Optimal response and covariate-adaptive biased-coin designs for clinical trials with continuous multivariate or longitudinal responses

Optimal Response and Covariate-Adaptive Biased-Coin Designs for Clinical Trials with Continuous Multivariate or Longitudinal Responses

Anthony C. Atkinson*

Department of Statistics, London School of Economics, London WC2A 2AE, UK

Atanu Biswas*

Applied Statistics Unit, Indian Statistical Institute, 203 Barrackpore Trunk Road, Kolkata 700 108, India

Abstract

Adaptive randomization of the sequential construction of optimum experimental designs is used to derive biased-coin designs for longitudinal clinical trials with continuous responses. The designs, coming from a very general rule, target pre-specified allocation proportions for the ranked treatment effects. Many of the properties of the designs are similar to those of well-understood designs for univariate responses. A numerical study illustrates this similarity in a comparison of four designs for longitudinal trials. Designs for multivariate responses can likewise be found, requiring only the appropriate information matrix. Some new results in the theory of optimum experimental design for multivariate responses are presented.

Keywords:
Biased-coin design; covariate balance; effective number of observations; ethical allocation; equivalence theorem; multivariate D_A-optimality; multivariate loss; power; skewed allocation.

*Corresponding Author

Email addresses: a.c.atkinson@lse.ac.uk (Anthony C. Atkinson), atanu@isical.ac.in (Atanu Biswas)

1. Introduction

Response-adaptive designs are becoming increasingly popular in phase III clinical trials with sequential entrance of patients. The ethical objective is to use the accumulating data to skew the allocation in favour of the better treatments, so ensuring that as few patients as possible receive bad treatments. The advantages of response-adaptive designs are extolled by Zelen and Wei (1995), Hu and Rosenberger (2003) and Rosenberger and Hu (2004). Gallo *et al.* (2006) provide a perspective from the pharmaceutical industry.

Our procedure is based on the adaptive randomization of treatment allocations from the sequential construction of optimum experimental designs. As a consequence, we require optimum designs for multivariate continuous responses that provide balance over the prognostic factors that may be included in the estimation of treatment effects. Unfortunately, the majority of the adaptive designs that have been developed are for a single binary response per patient in the absence of covariates. Examples include the play-the-winner (PW) design (Zelen, 1969), the randomized play-the-winner (RPW) design (Wei and Durham, 1978), the success driven design (Durham *et al.*, 1998) and the drop-the-loser (DL) rule (Ivanova, 2003). Related designs for continuous responses, using non-parametric methods to discretise the problem, include Rosenberger (1993) and Bandyopadhyay and Biswas (2004).

These designs work well in skewing the allocation in favour of the better treatment, although they are not derived from any optimality criterion. One form of optimality, for binary responses, consists of minimizing an aspect of behaviour, such as the total expected number of failures, for a given variance of the estimated treatment difference. Such designs include those of Rosenberger *et al.* (2001) and Biswas and Mandal (2007) for binary responses. Zhang and Rosenberger (2006, 2007) and Biswas *et al.* (2007) find optimum designs for continuous responses.

Several of these procedures have been extended to design in the presence of covariates, giving rise to Covariate Adjusted Response Adaptive (CARA) designs. For the randomized play-the-winner rule, Bandyopadhyay and Biswas (1999) combined polytomous covariates with binary responses and Bandyopadhyay and Biswas (2001) incorporated covariates in their design for continuous responses. Zhang *et al.* (2007) studied asymptotic properties of CARA designs under widely satisfied conditions. Optimum biased-coin designs for covariate balance, without response adaptivity, were introduced

by Atkinson (1982). This form of optimality was extended to response-adaptive designs for univariate responses by Atkinson and Biswas (2005a, 2005b). Rosenberger and Sverdlov (2008) discuss the arguments that have been advanced in the clinical trials literature for and against treatment allocations rules that provide some balance over covariates, as do Shao *et al.* (2010).

There is an appreciable literature on the analysis of data from clinical trials when the responses are observed at a series of monitoring times, for example Everitt and Pickles (2004, Chapters 5–7). Molenberghs *et al.* (2004) describe data from three clinical trials of anti-depressants in which the responses can be treated as continuous. Galbraith and Marschner (2002) provide guidelines for designing non-adaptive longitudinal clinical trials.

By comparison there is very limited literature on the design of adaptive longitudinal trials. Biswas and Dewanji (2004b) describe a trial of pulsed electro-magnetic field therapy in which each patient was monitored for about 16 weeks. The original responses in this trial had a complicated multivariate structure, which was ignored in the design. Instead a binary variable 'recurrence' was used. Biswas and Dewanji developed an urn design for longitudinal binary responses, which is a modification and simple extension of the RPW design where the covariates were ignored. See also Biswas and Dewanji (2004a,c). Sutradhar *et al.* (2005) used a similar urn model based design and allowed for the possibility of time-dependent covariates. Subsequently, Sutradhar and Jowaheer (2006) extended this approach for longitudinal count data. Biswas *et al.* (2012) provided an optimum response-adaptive design for longitudinal binary responses. Atkinson and Biswas (2014, Chapter 5) provide an account of work on response-adaptive designs for longitudinal responses. Further, Huang *et al.* (2013) proposed a general framework for longitudinal covariate-adjusted response-adaptive randomization procedures, and studied the related asymptotic properties.

In contrast, we obtain optimum biased-coin designs for multivariate and longitudinal responses by the extension of methods for univariate responses. The optimum designs in both cases are functions of the information matrix for the observations. The model for multivariate data is introduced in Section 2.1. In the rest of Section 2 we explore the consequences of a general formulation for randomized response-adaptive designs for univariate or multivariate responses. These designs use optimum design theory to provide covariate balance in a general adaptive rule that skews allocation to the better treatments, whilst maintaining a controllable degree of randomness. We

stress that these results are extremely general; to apply the rules we merely need to be able to provide the information matrix of the observations. Loss and bias, used to compare the designs, are presented in Section 3 with the information matrix for longitudinal designs explicitly presented in Section 4.

Four specific allocation rules are described in Section 5. These include the extension of the rule of Atkinson and Biswas (2005a), which achieves adaptivity through use of the link function of Bandyopadhyay and Biswas (2001), and our new rule. We take the particular form of this rule which targets specified proportional allocations to the ranked treatments: in our numerical example with two treatments, our target is that 80% of patients should be allocated to the unknown better treatment. This procedure overcomes the instability in early allocations with the link-function based rule that can lead to imbalance if the trial is stopped early. In our application we apply the results to the particular information pattern and covariance structure arising with longitudinal responses developed in Section 4. The numerical results are in Section 7.

Two main contributions of our paper are the provision of our general rule and its application to longitudinal trials. In the form we use here, the design ceases to be response-adaptive once the correct ordering of the treatments has been established. We can then extend standard results of the effect of randomization on inference (Burman, 1996; Atkinson, 2002) to multivariate designs. For longitudinal designs with correlated observations we define an effective number of observations that permits calculation of the loss from randomization. This important quantity indicates the average number of patients on whom information is lost due to a particular randomization rule. Simulations in Section 7 confirm the accuracy of this definition.

The methods of optimum experimental design are central to our construction of allocation rules. In Appendix A.1 we develop new results on multivariate D_A-optimality that allow us to estimate linear combinations of the treatment effects, such as differences, in the presence of the parameters associated with the prognostic factors over which we are balancing. An equivalence theorem satisfied by the optimum designs is in Appendix A.2. These contributions are discussed in Section 8.

2. Multivariate Data

2.1. Models

Patients arrive sequentially and are to be allocated one of t treatments. The particular treatment to be allocated to patient $n+1$ depends on a vector of prognostic factors x_{n+1}, on previous allocations and on the information available from the responses of previous patients. In a longitudinal study this information will increase as extra readings become available on earlier patients. But we start with the simpler multivariate case of n_h responses from each patient, all of which are available before the next patient arrives. In Section 4 we extend the model to the longitudinal case of incomplete time series of observations.

We assume that the results of the trial, perhaps after data transformation, will be analysed using a regression model which, for the n_h readings on patient i, is written

$$E(y_i) = F_i\beta = H_i\alpha + Z_i\theta + G\zeta, \tag{1}$$

where y_i is $n_h \times 1$. Although we take all observations to have the same dimension, there is no difficulty, other than notational, in y_i being of dimension $n_h(i)$.

Here α is the $t \times 1$ vector of treatment effects that are the focus of inferential interest and H_i is the $n_h \times t$ matrix of t indicator variables, the one non-zero column indicating which treatment the patient received. The $n_h \times v$ matrix Z_i contains those covariates, including any powers or interactions of the elements of x_i, which may be used to adjust the responses when estimating α. Because of the way we have parameterized the treatment effects, Z_i does not include a constant column. In the context of longitudinal data the $n_h - 1$ elements of ζ are arbitrary period effects, the same for all patients; the $n_h \times (n_h - 1)$ matrix G is the matrix of indicator variables for the period. These matrices are highly structured. All rows of H_i are the same for patient i, as are those of Z_i, whereas the rows of G are different.

