

# 3

## Adaptive Bayesian Designs for Dose-Ranging Drug Trials

Donald A. Berry, Peter Müller, Andy P. Grieve, Michael Smith, Tom Parke, Richard Blazek, Neil Mitchard, Michael Krams

### 1 Introduction

#### 1.1 *Stroke and Neuroprotective Agents*

If you hold this text in your right hand and read it aloud, parts of the left side of your brain would be very active. This activity would consume considerable amounts of energy. Arterial blood provides a constant source of fresh energy to the brain in the form of glucose and oxygen. Now imagine that there is a sudden interruption in the arterial flow to these parts of your brain. Brain cells would start to die within minutes. You would lose control and power in your right arm and the text would drop. You would want to say something but would be unable to speak, and you may be unable to understand what other people are saying to you. These symptoms would develop very rapidly. Each year in the United States some 600,000 patients experience such symptoms. A doctor would diagnose “stroke” or “brain attack”. In industrialized countries, stroke is the leading cause of disability, and it is the third leading cause of death, behind only heart disease and cancer. Stroke survivors can show either full or partial recovery. At about three months after the event most patients have reached their best recovery and at most small improvements occur after three months. The overall economic burden of stroke is enormous – approximately \$41 billion in the U.S. in 1997.

Until recently, there were few therapeutic options for stroke patients. This is now changing: both patients and medical teams have learned to deal with a brain attack as a medical emergency. The motto is “Time is brain”. There are two therapeutic approaches that have been shown to improve outcome (Hill, 1998):

1. Stroke units: Specialized semi-intensive care units, where blood pressure, blood sugar and other important factors relevant for outcome can be properly controlled and where patients receive physiotherapy.

2. Thrombolytic therapy: If a patient suffering from an ischemic stroke is admitted to a stroke unit within 3 hours of onset of symptoms, tPA (tissue plasminogen activator) can help to dissolve the blood clot that is occluding the cerebral artery. tPA can thus help to reinstall perfusion of the affected parts of the brain (that is, get the blood flowing through them again). However, tPA has risks: given to the wrong patient, it can lead to serious complications, such as secondary hemorrhage.

A third therapeutic concept involving neuroprotective agents is under development but is still experimental (Lees, 1998). The example considered throughout this report is the development of a particular neuroprotective agent. The logic of neuroprotective therapy in acute stroke is based on the pathophysiology of brain attacks. The following is a brief review of that pathophysiology (see also Dirnagl, 1999). Brain cells cannot function properly without a constant supply of oxygen and glucose. Stopping blood flow halts this supply. An interruption for longer than a few minutes will result in the death of affected brain cells. Once cells start to die, the fine equilibrium of intra- and extracellular ions spirals out of control. Neurotransmitters excite surrounding brain cells. This makes them hyperactive in an environment where the energy supply is already at dangerously low levels, which leads to further cell death. Complex inflammatory processes are triggered and complicate the situation further. The area of the brain that receives no or minimal blood flow is called the infarct core. Cells in the infarct core are destined to die. The penumbra is the area of the brain that surrounds the core and that receives some blood flow, but at a level insufficient to allow brain cells to function properly. Cells in the penumbra are able to survive. After a time they may recover some or all of their function. Think of a stroke as a rapidly expanding sphere inside a brain. The interior of the sphere (the infarct core), but not its surface (penumbra), is condemned to death. Now imagine that it is possible to intervene and interrupt the expansion of this sphere. Fewer brain cells would die. The goal of neuroprotective therapy in stroke is to stop or slow the expansion of the infarct core. To be effective it may have to be administered very soon after the onset of stroke.

To test whether a particular compound has neuroprotective properties, such a compound is first tested in animal stroke models. These tests are controlled for important factors such as:

- (i) Type of stroke (ischemic/hemorrhagic, cortical/subcortical).
- (ii) Stroke severity (localisation - proximal or distal occlusion of artery; duration over which a cerebral artery is occluded).
- (iii) Time delay between onset of stroke and administration of therapy (the shorter the delay, the more efficient the therapy).

- (iv) Administration of study drug (dose and duration over which drug is given).

If a reduction of infarct volume is observed in the treated group, this is taken as evidence for a compound to have neuroprotective properties. Whether these properties translate into a clinical benefit in human stroke patients, however, can only be seen by conducting clinical trials involving human stroke patients. Generally in drug development, the first clinical trials are in human volunteers and the goal is to assess the agent's safety. If it is proven safe then it is used in patients who have suffered stroke. The goal is to determine whether the agent is effective, and of course to continue assessing its safety.

Assessing neuroprotective drug efficacy in human stroke patients is not easy. An important question is: What is a relevant efficacy endpoint? Infarct volume is the endpoint often used in preclinical studies. In principle it is possible to assess infarct volume also in human stroke patients, by using sophisticated imaging techniques (diffusion/perfusion weighted magnetic resonance imaging). But it is as yet unclear how well reduction of infarct volume and clinical recovery correlate. The problem is that the neurological deficit depends on a) the stroke locale, b) the extent of the brain damage, and c) the stroke's etiology. A large cortical lesion in the primary motor cortex can produce a neurological deficit similar to a very small lesion along the corticospinal tract. And a rather large lesion in certain parts of the frontal lobes may produce only minor deficits, or least deficits that are difficult to measure. As regards etiology, approximately 80% of strokes are called "ischemic", often involving the obstruction of a cerebral artery. And about 15% of strokes are "hemorrhagic": if an arterial vessel bursts, blood will be pushed with high pressure into the delicate brain tissue, destroying the fine neuronal structures. Sometimes a hemorrhagic stroke can be a secondary complication of an ischemic stroke. This can happen when the lack of blood perfusion mollifies the brain tissue. When the occluded artery becomes cleared, blood may shoot with high pressure into the area of the brain that had been damaged, and potentially resulting in secondary hemorrhage.

In practice, most stroke trials take the principal endpoint to be the patient's condition 90 days after the stroke. The patient's condition is measured by one or more outcome rating scales. There is no general agreement regarding the appropriate way to assess recovery in stroke patients. Instruments are difficult to standardize across different types of strokes, different cultures, and different expectations. Some stroke scales for assessing recovery are dichotomous, with patients classified as "dependent" or "independent". Other scales have a finer rating of neurological deficit. We will use one of the latter, the Scandinavian Stroke Scale (SSS) (Scandinavian Stroke Study Group 1985). The SSS has nine components, each with possible scores ranging from 0 to some maximum number of points.

The maxima are shown in parentheses in the following list: consciousness (6), eye movement (4), upper extremity (6), hand (6), lower extremity (6), orientation (6), aphasia (10), facial (2), walking (12). Higher score means better performance. The maximum total SSS is 58.

Inclusion/exclusion criteria in stroke trials are very important, although exactly what they should be is not clear. Patients with mild stroke are easier to recruit into trials, in part because there are more of them. Such patients tend to improve even without therapy. On the other hand, perhaps they would improve even faster with neurological therapy. To help us understand the response over time of untreated patients we used an extensive longitudinal database over time for in excess of 1000 acute stroke patients, the Copenhagen Stroke Database (Jorgensen et al. 1994). Perhaps the most impressive aspect of this database is that it demonstrates the difficulty of developing drugs for treating stroke: the variability in response is enormous, even considering patients with similar baseline characteristics, including baseline SSS. Information from this database led us to focus on moderate to severe stroke in the trial described in this report. In addition, we use this database to simulate trials, as described later in this report.

Among the most important aspects of drug development programs is selecting the correct dose. Too many programs have focused too little attention on this aspect of drug development. There have been a number of failures in recent trials investigating neuroprotective agents in acute stroke and these may be due to an imperfect understanding of the dose-response relationship and a consequent incorrect selection of dose.

One of the reasons for the inadequate understanding about dose-response in development programs of neuroprotective agents is that proper learning requires trials with a very large number of patients – at least with traditional designs. Given the time and resource constraints, some programs may have extrapolated dose-response relationships from preclinical programs, sometimes assuming that from the perspective of efficacy, “more is better”.

Our approach described in the following section considers only single-shot IV therapy to be given as soon as possible after the stroke. We do not consider duration of therapy as an independent factor. A more comprehensive model would allow for learning about this dimension of dose.

We have suggested that acute stroke is a difficult indication for which to develop therapies. The design we propose in this report improves on the efficiency of drug development programs by applying Bayesian thinking and adaptive design. This methodology makes learning about novel therapies more efficient. It will help in the treatment of stroke patients by reducing the numbers of patients exposed to non-efficacious doses or drugs during clinical development of a new compound and by bringing efficacious therapies to the marketplace earlier. The approach is applicable to a wide range of drugs and indications, with stroke being but an example.

### 1.2 *Critical aspects of standard dose-selection trials*

In a standard dose-selection trial, patients are assigned randomly to a set of pre-defined doses, typically between 3 and 5 in number. A standard trial is balanced in the sense that equal or nearly equal numbers of patients are assigned to the various doses.

Such a trial may be inefficient. If the test drug is effective then the dose-response curve has a positive slope for some interval of doses. The sloping part of the curve may be located at doses greater than or less than those considered in the trial. In either case some of the observations at the opposite end from where the slope occurs may be wasted. Moreover, if the sloping part of the curve occurs within the range of doses considered and is steep relative to the interval between doses then observations at both ends of the range may be largely wasted. The case in which the drug is ineffective is similar to the one in which a positive slope occurs at doses larger than those considered in the trial; namely, many of the observations at lower doses may be wasted.

To illustrate some of these points, consider the sample of dose-response curves shown in Figure 1.1, (a)-(f). The true dose-response curves are logistic. Four active doses (labeled 1, 2, 3, 4) plus placebo (dose 0) are assigned in a balanced fashion so the sample sizes are equal. The true values of the response curves at these doses are shown as solid circles – these are the true mean responses for the various treatment arms in the trial. The vertical lines attached to the circles represent the set of likely values for the observed mean response at the corresponding dose. For example, this interval may be the one containing the middle 95% of the probability.

In Figure 1.1(a) the sloping part of the dose-response curve is between doses 0 and 1, with all four active doses having essentially the same underlying mean response. In retrospect, many (but not all) of the patients assigned to doses 2, 3 and 4 would have been more informative about the slope of the dose-response curve (and also about the ED<sub>95</sub>, the smallest dose at which 95% achieved) had they been assigned to doses between 0 and 1.

The accuracy with which a trial can determine the ED<sub>95</sub> is limited by the separation between doses. In Figure 1.1(a), for example, the ED<sub>95</sub> will likely be identified as being between doses 0 and 1. Such accuracy may be sufficient for some compounds, but for drugs that may be less safe at higher doses and for drugs that are costly to produce, it may be essential to identify the optimal dose with greater accuracy.

In Figure 1.1(b) the slope of the dose-response curve is the same as that in Figure 1.1(a) but the entire curve is shifted to the right. If the dose-response curve happens to be the one in Figure 1.1(b) then in retrospect it would have been more informative to have assigned patients to doses between 2 and 3 and also additional patients at doses 3 and 4. In addition, sampling variability present in the observed responses at doses 3 and 4

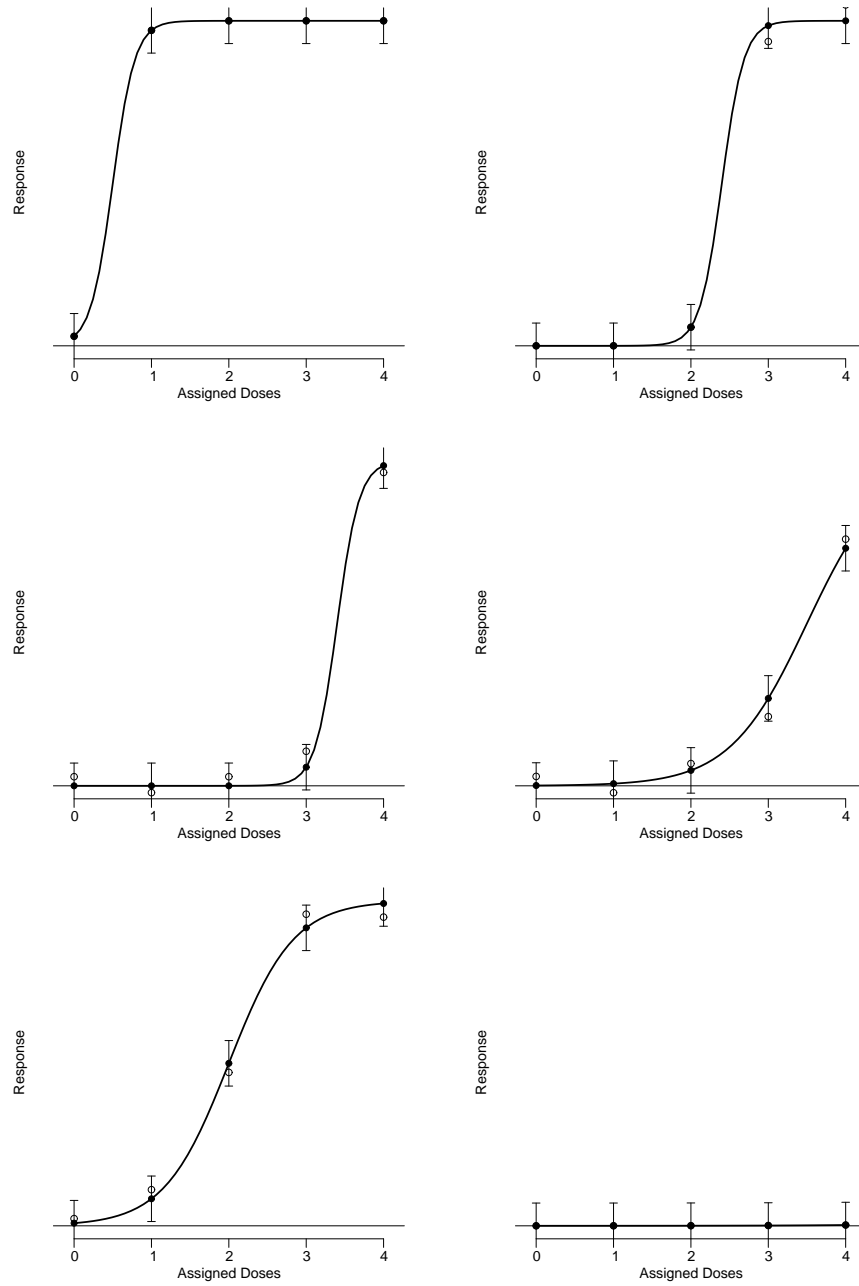


FIGURE 1.1.

will leave open the possibility that the asymptote of the curve has not been reached. For example, the observed means at doses 3 and 4 may be as indicated by the open circles. But regardless of the observations at these two doses, the value of the asymptote and the ED95 will be poorly known, and in retrospect assigning patients to doses higher than 4 (safety permitting) would have been more informative.

