BIOMETRICS 58, 560–568
September 2002

Adaptive Decision Making in a Lymphocyte Infusion Trial

Peter F. Thall,[1,*] Lurdes Y. T. Inoue,[1] and Thomas G. Martin[2]

[1]Department of Biostatistics, The University of Texas M. D. Anderson Cancer Center, Box 447, 1515 Holcombe Boulevard, Houston, Texas 77030, U.S.A.
[2]Department of Hematology/Oncology, University of California at San Francisco, Box 0324, San Francisco, California 94143, U.S.A.
*email: rex@mdanderson.org

SUMMARY. We describe an adaptive Bayesian design for a clinical trial of an experimental treatment for patients with hematologic malignancies who initially received an allogeneic bone marrow transplant but subsequently suffered a disease recurrence. Treatment consists of up to two courses of targeted immunotherapy followed by allogeneic donor lymphocyte infusion. The immunotherapy is a necessary precursor to the lymphocyte infusion, but it may cause severe liver toxicity and is certain to cause a low white blood cell count and low platelets. The primary scientific goal is to determine the infusion time that has the highest probability of treatment success, defined as the event that the patient does not suffer severe toxicity and is alive with recovered white blood cell count 50 days from the start of therapy. The method is based on a parametric model accounting for toxicity, time to white blood cell recovery, and survival time. The design includes an algorithm for between-patient immunotherapy dose de-escalation based on the toxicity data and an adaptive randomization among five possible infusion times according to their most recent posterior success probabilities. A simulation study shows that the design reliably selects the best infusion time while randomizing greater proportions of patients to superior infusion times.

KEY WORDS: Adaptive randomization; Bayesian design; Dose finding; Leukemia; Lymphocyte infusion.

1. Introduction

Patients with hematologic malignancies who have undergone allogeneic bone marrow transplantation (alloBMT) but whose disease has recurred have a very poor prognosis with conventional treatments. We describe an adaptive Bayesian design for a clinical trial of an experimental therapeutic strategy for patients who previously have received an alloBMT but who have suffered a disease recurrence more than 90 days after the transplant. Patients with acute myelogenous leukemia, advanced chronic myelogenous leukemia, or myelodysplastic syndrome will be included. Prognostically, the fact that the patient's disease has recurred after alloBMT supersedes the type of malignancy diagnosed initially. The 1-year survival rate of these relapsed patients is only about 5–10%, regardless of the type of malignancy. They also exhibit very similar response rates, of approximately 20–40%, with a variety of salvage chemotherapy regimens. The trial aims to evaluate the efficacy and safety of targeted immunotherapy followed by allogeneic donor lymphocyte infusion (DLI) while also attempting to optimize the interval between administration of targeted immunotherapy and DLI, hereafter referred to as DLI time. The duration of the DLI is only about 20 minutes. The rationale for this study stems from previous trials demonstrating that DLI can induce remissions in patients with acute and chronic leukemias relapsing after alloBMT (Antin, 1993; Collins et al., 1997). The donor T-cells can recognize antigens on the surface of the malignant cells and directly induce cell death. DLI, however, appears to be most successful in patients who have low tumor burden, known as minimal residual disease. The goal of the targeted immunotherapy is to induce minimal residual disease, thus providing the DLI with a better chance of success.

Mylotarg is an engineered monoclonal antibody attached to chalicheamicin, a potent cellular toxin (Bernstein, 2000). Mylotarg binds to the cell surface receptor CD33 located on immature blood progenitor cells, both normal and leukemia cells. Thus, an entry criterion for the trial is that the patient's cancer cells must express the CD33 cell surface antigen. Mylotarg has been administered to patients with relapsed acute myelogenous leukemia, with overall response rates of approximately 20–30% (Sievers et al., 1999). Its main side effects are liver dysfunction (hepatotoxicity, HT), infusion-related fevers and chills, prolonged neutropenia (low absolute neutrophil count, ANC, defined as <1000 cells/μl), and thrombocytopenia (low platelets). Severe HT is the adverse event of greatest concern with Mylotarg because it is potentially life threatening (Giles et al., 2001; Neumeister et al., 2001) and it is more likely to be severe in patients who have received prior stem cell transplantation than in those who have not. In previous clinical trials of Mylotarg, infusion-related fever and chills have been controlled through the administration of steroids prior to Mylotarg. Very few other severe adverse

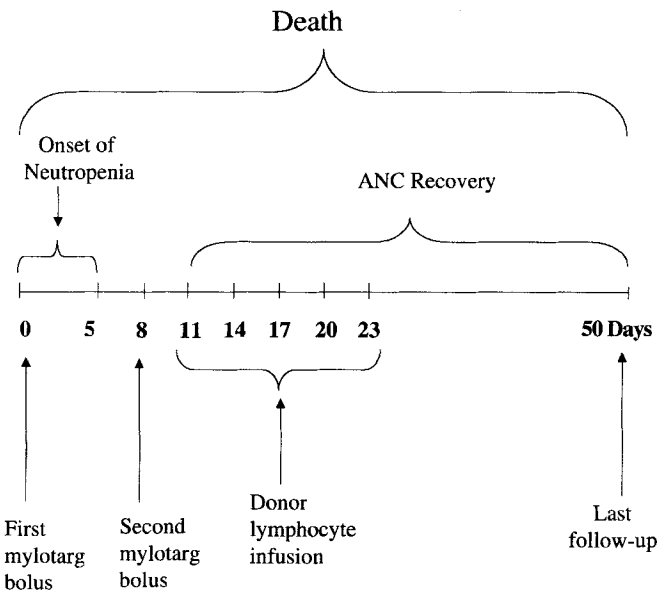

Figure 1. Treatment plan and patient outcomes for the DLI trial.

events were encountered, and almost all were resolved before day 50. Although HT is the most likely adverse event, in the sequel, we include any severe toxicity that cannot be resolved therapeutically, other than low blood counts, in the definition of HT.

The average time to achieve ANC >1000 cells/μl following Mylotarg is more than 40 days. In this study, mobilized DLI will be given following Mylotarg in an attempt to hasten neutrophil recovery. Usually, donor lymphocytes are collected without giving the donor white blood cell (WBC) growth factors. This unstimulated lymphocyte collection contains mostly lymphocytes and very few immature stem cells. However, if the donor is given WBC growth factors prior to lymphocyte collection, then this (mobilized) collection will contain a much larger proportion of immature stem cells (Lane et al., 1999). Mobilized DLI is usually given to patients who receive chemotherapy for disease relapse following alloBMT. The goals of mobilized DLI are to provide lymphocytes capable of fighting leukemia and immature stem cells capable of hastening bone marrow recovery.