Conditionally on the values of the x_i the additive errors of observation in (1) have an n_h dimensional normal distribution with covariance matrix

$$\text{cov } y_i = \sigma^2 V. \tag{2}$$

For longitudinal data V has a known structure; in our numerical example this comes from an AR(1) process of errors. Efficient estimation is by weighted

least squares, with a block diagonal weight matrix with diagonal elements V^{-1}. The conditional information matrix for all n patients is

$$\mathcal{I}(n) = \sum_{i=1}^{n} F_i^T V^{-1} F_i. \tag{3}$$

The design criterion does not depend on the value of σ^2, which is suppressed in (3). However, σ^2 is necessary in power calculations. In response adaptive designs the value of x_i is determined by previous values of y_i, so that, unconditionally, the observations are not independent. In Section 6 we give the argument that (3) is, however, the correct unconditional asymptotic information matrix.

2.2. Parameter Estimation and Randomization

To construct adaptive designs for the efficient estimation of α in (1) we employ randomized versions of the sequential construction of optimum experimental designs (Fedorov, 1972; Atkinson *et al.*, 2007; Fedorov and Leonov, 2014). The cost of randomization is that the design is likely to be unbalanced when it is stopped after a number of patients that is unknown at the planning stage; there is a consequent loss of efficiency in estimation.

In order to balance parameter estimation and randomization in general non-sequential designs, Ball *et al.* (1993) suggested finding designs to maximize the utility

$$U = U_V - \gamma U_R, \tag{4}$$

where the contribution of U_V is to provide estimates with low variance, whereas U_R provides randomness. The parameter γ provides a balance between the two. We extend their method to the sequential construction of longitudinal clinical trials.

With π_j the probability of allocating treatment j, let

$$U_V = \sum_{j=1}^{t} \pi_j \phi_j,$$

where ϕ_j is a measure of the information from applying treatment j. In the next section this is defined in terms of D_A-optimality.

Ball *et al.* (1993) are interested in randomness with equal allocation, when

$$U_R = \sum_{j=1}^{t} \pi_j \log \pi_j.$$

To combine randomness with greater allocation to preferred treatments we introduce a set of gains $G_1, \ldots, G_t$ for the allocation of the individual treatments. These gains can be quite general, although we require $G_i \geq 0 \ \forall \ i$. Then

$$U_R = \sum_{j=1}^{t} \pi_j(-G_j + \log \pi_j). \tag{5}$$

In Section 2.5 we associate the G_j with the target allocation proportions of the treatments ordered by desirability. When all G_j are zero, U_R reduces to the utility used by Ball *et al.* (1993).

To maximize the utility (4) subject to the constraint $\sum_{j=1}^{t} \pi_j = 1$ we introduce the Lagrange multiplier λ and maximize

$$U = \sum_{j=1}^{t} \pi_j \phi_j - \gamma \sum_{j=1}^{t} \pi_j(-G_j + \log \pi_j) + \lambda \left(\sum_{j=1}^{t} \pi_j - 1 \right). \tag{6}$$

Since the G_j occur in U with a positive coefficient, maximization of U gives large values of π_j for treatments with larger G_j. Differentiation of (6) with respect to π_j leads to the t relationships

$$\phi_j - \gamma(-G_j + 1 + \log \pi_j) + \lambda = 0,$$

so that all quantities

$$\phi_j/\gamma + G_j - \log \pi_j$$

must be constant. Since $\sum_{j=1}^{t} \pi_j = 1$, we obtain

$$\pi_j = \{\exp(\phi_j/\gamma + G_j)\}/S = \{\exp(\psi_j/\gamma)\}/S, \tag{7}$$

where

$$\psi_j = \phi_j + \gamma G_j$$

and

$$S = \sum_{j=1}^{t} \exp\{(\phi_j/\gamma) + G_j\} = \sum_{j=1}^{t} \exp(\psi_j/\gamma).$$

As $\gamma \to \infty$, emphasis is solely on randomization and now all quantities $G_j - \log \pi_j$ must be constant so that $\pi_j \propto \exp(G_j)$. When the gains G_j for the allocation of the individual treatments are equal, we obtain the equal randomization rule of Ball *et al.* (1993) with $\pi_j = 1/t$. This also follows directly from (5) since, with $G_j = G \ \forall \ j$, $\Sigma_j \pi_j G_j = G$, which does not affect the allocation.

2.3. Optimum and Sequential Designs

The probabilities of allocation π_j in (7) depend on the information measure ϕ_j and on the utility G_j from allocating treatment j. We first consider ϕ_j.

In the methods of optimum experimental design, treatments are allocated to make large some function of $\mathcal{I}(n)$. We allocate to minimize the variance of s linear combinations of the treatment estimates which are adaptively chosen to give the desired probability of allocation of each treatment. For this measure we follow Atkinson (1982) and use D_A-optimality (Atkinson *et al.*, 2007, §10.2).

The s linear combinations of the treatment effects are $L^T\alpha$, where L^T is $s \times t$, $s < t$. The s combinations of all parameters can be written

$$A^T\beta = L^T\alpha + W_1^T\theta + W_2^T\zeta, \tag{8}$$

where A^T is $s \times (t + v + n_h - 1)$, W_1^T is $s \times v$ and W_2^T is $s \times (n_h - 1)$. If the effects of the variables Z_i and G are not of interest in themselves, the parameters θ and ζ in (1) become nuisance parameters and the elements of W_1 and W_2 are zero. For any A the variance of the estimated combination of coefficients is

$$\text{var}\,\{A^T\widehat{\beta}\} = \sigma^2 A^T\{\mathcal{I}(n)\}^{-1}A, \tag{9}$$

where $\hat{\beta}$ is the least squares estimate of β.

The properties of the design depend on the number of treatments and on the dimension of the space of the nuisance parameters over which allocations are randomized. The exact relationship depends on the form of randomization. Atkinson (2014) provides many examples. There is no randomization over the values of the time profile G. Since, in (8), W_1^T is $s \times v$, the dimension of the nuisance parameters is $q = t + v - s$.

D_A-optimum designs minimize the logarithm of the determinant of the covariance matrix (9). Thus we seek designs to maximize the information measure

$$\phi = -\log|A^T\{\mathcal{I}(n)\}^{-1}A| = -\log\Psi. \tag{10}$$

If treatment j is allocated to the $(n+1)$st patient we extend the notation of (10) and obtain

$$\phi_j = -\log|A^T\{\mathcal{I}(n+1,j)\}^{-1}A| = -\log\Psi_j.$$

Once the allocation has been made, we can suppress the subscript j which is, however, required when we are comparing treatments for allocation. Substitution of this expression for ϕ_j in (7) yields

$$\pi_j = \Psi_j^{-1/\gamma} \exp(G_j)/S. \tag{11}$$

In the sequential generation of optimum designs we would make the allocation for which ϕ_j was a maximum.

For clinical trials with univariate responses Atkinson (1982) exploited the results on sequential generation of D_A-optimum designs from Silvey (1980). Here we have multivariate responses. The requisite extension of D_A-optimality is derived in Appendix A.1, with the Equivalence Theorem for multivariate D_A-optimality in Appendix A.2. It is clear from (A-1) that Ψ_j is the product of two terms, one of which is the same for all allocations. Substitution into (11) yields the allocation probability

$$\pi(j|x_{n+1}) = \frac{\{1 + d_A(j, n, x_{n+1})\}^{1/\gamma} \exp(G_j)}{\sum_{s=1}^{t} \{1 + d_A(s, n, x_{n+1})\}^{1/\gamma} \exp(G_s)}, \tag{12}$$

where $d_A(j, n, x_{n+1})$, the directional derivative of the D_A-optimality criterion, is given by (A-3) and x_{n+1} is the vector of covariates for the new patient that are included in $\mathcal{I}(n+1, j)$.

2.4. Gain and Allocation Probabilities

We have derived our very general allocation rule in terms of undefined gains G_j from allocation of the treatment j. We now find appropriate G_j for an allocation rule which targets proportions of adaptively ranked treatments.