On this same point, the conclusions drawn from trials in which the curves are as shown in Figures 1.1(c) and 1(d) will be similar. For example, the five observed mean responses may be as shown by the open circles, and these are identical in the two figures. In 1(c) but not in 1(d) the asymptote will have been reached (unknownst to the investigators). Moreover, if the responses at dose 4 in comparison to placebo are not clinically important then in 1(c) the drug should be abandoned while in 1(d) doses higher than 4 should be explored, safety permitting. The observed results of the trial will not help in deciding which strategy is better.

In Figure 1.1(e) the doses assigned happen to be ideally placed in the sense that they encompass the gamut of responses and they are spread out along the sloping part of the curve. But even here, because of uncertainty concerning the true slope of the curve between doses 3 and 4, the ED95 is not very well identified.

In Figure 1.1(f) there is no response over the range of doses considered. Assuming that the dose-response curve is known to be monotonic, information from patients assigned to doses 1, 2 and 3 will be largely wasted and would have been more informative had they been placed at either end. Either the drug is ineffective or its effect occurs at doses higher than 4. Explorations beyond dose 4 are essential (if safe) to decide which possibility is the correct one.

Another drawback to standard types of designs is the use of fixed sample sizes. Even if the doses considered are judged in retrospect to be appropriate, the variability in response across patients assigned to the same dose may be less or greater than anticipated. In the former case the sample size chosen may be unnecessarily large and in the latter case it may be too small to accomplish the study's goals. Figure 1.2 addresses this issue. In both 1.2(a) and 1.2(b) the dose-response curve is the same as in Figure 1.1(e). In Figure 1.2(a) the measurement standard deviation is half that of Figure 1.1(e) and in Figure 1.2(b) it is twice that of Figure 1.1(e). The standard errors at each dose will then be half as large and twice as large, respectively, as anticipated.

Not surprisingly, a retrospective view of conventional dose-finding trials suggests that assigning to doses in a different fashion from that actually used would have been more informative, or that a different sample size would have been more appropriate, or both.

At least two types of improvements to balanced designs with pre-selected doses are possible:




FIGURE 1.2.

- (i) Increase the number of doses considered. This can facilitate identifying critical aspects of the dose-response curve, but if the numbers of patients assigned at each dose is the same then the cost in terms of sample size may be prohibitively large.
- (ii) Avoid large sample numbers of observations at doses where the form of the curve can be reasonably well estimated from a smaller number by borrowing the information available from observations at nearby doses.

Both improvements are accomplished using the procedure we describe in this report. The result is that more accurate information is available about critical aspects of dose-response with, generally, smaller sample sizes. Moreover, when the total sample size is larger, this larger size is necessary to identify important aspects of the dose-response curve.

In this report we propose an adaptive scheme in which the doses assigned are allowed to depend on information from patients treated previously. The design is in two phases:

- (a) The focus of the first phase is dose finding – is there a dose that is sufficiently effective to be used in a confirmatory trial and thereafter in a clinical setting?
- (b) The second phase is confirmatory.

A second adaptive aspect of the design is that the duration of the dose-finding phase of the trial depends on the accumulating information about the dose-response curve: If the available information is sufficient to conclude that there is indeed a response to the drug and that the dose-response curve is well identified, assignment switches to a confirmatory mode in which a particular dose that has been identified as being effective – such as the

ED95 – is compared with placebo in a balanced, randomized fashion. The switch can take place seamlessly with no delay in patient accrual.

On the other hand, if the available evidence is sufficient to suggest that there is no dose-response or that the maximal effect is too small to make the drug worthwhile clinically, then dose finding will cease and the procedure will recommend terminating the trial.

The design is adaptive in both assigning dose and deciding when to stop dose finding. And it uses Bayesian updating and Bayesian decision analysis. However, the regulatory approval process is largely frequentist. Therefore, the decision analysis involves predicting statistical significance (in the usual frequentist sense) of the confirmatory trial, and its sample size is determined accordingly.

There is a regulatory tradition of requiring two trials, both with statistical significant results. The design we propose embodies two trials, one the dose-finding phase and the other the confirmatory phase. Statistical significance in the first phase can be for the slope of the dose-response curve or for the comparison of the ED95 with placebo, assuming a particular probability model. The current version of our program does not require statistical significance in the first phase (adjusted for the sequential nature of the design) before converting to the second phase. However, this option is being incorporated into the program so as to meet this requirement if it is imposed by regulatory agencies.

## 2 Design Process

The proposed design for dose selection is in four parts, which by the nature of sequential, adaptive designs are not necessarily ordered in time. The process is schematically represented in Figure 2.1, the 11 components of which will be described below and related to the four parts of the dose selection process.

### *2.1 Allocating patients to optimal dose*

On a particular day during the trial, a number of patients, typically fewer than 5, will be admitted to the trial [1 in Figure 2.1]. Consider the first patient. Based on the available information our knowledge concerning the ED95 will be represented by a posterior distribution with some variance. Assigning a particular dose to the next patient and observing the response will effect a change in the posterior distribution of the ED95, and hence in its variance. We can predict how the chosen dose and the corresponding response of the patient will change the consequent posterior variance. The most informative dose [11] – say  $\tilde{z}$  – is the one with the smallest average variance of the ED95, where the average is with respect to the predictive

probabilities of the various future observations, given the dose assigned. Before assigning the dose, however, some amount of randomization is employed [2]. First, if  $\bar{z}$  is not placebo then placebo is assigned with some minimal probability – such as 15%. Second, consider the set of doses for which the expected response is within 10% of that of  $\bar{z}$ . The dose assigned is selected randomly from within this set. Subsequent patients admitted on a particular day – up to 5 – are considered as above, but conditioning on the previously assigned doses and averaging with respect to the predictive distributions of the responses at these doses.

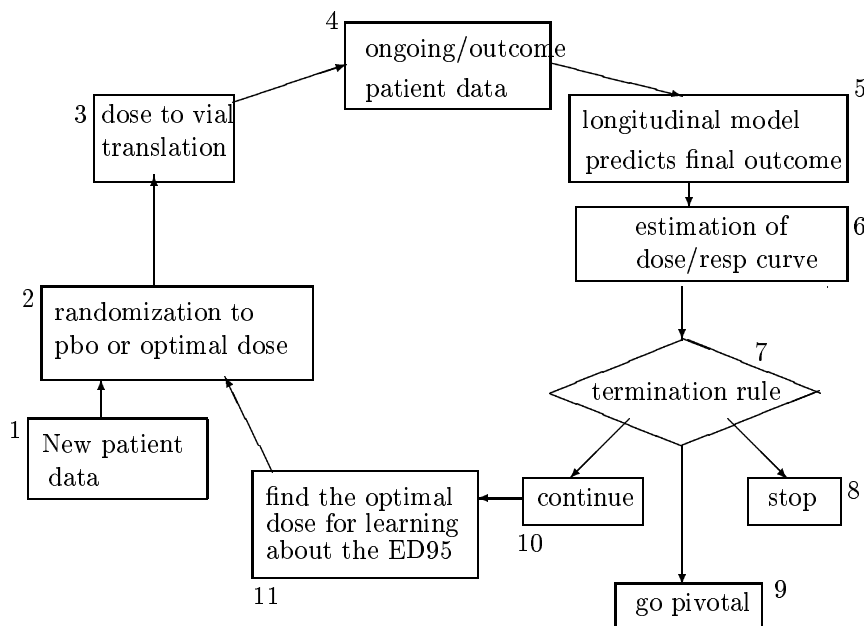


FIGURE 2.1.

## 2.2 Predicting final outcome: Longitudinal model and user interface

At any particular time, patients having entered the trial will provide information concerning their responses [4 in Figure 2.1]. Through the use of a longitudinal, predictive model to be described in Section 7, we impute the final response of the patient [5 in Figure 2.1]. This model allows an early update of knowledge about the dose-response relationship, which in turn allows for better choices concerning the doses to be allocated to future patients. In view of the critical importance of delivering neuroprotective therapy quickly, such updating requires rapid data processing. A rapid communications interface between the statistical system and the centers

where the study is carried out has been developed. Details of the interface are given in Section 8.

### *2.3 Estimating dose-response*

We use a dose-response model based on Normal Dynamic Linear Models (NDLM) as described, for example, in West and Harrison (1997). This is essentially a piecewise linear model. It provides the necessary flexibility to encompass both monotonic and non-monotonic dose-response relationships. An additional advantage of the NDLM is the existence of analytical results for the determination of the posterior distribution of the dose-response curve, and also for the ED95. Together with data from those patients who have completed the study, the imputed values described in the previous section are used to update the relevant posterior distributions. Details of the model and the updating procedure are given in Section 6.

### *2.4 Confirmatory decision*

On each day of the dose-finding phase the algorithm implements a decision rule [7 in Figure 2.1] that addresses the question of whether the process should continue in dose-finding [10 in Figure 2.1], end because of futility [8 in Figure 2.1], or shift to a confirmatory phase [9 in Figure 2.1]. These decisions are made based on the predictive probabilities of failure to show a benefit and successfully showing a benefit. (The decision to stop the trial is not made in the absence of human intervention. A Data Safety Monitoring Board reviews any decision concerning the conduct of the trial.) The shift from dose-finding to confirmatory may be seamless in that the phase of the trial may not be recognizable by physicians and others involved. If a shift is made in a seamless fashion then the follow-up information from patients treated most recently in the trial will not be complete.

The timing of the shift from dose-finding to confirmatory is important. Whether to shift is based on a Bayesian decision analysis. A component of this analysis is an assessment of the probability of showing an eventual benefit in a confirmatory trial at the current ED95. The decision to shift will depend on the available information about dose-response. The design of the confirmatory study will be based on the predictive power calculations, but whether to shift to the confirmatory phase of drug development is based on Bayesian calculations of the probability of showing a benefit over placebo.

### 3 The System

#### 3.1 Notation

Let  $y_i$  denote the final response (at week 13) for the  $i$ -th patient, measured as difference from baseline SSS (Scandinavian Stroke Score). Let  $y_{ij}, j = 1, \dots, M - 1$ , denote responses in weeks 1 through  $M - 1 = 12$ . We use  $y$  without an index to stand for week 13 response. Vectors  $y_i = (y_{ij}, j = 0, \dots, M)$  and  $y(i) = (y_1, \dots, y_i)$  stand for data for the  $i$ -th patient and for the first  $i$  patients, respectively. We will use  $D$  to denote all data collected in the study. Let  $x_i$  denote a vector of covariates for the  $i$ -th patient. In the current implementation the only covariate used is the baseline SSS score. But  $x_i$  could include other relevant patient-specific information, such as age, time between onset of stroke and start of treatment, etc. We use  $z$  to refer to a treatment dose,  $z_i$  to refer to the dose assigned to patient  $i$ , and  $\{Z_j; j = 0, \dots, J\}$  to refer to the set of allowable doses, where placebo  $Z_0 = 0$ . We use  $Y_{jk}, j = 0, \dots, J$ , to denote the response of the  $k$ -th patient assigned to dose  $Z_j$ . In Sections 5 and 6 we describe a sampling model for  $y_i$  and  $y_{ij}$ . In earlier sections the probability model is general. Our descriptions of dose assignment (Section 4) and stopping time (Section 5) apply for any underlying probability model. We will need generic notation to refer to the dose/response curve for  $y_i$  and the longitudinal data model for  $y_{ij}$ . We use  $f(z, \theta)$  to denote the dose/response curve  $E(y|\theta)$  as a function of dose  $z$ , parametrized by an unknown parameter vector  $\theta$ ,  $E(y_i|\theta) = f(z_i, \theta)$ . Given the mean curve  $f(z, \theta)$ , we assume normal errors:

$$y_i = f(z_i, \theta) + \epsilon_i \text{ with } \epsilon_i \sim N(0, \sigma^2). \quad (1)$$

Let  $df(z, \theta) = f(z, \theta) - f(0, \theta)$  denote the advantage over placebo. For a given parameter vector  $\theta$  let  $\hat{z}_\theta$  denote the dose with maximum expected response, i.e.,  $f(\hat{z}_\theta, \theta) = \max_j \{f(Z_j, \theta)\}$ . Denote with  $z95_\theta$  the ED95 of the unknown dose response curve, defined as the minimum dose  $Z_j$  with mean improvement greater or equal than 95% of the maximum possible improvement over placebo:

$$z95_\theta = \min\{Z_j : df(Z_j, \theta) \geq 0.95 df(\hat{z}_\theta, \theta)\}$$

We use  $f(z)$  for the posterior expected dose/response curve,  $f(z) = E[f(z, \theta)|D]$ , and  $z95$  for the ED95 of the posterior mean curve  $f(z)$ .

Also, we shall use  $\overline{z95}$  to denote the posterior mean of  $z95_\theta$ , i.e.,  $E(z95_\theta | D)$ . In view of the nonlinear nature of the ED95,  $\overline{z95} \neq z95$ . For technical reasons, in the program we sometimes consider conditional posterior means of  $z95_\theta$ , keeping some aspects of the model fixed. We indicate when this is the case.

### 3.2 Dose-finding study

The first phase of the trial is for dose-finding. Details of the dose allocation are discussed in Section 4. We will use  $D_{1t}$  to denote the data up to week  $t$  of the dose-finding phase, and  $n_{1t}$  to denote the number of patients enrolled by week  $t$ . Similarly  $D_1$  denotes the data from the complete dose-finding phase, and  $n_1$  denotes the total number of patients enrolled in the dose-finding phase.