Figure 1 illustrates the treatment plan and clinically relevant patient outcomes. Therapy begins with a bolus of Mylotarg aimed at reducing the patients' leukemic burden. The patient is then monitored for the following week for HT. If no HT is encountered, then a second dose of Mylotarg is administered. Mobilized DLI will be given at a defined interval following the second dose of Mylotarg. Patients who develop HT after the first dose of Mylotarg will not receive a second dose and will not receive scheduled DLI because one dose of Mylotarg likely will not result in prolonged neutropenia. These patients will be taken off study and may receive DLI at a later time once HT has been resolved. Patients who experience HT after the second course of Mylotarg will receive DLI as scheduled because, at that point, DLI is the only remaining therapeutic avenue and likely will hasten neutrophil recovery. Due to the effects of the first course of Mylotarg, all patients will suffer neutropenia, defined as an ANC falling below 1000, within 5 days of drug infusion. Neutrophils are white blood cells that play an important role in fighting infections. The dual therapeutic goals of the mobilized DLI are to enhance neutrophil recovery and kill remaining cancer cells. Restoring the patient's ANC is critical because, during the period of neutropenia, the patient is highly susceptible to life-threatening bacterial and fungal infections. Treatment success is defined as the event that the patient does not suffer HT and is alive with ANC $\geq$1000 at day 50.

The primary scientific goal of the trial is to determine the DLI time, among five predetermined times after the Mylotarg, having the highest treatment success probability. Our design is based on a parametric Bayesian model for the probability of HT as a function of Mylotarg dose and for the times to ANC recovery and death, each as a function of infusion time and whether the patient has experienced HT. We take a model-based approach because decisions must be made adaptively during the trial based on very limited information, and this problem is especially acute early in the trial. The design has two main components: an algorithm for between-patient de-escalation among six Mylotarg dose pairs if an unacceptably high HT rate is observed and a sequentially adaptive randomization of patients among the infusion times. Depending on the interim HT data and the dose de-escalation algorithm, each patient may receive any of five possible dose combinations in the two courses. Each patient's randomization probabilities for the five infusion times are based on their current posterior success probabilities. Thus, each patient is more likely to be infused at a time having higher current success rate, based on the most recent data from the trial. This method is used in place of balanced randomization because it yields unbiased comparisons among the infusion times while also providing each patient, with high probability, the current highest success rate. Because the therapeutic goal of the treatment being studied in this trial is to save patients for whom initial treatment has failed, it is an example of a salvage therapy.

We provide a detailed description of the patient outcomes and probability model in Section 2. The trial design is presented in Section 3, followed by a summary of a simulation study of the design in Section 4. We close with a discussion in Section 5.

2. Probability Model

2.1 *Patient Outcomes*

The indicator of HT in the jth Mylotarg course is denoted by Y_j for $j = 1, 2$. Because the patient's targeted immunotherapy is terminated if there is HT in the first course, Y_2 is defined only if $Y_1 = 0$, so that the three possible HT outcomes are $Y_1 = 1$, $(Y_1, Y_2) = (0, 1)$, and $(Y_1, Y_2) = (0, 0)$. We let Y_+ indicate HT in either course, which is either of the first two outcomes. We assume that the onset time of neutropenia, T_N, is uniformly distributed on the interval from 0 to 5 days. Denoting the infusion time by t_I and the time from infusion to ANC recovery by T_A, the patient's time from start of therapy to ANC recovery is $t_I + T_A$ and total duration of neutropenia is $t_I + T_A - T_N$.

The primary scientific goal of this study is to choose the best DLI time among 11, 14, 17, 20, and 23 days after the initial Mylotarg bolus. This is motivated by the potential

trade-off between giving the DLI as soon as possible after Mylotarg to hasten neutrophil recovery or potentially giving the DLI too soon after Mylotarg such that residual circulating immunotoxin can bind to the donor stem cells and further delay ANC recovery. Ideally, patients achieve ANC recovery within 14 days after DLI. Therefore, the neutropenic interval, $t_I + T_A - T_N$, begins within 5 days of the initial Mylotarg infusion and lasts until approximately 14 days following DLI. Because the half-life of Mylotarg is approximately 60–90 hours, giving DLI soon after Mylotarg may prolong neutrophil recovery to much longer than 14 days, the standard interval following DLI. Alternatively, waiting longer to infuse the patient to give more time for Mylotarg clearance may also prolong the neutropenia and thus increase the risk of death due to infection. This study is designed to identify the DLI time, t_I, that is associated with the shortest neutropenic interval, T_A, while maintaining survival. Formally, clinical success is defined as the event that the patient does not suffer HT within the first 8 days, recovers ANC $\geq$1000, and survives to day 50. Denoting the time of death by T_D, the success probability of a patient infused at time t_I is $\theta(t_I) = \Pr(Y_1 = 0$ and $t_I + T_A < 50 < T_D \mid t_I)$.

2.2 *Regression Models*

In addition to the fact that each patient receives either one or two courses of Mylotarg, the trial design includes an algorithm for Mylotarg dose de-escalation between patients, described in Section 4. Consequently, the Mylotarg doses given in courses 1 and 2, denoted d_1 and d_2, may vary between patients. We assume that the probability of HT depends on the cumulative dose of Mylotarg according to the logistic model

$$\pi_1(d_1, \gamma) = \Pr(Y_1 = 1 \mid d_1, \gamma) = \text{logit}^{-1}(\gamma_0 + \gamma_1 d_1) \quad (1)$$

and

$$\begin{aligned}\pi_2(d_+, \gamma) &= \Pr(Y_2 = 1 \mid Y_1 = 0, d_1, d_2, \gamma) \\ &= \text{logit}^{-1}(\gamma_0 + \gamma_1 d_+),\end{aligned} \quad (2)$$

where $\text{logit}(p) = \log\{p/(1-p)\}$, $d_+ = d_1 + d_2$, and $\gamma = (\gamma_0, \gamma_1)$. The linear components of π_1 and π_2 have identical parameters and differ only in terms of Mylotarg dose, reflecting the assumption that the risk of toxicity is a function of cumulative dose. Recalling that (Y_1, Y_2) take on three possible values, the overall probability that a patient has HT is $\pi_+ = \pi_1 + (1-\pi_1)\pi_2$, with $(1-\pi_+) = (1-\pi_1)(1-\pi_2)$.