Let the target proportion of patients receiving treatment ranked j be p_j^*. Then we require that

$$p_1^* \geq p_2^* \geq \ldots \geq p_j^* \geq \ldots \geq p_t^*, \tag{13}$$

with, to avoid uniform allocation, at least one inequality. With $p_1^* = 1$ (and all other $p_j^* = 0$) we obtain a rule in which only the most highly ranked treatment is allocated. Plausible rules allow allocation to all treatments that have not been eliminated from the study and have the p_j^* a decreasing function of j. The purpose of the rule is to ensure a specified ethical gain without going through possible extreme designs even if, for example for a two treatment design, α_1 is very much greater than α_2. By using ranks,

we ensure both a prefixed allocation which is ethically skewed and sufficient allocation to each treatment to ensure that the design is not too inferentially inefficient.

At the optimum design it follows from the Equivalence Theorem of Appendix A.2 that all $d_A(j, n, x_{n+1})$ are equal and the treatments are correctly ordered. Let the correct, but unknown, rank of treatment j be $R(j)$. Then, from (12)

$$\pi(j|x_{n+1}) = p^*_{R(j)} = \frac{\exp\{G_{R(j)}\}}{\sum_{s=1}^{t} \exp\{G_{R(s)}\}}. \tag{14}$$

The probabilities of allocation in (12) and (14) are unaltered if we replace $G_{R(j)}$ with

$$G^c_{R(j)} = G_{R(j)} + c.$$

We choose c so that $\sum_{i=1}^{t} \exp\{G^c_{R(i)}\} = 1$. Then (14) becomes

$$G^c_{R(j)} = \log p^*_{R(j)}$$

and the allocation probabilities (14) have the simple form

$$\pi_G(j|x_{n+1}) = \frac{\{1 + d_A(j, n, x_{n+1})\}^{1/\gamma} p^*_{R(j)}}{\sum_{s=1}^{t} \{1 + d_A(s, n, x_{n+1})\}^{1/\gamma} p^*_{R(s)}}, \tag{15}$$

provided the ranking of the treatments is known. In designing the trial, the p^*_j are the fundamental quantities which are to be specified, rather than the gains G_j.

2.5. Skewed Allocations

Replacement of the ranks $R(j)$ in (15) with the ranks $\widehat{R}(j)$ based on the estimated ordering of the treatment effects α gives an operational rule, but we have also to specify the coefficients L that give an efficient design. We develop the argument for $s = 1$, when the properties of a linear combination $l^T\alpha$ are of interest. First consider univariate responses, variance σ^2, with a proportion $r_j = n_j/n$ of the n patients receiving treatment j. When there are two treatments and in (8) we take $l^T = (1 \quad -1)$, or equivalently $(0.5 \quad -0.5)$, the inferential purpose is to estimate $\alpha_1 - \alpha_2$ with minimum variance. The design minimizing (9) provides balance over the covariates and equal allocation to the two treatments, so that $r_1 = r_2 = 0.5$.

To obtain skewed allocation for t treatments combined with efficient parameter estimation we find designs for estimation of the linear combination with

$$l^T\alpha = \pm l_1\alpha_1 \mp \ldots \pm l_t\alpha_t, \tag{16}$$

where the coefficients $l_j, j = 1, \ldots, t$ are such that $0 < l_j < 1$ and $\sum l_j = 1$. Then, in the absence of covariates,

$$\text{var}\,\{l^T\hat{\alpha}\} = (\sigma^2/n)\sum_{j=1}^{t} l_j^2/r_j \quad \text{with} \quad \sum_{j=1}^{t} r_j = 1. \tag{17}$$

Use of a Lagrange multiplier shows that this variance is minimized when the proportion of patients receiving treatment j is l_j, as it is when the design is balanced across treatments, in the sense of the covariates having the same conditional distribution for each treatment. The signs in (16) are a generalization to t treatments of the weights 0.5 and -0.5 that give efficient designs with $r_j = 0.5$ for the treatment difference.

To obtain an adaptive design targeting (13) we take weights

$$l_j = p^*_{\hat{R}(j)}.$$

2.6. Adaptive Design: Rule G

With estimated rankings the probability of allocation of treatment j for Rule G is

$$\pi(j|x_{n+1}) = \frac{\{1 + d_A(j, n, x_{n+1})\}^{1/\gamma} p^*_{\hat{R}(j)}}{\sum_{s=1}^{t}\{1 + d_A(s, n, x_{n+1})\}^{1/\gamma} p^*_{\hat{R}(s)}}. \tag{18}$$

The effect of different values of the parameter γ can be elucidated by simulation, as in Atkinson (2014) for univariate responses. In our example we take $\gamma = 0.1$. Distributional results for Rule G are in Section 6.

It is a characteristic of this scheme that the probability of allocating the treatments depends on the p^*_j and on the ordering of the α_j, but not on the differences between them. Suppose there are two treatments. Then, if $\alpha_1 > \alpha_2$, treatment 1 will eventually be allocated in a proportion p^*_1 of the trials regardless of the value of $\Delta = \alpha_1 - \alpha_2$. Of course, if Δ is small relative to the measurement error, in many of the initial trials, $\hat{\alpha}_1 < \hat{\alpha}_2$ and it will seem that treatment 2 is better. Then some individual allocations will be skewed in favour of treatment 2 with target p^*_1, that is $p_{\hat{R}(2)}$. When $\hat{a}_1 > \hat{\alpha}_2$, treatment 1 will be preferred. If the trial is terminated before a clear difference between the treatments has been established, each treatment may have been be allocated to around half the patients.

3. Loss, Bias and the Assessment of Designs

3.1. Loss

Our adaptive designs have some randomness in allocation. The effect of randomness is slightly to unbalance the designs while reducing the chance of guessing the next allocation. To compare designs we need measures of these two aspects.

We start with the effect of imbalance and extend the idea of loss (Burman, 1996) to multivariate responses. For univariate responses the variance of the estimated linear combination (9) has a minimum value of σ^2/n for the optimum design with proportions l_j and balance across the covariates. For multivariate data with n_h independent observations on each patient, the minimum variance is $\sigma^2/(nn_h)$. However, in general the variance depends on the correlation of the multivariate observations.

For observations with structure given by (1), the contribution of y_i to the total sum of squares is $y_i^T V^{-1} y_i$. For positively correlated data this is a smaller contribution than if the observations were independent. We define the effective number of observations as

$$n_{\text{effec}} = J^T V^{-1} J, \tag{19}$$

where J is a vector of ones.

The effective number of observations n_{effec} depends on the structure of V. For independent observations $n_{\text{effec}} = n_h$. For correlated observations it decreases with increasing correlation. With this definition the variance of the estimated linear combination for the optimum design has the minimum value

$$\text{var}\left\{l^T \hat{\alpha}_*\right\} = \sigma^2/(n \times n_{\text{effec}}), \tag{20}$$

where $\hat{\alpha}_*$ is the least squares estimate of α from the optimum design.

We can find from (9) the variance of the same linear combination for any other design. The efficiency of the design is then

$$E_n = 1/\left[nl^T\{\mathcal{I}(n)\}^{-1}l\right].$$

The loss L_n is defined on comparing the variance (9) with the minimum value given by (20) as

$$\text{var}\left\{l^T \widehat{\alpha}\right\} = \frac{\sigma^2}{n \times n_{\text{effec}} - L_n}.$$

Comparisons can use either the efficiency E_n, or the loss, calculated by Atkinson and Biswas (2014) for numerous rules for skewed and unskewed allocations when the responses are univariate.

For all reasonable designs the efficiency tends to one as $n \to \infty$. However, distinct limits of loss are known for several classes of design for univariate responses and comparisons of loss provide an incisive means of comparing designs for the quality of the estimates of the treatment effects. The number of such parameters does not depend on the dimension of the multivariate observations and, in the multivariate setting, the loss can be interpreted as n_{effec} times the number of patients on whom information is lost due to the lack of optimality of the design.

3.2. Selection Bias

The purpose of including randomization in these rules is to prevent various kinds of bias. Selection bias occurs when the clinician is able correctly to guess the next treatment to be allocated. For two treatments it can be written as

$$\begin{aligned} B_n \;=\; & (\text{probability of correctly guessing the allocation to patient } n \\ & -\text{probability of incorrectly guessing the same allocation}). \end{aligned} \quad (21)$$

As do Heritier *et al.* (2005), we directly use the allocation probability $\pi_n(j)$ for patient n. The probability of correctly guessing the allocation of treatment j when $\pi_n(j) \geq 0.5$ is $\pi_n(j)$ and of an incorrect guess is $1-\pi_n(j)$. The selection bias can then be estimated by simulation as the average value of $\{2\pi_n(j)-1\}$. In simple cases the bias can be calculated explicitly. For example, in the non-randomized sequential construction of optimum designs the next treatment to be allocated is known exactly and the value of B_n is one.

4. Longitudinal Designs

The main algebraic difference between multivariate designs and longitudinal designs comes from the reduced amount of information that is available from previous patients when allocation is made to patient $n+1$. As an example, in our calculations we assume that patients arrive, or are grouped to arrive, in cohorts of size n_g (n_g can equal one). In the general case of $n_g > 1$ the various members of each cohort can be allocated different treatments.