### 3.3 Confirmatory phase

In the confirmatory phase, patients are randomized equally to placebo ( $z = 0$ ), recommended treatment dose  $z^*$ , and possibly a third dose  $z^{**}$  chosen as alternative to  $z^*$ . As a recommended treatment dose we use the ED95 dose,  $z^* = \overline{z95}$ . As alternative dose  $z^{**}$  we use the ED50,  $z^{**} = \overline{z50}$ . We will use  $D_2$  to denote the data from the confirmatory phase, and  $n_2$  to denote the number of patients enrolled in this phase at each of the three doses (i.e., the total sample size of the confirmatory phase is  $3n_2$ ).

To choose  $n_2$  we use predictive power. Let  $df^*$  denote the advantage  $df(z^*, \theta)$  of treatment over placebo at the recommended dose  $z^*$ . Let  $(m, s)$  denote posterior mean,  $m = E(df^* | D_1)$ , and standard deviation,  $s^2 = Var(df^* | D_1)$ , at the end of the dose-finding phase. Details of this estimation depend on the underlying probability model (See Section 6 for a discussion of the actual probability model used). Let  $\hat{\sigma}^2$  denote the posterior mean  $E(\sigma^2 | D_1)$  conditional on all current data. Fixing  $\sigma^2$  at  $\hat{\sigma}^2$  and approximating  $p(df^* | D_1) \approx N(m, s)$ , we can find the minimum sample size  $n_2$  that gives a desired predictive power (90% in our implementation).

## 4 Adaptive Dosing: The Allocator

During the dose-finding phase of the trial we must choose a dose for each new patient. As indicated in the introduction, we consider a decision-theoretic approach. To introduce notation and to clarify the context, we briefly review the general set-up of a Bayesian decision problem. Decision making under uncertainty is choosing an action  $d$  to maximize expected utility

$$U(d) = \int u(d, \theta, y) dp_d(\theta, y).$$

Here,  $u(d, \theta, y)$  is the utility function modeling preferences of consequences and  $p_d(\theta, y)$  is a probability distribution of parameter  $\theta$  and observation  $y$ , which may depend on the chosen action  $d$ . Typically  $p_d(\theta, y)$  is specified as a prior  $p(\theta)$  on the parameters and a sampling model  $p_d(y|\theta)$ . Utility  $u(d, \theta, y)$  must be specified only for a specific realization  $(\theta, y)$  of the experiment. To compute expectation  $U(d)$  means averaging over the quantities that are

unknown at the time of decision making. Often some of the data, say  $y_0$ , is already known at the time of decision making. Assume  $y = (y_0, y_1)$ . Then  $p(\theta)$  and  $p_d(y|\theta)$  are replaced by  $p(\theta|y_0)$  and  $p_d(y|\theta, y_0)$ . See Chaloner and Verdinelli (1995) and Verdinelli (1992) for reviews of Bayesian approaches to decision problems traditionally known as optimal design. Spiegelhalter, Freedman and Parmar (1994), Berry (1993) and Berry and Stangl (1996) discuss general issues related to the use of Bayesian optimal design methods in medical decision problems.

#### 4.1 Dose allocation as a decision problem

Central to our proposal is a utility function that quantifies relative preferences over alternative outcomes. The proposed utility function is related to learning about the unknown dose/response curve. Learning is formalized as minimizing the posterior variance for some key parameter  $g(\theta)$  of the dose/response curve. In the current implementation we choose as key parameter  $g(\theta)$  the mean response  $f(z95_\theta, \theta)$  at the ED95. The posterior variance of  $g(\theta)$  includes uncertainty in the unknown ED95 dose, as well as the unknown response at that dose.

The proposed approach is myopic in the sense that when we consider the optimal dose for the next patient we proceed as if he or she was the last patient to be recruited into the trial. Assume we currently have  $N$  patients enrolled in the trial. Let  $\tilde{y}_k = y_{N+k}$ ,  $\tilde{x}_k = x_{N+k}$  and  $\tilde{z}_k = z_{N+k}$  denote the response, covariate and assigned dose for the next  $K$  patients,  $k = 1, \dots, K$ . Let  $\tilde{y}(k) = (\tilde{y}_1, \dots, \tilde{y}_k)$ ,  $\tilde{z}(k) = (\tilde{z}_1, \dots, \tilde{z}_k)$ , and  $\tilde{x}(k) = (\tilde{x}_1, \dots, \tilde{x}_k)$ , denote responses, assigned doses and covariates up to the  $k$ -th new patient. Let  $D_N$  denote the observed data for the first  $N$  patients, and let  $\tilde{D}$  denote the still missing final responses for already enrolled patients. We define the utility function for choosing the dose  $\tilde{z}_k$  as

$$u_k[\tilde{z}(k), \tilde{y}(k), \tilde{x}(k), \tilde{D}, D_N] = -Var[g(\theta) | D_N, \tilde{D}, \tilde{y}(k), \tilde{x}(k), \tilde{z}(k)].$$

Of course we have to decide upon  $\tilde{z}_k$  before observing  $\tilde{D}$ ,  $\tilde{y}(k)$  and  $\tilde{x}(k)$ . Thus we choose the dose  $\tilde{z}_k$  by maximizing the utility  $u_k(\cdot)$  in expectation, averaging with respect to  $\tilde{D}$ ,  $\tilde{y}(k)$  and  $\tilde{x}(k-1)$ . The relevant distributions for  $\tilde{D}$  and  $\tilde{y}(k)$  are the posterior predictive distributions given the current data. For the covariates  $\tilde{x}_1, \dots, \tilde{x}_{k-1}$  we use the empirical distribution  $p(x)$  from the Copenhagen Stroke Study data base, assuming independence, i.e.,

$$p[\tilde{x}(k-1)] = \prod_{h=1}^{k-1} p(\tilde{x}_h).$$

And  $\tilde{x}_k$  is fixed at “typical” covariate values  $x^o$ , i.e., we find the optimal dose for an average next patient. For  $\tilde{z}_1, \dots, \tilde{z}_{k-1}$  we substitute the optimal

values found by optimizing the expected utilities  $U_1(\cdot), \dots, U_{k-1}(\cdot)$ .

$$U_k[\tilde{z}_k, \tilde{z}(k-1), \tilde{x}_k, D_N] = \int u_k[\tilde{z}(k), \tilde{y}(k), \tilde{x}(k), \tilde{D}, D_N] \times \\ \times p(\tilde{D} | D) p[\tilde{y}(k) | D, \tilde{z}(k)] p[\tilde{x}(k-1)] d\tilde{D} d\tilde{y}(k) d\tilde{x}(k-1). \quad (2)$$

Maximizing  $U_k(\cdot)$  over  $\tilde{z}_k$  we find the optimal action.

#### 4.2 Evaluating Expected Utility

Critical for a successful implementation of the proposed decision theoretic dose allocation is the availability of analytical or efficient numerical integration to evaluate the integrals in (2). Key to our implementation strategy is to rewrite the expected utility integral (2) as an integral with respect to the posterior distribution  $p(\theta | D)$

$$U_k[\tilde{z}(k), \tilde{x}_k, D_N] = \int u_k[\tilde{z}(k), \tilde{y}(k), \tilde{x}(k), \tilde{D}, D_N] \times \\ \times p(\tilde{D} | \theta, D) p[\tilde{y}(k) | \theta, \tilde{x}(k), \tilde{z}(k)] p[\tilde{x}(k-1)] \\ d\tilde{D} d\tilde{y}(k) d\tilde{x}(k-1) p(\theta | D) d\theta. \quad (3)$$

Most models allow efficient random variate generation from the posterior distribution  $p(\theta | D)$  using Markov chain Monte Carlo (MCMC) simulation. See, for example, Tierney (1994) or Gilks et al. (1996) for a summary of MCMC methods. Details of implementing the appropriate MCMC scheme depend on the specific probability model  $p(D | \theta)$ . In Section 6 we will discuss how we implement simulations from the posterior distribution in our model. But we do not need these details in the following discussion. We assume only that by appropriate simulation techniques it is possible to generate an (approximate) posterior Monte Carlo sample  $\Theta = \{\theta^1, \dots, \theta^T\}$  with  $\theta^t \sim p(\theta | D)$ . Using the Monte Carlo sample  $\Theta$  we can evaluate expected utilities by replacing (3) with a corresponding Monte Carlo sample average. For each  $\theta^t$  we simulate covariates  $\tilde{x}_h^t \sim p(x)$ ,  $h = 1, \dots, k-1$ , responses  $\tilde{y}_h^t \sim p(y_h | \tilde{z}_h, \tilde{x}_h^t, \theta^t)$  and missing responses  $\tilde{D}^t$  of current patients  $\tilde{D} \sim p(\tilde{D} | \theta^t)$ , assuming that these models are all available for efficient random variate generation. For each simulated experiment we then evaluate observed utility  $u_k^t = u_k(\tilde{z}(k), \tilde{y}(k)^t, \tilde{x}(k)^t, \tilde{D}^t, D)$ , and replace expected utility by a Monte Carlo average

$$\hat{U}_k(\tilde{z}(k), \tilde{x}_k, D) = \frac{1}{M} \sum_{t=1}^M u_k^t$$

Evaluating expected utility  $\hat{U}_k$  for a grid of possible choices  $\tilde{z}_k$  we find the optimal dose as the dose with maximum  $\hat{U}_k(\cdot)$ . To reduce numerical error

in the approximation of  $U_k$  by  $\hat{U}_k$  we use common random numbers, i.e., whenever possible we use the same Monte Carlo sample  $\Theta$  and random variates  $\tilde{x}_h^t, \tilde{y}_h^t, \tilde{D}^t$  when we evaluate  $\hat{U}_k(\tilde{z}_k)$  for alternative choices of  $\tilde{z}_k$ .

Figure 4.1 shows a typical expected utility curve. Of course, the expected curve changes from week to week. Assume at one time it is optimal to allocate to high doses. As more patients are allocated in that part of the dose range it will usually become more advantageous to allocate new patients in other parts of the dose range to learn about, for example, the response at placebo. While it is difficult to intuitively understand the expected utility curve and the dose assignment at a given time, typical patterns of dose allocations over the course of a clinical trial do seem intuitively meaningful. Figure 9.2 (Section 9) shows an example of dose assignments in a simulated trial.




FIGURE 4.1. Approximate expected utilities  $\hat{U}_k$  plotted against  $\tilde{z}_k$ . The circles show the Monte Carlo estimates of expected utilities at the allowable doses. The curve shows a smooth fit through the Monte Carlo estimates. For comparison a second curve (almost indistinguishable from the other curve) shows a smooth fit through Monte Carlo estimates of expected utilities using half the Monte Carlo sample size.

### 4.3 Additional Randomization: The Recommender

Maximizing  $\hat{U}_k(\cdot)$  delivers the optimal doses  $\tilde{z}_k^*$  be assigned to the next patients,  $k = 1, \dots, K$ . Assume  $\tilde{z}_k^*$  equals  $Z_{j^*}$  in the list of allowable doses. Before actually assigning a dose to a new patient we use an additional randomization. First, we guard against the possibility of a drift in the patient population over the course of the trial by fixing a minimum proportion  $p_0$  of patients assigned to placebo. Second, because of safety concerns we want to avoid unnecessarily high doses. To achieve these two aims we allocate

with probabilities  $p_0$  at placebo and split the remaining probability  $(1 - p_0)$  uniformly over all doses  $Z_j$  within a neighborhood of  $Z_{j^*}$ , defined as the set of all doses less than or equal and having estimated mean response within 10% of the estimated mean response at  $Z_{j^*}$ .

## 5 Optimal Stopping: The Terminator

The dose-finding phase of the trial involves two important decision problems, dose assignment and termination. We discussed an adaptive dose allocation scheme in Section 4. In the present section we discuss the second problem: optimal stopping for the dose-finding phase. At each period  $t$  of the trial, say once a week, we have to decide ( $d_t$ ) whether to terminate the trial and abandon the drug ( $d_t = A0$ ), continue with the dose-finding phase ( $d_t = A1$ ), or terminate the dose-finding phase and switch to confirmatory mode ( $d_t = A2$ ).

Making an optimal decision is complicated. The decision to continue (A1) requires evaluating the expected utility of this action which in turn requires knowing the utility of next week's optimal termination decision for all possible interim results, i.e., we need to solve a sequential decision problem. A standard approach is backward induction: Consider the optimal solution under every possible scenario at the latest possible termination time  $T$  (determined by some maximum number of patients in the trial), and record a table of such decisions, starting from  $T$ . Then, proceeding backwards, we find optimal decisions and their utilities at the penultimate week by averaging those from the following week using the predictive probabilities based on the current week's information. Proceeding backwards in this way we arrive at the current time and the decisions that are optimal for the current. See, for example, Berger (1985, chapter 7).

For the *terminator* we implement a numerical solution of the backward induction problem. The approach is based on a dual strategy of using a reduced action space to constrain the number of scenarios which we need to consider in backward induction, and forward simulation to evaluate expected utility under all relevant scenarios. Central to our approach is a formulation of the problem as a decision problem with a probability model describing all relevant uncertainties and a utility function which describes the relative preferences of possible outcomes.

Alternative Bayesian approaches to optimal sequential design in phase I drug trials are discussed in, among other references, Thall, Simon and Estey (1995), who define stopping criteria based on posterior probabilities of clinically meaningful events. Thall and Russell (1998) define ad-hoc sequential procedures based on monitoring posterior probabilities and evaluate their frequentist properties. Vlachos and Gelfand (1998) follow a similar strategy. Whitehead and Brunier (1995) and Whitehead and Williamson (1998)

use a Bayesian  $m$ -step look-ahead procedure to find an optimal dose for the next  $m$  patients in a dose-finding study.

