized discontinuation design: Application to cytostatic antineoplastic agents. *Journal of Clinical Oncology* Vol. 20, No.22, 4478-4484

Schnatter, S., Kaufmann, S. (2008) Model-Based Clustering of Multiple Time Series. *Journal of Business & Economic Statistics*, Vol. 26, No.1, 78-89

Stadler W.M. et al. (2005) 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*, Vol.23, No 16, 3726-3732

Stadler, W.M. (2007) The randomized discontinuation trial: a phase II design to assess growth-inhibitory agents. *Molecular Cancer Therapeutics*, Vol 6, 1180-1185

Stadler, W.M. (2009) Other paradigms: randomized discontinuation trial design. *The Cancer Journal*, Vol 15, 431-434

Therasse, P., Arbuck, S.G., Eisenhauer, E.A., et al.(2000) New Guidelines to Evaluate the Response to Treatment in Solid Tumors. *Journal of the National Cancer Institute* 92, 205-216.

Waliszewski, P., (2005) A principle of fractal-stochastic dualism and Gompertzian dynamics of growth and self-organization. *Byosistems,* Vol.82, 61-73

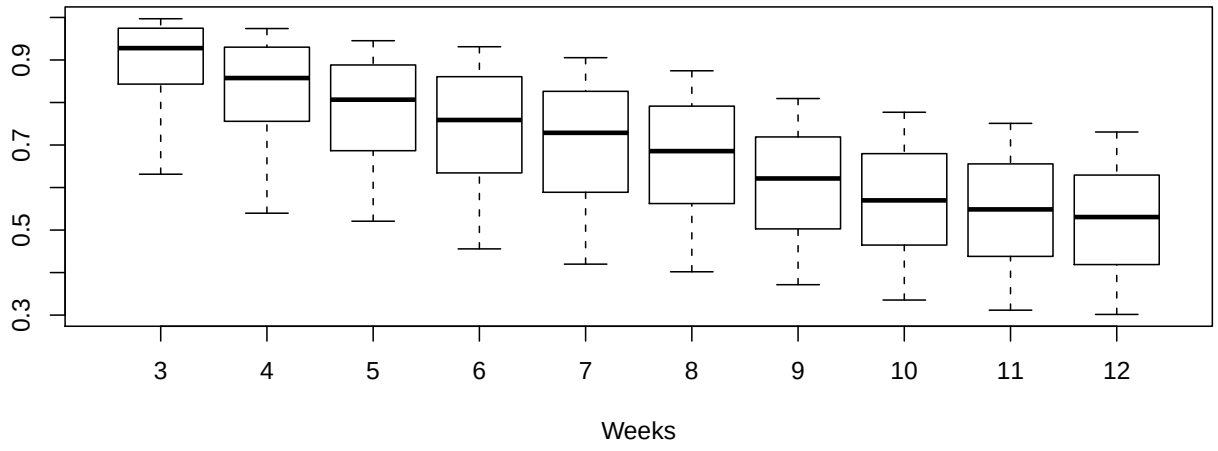

Figure 1. Boxplots of the *a priori* distribution of the proportion p_e of patients eligible for the second stage of the trial across discontinuation trials having alternative durations of the first stage. The *a priori* distributions are computed through a simple data driven procedure; the adopted procedure for specifying the prior is discussed in Section 4. The lines in boxplots indicate the 5-th ,25-th, 50-th, 75-th and 95-th percentiles.

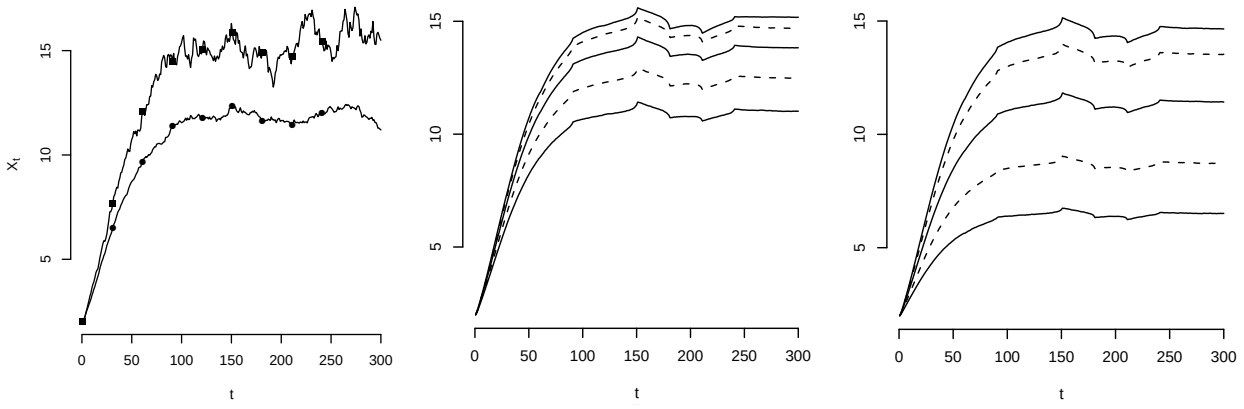

(i) (ii) (iii)

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 $X^1_{i,t} \sim \mathbf{GP}(\tilde{a}_j - \tilde{\psi}_j\tilde{b}_j, \tilde{b}_j, \tilde{\sigma}_j \mid X^1_{i,t_0} = \tilde{X}^1_{j,t_0}, \ldots, X_{i,t_k} = \tilde{X}^1_{j,t_k})$ under two alternative prior distributions for $\{\tilde{\psi}_j\}_{j=1}^M$.