Consider the first patient of cohort $k+1$, so that $n+1 = kn_g + 1$. We assume the responses are delayed, so that there is no information on the

responses from cohort k. Patients from cohort $k-1$ contribute one response, those from cohort $k-2$ two responses and so forth. Working backwards, the first cohort to contribute all n_h responses is number $k-n_h$. Let $S(i,n)$ denote this set of indexes for patient i as cohort $k+1$ starts. Then the information matrix for the allocation of patient $n+1$ is, in an extension of the notation of (3),

$$\mathcal{I}(n) = \sum_{i=1}^{n} F_{i,S(i,n)}^T V_{S(i)}^{-1} F_{i,S(i,n)}. \tag{22}$$

The same process of counting applies to the sufficient statistics and so to the estimates of α used in the adaptive allocation rules. For allocation of the remaining observations in cohort $k+1$ we increment the information matrix by the complete value of $F_i^T V^{-1} F_i$ for each allocated observation. This is a temporary measure to aid balance, so that we remove these contributions when recalculating (22) for the next cohort.

The other difference between longitudinal data and the multivariate data of Section 2.1 is the structure of the covariance matrix V. In our numerical example the errors form an AR(1) process. A stationary process can be simulated by generating $u_1 = \epsilon_1 \sim N\{0, 1/(1-\rho^2)\}$ and, for $i > 1$, $u_i = \rho u_{i-1} + \epsilon_i$ where $\epsilon_i \sim N(0,1)$. Then var $u_i = 1/(1-\rho^2)$ and cov $(u_i, u_{i-s}) = \rho^s/(1-\rho^2)$. The errors of observation will be σu_i, where σ is to be estimated.

For this error structure

$$V^{-1} = \begin{pmatrix} 1 & -\rho & 0 & 0 & \dots \\ -\rho & 1+\rho^2 & -\rho & 0 & \dots \\ 0 & -\rho & 1+\rho^2 & -\rho & \dots \\ \dots & \dots & \dots & \dots & \dots \end{pmatrix}$$

and, from (19)

$$n_{\text{effec}} = J^T V^{-1} J = n_h - 2(n_h - 1)\rho + (n_h - 2)\rho^2.$$

For $\rho = 0, n_{\text{effec}} = n_h$. The number however decreases appreciably as ρ increases; for $n_h = 4$ and $\rho = 0.5$, $n_{\text{effec}} = 1.5$.

5. Four Allocation Rules

We now present and compare four specific allocation rules, two of which depend directly on the estimated ranking of the treatments. As in §2.5, let

the ranking of treatment j be $R(j)$, estimated by $\widehat{R}(j)$ from the ranking of the estimated treatment effects $\hat{\alpha}$.

1. Rule D (Deterministic). For the purposes of comparison with adaptive rules, we assume that the correct ordering of the treatments is known. There is no randomization and the rule is that for the sequential construction of the optimum design. We allocate that treatment for which (A-2) is a maximum:

$$\pi_D(j|x_{n+1}) = \begin{cases} 1 & j = \arg\max\limits_{j\in 1,\ldots,t} d_A(j,n,x_{n+1}) \\ 0 & \text{otherwise} \end{cases} .$$

Simulations of designs for univariate responses mentioned above show that the loss L_n for this rule rapidly tends to zero and the bias B_n is one, since it is always known which treatment will be allocated next. These are extreme values; other rules have higher loss and lower bias.

2. Rule RA (Random and Response Adaptive). In this rule the treatments are allocated with probabilities p_j^* introduced in (13) and based on the estimated ranking of treatment effects; there is no attempt at covariate balance. Then (18) reduces to

$$\pi_{RA}(j) = p^*_{\widehat{R}j}.$$

The best guessing strategy is always to guess that the seemingly best treatment will be allocated. Asymptotically the probability of being correct is p_1^* (and of being wrong $1 - p_1^*$), so that the limit of B_n in (21) is $2p_1^* - 1$. For univariate responses the asymptotic value of the loss is q (Cox, 1951; Burman, 1996).

3. Rule G (General with ranks). This rule (18) extends Rule RA to include covariate balance. Since $d_A(j,n,x_{n+1})$ in (18) is not standardized for n, the rule tends asymptotically to Rule RA.

4. Rule L (Link). For two treatments the target probabilities depend on the estimated difference in treatment means $\widehat{\Delta} = \hat{\alpha}_1 - \hat{\alpha}_2$. Atkinson and Biswas (2005a) use a link function to relate the l_j to $\widehat{\Delta}$. Following Bandyopadhyay and Biswas (2001) they take $l_1 = \Phi(\widehat{\Delta}/T)$, where $\Phi(.)$ is the standard normal cumulative distribution function (c.d.f.). The value of

l_1 may be greater or less than 0.5 and the parameter T controls the degree of skewing of the allocation.

For generality we present the t treatment version of this rule. After n patients have been treated the estimated treatment parameters are $\hat{\alpha}_j$. To preserve the invariance of the procedure to the overall treatment mean let

$$\bar{\alpha} = \sum_{j=1}^{t} \hat{\alpha}_j / t \quad \text{and} \quad \widehat{\Delta}_j = \hat{\alpha}_j - \bar{\alpha}.$$

The cumulative normal distribution provides coefficients l_j by setting $l'_j = \Phi(\widehat{\Delta}_j / T)$ with

$$l_j = l'_j / S_l, \quad \text{where} \quad S_l = \sum_{k=1}^{t} l'_k. \tag{23}$$

For $t = 2$ this reduces to the design procedure of Bandyopadhyay and Biswas (2001) except that the standard deviation T is replaced by $2T$.

This design procedure does not depend on the values of the covariates. To provide a rule that is covariate adjusted, we use the values of l'_j in place of $p^*_{\widehat{R}(j)}$ in (18) to give allocation probabilities

$$\pi_L(j|x_{n+1}) = \frac{\{1 + d_A(j, n, x_{n+1})\}^{1/\gamma} l'_j}{\sum_{s=1}^{t} \{1 + d_A(s, n, x_{n+1})\}^{1/\gamma} l'_s}. \tag{24}$$

Use of l'_j, or the standardized l_j from (23), gives identical allocation probabilities $\pi_L(j|x_{n+1})$ since the summation S_l cancels in (24). As for Rule G, the emphasis on covariate balance reduces as n increases, the rule asymptotically reducing to that of Bandyopadhyay and Biswas (2001).

The result is a rule in which the targets vary more than in the three other rules of this section. When the estimated treatment difference is small the allocations are closer to 0.5 and when the treatment differences are overestimated, the allocation probabilities are more extreme.

6. Distributional Results

The conditional distribution of the responses in (1) is multivariate normal. But, for the adaptive design, the value of F_i depends on the preceding values of y_i, so that, unconditionally, the observations are no longer independent.

However, asymptotically, the least squares parameter estimates for this model have the same distribution as those from a non-adaptive design.

The crucial property is that the four rules considered in our paper are such that $\pi(j|x_{n+1}) > c$ for all n and j, with c some positive constant. Then, with $n_{j,n}$ the number of patients allocated to treatment j, $n_{j,n} \to \infty$ with n. Asymptotic results require that the information matrix (3) is such that $\mathcal{I}_n/n$ tends to a limit. For known covariance matrix V in (2) the form of the rules ensures that the information matrix for the treatment terms in (1) has a limit, as does the matrix for the non-stochastic period indicators G. The remaining requirement is that the distribution of the covariates Z_i is well behaved. It then follows, as in Lai and Wei (1982), that the asymptotic information matrix is that for least squares from the conditional model (1) and that the parameters are asymptotically normally distributed. If the value of ρ is not known, a consistent estimator of ρ is required in (2), together with an estimate of σ^2. Then the asymptotic information matrix is $\mathcal{I}(n)$ given by (22) for longitudinal designs.

As n increases, the treatment parameters α become increasingly precisely estimated and eventually the treatments are always correctly ordered. Then Rules D, RA and G are no longer response adaptive, although they still adapt to the values of the covariates. The rules become the analogue, for multivariate responses, of the covariate adaptive designs with skewed allocations included in Atkinson and Biswas (2005a,b). The losses for these univariate designs are similar to those of the rules for unskewed rules compared by Atkinson (2002), although convergence to asymptotic values is slower for more skewed allocations. Simulations in Table 1 for longitudinal responses show that these univariate results are also a good guide to the properties of the designs of this paper.