### 5.1 The Terminator

Before we discuss the underlying decision-theoretic framework, we present the form of the final implementation and attempt to make it intuitively appealing. Let  $z^* = \overline{z95}$  denote the ED95 dose. Let  $(m_t; s_t)$  denote the posterior mean and standard deviation of the advantage over placebo at the ED95 dose conditional on the data available at time  $t$ , i.e.,  $m_t = E[df(z^*, \theta) | D_{1t}]$  and  $s_t^2 = Var[df(z^*, \theta) | D_{1t}]$ . The proposed stopping rule is a function of  $(m_t, s_t)$ . We define cut-offs for  $(m_t, s_t)$  that partition the space into three subsets, corresponding to decisions  $A0$ ,  $A1$  and  $A2$ . Figure 5.1 provides an example partition. The pattern is intuitive. For very small and very large  $m$  the terminator recommends  $A0$  (abandon the drug) and  $A2$  (continue to the confirmatory phase), respectively. For intermediate values the terminator recommends continuing the dose-finding trial. And for smaller  $s$  there is less ambiguity about whether to stop.

		s[t]														
		1.40	1.50	1.60	1.70	1.80	1.90	2.01	2.11	2.21	2.31	2.41	2.51	2.61		
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	-0.36
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.10
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.57
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1.03
1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	1.49
1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	1.95
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2.41
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2.88
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	3.34
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	3.80
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4.26
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4.72
2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	5.18
2	2	2	2	2	2	2	2	2	2	2	1	1	1	1	1	5.65
2	2	2	2	2	2	2	2	2	2	2	2	1	1	1	1	6.11
2	2	2	2	2	2	2	2	2	2	2	2	2	1	1	1	6.57
2	2	2	2	2	2	2	2	2	2	2	2	2	2	1	1	7.03
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	7.49
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	7.96
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	8.42

FIGURE 5.1. Example terminator. The optimal decision for each pair  $(m, s)$  is either  $A0$ ,  $A1$ , or  $A2$ , indicated by entries 0, 1, and 2, respectively.

To decide whether to terminate at time  $t$ , we compute  $(m_t, s_t)$  and use the entry shown in Figure 5.1. For any two scenarios with equal summary  $(m_t, s_t)$  we use the same decision and ignore any additional information

beyond  $(m_t, s_t)$  which might be contained in the data. Thus the decision depends on the current data only indirectly through  $(m_t, s_t)$  which we use to look up the decision table. The main reason for this constraint is to simplify the resulting computations.

## 5.2 A decision-theoretic stopping rule

To find the boundaries in the  $(m_t, s_t)$  table we use a decision-theoretic argument. Specifying a utility function for each outcome of the experiment allows us to find the worth of a given sequential strategy. Maximizing this utility function in expectation defines the optimal termination decision.

### 5.2.1. The Utility Function

We start the process by defining the utilities for all possible outcomes assuming that the whole experiment is known. Expected utilities of decisions can then be found by finding expectations over all unknowns. For a given realization the utility includes a sampling cost for the number of patients recruited into the trial, and a payoff for successfully developing and marketing the drug. Of course that payoff is included only if the drug is approved for marketing by the regulatory authorities. And it depends on the magnitude of the drug's benefit in comparison with placebo as observed in the two phases of the trial taken together – a better drug has a greater payoff.

*Utility under  $d_t = A2$ .*

If we were to make decision  $d_t = A2$  then the resulting utility will include the sampling cost (negative) for the confirmatory phase of the study plus the payoff if the confirmatory study eventually concludes a significant treatment effect.

Let  $B$  denote the event in which the null hypothesis of no treatment effect is rejected in the confirmatory phase at some fixed significance level  $\alpha$ . Let  $\bar{y}_0$  and  $\bar{y}^*$  denote the sample average of the responses under placebo and treatment ( $z^*$ ), respectively, in the confirmatory phase.  $B$  is the event

$$B = \{(\bar{y}^* - \bar{y}_0)/\sqrt{2\hat{\sigma}^2/n_2} > q_\alpha\}$$

where  $q_\alpha$  is the  $(1 - \alpha)$  standard normal quantile. As in Section 3.3,  $n_2$  is the sample size in the confirmatory study, and  $\hat{\sigma}^2$  is the posterior mean of the measurement variance.

Let  $c_1$  denote the sampling cost per patient, and let  $c_2$  denote the payoff for a successful drug. Assuming that this payoff is proportional to the size of the effect we specify  $c_2$  as payoff per point advantage over placebo. Let  $m_2 = E(df^* | D_1, D_2)$ .

$$u(d_t = A2, D_1, D_2) = \begin{cases} -3 n_2 c_1 & \text{if not } B \\ -3 n_2 c_1 + c_2 m_2 & \text{if } B \end{cases}$$

The available data enters into the definition of  $u(\cdot)$  implicitly through  $n_2$ , which depends on  $D_1, m_2$ , which is a statistic depending on  $D_2$ , and  $B_2$ , which is an event in  $D_2$ . The sampling cost includes only those patients in the confirmatory study and omits the sampling cost for the first  $n_1$  patients in the dose-finding phase. Including the latter would add the same term to the utilities of all three alternative actions ( $A_0, A_1$  and  $A_2$ ) and would thus not change the decision problem. Of course, decision  $d_t$  needs to be made before observing  $D_2$ . The relevant expected utility is an average over  $D_2$ :

$$U(d_2 = A_2, D_1) = \int u(d_t = A_2, D_1, D_2) dp(D_2 | D_1). \quad (4)$$

As indicated in Section 3.3,  $n_2$  is a function of  $(m, s)$ . It remains to compute  $P(B|D_1)$  and  $E(m_2|D_1, B)$ . Both can be computed analytically if we use the following approximation:  $p(df^* | D_1) \approx N(m, s)$ .

*Utility under  $d_t = A_0$ .*

If we were to decide  $d_t = A_0$  then the trial is over and thus  $U(d_t = A_0) = 0$ , again omitting the sampling cost for the first  $n_1$  patients.

*Utility under  $d_t = A_1$ .*

Finally, if we were to decide  $d_t = A_1$ , then the utility depends on what we will decide in the next period  $t + 1$ . Let  $n^+$  denote the number of patients enrolled in week  $(t + 1)$  and let  $D^+$  denote the data collected in week  $(t + 1)$ .

$$u(d_t = A_1, D_{1t}, D^+) = -n^+ \cdot c_1 + U_{t+1}^*(D_{1t} \cup D^+),$$

where  $U_{t+1}^*(D_{1,t+1})$  is the expected utility under the optimal action at time  $(t + 1)$ . By taking the expectation with respect to  $D^+$  we obtain

$$U(d_t = A_1, D_{1t}) = \int u(d_t = A_1, D_{1t}, D^+) dp(D^+ | D_{1t}). \quad (5)$$

We are assuming that  $n^+$  is fixed or known. Extension to random  $n^+$  is straightforward. Evaluation of the integral (5) is discussed below in Section 5.2.3.

The optimal decision  $d_t^*(D_{1t})$  at time  $t$  and its value  $U_t^*(\cdot)$  can, in principle, be derived as

$$U_t^*(D_{1t}) = U(d_t = d_t^*, D_{1t}) = \max_{d_t \in \{A_0, A_1, A_2\}} \{U(d_t, D_{1t})\}. \quad (6)$$

There are at least two impediments to implementing (6). First, the solution requires backward induction in the sense that the definition of  $U(d_t = A_1, \dots)$  entails solving (6) for period  $t + 1$ . Second, the definitions of  $U(d_t = A_2, \dots)$  and  $U(d_t = A_1, \dots)$  involve integrals that are typically

analytically intractable. In the following two subsections we outline an implementation strategy based on constrained backward induction and forward simulation. See Müller, Berry, Grieve, Smith and Krams (1999) for a detailed discussion.

### 5.2.2. Constrained backward induction

To allow a practical solution to the backward induction problem we constrain decisions  $d_t$  to depend on the current data only indirectly through  $(m_t, s_t)$ , with  $(m_t, s_t)$  given for a finite grid, i.e., we use  $d^*(m_t, s_t)$ . This amounts to constraining the action space. Instead of allowing decisions to depend on the full information set  $D_{1t}$  they are allowed to depend on the current data only indirectly through  $(m_t, s_t)$ . Even without this constraint, expression (4) for  $U(d = A2, D_{1t})$  depends on  $D_{1t}$  only indirectly through  $(m, s)$ . We write  $U(d = A2, m, s)$  to emphasize this. However,  $U(d_t = A1, D_{1t})$  depends on  $D_{1t}$  beyond the statistic  $(m, s)$ . See Section 5.2.3 for the definition and evaluation of  $U(d_t = A1, m, s)$ . We will use  $U^*(m, s) = U(d_t^*(m, s), m, s)$  for the expected utility when making an optimal decision.

Computation proceeds as follows. Starting with an initial guess for  $d^*(m, s)$  and  $U^*(m, s)$  we update the tables using (4) through (6). We continue updating until the table remains unchanged over a complete cycle of updates.

### 5.2.3 Forward Simulation

There remains the problem of evaluating expected utility in (5). We use forward simulation (Carlin, Kadane and Gelfand, 1998). The idea is to simulate a trial including up to the maximum number of possible patients in the dose-finding phase. Then, many such trials, say  $T$ , are simulated and the results averaged.

Figure 5.2 illustrates how these forward simulations are used to compute posterior integrals. Assume we need a posterior integral conditional on data  $D_{1t}$ . We first compute the corresponding summary statistic  $(m_t, s_t)$ , and then approximate the desired posterior integral by a sample average over all simulated trials whose trajectories pass through  $(m_t, s_t)$ . If  $(m_t, s_t)$  were in fact a sufficient statistic for the unknown parameter vector then the sample average would provide a (simulation) consistent and unbiased estimate of the desired expectation. Carlin et al. (1998) use forward simulation in such a set-up. In general, some approximation is involved.

Let  $D_{1t}$  and  $n_{1t}$  denote the data and number of enrolled patients at the time of making the decision  $d_t$ . Let  $\tilde{D}$  denote the future data, including the remaining  $n_1 - n_{1t}$  patients, as well as still missing final measurements on current patients who have been recruited during the last 13 weeks. The  $T$  trials for the forward simulation are generated by simulating from the posterior predictive distribution  $p(\tilde{D}|D_{1t})$ . Denote the simulated values as




FIGURE 5.2. Forward simulation. The figure shows trajectories in the  $(m_t; s_t)$  space for some simulated trials. See the text for how these simulations are used to compute a posterior integral.

$\tilde{D}^{(j)}$ . For each week of each simulation we record the moments  $(m_t^{(j)}, s_t^{(j)})$ .

Assume in the constrained backward induction we have current estimates  $d^*(m, s)$  and  $U^*(m, s)$ . Assume we are updating a specific cell  $(m, s)$  and want to find the optimal decision for that cell. Expected utility  $U(d = A2, m, s)$  is easily computed using (4), and  $U(d = A0, m, s) = 0$  is fixed. To compute integral (5) for  $U(d = A1, \dots)$  we consider the subset  $A$  of simulations which pass through cell  $(m, s)$  at some time, i.e.,  $j \in A$  if  $(m_{t_j}, s_{t_j}) = (m, s)$  in some week  $t_j$ . Let  $T'$  denote the number of indices in  $A$ . Approximate (5) by the sample average

$$\hat{U}(d = A1, m, s) = -n^+ c_1 + \frac{1}{T'} \sum_{j \in A} U^*(m_{t_j+1}^{(j)}, s_{t_j+1}^{(j)}),$$

using the current estimates for  $U^*(m, s)$ . The approximation for  $U(d = A1, \dots)$  is the same for all  $D_{1t}$  with the same summary statistic  $(m, s)$ . Thus we write  $U(d = A1, m, s)$ . Compare with  $U(d = A2, m, s)$  and  $U(d = A0, m, s)$  to find the optimal decision  $d^*(m, s)$  and its value  $U^*(m, s)$ . Repeating the same process for all cells  $(m, s)$  updates the currently imputed values for  $d^*(m, s)$  and  $U^*(m, s)$ . We repeat updating until none of these values change after a cycle of updates.

### 5.3 From stopping rule to terminator

The process described in the previous two sections is computationally intensive. It is impractical to repeat this computation every week, each time with slightly different data  $D_{1t}$ . Instead, we build a single decision table such as the one in Figure 5.1 by combining estimated utilities  $\hat{U}(d_t = a, m, s)$ ,

$a \in \{A0, A1, A2\}$ , computed under a set of “typical” dose/response curves. This gives rise to a “static terminator.” In the current implementation we use four dose/response curves. For each cell in the  $(m, s)$  we compute the average expected utility approximation under the four curves. For A2 we use equal weights. For A1 we use weights proportional to the number of simulated trajectories in each experiment that passed through the specific  $(m, s)$  cell.

## 6 Dose/Response Curve

### 6.1 Probability model

The choice of the probability model for  $f(z)$  is guided by the following considerations. First, the model should allow for posterior inferences that are analytically tractable to facilitate efficient computation of expected utilities. Second, we want a flexible model which supports a priori a wide range of dose/response curves. Although a monotonic curve with asymptotes near 0 and for large doses is likely, the model should allow for non-monotonicity and other irregular features. This is particularly important since death is a possible consequence of stroke. The drug may have a different dose-effect as regards rehabilitating patients and mortality.