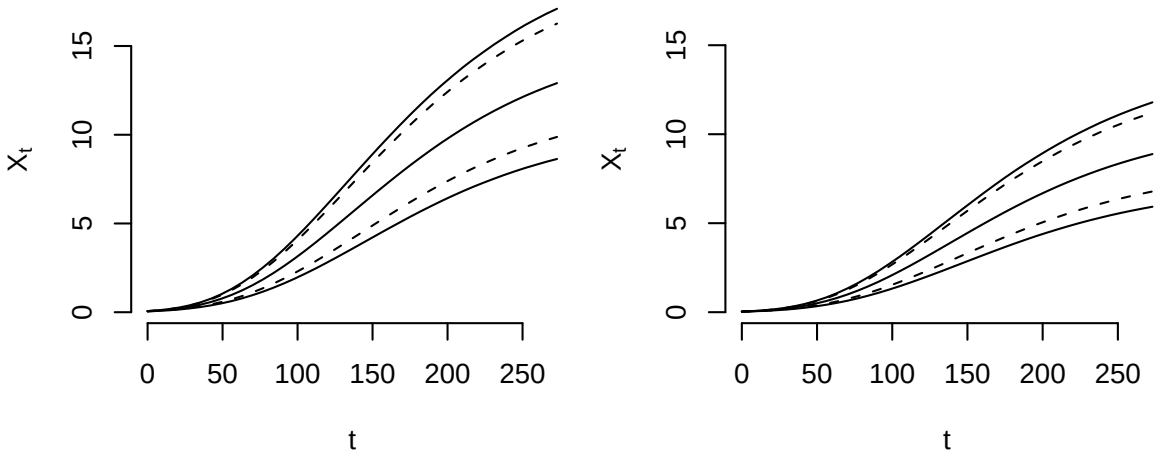

Control, $E = 0$ Effective treatment, $E = 1$

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 $E(X_t \mid \theta)$.

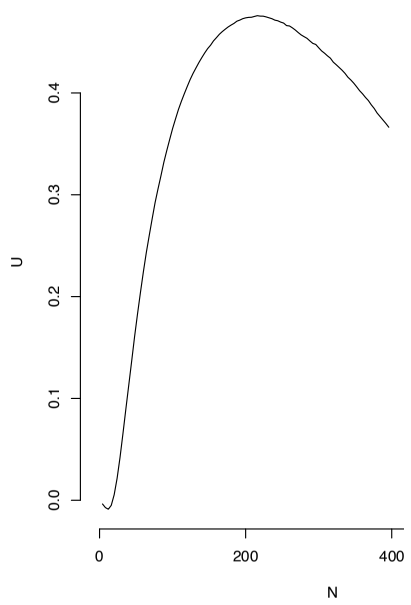

$c_1 = c_2 = 0.001, c_3 = 3.3\,10^{-5}$

$d^\star = (216, 60, 102)$

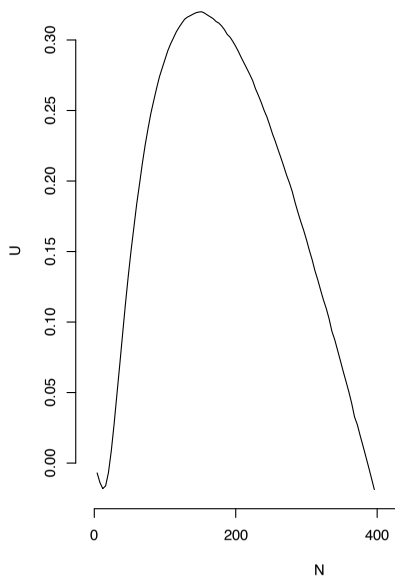

$c_1 = c_2 = 0.002, c_3 = 6.6\,10^{-5}$

$d^\star = (152, 56, 105)$

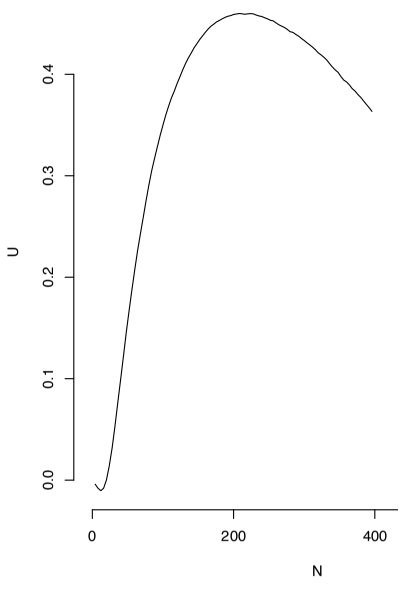

$c_1 = 0.004, c_2 = c_3 = 0$

$d^\star = (204, 62, 108)$

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

	$p_e = 0.2$			$p_e = 0.4$			$p_e = 0.6$			$p_e = 0.8$		
		p_0			p_0			p_0			p_0	
p_1	0.2	0.4	0.6	0.2	0.4	0.6	0.2	0.4	0.6	0.2	0.4	0.6
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 1

Probabilities of rejecting the null hypothesis of ineffectiveness of the novel treatment for hypothetical values of p_e, p_0 and p_1 at the optimal design $d^* = (N^* = 216, T_1^* = 66, T_2^* = 102)$.

	$N = 100$			$N = 150$			$N = 200$			$N = 250$		
		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

Table 2

Probabilities of rejecting the null hypothesis of ineffectiveness of the novel treatment for hypothetical values of (p_0, p_1) *for alternative upfront randomize trials by sample size per treatment.*