The value of loss indicates the effective reduction in sample size due to randomization. This value can then be used in power calculations to assess the effect of the randomization on average power. A different approach to calculation of the effect of randomization on power is that of Hu and Rosenberger (2003) who derive a relationship between power and the variability in the proportions of observations allocated to each treatment in a design targeting maximum power. Under such conditions, a procedure with higher loss will in general have a higher variance of $n_{j,n}$. For random allocation to treatment j with probability p_j^*,

$$\sqrt{n}\{\mathrm{var}(n_{j,n}/n) - p_j^*\} = p_j^*(1 - p_j^*), \tag{25}$$

which equals 1/4 for unskewed randomization. Zhang *et al.* (2007, Remark 3.3) give the term for the additional variance above that of (25) for the rule of Bandyopadhyay and Biswas (2001) which depends on α_j and on T in (23). We can therefore expect that asymptotically Rule L will have a higher loss than the other rules of Section 5.

7. Numerical Results

We now present simulation results for the four rules of Section 5 when there are two treatments. The ethical goal is to allocate 80% of the patients to the better treatment, taken as treatment 1. With $p_1^* = 0.8$, the value of the bias B_n tends to $2p_1^* - 1 = 0.6$ for large n. In these simulations we take cohorts of size 3 ($n_g = 3$) with $n_h = 4$. In (1) $\sigma^2 = 1, \Delta = \alpha_1 - \alpha_2 = 0.5$ and there are four prognostic factors with independent standard normal distributions, so that $q = 5$. In (18) three values of γ, 0.1, 0.01 and 0.001 are taken, both for Rules L and G, so that in all eight rules are compared. For designs with univariate responses which are not response adaptive, smaller values of γ initially force appreciable balance like Rule D, taking longer to move towards random allocation as n increases. The same pattern is shown here. For Rule L we take $T = 0.5941$, giving the required value of 0.8 for $\Phi(\Delta/T)$. Each rule was simulated 10,000 times. All designs were regularized to avoid extreme designs by ensuring that the minimum number of allocations of each treatment was never below $\sqrt{n}$ nor the maximum above $n - \sqrt{n}$.

Table 1 gives results for $n = 48$ and 192, both multiples of n_g, for three values of ρ. The behaviour of loss and bias as n increases is shown in Figures 1 and 2.

The left-hand third of the table contains results for the average losses $\bar{L}_{48}$ and $\bar{L}_{192}$ for the three values of ρ. When $n = 48$ the values for Rule D (non-adaptive sequential design construction) are little affected by ρ. However, when $n = 192$ the loss for $\rho = 0$ decreases to 0.27 in line with the known limit of 0 as $n \to \infty$. The decrease with n is less dramatic for the other values of ρ. For Rule RA (random allocation) there is no balancing over covariates. For $n = 48$ the loss is above 5 as it is for $n = 192$; for $\rho = 0$, the value of 5.40 is closest to that of q, which would be the loss if the ordering of the treatments were known. The loss increases with ρ, since the effective number of observations decreases, making the ordering of treatments more variable.

Table 1: The four rules of Section 5: average loss, bias and proportion of total allocations to treatment 1 when (a) $n = 48$ and (b) $n = 192$. Three values of γ each for rules G and L

$n = 48$	Loss $\bar{L}_{48}$			Bias $\bar{B}_{48}$			Total Proportion $\bar{r}_{1,48}$		
ρ	0.0	0.3	0.5	0.0	0.3	0.5	0.0	0.3	0.5
D	2.42	2.49	2.73	1.00	1.00	1.00	0.756	0.756	0.756
RA	6.63	7.53	8.95	0.59	0.59	0.60	0.726	0.711	0.686
G(0.1)	3.94	4.88	6.45	0.63	0.63	0.64	0.752	0.738	0.712
G(0.01)	2.54	3.51	5.14	0.76	0.78	0.79	0.765	0.752	0.726
G(0.001)	2.55	3.47	5.08	0.96	0.96	0.96	0.754	0.743	0.718
L(0.1)	4.58	5.45	6.79	0.64	0.64	0.63	0.749	0.741	0.727
L(0.01)	3.68	4.17	5.42	0.82	0.79	0.81	0.741	0.750	0.736
L(0.001)	3.22	4.08	5.24	0.96	0.97	0.97	0.747	0.739	0.727

$n = 192$	Loss $\bar{L}_{192}$			Bias $\bar{B}_{192}$			Total Proportion $\bar{r}_{1,192}$		
ρ	0.0	0.3	0.5	0.0	0.3	0.5	0.0	0.3	0.5
D	0.27	1.35	2.40	1.00	1.00	1.00	0.792	0.792	0.792
RA	5.40	6.77	8.90	0.61	0.61	0.61	0.782	0.778	0.769
G(0.1)	3.96	5.31	7.40	0.61	0.61	0.61	0.791	0.788	0.779
G(0.01)	1.74	2.96	4.73	0.61	0.61	0.62	0.800	0.796	0.790
G(0.001)	0.43	1.60	3.24	0.64	0.64	0.64	0.796	0.795	0.792
L(0.1)	6.87	9.90	13.92	0.61	0.62	0.63	0.793	0.792	0.788
L (0.01)	4.85	8.06	12.12	0.62	0.63	0.64	0.800	0.800	0.799
L(0.001)	3.20	6.46	10.86	0.67	0.70	0.74	0.798	0.800	0.802

The losses for Rule G show the strong effect of the value of γ, particularly when $n = 192$ and $\gamma = 0.001$. However, for $\gamma = 0.1$, there is a slight increase in loss with n as allocation becomes closer to that for Rule RA. For Rule L the loss decreases with γ although it increases going from $n = 48$ to 192 giving values significantly larger than those for Rule G with the same value of γ, with a maximum value of 13.92 when $\rho = 0.5$ and $\gamma = 0.1$. As we described in Section 5, this design remains response adaptive and the larger values of loss reflect the effect of the estimated treatment difference $\widehat{\Delta}$ on the values of l_1 and l_2.

The central third of the table gives the average biases for $n = 48$ and 192. The values agree with the theory of Section 3.2. For Rule D the bias is one for both values of n; for Rule RA the values are 0.59, 0.60 or 0.61, close to $2p_1^* - 1$. When $n = 48$ and $\gamma = 0.1$ the biases for Rule G and L are a little higher at 0.63 or 0.64. Both rules show that the effect of decreasing γ is to increase bias, strongly for $n = 48$, although much less so for $n = 192$, a feature clearly shown in the plots of Figures 1 and 2.

The right-hand side of the table gives the average allocation to treatment 1, $\bar{r}_{1,48}$. The regularization of the total number allocated to any treatment to be at most $n - \sqrt{n}$ precludes extreme allocations for n small so that all rules initially approach the overall target of 0.8 from below; Rule D is closest to the target when $n = 48$ and Rule RA the furthest from it. For Rule L the calculated values of $\hat{p}_1$ can be greater than 0.8. The entries in the table for $n = 192$ show this effect for small γ, when convergence to a proportion of 0.8 is even faster than for Rule D.

Figures 1 and 2 provide further insight into the dependence of loss and bias on the value of n. Results for $\rho = 0$ are in Figure 1. Apart from L all rules ultimately target 0.8. The upper limit of loss, as n increases, is that for Rule RA which gradually decreases to $q = 5$. The Loss is bounded below by results for Rule D, which gradually decrease to zero. The losses for Rule G are smallest for $\gamma = 0.001$.

Rules with small loss generally have high bias, a phenomenon shown in the right-hand panel of the figure. The highest bias is one for deterministic design construction, Rule D, and the lowest close to 0.6 for Rule RA. The bias for Rule G is highest for the smallest value of γ, which is the value giving the lowest loss. Rule L behaves relatively poorly in these comparisons; loss is always higher than that for Rule G with the same value of γ, as is the bias, except for $\gamma = 0.1$, when there is little to choose between the rules.