To reduce collinearity among the parameter estimates in the following models for regression of the means of T_A and T_D on t_I, we use the standardized infusion time $x_I = (t_I - 17)/3$, which takes on the five possible values $\{-2, -1, 0, 1, 2\}$. For time from infusion to ANC recovery, T_A, denoting $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2)$, we assume an exponential distribution with mean

$$\mu_A(t_I, \boldsymbol{\beta}) = \exp(\beta_0 + \beta_1 x_I + \beta_2 x_I^2), \quad (3)$$

so that T_A has probability density function (p.d.f.)

$$f_A(t \mid t_I, \boldsymbol{\beta}) = {\mu_A}^{-1} \exp\{-(t - t_I)/\mu_A\}, \qquad t > t_I. \quad (4)$$

We define this distribution conditional on $Y_1 = 0$ to limit attention to patients who are randomized among the five infusion times because this is the methodological and scientific focus. Moreover, patients are unlikely to suffer HT in course 1.

To account for the effects of toxicity and lymphocyte infusion time on survival, we assume that T_D follows a piecewise exponential distribution with mean

$$\mu_{D,1}(Y_+, \boldsymbol{\alpha}) = \exp(\alpha_0 + \alpha_1 Y_+) \quad (5)$$

before infusion and, if the patient survives long enough to be infused, mean

$$\mu_{D,2}(t_I, Y_+, \boldsymbol{\alpha}) = \exp(\alpha_2 + \alpha_3 x_I + \alpha_4 x_I^2 + \alpha_1 Y_+) \quad (6)$$

thereafter, denoting $\boldsymbol{\alpha} = (\alpha_0, \cdots, \alpha_4)$, i.e., the preinfusion baseline death rate α_0 is replaced by the quadratic $\alpha_2 + \alpha_3 x_I + \alpha_4 x_I^2$ after the patient is infused. Thus, T_D has p.d.f.

$$\begin{aligned} &f_D(t \mid t_I, \boldsymbol{\alpha}) \\ &\quad = \begin{cases} \mu_{D,1}^{-1} \exp(-t/\mu_{D,1}) & 0 < t \leq t_I, \\ \exp(-t_I/\mu_{D,1})\mu_{D,2}^{-1} & \\ \quad \times \exp\{-(t - t_I)/\mu_{D,2}\} & t > t_I. \end{cases} \end{aligned} \quad (7)$$

In particular, patients who suffer HT in the first course are subject to the death rate $\mu_{D,1}$ thereafter because they are not infused. For convenience, we will denote the p.d.f. with parameter $\mu_{D,j}$ by $f_{D,j}$ for $j = 1, 2$.

We assume log quadratic functions of x_I for both $\mu_A(t_I, \boldsymbol{\beta})$ and $\mu_{D,2}(t_I, Y_+, \boldsymbol{\alpha})$. This is a flexible family of functions allowing a broad range of possible ways in which T_A and T_D each may depend on infusion time. We assume that T_A and T_D are independent because any probability model including a parameter characterizing their association is not identifiable (Prentice et al., 1978). The piecewise exponential model specifies mean postinfusion survival time as an explicit function of infusion time, however.

2.3 *Likelihood*

Let δ_A denote the indicator that T_A is not right-censored either administratively or by death, with T_A^o the observed value of T_A or its censoring time, and let δ_D and T_D^o be the analogous quantities for survival time. Denoting $\mathcal{F}(t) = \Pr(T > t)$, the likelihood is given by

$$\begin{aligned} &\mathcal{L}(Y_1, Y_2, T_A, \delta_A, T_D, \delta_D \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) \\ &\quad = \left\{ \pi_1 f_{D,1}(T_D)^{\delta_D} \mathcal{F}_{D,1}(T_D^o)^{1-\delta_D} \right\}^{Y_1} \\ &\qquad \times \left\{ (1-\pi_1)\pi_2^{Y_2}(1-\pi_2)^{1-Y_2} f_A(T_A \mid t_I)^{\delta_A} \right. \\ &\qquad\quad \left. \times \mathcal{F}_A(T_A^o \mid t_I)^{1-\delta_A} f_D(T_D)^{\delta_D} \mathcal{F}_D(T_D^o)^{1-\delta_D} \right\}^{1-Y_1}. \end{aligned}$$

Because administrative censoring takes place only at $t = 50$, when $Y_1 = 1$ the event time portion of the likelihood takes one of two forms: $f_{D,1}(T_D)$ if $T_D < 50$ or $\mathcal{F}_{D,1}(50)$ if $T_D > 50$. Because f_D is piecewise exponential, when $Y_1 = 0$, the event time portion of the likelihood takes one of the following five forms:

$$\begin{cases} f_{D,1}(T_D) & \text{if } T_D < t_I, \\ f_A(T_A \mid t_I)\mathcal{F}_{D,1}(t_I) & \\ \quad \times \mathcal{F}_{D,2}(50 - t_I) & \text{if } t_I + T_A < 50 < T_D, \\ f_A(T_A \mid t_I)\mathcal{F}_{D,1}(t_I) & \\ \quad \times f_{D,2}(T_D - t_I) & \text{if } t_I + T_A < T_D < 50, \\ \mathcal{F}_A(50 - t_I \mid t_I) & \\ \quad \times \mathcal{F}_{D,1}(t_I)\mathcal{F}_{D,2}(50 - t_I) & \text{if } 50 < \min\{t_I + T_A, T_D\} \end{cases}$$

or

$$\mathcal{F}_A(T_D - t_I \mid t_I)\mathcal{F}_{D,1}(t_I)f_{D,2}(T_D - t_I) \quad \text{if } t_I < T_D < \min\{t_I + T_A, 50\}. \quad (9)$$

The last four rows correspond to cases in which the patient survives long enough to be infused at t_I, with $f_D(t) = \mathcal{F}_{D,1}(t_I)f_{D,2}(t - t_I)$ and $\mathcal{F}_D(t) = \mathcal{F}_{D,1}(t_I)\mathcal{F}_{D,2}(t - t_I)$. Thus, for Mylotarg dose schedule (d_1, d_2), the 50-day success probability for a patient infused at t_I is

$$\begin{aligned}\theta(t_I) &= \{1 - \pi_1(d_1)\}F_A(50 - t_I \mid t_I) \\ &\quad \times \sum_{y_2=0}^{1} \pi_2(d_+)^{y_2}\{1 - \pi_2(d_+)\}^{1-y_2} \\ &\qquad \times \mathcal{F}_{D,1}(t_I, y_2)\mathcal{F}_{D,2}(50 - t_I, y_2 \mid t_I) \\ &= \{1 - \pi_1(d_1)\}F_A(50 - t_I \mid t_I) \\ &\quad \times \Bigg[\pi_2(d_+)\exp\Bigg\{-t_I e^{-(\alpha_0+\alpha_1)} \\ &\qquad - (50 - t_I)e^{-(\alpha_2+\alpha_3 x_I+\alpha_4 x_I^2+\alpha_1)}\Bigg\} \\ &\quad + \{1 - \pi_2(d_+)\} \\ &\quad \times \exp\Bigg\{-t_I e^{-\alpha_0} - (50 - t_I)e^{-(\alpha_2+\alpha_3 x_I+\alpha_4 x_I^2)}\Bigg\}\Bigg], \end{aligned} \quad (10)$$

where

$$F_A(50 - t_I \mid t_I) = 1 - \exp\{-(50 - t_I)\exp(-\beta_0 - \beta_1 x_I - \beta_2 x_I^2)\}$$

and $\pi_j(d) = \text{logit}^{-1}(\gamma_0 + \gamma_1 d)$. The quantities $\{\theta(t_I), t_I = 11, 14, 17, 20, 23\}$, will provide the basis for adaptively randomizing patients among the five infusion times throughout the trial.