Based on these considerations we chose a normal dynamic linear model (NDLM). West and Harrison (1997) give a formal definition and discussion of NDLMs. Before we describe details of the model, we outline some important features. Denote by  $Z_j, j = 0, \dots, J$ , the range of allowable doses, and by  $\theta_j = f(Z_j, \theta), j = 0, \dots, J$ , the vector of mean responses at the allowable doses. The underlying idea is to formalize a model which locally, for  $z$  close to  $Z_j$ , fits a straight line  $\theta_j + (z - Z_j) \delta_j$ , with level  $\theta_j$  and slope  $\delta_j$ . This is illustrated in Figure 6.1. When moving from dose  $Z_{j-1}$  to  $Z_j$  the parameters  $\alpha_j = (\theta_j, \delta_j)$  change by adding a (small) so-called evolution noise  $e_j$  and adjusting the level  $\theta_j = \theta_{j-1} + \delta_{j-1}$ . Let  $Y_{jk}, k = 1, \dots, \nu_j$ , denote the  $k$ -th response observed at dose  $Z_j$ . Therefore,  $Y_j = (Y_{jk}, k = 1, \dots, \nu_j)$  is the vector of responses  $y_i$  of all patients with assigned dose  $z_i = Z_j$ . The resulting model is

$$Y_{jk} = \theta_j + \epsilon_{jk}, \quad j = 1, \dots, n, \quad k = 1, \dots, \nu_j$$

and

$$\begin{pmatrix} \theta_j \\ \delta_j \end{pmatrix} = \begin{pmatrix} \theta_{j-1} + \delta_{j-1} \\ \delta_{j-1} \end{pmatrix} + e_j, \quad (7)$$

with independent errors  $\epsilon_{jk} \sim N(0, V\sigma^2)$  and  $e_j \sim N(w_j, W_j\sigma^2)$ . The first equation describes the distribution of  $Y_{jk}$  conditional on the state parameters and is referred to as the “observation equation.” The second

equation formalizes the change of  $\alpha_j$  between doses and is referred to as the “evolution equation”. For a given specification of  $\{V, W_j, j = 1, \dots, n\}$  and prior

$$p(\alpha_0) = N(m_0, C_0), p(\sigma^{-2}) = \text{Gamma}(n_0/2, S_0/2)$$

with given hyperparameters  $m_0, C_0, S_0$ , and  $n_0$ , there exists a straightforward recursive algorithm to compute posterior distributions  $p(\alpha_j|Y_0, \dots, Y_j)$  and therefore any other desired posterior inference. It can be shown that  $p(\alpha_j|Y_0, \dots, Y_j)$  is bivariate normal  $N(m_j, C_j)$  with some moments  $m_j, C_j$ . For later reference: the predictive distributions  $p(Y_j|Y_0, \dots, Y_{j-1})$  are normal distributions with moments  $f_j, Q_j$ , and the posterior distributions  $p(\alpha_j|Y_0, \dots, Y_j)$  are bivariate  $N(m_j^*, C_j^*)$ . West and Harrison (1997) give the recursive equations to compute  $m_j, C_j, f_j, Q_j, m_j^*, C_j^*$  and other posterior moments. An algorithm known as Forward Filtering Backward Sampling (FFBS) allows efficient random variate generation from the full posterior distribution. It is described in Frühwirth-Schnatter (1994) and Carter and Kohn (1994).

Specifying the evolution variances  $W_j$  is not easy. As an alternative to specifying  $W_j$  *a priori*, West and Harrison (1997) propose to define  $W_j$  as a scalar multiple of the posterior variance  $C_j$ , i.e.,  $W_j = (1 - r)/r C_j$  for some scalar  $r \in (0, 1)$ . This corresponds to thinking of the evolution noise  $e_j$  as discounting some of the current information on  $\alpha_j$  as represented by the posterior variance  $C_j$ . The scalar factor  $r$  is the “discount factor”. A large discount factor implies a small variance  $W_j$ . This in turn implies that a small change  $e_j$  between adjacent doses is likely and there is greater smoothing. On the other hand, a smaller  $r$  implies larger  $e_j$  and therefore allows for larger changes between adjacent doses.

## 6.2 Prior specification for the NDLM

A minor shortcoming of the NDLM in the present application is that the prior specification with the hyperparameters  $m_0, C_0$  and  $r$  does not naturally allow us to fix arbitrary desired prior moments for  $\theta_j = f(Z_j)$ . Only  $E(\theta_0)$  and  $Var(\theta_0)$  are fixed as  $m_0$  and  $C_0$ . Prior expectation and variance for  $\theta_j, j > 0$ , are then implied by the evolution equation. To increase the number of prior parameters and allow for essentially arbitrary prior moments  $E(\theta_j)$  and  $Var(\theta_j)$  we augment the model by introducing dummy observations  $\tilde{Y}_j, j = 1, \dots, J$ , with associated observation variance  $\tilde{\sigma}_j^2$ . When going through the FFBS scheme for posterior inference in the NDLM we proceed then as if  $\tilde{Y}_j$  were data sampled from  $\tilde{Y}_j \sim N(\theta_j, \tilde{\sigma}_j^2)$ . By appropriate choice of  $\tilde{Y}_j$  and  $\tilde{\sigma}_j^2$  we can achieve any prior moments for  $\theta_j$ , subject only to technical constraints (for example, the marginal prior variances  $Var(\theta_j)$  can not be larger than those implied without the dummy observations).




FIGURE 6.1. The NDLM fits a smooth curve to the data by defining for each dose  $j$ ,  $j = 0, \dots, J$ , a locally straight line, parametrized by  $\alpha_j$ . The plot shows the local lines fit at doses  $z_j = 0.4$  and  $z_{j0} = 0.9$ . The solid circle marks the point and a solid line segment indicates the fitted line. Between dose levels the slope of the line changes by adding an evolution noise  $e_j$ .

## 7 Longitudinal Model

For many patients with missing final score  $y_{ij}$  we have earlier scores,  $y_{ij}$  available, where  $j < M$ . From historical data (Copenhagen Stroke Database) we expect  $y_{ij}$  to be a good predictor of  $y_i$ , even as early as week 4 ( $j = 4$ ). To allow formal imputation we extend the probability model with a longitudinal data model for  $(y_{ij}, j = 1, \dots, 13)$ .

$$y_{ij}|y_i \sim N(m_j + a_j y_i, s_j^2)$$

We assume that given a patient's current status, the patient's future status does not depend further on the patient's previous history. That is, given  $y_{ij}$ ,  $y_i$  and any earlier response  $y'_{ij}$ ,  $j' < j$ , are conditionally independent. This implies

$$p(y_{ij}|y_{i,j+1}, \dots, y_{iM}) = p(y_{ij}|y_{i,j+1}) = N(m'_j + a'_j y_{i,j+1}, s'^2_j)$$

with parameters  $m'$ ,  $a'$  and  $s'_j$  which are easily derived from  $m$ ,  $a$  and  $s_j$ .

In the following discussion we will use  $\theta = (\theta_1, \delta_1, \dots, \theta_K, \delta_K, \sigma, m_1, a_1, s_1, \dots, m_M, a_M, s_M)$  to denote the vector of all unknown model parameters.

We use the Copenhagen Stroke Database (Jorgensen et al. 1994) to derive informative (conjugate) priors for  $m'$ ,  $a'$  and  $s'_j$ , allowing us to obtain conditional posterior distributions in closed form. All data, including that from patients who have not yet completed the full 13-week study period, are used in updating the longitudinal model parameters. To simplify computations, in the current implementation we start updating the longitudinal

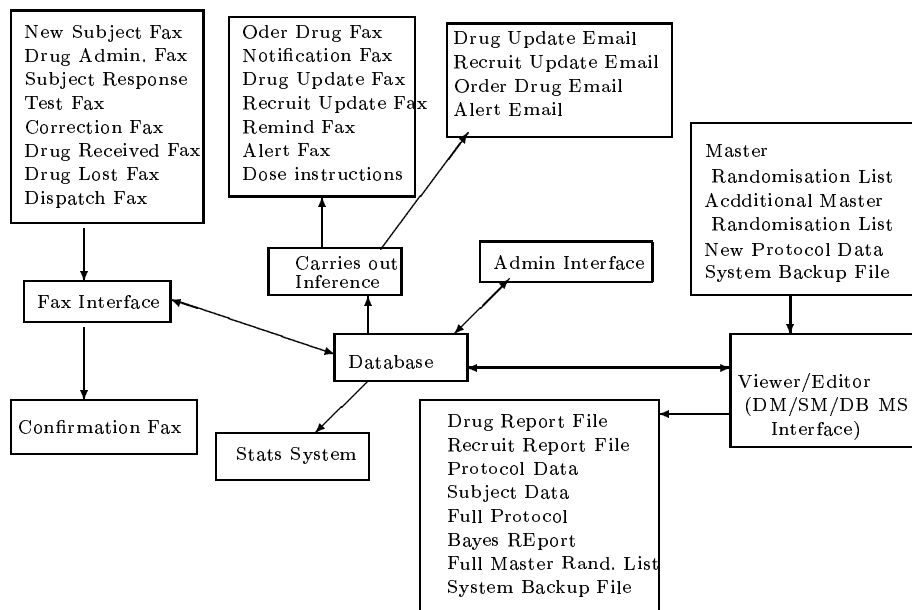
data model only when some minimum fraction of the patients have completed final observations (currently 35% of the data set).

## 8 Investigator/System Interface

### 8.1 Introduction

The purpose of the present section is to describe how the program described in the previous sections can be used in an actual clinical trial. The investigator/system interface incorporates the information received from the investigators, orchestrates the running of the statistics programs, transmits to the investigators the details of what dose to administer the patient, and provides a means of, monitoring the progress of the study.

#### System Overview



### 8.2 Using the System: The Investigator

Upon admitting a patient suitable for the study, the investigator:

1. Assesses their weight and their baseline score on the Scandinavian Stroke Scale,

2. Enters this information and the patient's initials and date of birth on a single page pre-printed form,
3. Faxes the form to the data system,
4. Receives a confirmation of her fax, with the number for the patient within the study, and the details as read from the investigator's form by the system so they can be checked, and
5. A minute or two later, receives a fax with instructions on how to make up the dose to give the patient.

The dose assigned to a patient is specified to the investigator, or the investigator's pharmacy, by identifying two vials and instructing how much is to be taken from each and combined to make the dose. A unique number identifies each vial. Vials contain one of three concentrations of the drug or placebo.

This allows a wide range of doses to be used and a larger than usual number of different dose levels within the range is available for the dose allocation algorithm. The system is also able to scale the dose according to the patient's weight.

As well as the fax forms with their minimum of information, the investigators complete more detailed study forms and pass them to the CRAs in the usual way.

To enable updating the current distribution of the relevant parameters, the investigators complete a form indicating the results of follow-up visits for all patients in the study. They also fax these to the system. Should a patient develop complications the investigator can request being unblinding as regards that patient, again by completing a form and faxing it to the system. The system replies with details of the dose assigned to the patient.

Faxed forms are being used rather than an internet and Web browser-based interface because of the lack of availability of such facilities in some participating centers, and the fear that if the study sponsor provided the necessary facilities they would be abused or stolen. Faxed forms are slower and less reliable than an internet-based solution, and implementing an interface would be easier if internet access could be effected.

### *8.3 Using the System: The Study Administrators*

The system will hold a large amount of relevant and up-to-date information of interest to the Study Administrators (Statisticians, Data Managers, Study Manager and the Data Safety Monitoring Board). This is made available to them through a conventional PC-based program that allows the data to be navigated and viewed.

As well as the data, all the fax images can be viewed. This allows the Data Manager to check the data on the form against that read in by the

28 Berry, Müller, Grieve, Smith, Parke, Blazek, Mitchard and Krams

Screen shot of program

Full page Figure.  
Figure not available in postscript

system and correct it where necessary. A number of events, e.g. invalid faxes, receipt of correction forms, and requests to unblind patients, are flagged to the Data Manger to check and make corrections if necessary.

The data are available more quickly and more accessibly than for conventional trials. An additional advantage of having such a system is that it can proactively send reminders to investigators. For example, it can remind them when follow-up checks are due on patients. And it can let them know when their pharmacy is running low so they can request further packs of vials. The information on doses, effective doses and the status of the study as determined by the terminator algorithm are blinded to all but the Data Safety Monitoring Board.

#### *8.4 Using the system: The Pfizer pharmacy*

The pharmacy at Pfizer is responsible for providing a Master Randomization List of all the vials that may be used in the study, their number, the concentration they contain and their expiration date. Whenever the pharmacy dispatches a shipment of vials to a center they send a fax to the system indicating the vials sent.

Thus, the system has the information it requires to specify doses to investigators in terms of the blinded vial numbers.

The system is informed that the vials are available at a center by the investigator faxing in a form acknowledging receipt of a shipment. The investigators also have a form they can fax should a dose administered be other than that directed and a form indicating that vials have been lost or broken.

The system sends a fax or e-mail to the Pfizer pharmacy when a center's supply of any particular concentration of drug is getting low. It also provides regular reports to the Pfizer pharmacy indicating what vials are at which center and which vials are in transit.

#### *8.5 How the system is constructed*

The system comprises two servers and a number of connected client PCs. There is one Fax server and one Statistics server. The Fax server runs the Data System and the Statistics server runs the Statistics Programs. The Data System comprises:

- A database built using Oracle 8,
- A fax reception package Zetafax,
- The form recognition package Teleforms with the forms designed specifically for this system and the associated Basic scripts to validate the data and enter into the database,

An example fax:

Full page Figure.  
Figure not available in postscript

- The Study Administrators client program, written in Visual Basic,
- The Fax-Statistics interface programs written in C and UNIX shell script.

### 8.6 *The development process*

The system has been developed in two phases.

1. The first phase concentrated on the core areas of:
  - (a) Receiving patient data on faxes,
  - (b) Storing the data in a database,
  - (c) Passing the information to the statistics programs, running the programs and storing the results,
  - (d) Providing facilities to browse and edit the data.
2. The second phase extends the first phase system by:
  - (a) Sending dose information to the investigator,
  - (b) Managing the vial data,
  - (c) Extending the Study Administration facilities,
  - (d) Providing proactive reports, warnings and reminders.

For each phase:

- Detailed Use Cases and User Requirements were drawn up that detailed what the system had to achieve.
- These were refined into System Requirements that detailed the functionality and properties of a System.
- A Design was then drawn up for the system, including a definition of the interface between the Data System and the Statistics
- The components identified in the Design written.
- Individual components of the system were tested.
- The whole system was tested.
- The system was then handed over for acceptance testing then user testing.

The System dataflow model

Full page Figure.  
Figure not available in postscript

In the first phase it soon became clear that the Data System would be written before the fax form designs were finalized or the statistics programs were finished. So we decided to develop two versions during in the first phase, the first using early versions of the fax forms and statistics programs and the second only after the fax form and statistics programs for the first phase had been finalized. The first version was released after the component testing was complete but holding back some development effort for later rework and the system testing was held over until the second version was complete.

This had the advantage of allowing a thorough evaluation of the Study Administration client program in the first version and for the comments to be taken into account in the rework for the second version and in not duplicating any testing effort.

The first phase took about 18 person-months to develop. At the time of this writing the second phase is not yet complete.

### 8.7 *Validation process*

The system was validated by:

- Project reviews to ensure consistency of the requirements documents and the design.
- Code reviews.
- Component testing, a component test plan was drawn up and light weight component test scripts drawn up stating what was to be tested but not how it was to be tested. The tests were then performed by the developers themselves.
- System testing, a system test plan was drawn up, enumerating the test scripts to be written and cross referring the test scripts to the system requirements. The system test scripts contained more detailed instructions on what steps to take to perform the tests and on the expected outcomes. Independent testers ran the system tests. The system testing also included some unscripted testing of the user interfaces by an independent tester, this proved particularly effective at finding faults.