These results for uncorrelated observations are close to those for individ-

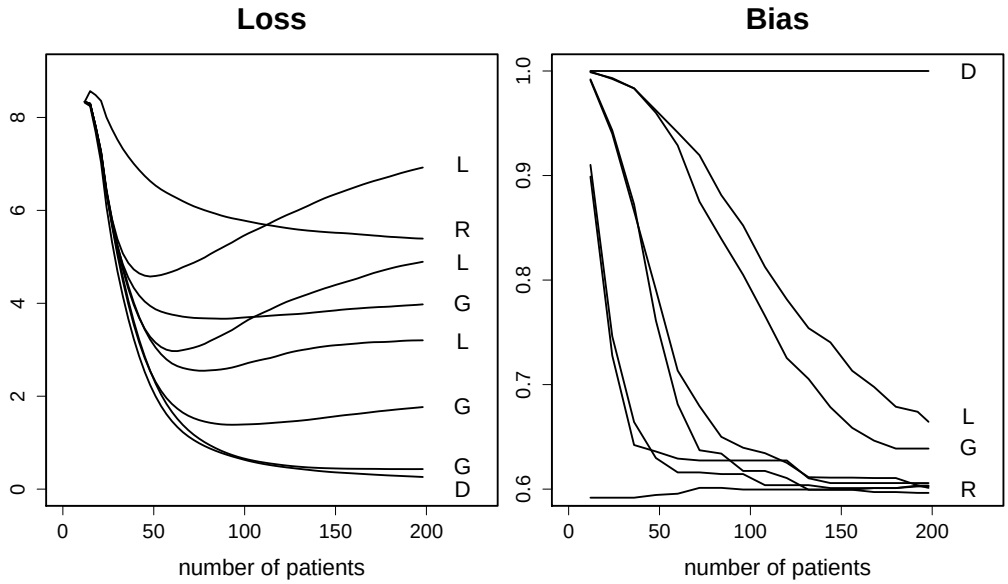

Figure 1: Correlation $\rho = 0$. Left-hand panel: Average loss $\bar{L}_n$ as a function of patient number n. Reading down at $n = 200$, Rules L(0.1), RA, L(0.01), G(0.1), L(0.001), G(0.01), G(0.001), D. Right-hand panel: Average smoothed bias $\bar{B}_n$. Reading down at $n = 70$, Rules D, L(0.001), G(0.001), L(0.01), G(0.01), G(0.1), L(0.1), R.

ual patients exhibited in Atkinson (2014). The results in Figure 2 for $\rho = 0.5$ have a similar structure to those of Figure 1 but the numbers are slightly different. The values of loss for Rules RA, G and D do not decline so fast from the initial value and, at $n = 200$ are 2 to 3 higher than those in the uncorrelated case. The loss for Rule L increases more rapidly with n, again always being above that for Rule G with the same value of γ. The bias in the right-hand panel, apart from that for Rule D, again decreases to 0.6, but more gradually than in the uncorrelated case.

The advantage of Rule L, as is shown in Table 1 is that the proportion of patients receiving the better treatment converges more rapidly to 0.8 than for the other rules. The cost is the higher loss arising from the estimation of the target proportion, rather than its convergence to the given value of 0.8. In practice, an exact allocation of 0.8 is not likely to be important and the choice of rule should be based on loss and bias. Atkinson (2014) gives plots, for rules for individual patients, of loss against bias as n increases. A rule with values of both quantities below those for all other rules at a particular n is called "admissible". The results of our comparisons show that Rule G is

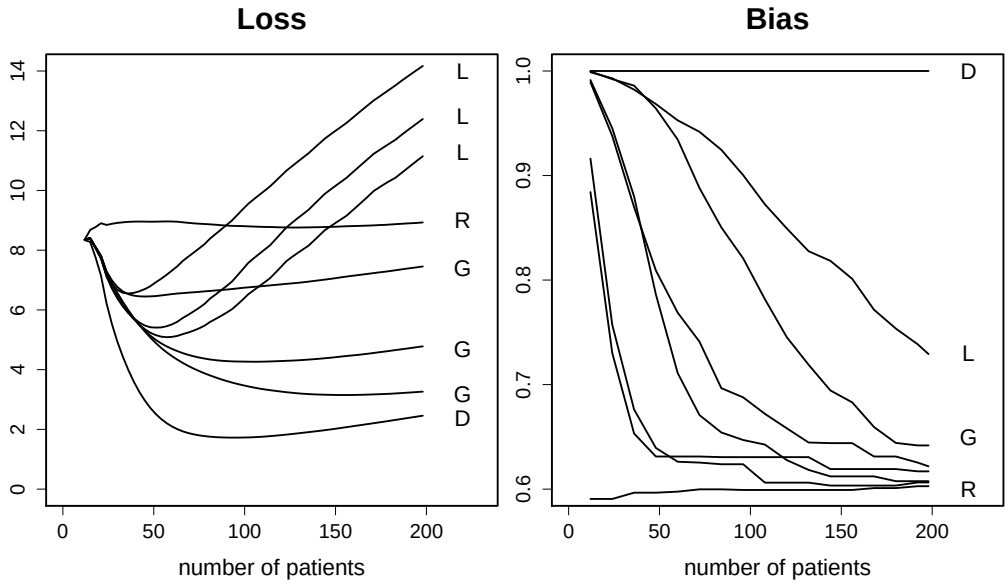

Figure 2: Correlation $\rho = 0.5$. Left-hand panel: Average loss $\bar{L}_n$ as a function of patient number n. Reading down at $n = 200$, Rules L(0.1), L(0.01), L(0.001), RA, G(0.1), G(0.01), G(0.001), D. Right-hand panel: Average smoothed bias $\bar{B}_n$. Reading down at $n = 100$, Rules D, L(0.001), G(0.001), L(0.01), G(0.01), L(0.1), G(0.1), RA.

admissible when compared to Rule L. The choice of γ depends on the relative importance of bias and loss.

Finally we consider the power of the t-test for equality of the treatment means. This is a maximum for equal allocation to the two treatments with balance over the covariates; skewing the design causes a reduction in power, as do increasing values of ρ, which effectively reduce the number of observations. We calculate the t-statistic using the elements of $\mathcal{I}_n$ to allow for any correlation in the estimates $\hat{\alpha}_1$ and $\hat{\alpha}_2$. Since very small numbers of patients are not of interest in these simulations, power is assessed by counting the number of t-statistics that are greater than 1.96. Figure 3 shows the logits of the powers for Rules D and L(0.1) as a function of ρ for three values of n. The values for the six other rules of Table 1 lie between these two and are not shown. For lower values of ρ Rule L has higher power than Rule D as it often gives a more nearly balanced allocation. As ρ increases the performance of Rule L is sometimes degraded by poor estimates of the treatment difference leading to occasional designs which are very unbalanced over treatments.

If randomization is required, Rule D is not appropriate. Then Rules G and

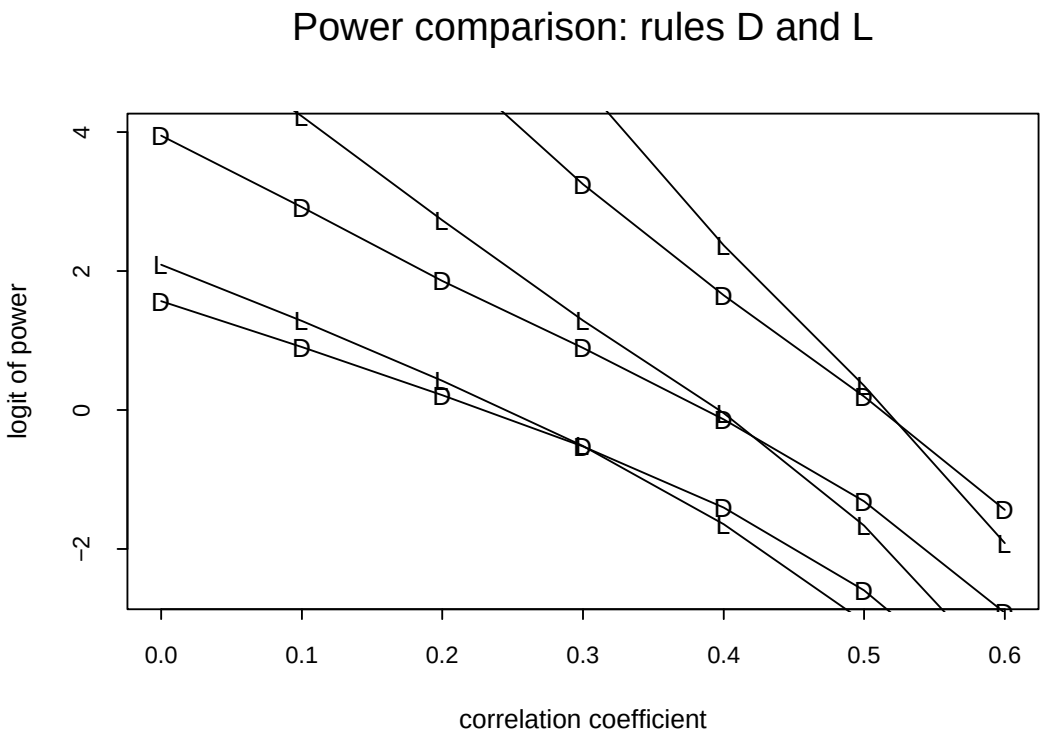

Figure 3: Power of Rules D and L(0.1) for testing treatment difference; logit of proportion significant for the t test of size 0.05. The three pairs of lines, reading down, are for $n = 198$, 99 and 48. The upper limit of power shown in the plot is approximately 0.95. The effect of the correlation ρ on the effective sample size is evident

L are to be preferred to Rule RA. However, the choice between the two sets of rules does not only depend on statistical properties. In our comparison we chose Δ/T such that $l_1 = 0.8$. If this is the clinician's ideal skewing towards the better treatment, Rule G should be used. If, on the other hand, greater skewing is required for larger values of Δ, then Rule L is appropriate. In either case the value of γ will have to be chosen to balance loss against bias.