2.4 *Prior Distributions*

When using Bayesian adaptive decision rules in small-scale clinical trials, the priors may have a substantive effect on the decisions early in the trial when relatively little data are available, i.e., in clinical trials including decisions based on small amounts of data, no prior is truly noninformative. In the present setting, assuming a highly dispersed prior on a given parameter will have particular consequences with regard to the doses assigned to the patients earlier in the trial as well as on the decision of whether to stop the trial early if unacceptably high rates of either toxicity or death are observed.

We assume *a priori* that each of the parameter vectors $\boldsymbol{\beta}$, $\boldsymbol{\alpha}$, and $\boldsymbol{\gamma}$ is multivariate normal. We apply a slightly modified version of the method of Bedrick, Christensen, and Johnson (1996) for specifying priors in a generalized linear regression model setting. The Bedrick et al. method begins with specification of a multivariate prior having dimension p equal to that of the parameter vector, based on p distinct covariate vectors, in the natural domain corresponding to the phenomenon described by the parameter. This distribution is then transformed to obtain the prior on the parameter vector. Here, the natural domains corresponding to $\boldsymbol{\beta}$, $\boldsymbol{\alpha}$, and $\boldsymbol{\gamma}$ that we employ are, respectively, the mean time to ANC recovery, mean probability of surviving 50 days, and HT probability.

Table 1

Prior means and 90% credibility intervals (CI)

a. Prior mean time to ANC recovery: $\mu_A(t_I, \boldsymbol{\beta}) = \exp(\beta_0 + \beta_1 x_I + \beta_2 x_I^2)$

t_I	x_I	Mean	90% CI
11	−2	20	10–30
17	0	10	8–20
23	2	10	8–15

b. Prior probability of death before day 50 given survival to infusion at time t_I: $1 - \exp\{-(50 - t_I) \times e^{-(\alpha_2+\alpha_3 x_I+\alpha_4 x_I^2+\alpha_1 Y_+)}\}$

t_I	x_I	Y_+	Mean	90% CI
11	−2	0	0.30	0.10–0.50
11	−2	1	0.50	0.20–0.80
17	0	0	0.25	0.10–0.40
23	2	0	0.30	0.10–0.50

c. Prior mean hepatoxicity probability: $\text{E}\{\text{logit}^{-1}(\gamma_0 + \gamma_1 d_+)\}$

Dose	d_+	Mean	90% CI
$d_1 = 4$	4	0.025	0.00–0.20
$(d_1, d_2) = (4, 4)$	8	0.10	0.00–0.60

Because the third author of this paper (TM) is a physician specializing in BMT and is a coinvestigator in the DLI trial, the numerical values of the priors are based on his clinical experience and prior beliefs regarding the patient outcomes in the trial.

We denote the mean and variance of α_j by μ_{α_j} and $\sigma^2_{\alpha_j}$ and $\text{cov}(\alpha_j, \alpha_k) = \sigma_{\alpha_j,\alpha_k}$, with the prior parameters of the β_j's and γ_j's denoted similarly. Because $\boldsymbol{\beta}$ is three dimensional, we first specified the prior mean and a 90% credibility interval (CI) for the mean time to ANC recovery, $\mu_A(t_I, \boldsymbol{\beta})$, at each of three distinct infusion times. These are summarized in Table 1, part a. The rationale for these numerical prior parameters is that the earliest possible time to neutrophil recovery following DLI is approximately 8 days and the typical range is between 8 and 15 days. By $t_I = 23$ days, the Mylotarg should clear the patient's bloodstream and standard neutrophil recovery would be expected following DLI administration, i.e., recovery in 8–15 days. At $t_I = 17$ days, there is greater uncertainty about the amount of residual circulating Mylotarg, with greater amounts associated with a delay in neutrophil recovery. This motivates the larger, 8–20 day, 90% CI. At $t_I = 11$ days, there certainly will be circulating Mylotarg and neutrophil recovery may be delayed slightly or significantly, hence the wider CI.

We used this prior information to determine the means of the linear terms $\eta_A(x_I, \boldsymbol{\beta}) = \beta_0 + \beta_1 x_I + \beta_2 x_I^2$ algebraically for $x_I = -2$, 0, and 2 and then used the credibility intervals to solve for the variance of each entry of $\boldsymbol{\eta}_A(x_I, \boldsymbol{\beta}) = (\eta_A(-2, \boldsymbol{\beta}), \eta_A(0, \boldsymbol{\beta}), \eta_A(2, \boldsymbol{\beta}))$. Because $\boldsymbol{\beta}$ is a one-to-one linear function of $\boldsymbol{\eta}_A(x_I, \boldsymbol{\beta})$, applying the usual transformation theorem for multivariate normals then yielded the three-dimensional normal distribution of $\boldsymbol{\beta}$.