## 9 Simulation Results

### 9.1 *Dose range, data properties and underlying curves*

The number of doses used is arbitrary, with a greater number being better from the perspective of learning about the dose-response curve. But logistical considerations limit the number in practice. We used sixteen possible

doses in all simulations: 0, 0.1, 0.2, . . . , 1.5. These doses are realistic for the actual trial; but the set of doses to be used is not yet final. Equal spacing is natural for the NDLM but unequally spaced doses can be accommodated.

We investigate the properties of our design by assuming a variety of dose-response curves. These were chosen to reflect different scenarios that might be expected in a dose-response study (Figure 9.1).

- Curve 1: Null or flat dose-response, corresponding to no drug effect.

The shape of the next three curves is logistic with the lower 4 doses having no improvement over placebo, and the highest 5 doses forming a plateau at the maximal drug effect.

- Curve 2: Minimal clinically relevant benefit over placebo (2 points on SSS) at the highest dose.
- Curve 3: Modest benefit over placebo (4 points) at the highest dose.
- Curve 4: Very large benefit over placebo (8 points) at the highest dose.

The next two curves reach their maxima earlier (Curve 5) and later (Curve 6) than those above.

- Curve 5: Modest benefit over placebo (3.3 points) but none until the last 6 doses.
- Curve 6: Modest benefit over placebo (4 points) where the increasing is over the range of the first 6 doses similar to Curve 3 but with increase occurring 0.5 sooner in terms of dose.

The last curve we consider is the only one that is not monotonic.

- Curve 7: Bell-shape with increasing effect up to the middle doses and then the effect drops back to placebo level.

In generating weekly patient data we use the Copenhagen Stroke Database (Jorgensen et al. 1994). The SSS data for patients assigned to placebo are as in the CSD. A patient assigned to a positive dose responds as in the CSD, but with the improvement depending on the response at that dose for the curve assumed. For example, if Curve 3 or Curve 6 is assumed and the patient is assigned to dose 1.5 then the patient's change from baseline SSS is 4 points greater than actual.

## 9.2 Process and goals

We simulate a number of trials from a given curve. The statistical system incorporates software to enable us to generate patient data from the underlying dose-response curve, simulate the adaptive allocation of patients

to doses, estimate the dose-response curve given the current patient data and make a decision regarding stopping the dose-finding phase of the trial.

We examine general properties of the algorithm, as follows.

- Doses allocated, depending on the underlying dose-response curve.
- Accuracy of estimate of dose-response curve.
- Sample size in comparison with standard designs.

We also wish to learn about the effect of different settings of the initialization files on the time taken to run the allocator and terminator functions. To use the system in an actual trial it must allocate a dose to the next patient in a matter of a few minutes.

There are a great deal of variable parameters in this approach and it is the aim of simulation to find out the effect of these parameters and identify optimal settings which allow best inferences to be made about the majority of the underlying curves.




FIGURE 9.1. Nominal dose-response curves used in simulations

### 9.3 Initialization file settings

An initialization file controls the critical aspects of running the algorithm. For example, the default values for the prior estimates and longitudinal model analysis are from the Copenhagen Stroke Database. Default settings for Markov chain Monte Carlo (MCMC) iterations have been chosen to provide reasonable estimates of the parameters of interest but also to keep the calculation time within reasonable limits. Many options that can be set in the initialization file allow fallback positions should the software not

function correctly. Default allocations include allocation at given posterior quantiles and random allocation at prespecified doses.

Other settings are invoked in generating data for simulations. In running a “trial” a patient database is created. If the settings indicate that the database should include deaths then this is incorporated. The patient data file is then read in sequentially (one week’s worth of responses at a time) and the doses are adaptively allocated to the patients in trial. The adaptive nature of the design means the simulation must be run with patients being allocated to doses and patient responses being generated sequentially.

The simulation process allows for setting the dose range to be used, the inclusion and exclusion criteria for baseline scores, the patient horizon (maximum number of patients), the maximum number of observations per patient and properties of the distribution and scoring method for deaths in the generated patient database. The initialization file also controls the specification of the prior knowledge about the form of the dose-response curve and controls some aspects of the eventual design of the study – for example, the proportion of patients allocated to placebo.

In all simulations reported here we have set a lower limit of the sample size to be 250. This lower bound allows for getting at least a modest amount of data about the safety of the drug. In addition, we set the upper limit on the sample size to be 1000. Some upper limit is helpful in carrying out the necessary calculations. Also, an upper limit is appealing to drug company management to facilitate budgeting and planning. The limit of 1000 is moderately large, as verified using simulation. When we set the limit to 5000, only occasionally did the terminator go beyond 1000 patients, irrespective of the dose-response curve assumed.

#### 9.4 *Simulation process*

In most of the simulations reported here, 100 trials were generated for each curve. When initialization settings were changed, these were made one item at a time in order to investigate the effects and to keep the number of sources of change to a minimum at any one step. For example for any one nominal curve we simulated 100 trials with no deaths and with one setting of minimum and maximum sample sizes, one setting of placebo proportion and one setting of the different NDLM parameters. Then we introduce deaths with a particular method of scoring death, holding all other factors constant. In this way we have attempted to build up a picture of the influence of different factors on our estimation of the underlying curve and on the inferences derived when we stop the trial.

To evaluate the overall and long-term performance of the algorithms and for comparison with standard designs we performed 1000 additional trial simulations on selected curves and for selected settings of the initialization file. These larger simulations are used to calculate Type I error probabilities (significant dose-response when the underlying dose-response curve is flat)

and power of the procedure (proportion of trials with statistical significance for a nominal dose-response curve with the hypothesized clinically relevant difference).

The statistical package S-Plus was used to display the data from each simulation run – we examined the overall picture across the 100 trials and also observed the results for the individual trials.

Although we can evaluate the Type I error and power of our procedure there are very few metrics which will allow us to explicitly define its performance. As a result most of the evaluation from simulations has involved graphical presentation of the simulation results and comparing the dose allocation, estimation of the curve and stopping decision against what we would expect given knowledge of the underlying curve. We also compare the results of simulations against each other in order to identify any “oddities.”

### 9.5 *Conventional dose-finding study*

We compare our results against conventional dose-finding study, one with three doses plus placebo. For a clinically relevant difference we wish to detect, for an estimate of variability, and for appropriate power and significance levels, we calculate sample sizes for each group. There are many assumptions underlying this calculation. The variability in the actual data may be greater than that assumed (taken from the Copenhagen Stroke Database). The effect of treatment may not be as great as assumed, or it may be greater. The location of the best effect may be outside the dose range assumed. And characteristics of the patient population may be different from those in other studies.

Any conventional study is saddled with such assumptions. If any are wrong then it may not be possible to draw a firm conclusion or the sample size may in retrospect be too large. Conventional studies may be able to recalculate variability and adjust sample sizes accordingly, or there may be other ways to adjust for some of the effects above while the study is ongoing, but in principle many factors are fixed at the outset of the study.

Table 1 shows sample sizes for a conventional study powered at 90% to detect a difference of 2 points or 4 points improvement over placebo. This assumes a 5% significance and testing using Dunnett’s correction for three active doses and placebo. Estimates of variance in change from baseline have been drawn from our analysis of the Copenhagen Stroke Database and reflect variability seen in an untreated patient population.

We polled three clinicians about their choice of doses for the conventional dose-response study described above in the absence of any prior information about our underlying dose-response curves. Given those doses, the sample sizes above, an estimate of variability in the responses and the known responses at each of the doses we were able to construct true power calculations of Dunnett’s test for the study. These are shown in Table 2.

It is quite clear from these results that the choice of doses has a large

TABLE 1. Sample size estimates for a conventional dose-response study design.

Improvement over placebo to detect	Assumed variance; change from baseline	Total sample size (in 4 groups)
2 pts.	150	3216
4 pts.	150	804

TABLE 2. Power calculations for the conventional dose-response study design. Curve 3 – Maximum improvement over placebo = 4pts.

Improvement over placebo to detect	Assumed variance	Clinician 1			Clinician 2			Clinician 3		
		0	0.06	0.2 0.5	0	0.5	0.75 1.0	0	0.33	0.66 1.0
2	150			0.13			>0.99			>0.99
2	50			0.27			>0.99			>0.99
2	250			0.09			>0.99			>0.99
4	150			0.07			0.88			0.86
4	50			0.11			>0.99			>0.99
4	250			0.07			0.66			0.63

impact on the true power of the test. Clinician 1 adopted a conservative approach in his choice of doses and as a result ends up being underpowered for most of the underlying dose-response curves which we looked at. Clinicians 2 and 3 chose very similar doses and so on the whole their true powers are similar across curves. The overall picture changes for different curves. For curve 6, which changes early in the dose range, clinician 1's choice of doses picks up the critical part of the dose-response curve and so his true power is approximately 90%. In this curve clinicians 2 and 3 are overpowered. In curve 5, which changes late in the dose range, none of the clinicians have adequate power.

For curve 3 it can be seen quite easily that if we had powered up to detect a difference of 2 points improvement over placebo then we would have been grossly overpowered under the choices of doses by clinicians 2 and 3. Interim analyses may have been performed in order to stop the study early in this trial. However, it is still true that even when the variability is substantially larger than we originally anticipated we have a total sample size which is too big given the observed underlying dose-response curve. This seems wasteful of resources and time.

In the case where we sample sized to detect a difference of 4 points improvement over placebo we would hope that our estimate of the true power is around the 90% used in the sample size calculation. Unfortunately clinician 1 misses most of the ascending part of the dose-response curve and so is underpowered. Clinicians 2 and 3 almost achieve the 90% power. In both cases though their choice of doses would mean that if there were any safety concerns about the top dose (at which we have a significant

difference from placebo) then neither of their lower doses would achieve significance. This may lead to the study being repeated. We can also see that if the observed variability is different from our prior estimate then their true power is affected.

This illustrates that even in the “objective” world of frequentist sample sizing our judgements about likely effect size, estimate of variability and choice of doses can have a substantial impact on the inferences we make at the end of the study.

### 9.6 Results of simulations

Consider Curve 3, where the maximal improvement over placebo is 4 points. The ED95 for this curve is approximately 1.0. Placebo is allocated with a fixed minimum proportion, in this case 0.15. One patient accrues per day. Figure 9.2 shows the frequency distribution of doses assigned in a single simulation (left panel) and the actual doses allocated by time over the course of the trial (right panel). The trial lasted 36 weeks and 250 patients were accrued. Therefore, at the time of stopping, complete information concerning the 13 weeks of the study was available on about 2/3 of the 250 patients.

Early in the trial the algorithm seeks to identify the minimum and maximum effects for these are critical references against which the ED95 must be evaluated. In the latter phases of the study the algorithm homes in on the ED95 value, while also getting information about the minimum and maximum. The histogram shows that a small fraction of doses are allocated in the dose range from 0.1 to 0.7. The nominal dose-response curve in this range is nearly flat and these doses are away from the area of principal interest, the region containing the ED95.

Figure 9.3 shows the distribution of doses across 100 simulations. Superimposed on this is a summary of the fitted dose-response curves across these simulations. The diamonds represent the true (assumed) dose-response curve. The middle line shows the median of the fitted values with the quartiles of the distribution of the fitted values at each dose. This figure shows that the low doses have a low probability of being allocated. It also shows that pre-specified minimal proportion of placebo doses of 15% is achieved. The fitted curves accurately estimate the underlying dose-response curve.

Figure 9.4 (left panel) shows the distribution of observed effect sizes ( $m$ ) at the ED95 – measured as the number of points improvement over placebo on the Scandinavian Stroke Scale – against the standard error ( $s$ ) of this estimate. The triangles represent those simulations in which the decision was made to go into a confirmatory phase and the circles correspond to simulations in which the decision is made to not go confirmatory. The sizes of the symbols are proportional to the trial sample sizes. Recall that the terminator function is called only after 250 patients have accrued and that the maximum study size is 1000. The figure shows that when the trial is

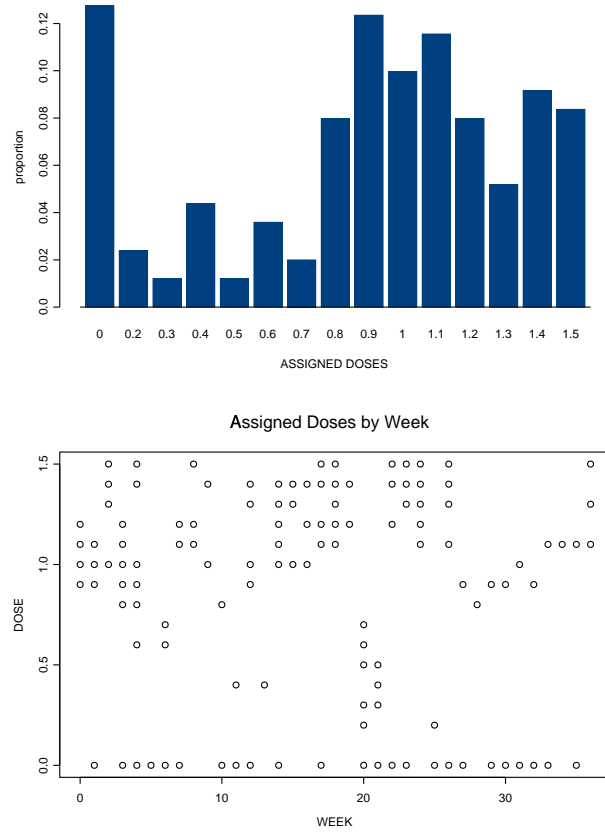


FIGURE 9.2. Results of a single simulation assuming Curve 3. Distribution of doses (top panel) and doses allocated over time (bottom panel).