8. Discussion

For homoscedastic regression models with t treatments giving univariate responses, the optimum allocation for testing hypotheses about the equality of the means of the treatments is to allocate a proportion $1/t$ of the patients to each treatment. However equal allocation is not always required. Dumville

et al. (2006) and Peckham *et al.* (2015) review the use of unequal allocation ratios in clinical trials.

In such work it is assumed that both the target weights, and the treatments to which they apply, are known. Unequal allocation targets arise naturally for models in which the variances of the responses to the treatments are not the same. For example, Wong and Zhu (2008) extend the D_A-optimum designs of Atkinson (1982) to heteroscedastic models in which the variances differ between treatments. Atkinson (2015) gives details for two-treatment designs in the presence of covariates.

Baldi Antognini and Giovagnoli (2015) describe compound optimum designs balancing between inference and allocation of as many patients as possible to the better treatment; the allocation targets for the various treatments may depend adaptively on parameters estimated from the responses to earlier allocations; longitudinal responses are not considered.

The linear contrast (16) provides a mechanism for skewed allocation for which var $\{A^T\widehat{\beta}\}$ is a scalar. Although the exposition here is in terms of D_A-optimality, criteria such as A- or E- optimality will yield the same optimum design minimising var $\{A^T\widehat{\beta}\}$.

Our general rule (12) provides a family of potential treatment allocation schemes which are covariate adaptive. The particular choice of G_j that led to (18) provides a rule for longitudinal responses that has many properties in common with better understood rules for univariate responses. Despite the relative computational complexity of the counting and consequent treatment allocation algorithm described in Section 4, the loss and bias of the rule are straightforward analogues of those for univariate responses. In particular, the analogy of the values of loss is striking, but depends on the correct definition of n_{effec}.

For small n Rule G forces skewed allocation and the loss is close to that of Rule D. But, as n increases the rule becomes increasingly like skewed random allocation, with a higher loss but with bias tending to zero. Rules for smaller values of the tuning constant γ have a higher initial emphasis on targeting the target skewing proportions p_j^*. In the selection of a biased-coin design for clinical trials the emphasis in the statistical literature is often on trials that provide allocations very close to these targets. Atkinson (2012) stresses the importance of considering both loss and selection bias. Rule G is such that the bias decreases as n increases, but in such a way that the efficiency of estimation goes to 100%. A rule with constant loss of $q/5$ can be obtained

by simplifying (18) and taking the allocation probabilities proportional to $\widehat{R}(j)d_A(j, n, x_{n+1})$ (Burman, 1996).

Two final points. First we note that we have assumed a known value of ρ. Indeed, the evaluation of the inferential properties of the design, particularly n_{effec} and the loss will depend strongly on ρ. However, the dependence of the design itself on ρ is slight. We would recommend designing for an arbitrary value, such as 0.3, whilst sequentially estimating ρ. Of course, the value of the estimate $\hat{\rho}$ should be used in any inferences drawn from the results of the trial. Secondly, the methods may be extended to regression models with distributions other than the normal through the use of elemental information matrices as described in Atkinson *et al.* (2014).

Acknowledgements

We are grateful to the referees for comments which helped clarify the exposition of our results.

Appendix A. Multivariate $\mathbf{D_A}$-Optimum Designs

Appendix A.1. Sequential Design Construction

For the sake of generality we extend the single linear combination of the parameters $a^T\beta$ in (8) to the set of s combinations $A^T\beta$, where A is a $(t + v + n_h - 1) \times s$ matrix of known constants. D_A-optimum experimental designs for the linear regression model (1) maximize $|A^T\{\mathcal{I}(n)\}^{-1}A|^{-1}$ and so minimize the generalized variance of these linear combinations, providing a normal theory confidence region of minimum volume. Such optimum designs can be constructed sequentially.

We first consider univariate responses when F_i in (1) becomes the vector f_i^T, which includes the vectors of allocation and prognostic factors for the ith patient. When allocation is made to patient $n + 1$, all other allocations are known. A useful matrix result for D-optimum designs maximizing $|\mathcal{I}(n+1)|$ is that

$$|\mathcal{I}(n+1)| = [1 + f_{n+1}^T\{\mathcal{I}(n)\}^{-1}f_{n+1}]|\mathcal{I}(n)| = \{1 + d(j, n, x_{n+1})\}|\mathcal{I}(n)|. \quad \text{(A-1)}$$

That treatment is allocated for which $d(j, n, x_{n+1})$ is a maximum.

In the iterative construction of D_A-optimum designs for a univariate response,

$$d_A(j,n,x_{n+1}) = f_{n+1}^T\{\mathcal{I}(n)\}^{-1}A[A^T\{\mathcal{I}(n)\}^{-1}A]^{-1}A^T\{\mathcal{I}(n)\}^{-1}f_{n+1}, \quad (j=1,\ldots,t). \tag{A-2}$$

In the absence of randomization, patient $n+1$ would receive the treatment for which $d_A(j,n,x_{n+1})$ is a maximum.

We now turn to multivariate data and find the equivalent of (A-2). Let the uth row of F_{n+1} be denoted $f_{u,n+1}^T$. We extend (A-2) to

$$d_{A,uv}(j,n,x_{n+1}) = f_{u,n+1}^T\{\mathcal{I}(n)\}^{-1}A[A^T\{\mathcal{I}(n)\}^{-1}A]^{-1}A^T\{\mathcal{I}(n)\}^{-1}f_{v,n+1}.$$

With element u,v of V^{-1} written V^{uv}, the equivalent of (A-2) is

$$d_A(j,n,x_{n+1}) = \sum_{u=1}^{n_h}\sum_{v=1}^{n_h} V^{uv}d_{A,uv}(j,n,x_{n+1}). \tag{A-3}$$

This is the function in our generic rule (18).

Appendix A.2. An Equivalence Theorem

As $n\to\infty$, $d_A(j,n,x_{n+1})\to 0$. If we replace the number of patients allocated to each treatment by the continuous distribution of asymptotic proportions of allocation we obtain a design measure ξ and an information matrix $\mathcal{I}(\xi)$. Also $d_A(j,n,x_{n+1})$ tends to the directional derivative $d_A(j,\xi,x)$. In the case of non-sequential design we can consider choosing treatments from a space $\mathcal{J}$ and covariates from a space $\mathcal{Z}$. The design region is then $\mathcal{X}=\mathcal{J}\times\mathcal{Z}$. We now extend the General Equivalence Theorem of optimum design theory (Kiefer and Wolfowitz 1960, Whittle 1973) to multivariate D_A-optimality.

If we let

$$\bar{d}_A(\xi) = \sup_{j,x\in\mathcal{X}} d_A(x,\xi),$$

the equivalence theorem states that the D_A-optimum design, denoted ξ^*, is such that

$$\bar{d}_A(\xi_{DA}^*) = s.$$

Here s is the number of independent linear combinations of the parameters specified by A.

In the sequential construction of optimum designs for clinical trials there will asymptotically be balance for each treatment over the prognostic factors. With such balance we ignore x_{n+1} and write

$$d_A(j, \xi^*) = s, \quad (j = 1, ..., t).$$

The balanced design used in (20) to derive the expression for loss is $n_{\text{effec}} n\xi^*$. The results in Table 1 for Rule D show how fast the sequential construction of the optimum design converges.

References

Atkinson, A. C. (1982). Optimum biased coin designs for sequential clinical trials with prognostic factors. *Biometrika*, **69**, 61–67.

Atkinson, A. C. (2002). The comparison of designs for sequential clinical trials with covariate information. *Journal of the Royal Statistical Society, Series A*, **165**, 349–373.

Atkinson, A. C. (2012). Bias and loss: the two sides of a biased coin. *Statistics in Medicine*, **31**, 3494–3503.

Atkinson, A. C. (2014). Selecting a biased-coin design. *Statistical Science*, **29**. 144–163.

Atkinson, A. C. (2015). Optimum designs for two treatments with unequal variances in the presence of covariates. *Biometrika*, **102**, 494–499.

Atkinson, A. C. and Biswas, A. (2005a). Bayesian adaptive biased-coin designs for clinical trials with normal responses. *Biometrics*, **61**, 118–125.

Atkinson, A. C. and Biswas, A. (2005b). Optimum design theory and adaptive-biased coin designs for skewing the allocation proportion in clinical trials. *Statistics in Medicine*, **24**, 2477–2492.