For $\boldsymbol{\alpha}$, we first determined the prior of α_0 from the mean probability 0.02 of death without HT within the first 11 days and the probability 0.90 that the probability of this event will be less than 10%. Formally, we assumed that

$$1 - \exp(-11e^{-\mu_{\alpha_0}}) = 0.02,$$

which yields μ_{α_0} and

$$Pr\{1 - \exp(-11e^{-\alpha_0}) < 0.10\} = 0.90,$$

which then gives σ_{α_0} under normality. The prior of $\boldsymbol{\alpha}_{-0} = (\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ was then determined analogously to that of $\boldsymbol{\beta}$. Because $\boldsymbol{\alpha}_{-0}$ has four elements, the mean and variance of the linear term $\eta_D(x_I, Y_+, \boldsymbol{\alpha}_{-0}) = \alpha_2 + \alpha_3 x_I + \alpha_4 x_I^2 + \alpha_1 Y_+$ characterizing the survival distribution of a patient infused at t_I is required for each of four distinct values of the vector (x_I, Y_+). We thus elicited the conditional probability of death before day 50 given that the patient survived to be infused, which is

$$\Pr(T_D < 50 \mid T_D > t_I) = 1 - \exp\{-(50 - t_I)e^{-(\alpha_2 + \alpha_3 x_I + \alpha_4 x_I^2 + \alpha_1 Y_+)}\} \quad (11)$$

under the piecewise exponential model. This prior is summarized in Table 1, part b. The rationale for these numerical prior parameters is that, in previous DLI studies, the mortality rates at 60–90 days post-DLI have been approximately 30%, with deaths due to infection and disease recurrence. Mylotarg may cause more neutropenia and more infection, hence higher mortality, but it also may decrease mortality by preventing disease relapse. At $t_I = 11$ or 23 days, the total duration of neutropenia may be longer than for $t_I = 17$ days. Therefore, estimated mortality will be lowest for a t_I of 17 days versus 11 or 23 days. Given this, if HT occurs, then mortality is expected to increase.

For the prior on $\boldsymbol{\gamma}$, the mean HT probabilities and 90% CIs, summarized in Table 1, part c, correspond to Mylotarg doses of 9 mg/m^2 in each of two courses. The prior corresponds to this dose level since it has been used in previous trials not involving DLI. As above, we began with the assumption that $\eta_{HT}(9, \boldsymbol{\gamma}) = \gamma_0 + \gamma_1 9$ has mean logit(0.05) and $\eta_{HT}(18, \boldsymbol{\gamma}) = \gamma_0 + \gamma_1 18$ has mean logit(0.10), used the credibility intervals to solve for the variances of these two linear terms, and applied the transformation theorem to obtain the bivariate normal distribution of (γ_0, γ_1). The rationale for the numerical prior in Table 1, part c, is that, in the phase I and II trials utilizing Mylotarg at a total dose of 18 mg given in two courses, the incidence of HT in patients with relapsed AML following allogeneic transplantation was 16%, with half of these cases fatal. This motivated the starting dose here of 8 mg in two courses. The prior mean probabilities of HT of 2.5% after 4 mg and 10% after 8 mg are based on the above experience at the higher dose. The larger CI for 8 mg at $(d_1, d_2) = (4, 4)$ reflects the greater uncertainty over two courses of therapy.

2.5 *Posterior Distributions*

Based on the assumed model and priors, neither the joint posterior distribution of all the parameters nor the full conditionals are available analytically. Moreover, numerical evaluation of the posterior distribution is not feasible given the high dimension of the parameter vector. Estimation, however, can proceed using Markov chain Monte Carlo (MCMC) methods, which construct a Markov chain having stationary distribution that is the posterior distribution of interest. After an initial burn-in to reach the stationary distribution, sampled values of the chain are used to estimate the posterior. One of the key elements for using this approach is the algorithm used for constructing the chain. We used the Gibbs sampler algorithm (Geman and Geman, 1984), in which samples are taken from all univariate full conditional distributions. While not available analytically, the full conditionals are log concave so that sampling can proceed by using the adaptive rejection sampling (ARS) algorithm (Gilks and Wild, 1992). This algorithm is based on the fact that any concave function can be bounded from above and from below by piecewise hulls constructed by using tangents and chords between points that are evaluated in the density's domain. Because the rejection probability decreases with the number of sampled values, this algorithm is quite efficient.

To evaluate the performance of the ARS algorithm within the Gibbs sampler, we initially ran parallel chains with different starting points and assessed convergence of the chains using the software CODA (Best, Cowles, and Vines, 1995). This preliminary analysis indicated that the chains converged when using a burn-in of 500 iterations and an additional 1000 iterations for computing the randomization probabilities defined in equation (15). A discussion of implementation aspects of MCMC methods is given by Gilks, Richardson, and Spiegelhalter (1996). For actual trial conduct, a larger posterior sample size may be used to improve precision of the posterior estimates.

3. Trial Conduct

The Mylotarg dose de-escalation algorithm is based on the following definitions. Based on the clinical judgment that toxicity rates of 25% in the first course and 30% in the second course would be acceptable, we defined the course 1 Mylotarg dose d_1 to be unacceptably toxic if

$$\Pr(\pi_1(d_1) > 0.25 \mid data) > 0.95 \quad (12)$$

and the two-course pair (d_1, d_2) to be unacceptably toxic if

$$\Pr(\pi_2(d_+) > 0.30 \mid data) > 0.95. \quad (13)$$

The dose de-escalation algorithm is given in Figure 2, where H_1 is the event that d_1 is unacceptably toxic, regardless of

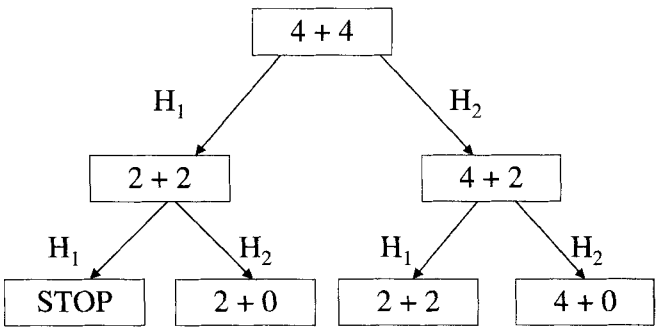

Figure 2. The Mylotarg between-patient dose de-escalation algorithm. "4+2" denotes 4 mg/m^2 in course 1 and 2 mg/m^2 in course 2, etc., H_1 = [d_1 unacceptably toxic], H_2 = [(d_1, d_2) unacceptably toxic but d_1 acceptable]. If either H_1 or H_2 occurs after the third level, then the trial is stopped.

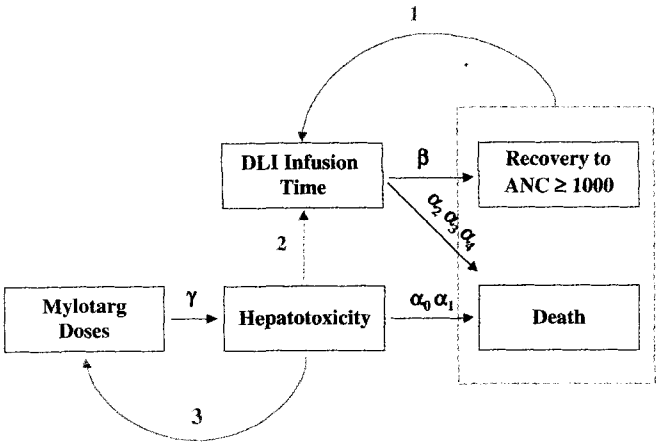

Figure 3. Patient outcomes and adaptive decisions. Arrow 1 denotes the adaptive randomization probabilities for infusion time, arrow 2 the fact that no DLI is given if there is HT in course 1, and arrow 3 the use of HT data by the dose de-escalation algorithm.

d_2, and H_2 is the event that (d_1, d_2) is unacceptably toxic but d_1 is acceptable.