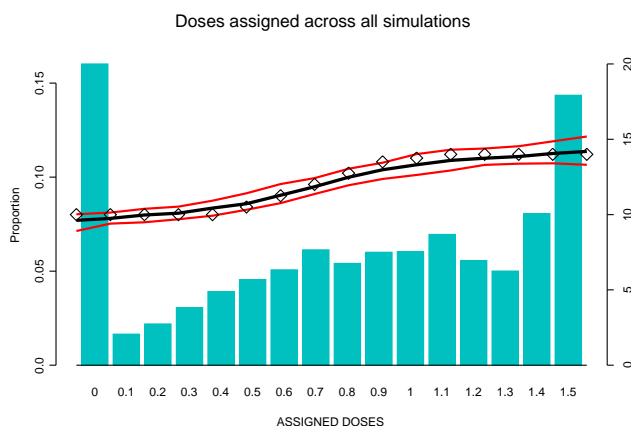


FIGURE 9.3. Allocation of doses across 100 simulations, assuming Curve 3 for the dose-response relationship.

stopped early it is because the effect is probably large. Naturally, for smaller sample sizes there is a bias in the observed response at the ED95.

For Curve 3, with its maximum of 4 points improvement over placebo the algorithm makes very few decisions to terminate the study because of lack of an observed benefit. In Figure 9.4. we see a “funnel” shape. This may be because in order to identify a small effect we will need to continue the trial to a larger number of patients and so we will end up with better precision in this estimate (smaller “s”).

Figure 9.4 (bottom panel) shows distributions of the sample size taken to reach a decision using the stopping rule, the observed effect size, the standard error of the effect size and the posterior estimate of sigma. The bias in observed effect size at the ED95 can be seen from the distribution here, as can the fact that in the majority of cases the decision to stop has been taken comparatively early. The fact that some simulation runs have continued on to 1000 patients may be influencing the standard error of the estimated effect size causing the bimodality we see in this distribution – the longer we continue the more certain we can be about the estimate of the effect size.

Figure 9.5 shows similar results for 200 simulations assuming the null or flat dose-response curve (Curve 1). The proportion of simulations in which the decision is to go into a confirmatory study is made is very low (2.2%). In the terminator function if the observed effect is less than 2 points improvement over placebo then we conclude that the drug has insufficient clinical effect to warrant further development. The bottom panel of Figure 9.5 (upper left histogram) shows that the decision to stop was made after relatively few patients (median of 258 where the minimum is 250). As indicated in the introduction, an important benefit of our adaptive scheme

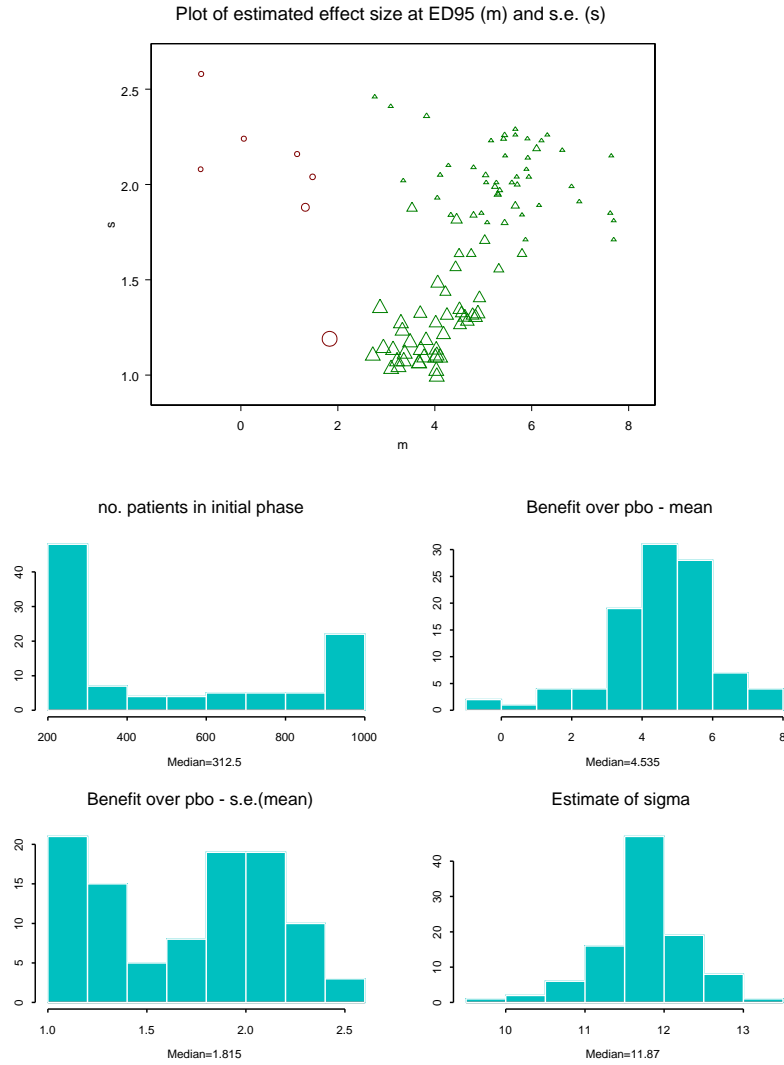


FIGURE 9.4. Distribution of observed effect sizes at the ED95 with corresponding precision. Dose-response is that of Curve 3. In the top panel symbol size is proportional to trial sample size at stopping. The bottom four panels show frequency distributions of sample size, posterior mean benefit over placebo at the ED95, its standard error, and the error standard deviation  $\sigma$ .

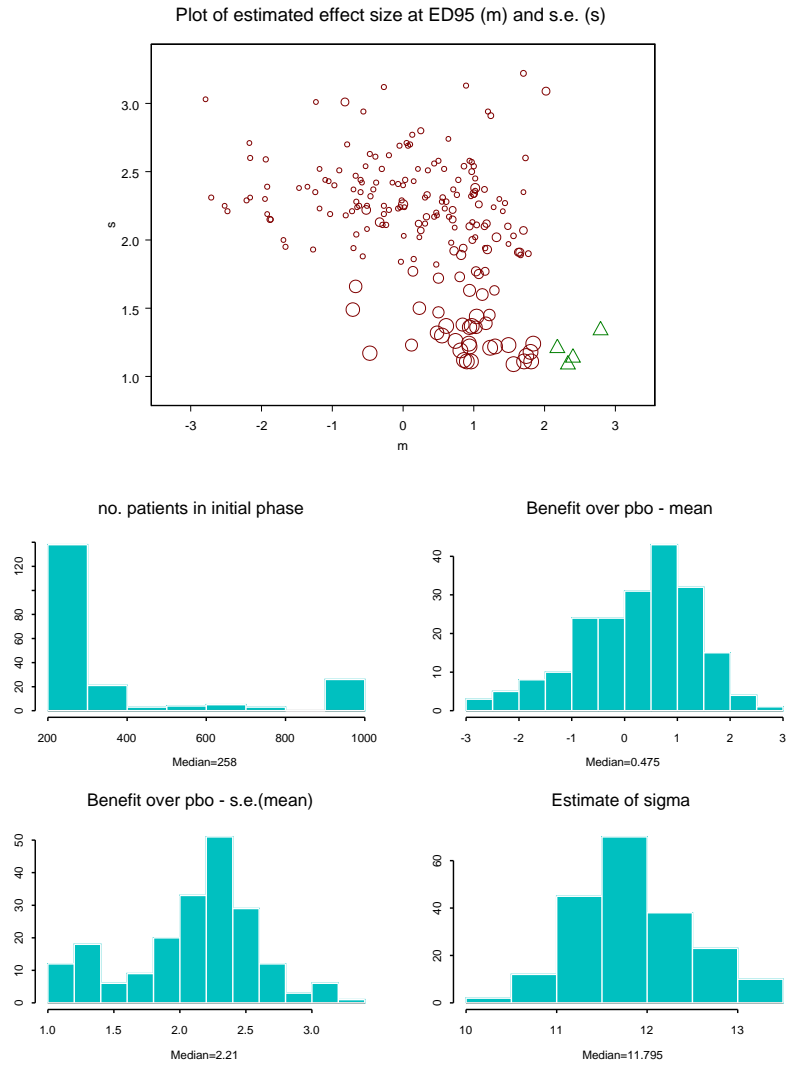


FIGURE 9.5. Distribution of observed effect sizes at ED95 with corresponding precision. The underlying dose-response is null, that of Curve 1. The symbol size in the top panel is proportional to sample size in the dose-finding phase.

is that when there is no drug effect the trial can be curtailed relatively early.

Figures 9.6 and 9.7 show the distribution of doses for the remaining dose-response curves considered. The posterior estimates of the dose-response curves are superimposed on the histograms of doses. Evidently, our adaptive allocation scheme is capable of dealing with a variety of underlying dose-response relationships. In most cases it is effective in learning about the form of the underlying curve.

To find the statistical power of our adaptive scheme we calculate a minimal slope in the observed dose-response. “Minimal slope” is the slope of a line spanning the whole dose range from placebo to the size of effect at the dose that will be selected to take into the confirmatory trial. This is minimal in the sense that if the maximal effect occurs anywhere other than at the maximum dose then the slope of the line will be greater. The slope is calculated only for those simulations that call for going forward into a confirmatory study. We can test this minimal slope for significance.

TABLE 3. Power of the adaptive approach assuming Curve 3 in which the maximal effect is 4 points improvement over placebo.

		Terminator decision	
		Stop (no effect)	Continue to confirmatory trial
Significance of slope*	Not significant	7	3
	Significant	0	90
	Total	7	93
Median sample size at stopping		255	343
Average power of test in confirmatory trial (median sample size) <sup>+</sup>	Posterior estimates of effect size, variability	0.76 (108)	
	Minimal clinically relevant difference (=2 pts)	>0.99 (603)	
	Predictive power	0.92 (174)	

\* Test of conservative estimate of slope = 0 at  $\alpha = 0.05$ .

<sup>+</sup> Average power of confirmatory trial with  $z^*$  (ED95 dose) and placebo to detect effect from known underlying dose-response curve at the chosen dose given that the dose-finding phase recommends going to confirmatory study. Nominal power=90%.

Table 3 gives a summary of the results from 100 simulations assuming Curve 3. In only 7 of 100 trials was the decision against going to confirmatory phase. As regards statistical significance of the minimal slope, only 3

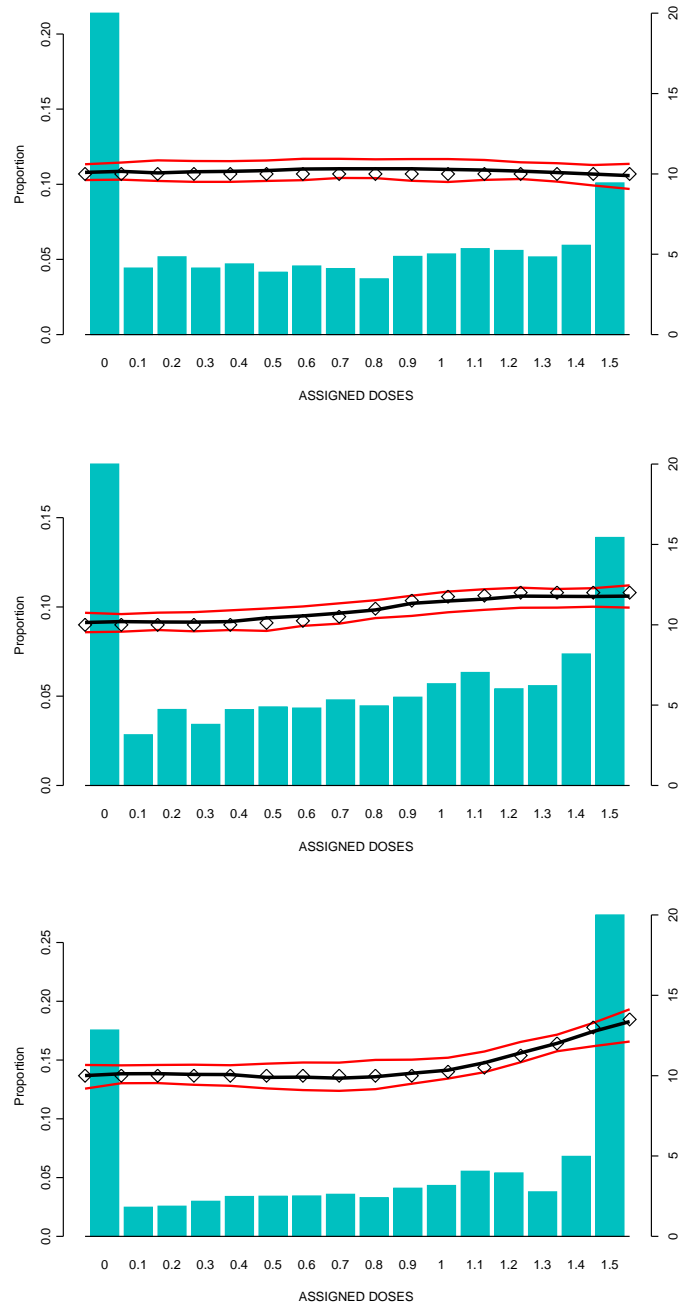


FIGURE 9.6. Distribution of doses and estimated dose-response functions for dose-response Curves 1, 2, and 5 from Figure 9.1.

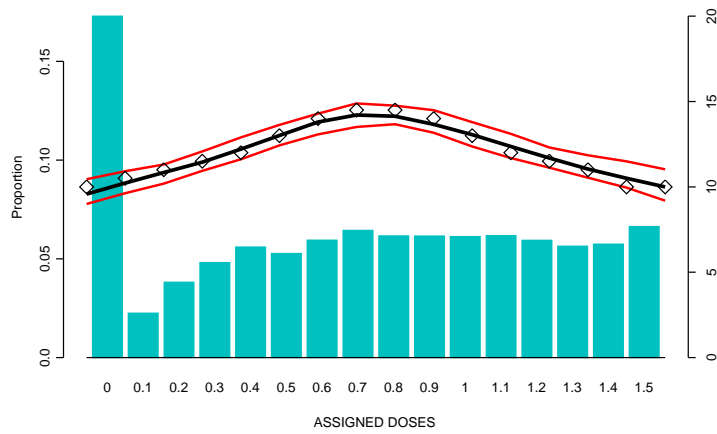
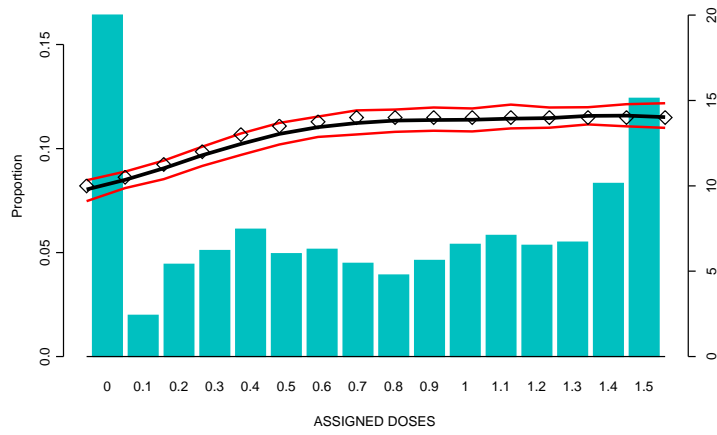


FIGURE 9.7. Distribution of doses and estimated dose-response functions for dose-response Curves 6 and 7 from Figure 9.1.

out of the 93 simulated trials that go forward into the confirmatory phase did not show a significant slope.