Atkinson, A. C. and Biswas, A. (2014). *Randomised Response-Adaptive Designs in Clinical Trials*. Chapman and Hall/ CRC Press, Boca Raton.

Atkinson, A. C., Donev, A. N., and Tobias, R. D. (2007). *Optimum Experimental Designs, with SAS*. Oxford University Press, Oxford.

Atkinson, A. C., Fedorov, V. V., Herzberg, A. M., and Zhang, R. (2014). Elemental information matrices and optimal experimental design for generalized regression models. *Journal of Statistical Planning and Inference*, **144**, 81–91.

Baldi Antognini, A. and Giovagnoli, A. (2015). *Adaptive Designs for Sequential Treatment Allocation*. Chapman and Hall/ CRC Press, Boca Raton.

Ball, F. G., Smith, A. F. M., and Verdinelli, I. (1993). Biased coin designs with a Bayesian bias. *Journal of Statistical Planning and Inference*, **34**, 403–421.

Bandyopadhyay, U. and Biswas, A. (1999). Allocation by randomized play-the-winner rule in the presence of prognostic factors. *Sankhya, Series B*, **61**, 397–412.

Bandyopadhyay, U. and Biswas, A. (2001). Adaptive designs for normal responses with prognostic factors. *Biometrika*, **88**, 409–419.

Bandyopadhyay, U. and Biswas, A. (2004). An adaptive allocation for continuous response using wilcoxonmannwhitney score. *Journal of Statistical Planning and Inference*, **123**, 207–224.

Biswas, A. and Dewanji, A. (2004a). Inference for a RPW-type clinical trial with repeated monitoring for the treatment of rheumatoid arthritis. *Biometrical Journal*, **46**, 769–779.

Biswas, A. and Dewanji, A. (2004b). A randomized longitudinal play-the-winner design for repeated binary data. *Australian and New Zealand Journal of Statistics*, **46**, 675–684.

Biswas, A. and Dewanji, A. (2004c). Sequential adaptive designs for clinical trials with longitudinal response. In N. Mukhopadhyay, S. Chattopadhyay, and S. Datta, editors, *Applied Sequential Methodologies*, pages 69–84. Marcel Dekker, New York.

Biswas, A. and Mandal, S. (2007). Optimal three-treatment response-adaptive designs for phase iii clinical trials with binary responses. In J. Lopez-Fidalgo, J. M. Rodriguez-Diaz, and B. Torsney, editors, *MODA 8 – Advances in Model-Oriented Design and Analysis*, pages 32–40. Physica-Verlag, Heidelberg.

Biswas, A., Bhattacharya, R., and Zhang, L. (2007). Optimal response-adaptive designs for continuous responses in phase III trials. *Biometrical Journal*, **49**, 928–940.

Biswas, A., Park, E., and Bhattacharya, R. (2012). Covariate-adjusted response-adaptive designs for longitudinal treatment responses: PEMF trial revisited. *Statistical Methods in Medical Research*, **21**, 379–392.

Burman, C.-F. (1996). *On Sequential Treatment Allocations in Clinical Trials*. Department of Mathematics, Göteborg.

Cox, D. R. (1951). Some systematic experimental designs. *Biometrika*, **38**, 312–323.

Dumville, J. C., Hahn, S., Miles, J. N. V., and Torgerson, D. J. (2006). The use of unequal randomisation ratios in clinical trials: a review. *Contemporary Clinical Trials*, **27**, 1–12. doi:10.1016/j.cct.2005.08.003.

Durham, S., Fluornoy, N., and Li, W. (1998). A sequential design for maximizing the probability of a favourable response. *Canadian Journal of Statistics*, **26**, 479–495.

Everitt, B. S. and Pickles, A. (2004). *Statistical Aspects of the Design and Analysis of Clinical Trials, Revised edition*. Imperial College Press, London.

Fedorov, V. V. (1972). *Theory of Optimal Experiments*. Academic Press, New York.

Fedorov, V. V. and Leonov, S. L. (2014). *Optimal Design for Nonlinear Response Models*. Chapman and Hall/ CRC Press, Boca Raton.

Galbraith, S. and Marschner, I. C. (2002). Guidelines for the design of clinical trials with longitudinal outcomes. *Controlled Clinical Trials*, **23**, 257–273.

Gallo, P., Chuang-Stein, C., Dragalin, V., Gaydos, B., Krams, M., and Pinheiro, J. (2006). Adaptive designs in clinical drug development: An executive summary of the PhRMA working group. *Journal of Biopharmaceutical Statistics*, **16**, 275–283.

Heritier, S., Gebski, V., and Pillai, A. (2005). Dynamic balancing randomization in controlled clinical trials. *Statistics in Medicine*, **24**, 3729–3741.

Hu, F. and Rosenberger, W. F. (2003). Optimality, variability, power: Evaluating response-adaptive randomization procedures for treatment comparisons. *Journal of the American Statistical Association*, **98**, 671–678.

Huang, T., Liu, Z., and Hu, F. (2013). Longitudinal covariate-adjusted response-adaptive randomized designs. *Journal of Statistical Planning and Inference*, **143**, 1816–1827.

Ivanova, A. (2003). A play-the-winner-type urn design with reduced variability. *Metrika*, **58**, 1–13.

Kiefer, J. and Wolfowitz, J. (1960). The equivalence of two extremum problems. *Canadian Journal of Mathematics*, **12**, 363–366.

Lai, T. L. and Wei, C. Z. (1982). Least squares estimates in stochastic regression models with applications to identification and control of dynamic systems. *Annals of Statistics*, **10**, 154–166.

Molenberghs, G., Thijs, H., Jansen, I., Beunckens, C., Kenward, M. G., Mallinckrodt, C., and Carroll, R. J. (2004). Analyzing incomplete longitudinal clinical trial data. *Biostatistics*, **5**, 445–464.

Peckham, E., Brabyn, S., Cook, L., Devlin, T., Dumville, J., and Torgerson, D. J. (2015). The use of unequal randomisation in clinical trials An update. *Contemporary Clinical Trials*, **45**, 113 –122.

Rosenberger, W. F. (1993). Asymptotic inference with response-adaptive treatment allocation designs. *Annals of Statistics*, **21**, 2098–2107.

Rosenberger, W. F. and Hu, F. (2004). Maximizing power and minimizing treatment failures in clinical trials. *Clinical Trials*, **1**, 141–147.

Rosenberger, W. F. and Sverdlov, O. (2008). Handling covariates in the design of clinical trials. *Statistical Science*, **23**, 404–419.

Rosenberger, W. F., Stallard, N., Ivanova, A., Harper, C. N., and Ricks, M. L. (2001). Optimal adaptive designs for binary response trials. *Biometrics*, **57**, 909–913.

Shao, J., Yu, X., Y., and Zhong, B. (2010). A theory for testing hypotheses under covariate-adaptive randomization. *Biometrika*, **97**, 347–360.

Silvey, S. D. (1980). *Optimum Design*. Chapman and Hall, London.

Sutradhar, B. C. and Jowaheer, V. (2006). Analyzing longitudinal count data from adaptive clinical trials: a weighted generalized quasi-likelihood approach. *Journal of Statistical Computation and Simulation*, **76**, 1079–1093.

Sutradhar, B. C., Biswas, A., and Bari, W. (2005). Marginal regression for binary longitudinal data in adaptive clinical trials. *Scandinavian Journal of Statistics*, **32**, 93–114.

Wei, L. and Durham, S. D. (1978). The randomized play-the-winner rule in medical trials. *Journal of the American Statistical Association*, **73**, 840–843.

Whittle, P. (1973). Some general points in the theory of optimal experimental design. *Journal of the Royal Statistical Society, Series B*, **35**, 123–130.

Wong, W. K. and Zhu, W. (2008). Optimum treatment allocation rules under a variance heterogeneity model. *Statistics in Medicine*, **27**, 4581–4595.

Zelen, M. (1969). Play the winner and the controlled clinical trial. *Journal of the American Statistical Association*, **64**, 131–146.

Zelen, M. and Wei, L. J. (1995). Foreword. In N. Flournoy and W. F. Rosenberger, editors, *Adaptive Designs*. IMS, Hayward, CA.

Zhang, L. and Rosenberger, W. F. (2006). Response-adaptive randomization for clinical trials with continuous outcomes. *Biometrics*, **62**, 562–569.

Zhang, L. and Rosenberger, W. F. (2007). Response-adaptive randomization for survival trials: the parametric approach. *Applied Statistics*, **56**, 153–165.

Zhang, L.-X., Hu, F., Cheung, S. H., and Chan, W. S. (2007). Asymptotic properties of covariate-adjusted response-adaptive designs. *The Annals of Statistics*, **35**, 1166–1182.