The adaptive randomization is as follows: At the time a patient completes a second course of Mylotarg, the posterior of (α, β, γ) is updated based on the current data and the patient is randomized to infusion time t_j with probability

$$\rho_j(\text{data}) = \Pr\{\theta(t_j) = \max_{1 \le k \le 5} \theta(t_k) \mid \text{data}\}, \quad j = 1,2,3,4,5. \quad (14)$$

This criterion is a generalization of that given by Thompson (1933) for the case of two treatments A and B with binary outcomes having success probabilities θ_A and θ_B that follow beta distributions. Thompson proposed randomizing patients to treatment A with probability $p_A = \Pr(\theta_A > \theta_B \mid \text{data})$ and to B with probability $1 - p_A$. A review of adaptive randomization methods, including comparison with balanced randomization, is given by Berry and Eick (1995). Our approach may be considered a compromise between balanced randomization and dropping any infusion time t_j that, based on interim data, has an unacceptably small posterior success probability, say in terms of $\mathrm{E}\{\theta(t_j) \mid \text{data}\}$ falling below some predetermined fixed cutoff. The latter approach risks incorrectly dropping a superior infusion time that happens to have inferior interim outcomes, however, while balanced randomization is less clinically attractive because it ignores data favoring some infusion times over others. Another approach would be to simply infuse each patient at the time having the maximum mean posterior success probability. This strategy is well known to be inferior from the theory of bandit problems (cf., Berry and Fristedt, 1985), however. To see this, suppose an infusion time, t^S, that is in fact superior has poor results early in the trial due to the play of chance. Under the above play-the-maximum strategy, it could easily follow that no additional patients would be assigned to t^S because the posterior of $\theta(t^S)$, based on a small amount of early unlucky data, thereafter makes t^S appear to be inferior.

Figure 3 illustrates the relationship between the immunotherapy and DLI, the possible patient outcomes, and the adaptive decision rules. Arrow 1 in the figure refers to the reliance of the adaptive randomization probabilities for infusion time on the posterior success probabilities, which in turn incorporate both the ANC recovery time and survival time data. Arrow 2 refers to the fact that the patient's therapy is terminated without DLI if HT occurs in course 1. Arrow 3 refers to the use of the HT data by the dose de-escalation algorithm.

Based on historical experience at the M. D. Anderson Cancer Center, an accrual rate of about two patients per month is anticipated. The trial has a maximum sample size of 60 patients and a 30-month maximum duration, allowing the possibility that the trial may be stopped early due to an excessively high toxicity rate. In contrast with the extensive time spent simulating the trial as part of the design process, during actual trial conduct, only a few minutes are required to update the posterior and compute the probabilities (12)–(14) used as decision criteria. Thus, each new patient's assigned Mylotarg doses and DLI time or the decision to terminate the trial entirely are available immediately at the patient's accrual time.

4. Simulation Study

Because interim decisions typically have important ethical, scientific, and economic consequences, while developing a new clinical trial design, it is important for the statisticians and physicians to closely study both the properties of the design and the particular decisions that will be made under specific clinical scenarios. If any aspect of the design is undesirable, then the design parameters or model components should be modified appropriately before the design is actually used to conduct the trial. The adaptive decisions made during the DLI trial rely on the posterior distributions of the probabilities $\pi_1(d_1)$ and $\pi_2(d_+)$ of HT, the mean time to ANC recovery, $\mu_A(t_I, Y_+)$, and mean survival time, $\mu_D(t_I, Y_+)$. The aim of the simulation study was to evaluate the behavior of the design under different possible qualitative and quantitative forms of these functions. Designing such a simulation study for the DLI trial is not straightforward, however, because the number of possible shapes of μ_A and μ_D as functions of t_I is quite large. In order to evaluate the design under a reasonable array of possible combinations of toxicity, ANC recovery, and survival time, we proceeded as follows.

We first evaluated only the dose de-escalation algorithm in terms of HT while ignoring the rest of the trial's structure. The aim was to first calibrate the probability cutoffs used in the toxicity criteria (12) and (13) and the overall maximum sample size, based on the behavior of the dose de-escalation algorithm. Initially, an algorithm substantially more complex than that given in Figure 2 was considered, beginning with $(d_1, d_2) = (9,9)$ and involving up to five possible dose changes. We simulated this algorithm with a maximum of 50 patients under five cases with $(\pi_1(9), \pi_2(18))$ varying from $(0.05, 0.10)$ to $(0.50, 0.55)$. For each case, we solved for γ_0 and γ_1 and used these as fixed values in the simulations to accommodate the varying doses resulting from the de-escalation. These preliminary simulations showed that, even with the acceptable toxicity probabilities $\pi_1(9) = 0.25$ and $\pi_2(18) = 0.30$, on average, 16 patients suffered HT in the first course and so

only 34 remained to be infused. Based on these numerical results and the additional medical consideration that starting with $(d_1, d_2) = (9, 9)$ might be overly risky, it was decided to modify the design. Accordingly, it began with $(d_1, d_2) = (4, 4)$ rather than (9, 9), the Mylotarg dose de-escalation algorithm was simplified to that described in Figure 1, and the maximum sample size was increased from 50 to 60.

We simulated the entire trial under this design, as follows. From the prior on γ (Table 1, part c), the fixed toxicity probabilities used for the simulations corresponded to $\pi_1(4) = 0.025$ and $\pi_2(8) = 0.10$. We solved these two equations for γ_0 and γ_1 and simulated $\mathbf{Y}$ using the toxicity probabilities determined by these values throughout. To obtain fixed values of the baseline death rate $\mu_D(t_I, Y_+)$, we solved for fixed values of α_0 and α_1 based on their prior means. Recall that $\Pr(T_D < 11 \mid Y_+ = 0) = 1 - \exp(-11e^{-\alpha_0})$, so the prior mean 0.02 for this probability gives fixed $\alpha_0 = 6.3$ for the simulation. Similarly, using the means in the first two rows of Table 1, part b, give fixed value $\alpha_1 = -0.6644$. These values give death rates of 0.0178 for patients with HT and 0.0091 for patients without HT. The simulations assumed these two death rates throughout for a patient never infused, and for infused patients the postinfusion mean survival time was multiplied by $\exp(\alpha_1) = 0.51$ if the patient had HT.