The dose chosen from the adaptive scheme will be taken forward into a confirmatory trial. Since we know the true response at that dose we can easily calculate the power of the confirmatory study to show a significant difference from placebo. We can select a sample size for this confirmatory trial in a number of ways: using the effect size and estimate of variability from the dose-finding study, using a minimal clinically relevant difference along with the estimate of variability, and using predictive power to account for uncertainty in the estimate of effect from the dose-finding study. Estimates of the average power of the test for the confirmatory study using these three methods of sample sizing is also shown in Table 3.

The power to show a difference from placebo in the confirmatory trial shows that if we use the minimal clinically relevant difference of 2 points improvement over placebo then we will overpower the study since the true effect is approximately 4 points improvement. Using the posterior estimate of the effect size at the ED95 leads to less than the nominal 90% power since we are not allowing for uncertainty in our estimate of effect size in the calculation of sample size. This is incorporated using the standard error around the estimate of effect size (as shown previously) and using predictive power methods. We can see that using this approach our average power is much closer to the nominal 90% power. Of course in practice all of these sample size estimates may be overridden by the need to have a substantial amount of safety data for regulatory submission.

TABLE 4. Type I error of the adaptive approach compared to the conventional design. Curve 1—no effect.

		Terminator decision	
		Stop (no effect)	Continue to confirmatory trial
Significance of slope*	Not significant	196 (98%)	0
	Significant	0	4 (2%)
	Total	196 (98%)	4 (2%)
Median sample size at stopping		257	1000

\* Test of conservative estimate of slope = 0 at  $\alpha = 0.05$ .

Table 4 shows the same summary for the null curve (Curve 1) in which there is no drug effect. This analysis provides the Type I error of this approach: any significant difference is a false-positive. In only 2% of the 200 simulated trials do we make the (incorrect) decision to go forward into a confirmatory study.

Table 5 shows a summary similar to that shown above for Curve 3 but for different underlying variance. A major benefit of the adaptive design is that

TABLE 5. Power of the adaptive approach vs. conventional design. Curve 3 – maximal effect = 4 points improvement over placebo, different observed variances.

		Terminator decision	
		Stop (no effect)	Continue to confirmatory trial
Significance of slope*	Not significant	8	92
	Significant	0	0
	Total	8	92
Median n		1000	257
Power of test in confirmatory trial <sup>+</sup>	Posterior estimates; effect size, variance		0.84 (42)
(median n)+	Min. clin. relevant difference (=2 pts)		>0.99 (213)
	Predictive power		0.93 (60)
Significance of slope *	Not significant	14	3
	Significant	0	83
	Total	14	86
Median n		277	413
Power of test in confirmatory trial <sup>+</sup>	Posterior estimates; effect size, variance		0.67 (150)
(median n)+	Min. clin. relevant difference (=2 pts)		>0.99 (847)
	Predictive power		0.84 (238)

\* Test of conservative estimate of slope = 0 at  $\alpha = 0.05$ <sup>+</sup> Average power of confirmatory trial with  $z^*$  (ED95 dose) and placebo to detect effect from known underlying dose-response curve at the chosen dose given that the dose-finding phase recommends going to confirmatory study. Nominal power=90%.

it adapts to the variability in the data. As a result, the confirmatory study will be powered appropriately based on the best estimate of variability from the data observed in the dose-finding phase. If variability in the data is substantially less than we would expect a priori then our inferences can also be made much sooner and with fewer patients.

However, Table 5 shows that when the variance is smaller, statistical power is not greatly increased. Rather, upon learning that the variance is small the algorithm stops the trial more quickly. The median sample size required to reach the decision to go forward into the confirmatory trial is smaller, and in fact is close to the minimum possible sample size of 250.

These power calculations show that again the adaptive design achieves close to the nominal 90% power for the confirmatory trial for either assumed variance. The predictive power approach takes into consideration the uncertainty in the estimate of the treatment difference and so achieves better power overall, albeit with larger average sample size. In the adaptive design the impact on our study of different observed variability is in the sample size taken to make a decision and not on the power of the test – which is shown in the calculations of true power shown for the conventional study.

From these simulations we conclude that our sample sizing and final inference across both trials are appropriate. The algorithm does what it is supposed to do: adapt to the underlying dose-response curve and to the variability in accumulating data. The dose-finding trial allows for stopping the drug's development early – either because it has no effect, or because the evidence is sufficiently clear to conclude that the treatment is effective. The simulations also show that when using the results of the dose-ranging phase to sample size a confirmatory trial, the latter will very likely confirm that the dose effect seen in the earlier trial is significant.

The simulations summarized above are (loosely) based on information from the Copenhagen Stroke Database but ignoring deaths. There are a number of issues we face in incorporating deaths into the generated data set, and we must deal with these as well when handling deaths in our analysis of the primary endpoint, change in SSS. Death must be included in some way.

To include mortality in the generated data we specify the distribution of deaths for different baseline scores – patients with mild initial stroke are less likely to die. We must also address the relationship between dose and death. The former we may evaluate from our analysis of the Copenhagen Stroke Database assuming that dose levels do not interact with baseline severity. The latter we must hypothesize at present because we have no information about whether and how the drug may affect the incidence of death. It may be that the effect of dose on death is not monotone, especially when averaging over severity of stroke. Perhaps low doses have no effect on mortality, medium doses offer protection against death and higher doses are linked to a mechanism that hastens death. This would give us a U-shaped

incidence of deaths across the dose range.

We want to incorporate death into the SSS endpoint, down-weighting any benefit of the treatment at doses should there be an increase in mortality rate. This will allow for addressing the role of mortality in the dose-response relationship. An alternative would be to evaluate the dose-response among survivors only and analyze deaths separately. But the two analyses – SSS improvement by dose and death by dose – must be combined eventually in some fashion. We feel it best to incorporate mortality at the earliest opportunity, but of course this does not prohibit considering the two measures separately after the trial is complete.

It is not clear how to best handle deaths as part of the principal endpoint and we have not yet settled on a final strategy. We consider two approaches here. In terms of neurological function one obvious solution would be to set  $SSS = 0$  should the patient die. This would give a negative change from baseline, which seems appropriate. A patient who enters the trial with an SSS of 20 and dies has a change from baseline of -20. However, a patient who arrives with  $SSS = 40$  and at the end of the study has a score of 10 has a worse change from baseline (-30) and therefore a greater negative impact on the effect at the dose assigned. Perhaps this is appropriate and perhaps it is not.

An alternative approach is to regard the score at death as the minimum SSS over the period of the trial, including baseline. A patient who improves but then dies nets a 0. A patient who declines by 10 points and then dies scores -10. But it might make more sense to count them the same. Indeed, the first patient may have passed through a change of -10 while dying but happened not to be tested at that time. Again, survivors may have lower scores than those who do not survive.

The plots of the data (Figure 9.8, top two panels) show the distribution of deaths across doses. The top panel shows that scoring death as zero separates the distribution of scores for patients who die from that those who survive. The second panel (center) shows slightly more overlap by setting the score for death to the minimum SSS over time.

The assumed dose-response curve in this example is the flat dose-response curve (null Curve 1) and so the effect of including deaths in the data can easily be seen. For this reason we have also investigated dose-response curves that are not monotonic (Curve 7). Figure 9.6 shows results of our simulations (without deaths) in this curve. Doses assigned have adapted to the form of the dose-response curve and there are fewer patients allocated in the extremes of the dose range.




FIGURE 9.8. Example of generated data including deaths. Upper panel: death scored as zero. Center panel: death scored as minimum score observed in the trial. Deaths are indicated as circled pluses. Bottom panel: estimated dose-response curve for Curve 1 (flat dose-response curve) but including deaths.

## 10 Discussion

Drug development is an extremely expensive enterprise. Part of the expense is due to the sluggardly nature of modern clinical trials. A fixed (and generally) large number of patients are accrued in trials that have deterministic designs. Although many trials employ optional stopping via boundaries such as those proposed by O'Brien and Fleming (1979), these stopping rules are very conservative. The reason is that they must diligently protect the overall Type I error rate of the frequentist tradition.

There is no such requirement in Bayesian approach. But Bayesians must live in a world still dominated by frequentist attitudes. In this report we use the flexibility and naturally adaptive nature of the Bayesian approach but with an overall goal that is consistent with frequentist tradition. It is a frequentist cake with Bayesian icing.

A ubiquitous problem in drug development is that the correct dose of a drug may never be adequately established. Our design takes the dose-finding issue very seriously, while deciding whether any dose of the drug is effective.

The first phase of our study uses a normal dynamic linear model and Bayesian updating for the relationship between dose and response. Should the results be sufficiently compelling (as judged by a decision analysis) and should a highly effective dose be identified then development of the drug enters a standard randomized confirmatory trial in which drug is compared with placebo. If the first phase establishes the drug as effective and if the confirmatory trial verifies its effectiveness then this combination would be very powerful evidence submitted to a regulatory agency in seeking marketing approval.

There are several virtues to our approach:

- (i) It avoids wasting patient resources investigating uninteresting doses.
- (ii) It stops the trial as soon as it becomes sufficiently clear that the drug is not effective.
- (iii) It moves into confirmatory phase when the evidence suggests that such is appropriate.
- (iv) It adapts to the actual variability in the accumulating data, perhaps stopping earlier than would be predicted (small variance) and perhaps continuing longer than would be predicted (large variance).
- (v) It uses a longitudinal model for the course of stroke in predicting final results for those patients having incomplete data, and learns about this model on the basis of the patients accrued to the trial.

The only disadvantage in our approach as opposed to conventional designs is that it requires a reliable and rapid system for communicating with participating centers. This system must be able to update the statistical

software based on recent information and it must be able to relay the doses assigned in very short order. Such a system is expensive, but its cost is tiny in comparison with the potential savings that accrue from using an efficient design.

## REFERENCES

- Berger, J.O. (1985). *Statistical Decision Theory and Bayesian Analysis*, New York: Springer-Verlag.
- Berry, D.A. (1993). A case for Bayesianism in clinical trials (with discussion), *Statistics in Medicine*, 12, 1377-1404.
- Berry, D.A. and Stangl, D.K. (1996). Bayesian methods in health-related research. In *Bayesian Biostatistics*, (D.A. Berry and D.K. Stangl, eds.), New-York: Dekker, 3-66.
- Carter, C. K., and Kohn, R. (1994). On Gibbs sampling for state space models, *Biometrika*, 81, 541-553.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: a review, *Statistical Science*, 10, 273-304.
- DeGroot, M. (1970). *Optimal Statistical Decisions*. New York: McGraw Hill.
- DeHaan, R., Horn, J. et al.(1993). A comparison of five stroke scales with measures of disability, handicap, and quality of life. *Stroke*, 24, 1178-1181.
- Dirnagl, U., Iadecola, C., and Moskowitz, M.A. (1999). Pathobiology of ischaemic stroke: an integrated view. *Trends Neurosci.*, 22, 391-397
- Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models, *Journal of Time Series Analysis*, 15, 183-202.
- Gilks, W.R., Richardson, S., and Spiegelhalter, D.J. (eds.) (1996). *Markov Chain Monte Carlo in Practice*, Chapman and Hall.
- Hill, M.D. (1998). Stroke treatment: Time is brain, *Lancet*, 352(SIII), 10-14.
- Jorgensen, H.S., Nakayama, H., Raaschou, H.O., Vive-Larsen, J., Stoier, M. and Olsen, T.S. (1994). Outcome and time course of recovery in stroke. Part I: Outcome. The Copenhagen Stroke Study, *Archives Physical and Medical Rehabilitation*, 76, 399-405.

- Lees, K.R. (1998). Does neuroprotection improve stroke outcome? *Lancet*, 351, 1447-1448
- Lindenstrom, E., Boysen, G, Christiansen, L.W., Rogvi Hansen, B. and Nielsen, B.W. (1991). Reliability of Scandinavian Stroke Scale. *Cerebrovascular Disease*, 1, 103-107.
- O'Brien, P.C., Fleming, T.R. (1979). A multiple testing procedure for clinical trials, *Biometrics*, 35, 549-556.
- Scandinavian Stroke Study Group (1985). Multicenter trial of hemodilution in ischemic stroke. Background and study protocol. *Stroke*, 16, 885-890.
- Thall, P.F., Simon, R.M., and Estey, E.H. (1995). Bayesian sequential monitoring designs for single-arm clinical trials with multiple outcomes, *Statistics in Medicine*, 14, 357-379.
- Thall, P.F., and Russell, K.E. (1999). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. to appear in *Biometrics*.
- Tierney, L., (1994). Markov chains for exploring posterior distributions, *Annals of Statistics*, 22, 1701-1762.
- Verdinelli, I. (1992). Advances in Bayesian experimental design. In *Bayesian Statistics IV*, (J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith, eds), pp. 467-48, Oxford: Oxford University Press.
- Warlow, C.P., Dennis M.S. et al. (1996). *Stroke: A practical guide to management*. Oxford, Blackwell Scientific.
- West, M. and Harrison, P.J. (1997). *Bayesian Forecasting and Dynamic Models*. Second Edition. New York: Springer-Verlag.
- Whitehead, J. and Brunier, H. (1995). Bayesian decision procedures for dose determining experiments, *Statistics in Medicine*, 14, 885-893.
- Whitehead, J., and Williamson, D. (1998). Bayesian decision procedures based on logistic regression models for dose-finding studies, *Journal of Biopharmaceutical Statistics*, 8, 445-467.