The remaining fixed parameter values needed for the simulations are those determining how μ_A and the 50-day survival probability each vary as functions of t_I. Rather than specifying these in terms of fixed values of $\alpha_2, \alpha_3, \alpha_4$, and β, we specified three non–model-based cases of each function in terms of their numerical values at each t_I. These are given in Table 2. Cases A_1, A_2, and A_3 for μ_A are functions that decrease, increase, and form a U shape with t_I, respectively. Similarly, cases D_1, D_2, and D_3 for the probability of death before day 50 are initially flat and then increasing, U shaped, and decreasing with t_I, respectively. Each scenario thus consists of a combination (A_i, D_j) from Table 2, with the fixed values of γ, α_0, and α_1 as described above. The cases A_1 and D_3 with decreasing functions correspond to the case where Mylotarg is unlikely to destroy infused donor cells and hence neutropenia increases with t_I. The cases A_2 and D_1 with increasing functions correspond to the case where Mylotarg kills a substantial number of donor cells, hence neutropenia decreases with t_I. The U-shaped functions A_3 and D_2 correspond to the case where Mylotarg kills donor cells but increasing t_I prolongs neutropenia; hence, there is an optimal intermediate infusion time.

Initial exploratory simulations with 100 repetitions of the trial took roughly 2 days for one scenario on a personal computer (PC) with a Pentium III processor. In order to simulate 1000 trials under each scenario, we used parallel processing that distributed the computations for each simulation over 35 PCs. Using this approach, on average, it took roughly 12 hours to simulate each scenario. Because there was some redundancy among the nine scenarios in terms of their 50-day survival probability curve $(t_I, \theta(t_I))$, in Table 3, we summarize results for the balanced incomplete block subdesign consisting of the three cases (A_1, D_3), (A_2, D_1), and (A_3, D_2). Results for the other six cases were substantively very similar to those in Table 3.

Table 2
Simulation cases for varying mean time to ANC recovery and probability of death before day 50 if infused at t_I and no HT

	Day of infusion				
Case	11	14	17	20	23
	Mean Number of Days to ANC ≥1000				
A_1	25	21.25	17.50	13.75	10
A_2	10	13.75	17.50	21.25	25
A_3	25	20	15	20	25
	Pr(Death Before Day 50 | No HT)				
D_1	0.25	0.25	0.25	0.30	0.35
D_2	0.25	0.20	0.15	0.20	0.25
D_3	0.40	0.35	0.30	0.25	0.20

Under scenario (A_1, D_3), $\theta(t_I)$ increases monotonically from 0.43 at $t_I = 11$ to 0.69 at $t_I = 23$. The selection probability, $SP(t_I)$, is very nearly monotone increasing in t_I, with $SP(23) = 0.56$ for the optimal day $t_I = 23$. Moreover, on average, the adaptive randomization treats more patients at the more desirable infusion times. The success probabilities under scenario (A_2, D_1) decrease monotonically with infusion day, from 0.67 at $t_I = 11$ to 0.39 at $t_I = 23$. Again, in this case, $SP(t_I)$ follows this monotone pattern, aside from a very small increase at $t_I = 23$. Given the small difference in the two largest success probabilities, $\theta(11) = 0.67$ and $\theta(14) = 0.64$, under this scenario, the fact that $t_I = 11$ and 14 are selected with respective probabilities 0.50 and 0.26 is quite encouraging. Under scenario (A_3, D_2), $SP(t_I)$ increases to a maximum of 0.70 at $t_I = 17$ and then decreases. Both $SP(t_I)$ and the numbers of patients infused follow this pattern, with $SP(17) = 0.54$ for the optimal day $t_I = 17$. These results indicate that the quadratic functions of infusion time in the log-mean ANC recovery rate and the log-mean death rate do a very effective job of recognizing nonlinear rates in terms of the Bayesian adaptive randomization based on the posteriors of the $\theta(t_I)$'s. Given that the adaptive randomization must rely on very small amounts of data early in the trial, the method behaves remarkably well.

Our elicited prior is informative. While a thorough analysis of the design's sensitivity to the prior is beyond the scope of this article, it is useful to examine how the operating characteristics change in a specific case. To do this, we (1) decreased the mean and 95% CI of $\mu_A(11, \beta)$ from (20, 10–30) to the more optimistic values (15, 5–25), (2) increased the mean and 95% CI of $\mu_A(17, \beta)$ from (10, 8–20) to the less optimistic values (15, 13–25), and (3) increased the mean and 95% CI of the probability of death before day 50, given survival to infusion, for $t_I = 11$ to (0.40, 0.10–0.60) when $Y_+ = 0$ and to (0.60, 0.25–0.90) when $Y_+ = 1$. The effects of these changes are to shorten both mean ANC recovery time and survival time for patients with the earliest infusion day, $t_I = 11$. With these changes, under scenario (A_1, D_3), the mean selection probabilities are $\mathbf{SP} = (SP(11), \ldots, SP(23)) = (0.07, 0.02, 0.11, 0.17, 0.63)$ and the mean numbers of patients infused are $\mathbf{N} = (N(11), \ldots, N(23)) = (7.9, 3.3, 12.1,$

Table 3
Design operating characteristics under three scenarios

Scenario		Not infused	Day of Infusion 11	14	17	20	23	Overall
A_1, D_3	Pr(50-day success)		0.43	0.48	0.54	0.61	0.69	
	Pr(selected)	—	0.06	0.03	0.15	0.20	0.56	—
	No. patients treated	1.6	5.7	4.2	15.1	10.3	19.3	56.3
A_2, D_1	Pr(50-day success)		0.67	0.64	0.58	0.48	0.39	
	Pr(selected)	—	0.50	0.26	0.16	0.03	0.05	—
	No. patients treated	1.3	23.5	11.6	12.4	3.3	5.2	57.3
A_3, D_2	Pr(50-day success)		0.54	0.61	0.70	0.57	0.45	
	Pr(selected)	—	0.13	0.17	0.54	0.11	0.05	—
	No. patients treated	1.5	7.5	9.8	26.6	7.6	4.6	57.6

9.7, 23.8); under (A_2, D_1), $\mathbf{SP} = (0.48, 0.22, 0.14, 0.06, 0.11)$ and $\mathbf{N} = (25.5, 9.4, 10.4, 3.6, 7.1)$; under (A_3, D_2), $\mathbf{SP} = (0.14, 0.16, 0.49, 0.13, 0.08)$ and $\mathbf{N} = (10.4, 9.9, 22.0, 7.0, 6.5)$. In all three scenarios, both $SP(17)$ and $N(17)$ decrease slightly while $SP(23)$ and $N(23)$ increase compared with the values in Table 3 for the actual prior. The general behavior of the design persists under all three scenarios, however, in that the best DLI time still has the highest values of SP and N. It thus appears that the method is sensitive to the prior, at least for changes of the magnitude considered, but that it still behaves as desired.

5. Discussion

We have proposed an adaptive statistical design for conducting a clinical trial of an innovative therapeutic strategy in a group of cancer patients with very poor prognosis. Each patient's outcome is the vector $(\mathbf{Y}, T_A^o, \delta_A, T_D^o, \delta_D)$, with several possible configurations of these variables observable depending on whether $Y_1 = 0$ or 1 and the values of T_A and T_D. The therapeutic strategy is adaptive within each patient in that no further therapy is given if $Y_1 = 1$. The design is adaptive between patients in that each patient's Mylotarg doses and infusion time are determined based on the data from previous patients in the trial. This may be considered a phase I/II trial (Thall and Russell, 1998) in the sense that it includes dose finding and evaluation of both adverse and desirable outcomes. While the design described here has been tailored to the particular application at hand, the dose-finding and adaptive randomization methods could be applied, with suitable modification, to design future trials having similar goals.

Acknowledgements

The authors thank two referees and an associate editor for their constructive comments. Dr Thall's work was partially supported by NIH grant RO1-CA-83932.

Résumé

Cet article décrit un plan adaptatif bayésien, conçu pour l'essai clinique d'un traitement expérimental destiné à des patients atteints d'affections tumorales hématologiques et accusant, après une greffe de moëlle osseuse allogénique, une récidive. Le traitement consiste en une ou deux immunothérapies ciblées, suivies par une transfusion de lymphocytes. Si l'on sait que l'immunothérapie est un préalable nécessaire à la transfusion, on sait aussi qu'elle peut produire des toxicités hépatiques sévères, et qu'elle provoque à coup sûr des diminutions importantes des globules blancs et des plaquettes. L'objectif scientifique primordial est ici de déterminer le délai optimal (mesuré à partir de l'immunothérapie) pour procéder à la transfusion, afin que celle-ci ait les meilleures chances de succès (on définit un succès par le fait que le patient, 50 jours après le début de la thérapie, ait non seulement survécu, mais qu'il n'ait subi aucune atteinte hépatique sévère, tout en ayant par ailleurs recouvré une quantité normale de globules blancs). Notre méthode utilise un modèle paramétrique prenant en compte la survie, la toxicité hépatique et le temps de recouvrement d'une quantité normale de globules blancs. Le plan expérimental repose d'une part sur un algorithme comparant, à partir de données de toxicité hépatique, la désescalade des doses de l'immunothérapie d'un patient à l'autre; il repose d'autre part, sur une randomisation adaptative de cinq délais possibles pour la transfusion, randomisation effectuée en fonction des dernières estimations a posteriori des taux de succès liés à chacun des cinq délais. Une simulation démontre que c'est avec une certaine fiabilité que ce plan sélectionne le meilleur délai après immunothérapie; de surcroî t, la proportion des patients randomisés sur les délais les moins bons s'avère relativement faible.

References

Antin, J. (1993). Graft-versus-leukemia: No longer an epiphenomenon. *Blood* **82,** 2273–2277.

Bedrick, E. J., Christensen, R., and Johnson, W. (1996). A new perspective on priors for generalized linear models. *Journal of the American Statistical Association* **91,** 1450–1460.

Bernstein, I. (2000). Monoclonal antibodies to the myeloid stem cells: Therapeutic implications of CMA-676, a humanized anti-CD33 antibody chalicheamicin conjugate. *Leukemia* **14,** 474–475.

Berry, D. A. and Eick, S. G. (1995). Adaptive assignment versus balanced randomization in clinical trials: A decision analysis. *Statistics in Medicine* **14,** 231–224.

Berry, D. A. and Fristedt, B. (1985). *Bandit Problems.* New York: Chapman and Hall.

Best, N., Cowles, M. K., and Vines, K. (1995). *CODA—Convergence Diagnosis and Output Analysis Software for Gibbs Sampling Output.* Cambridge: Medical Research Council, Biostatistics Unit.

Collins, R., Shpilberg, O., Drobyski, W., et al. (1997). Donor leukocyte infusions in 140 patients with relapsed malignancy after allogeneic bone marrow transplantation. *Journal of Clinical Oncology* **15,** 433–444.

Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6,** 721–741.

Giles, F., Kantarjian, H., Kornblau, S., Thomas, D. A., Garcia-Manero, G., Waddelo, T. A., David, C. L., Phan, A. T., Colburn, D. E., Rashid, A., and Estey, E. A. (2001). Mylotarg (Gemtuzumab Ozogomicin) therapy is associated with hepatic veno-occlusive disease in patients who have not received stem cell transplant. *Cancer* **92,** 406–413.

Gilks, W. R. and Wild, P. (1992). Adaptive rejection sampling for Gibbs sampling. *Applied Statistics* **41,** 337–348.

Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (1996). *Markov Chain Monte Carlo in Practice.* London: Chapman and Hall.

Lane, T., Ho, A., Bashey, A., Peterson, S., Young, D., and Law, P. (1999). Mobilization of blood-derived stem and progenitor cells in normal subjects by granulocyte-macrophage and granulocyte-colony-stimulating factors. *Transfusion* **39,** 39–47.

Neumeister, P., Eibl, M., Zinke-Cerwenka, W., Scarpetti, M., Sill, H., and Linkesch, W. (2001). Hepatic veno-occlusive disease in two patients with relapsed acute myeloid leukemia treated with anti-CD33 calicheamicin (CMA-676) immunoconjugate. *Annals of Hematology* **80,** 119–120.

Prentice, R. L., Kalbfleish, J. D., Peterson, A. V., Jr., Flournoy, N., Farewell, V. T., and Breslow, N. E. (1978). The analysis of failure times in the presence of competing risks. *Biometrics* **34,** 541–554.

Sievers, E. L., Appelbaum, F. R., Spielberger, R. T., et al. (1999). Selective ablation of acute myeloid leukemia using antibody-targeted chemotherapy: A phase I study of an anti-CD33 calicheamicin immunoconjugate. *Blood* **93,** 3678–3684.

Thall, P. F. and Russell, K. T. (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* **54,** 251–264.

Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* **25,** 275–294.

Received June 2001. *Revised January* 2002.
Accepted February 2002.