

***Jointly optimal design of Phase II and Phase III
clinical trials: an over-arching approach***

Christopher Jennison

Department of Mathematical Sciences,

University of Bath, UK

<http://people.bath.ac.uk/mascj>

PSI Conference, The Belfry

May 2012

The Phase II / Phase III problem

Phases of drug development occur sequentially.

These phases can be “optimised” individually but it is also sensible to consider design of the overall development process.

Dose finding in Phase IIb is often based on a non-linear, parametric dose response model with 3 or 4 parameters.

Thus, a model for the two Phases is quite different from that for two groups in a group sequential Phase III trial.

However, the fundamentals of problem formulation and methods of solution should have much in common with designing an optimal group sequential trial.

Problem formulation will require many assumptions: investigators may not find it easy to provide the required probabilities or costs and benefits — but they do currently make decisions that appear to depend on these unknowns!

DIA (formerly PhRMA) Working Group

The Adaptive Programs stream of the Adaptive Design Scientific Working Group has been studying the joint design of Phase II and Phase III trials.

Members of the “Main model” team studying generic methods are:

Carl-Fredrik Burman (leader)

Zoran Antonijevic

Christy Chuang-Stein

Chris Jennison

Fredrik Öhrn

Nitin Patel

José Pinheiro

Alun Bedding

Other teams are working on specific application areas: Diabetes, Neuropathic pain, and Oncology.

Outline of talk

1. Elements of the Phase IIb / Phase III decision process

Dose response model and prior for Bayesian analysis

Phase IIb responses, Phase III responses and the final decision

Gain function and sampling costs

Risk of losing a drug due to poor safety results

2. Optimising Phase IIb / Phase III

Formulation of the problem

An optimisation algorithm

Some preliminary results

3. Extensions and methods for improving computational speed


4. Results when the Phase III trial is group sequential

The Emax dose response model

We assume a 4 parameter Emax model with mean response at dose d_j

$$\mu(d_j) = \theta_1 + \theta_2 \frac{d_j^{\theta_4}}{\theta_3^{\theta_4} + d_j^{\theta_4}}.$$

Emax dose response curve



Dose response model and prior for Bayesian analysis

Doses $d_0 = 0$ (control) and $d_j, j = 1, \dots, J$, (active) are to be tested.

We shall have $J = 7$ active doses in our example, with Emax model

$$\mu(d_j) = \theta_1 + \theta_2 \frac{d_j^{\theta_4}}{\theta_3^{\theta_4} + d_j^{\theta_4}}.$$

In the prior, we suppose the four parameters are independent and

$$\theta_1 \sim N(5, 10^2),$$

$$\theta_2 \sim N(5, 10^2),$$

$$\theta_3 \sim N^+((d_J - d_0)/2, (d_J - d_0)^2),$$


$$\theta_4 \sim N^+(1, 1).$$

Here, N^+ denotes a normal distribution restricted to positive values.

Assessing the prior for Emax model parameters

The following plot shows a sample of 20 Emax dose response curves using parameter values generated from the chosen prior.

The curves with a negative treatment effect are coloured blue.



Phase IIb responses

We assume a fixed Phase IIb design with n_{2j} subjects on dose j , $j = 0, \dots, J$.

In our example, we suppose patients are allocated equally to each active dose and at 3 times this rate to dose zero. Thus, with a total of n_2 subjects in Phase IIb,

$$n_{20} = 0.3 n_2 \quad \text{on dose zero,}$$

$$n_{2j} = (0.7/J) n_2 \quad \text{on each active dose } j = 1, \dots, J.$$

Given $(\theta_1, \theta_2, \theta_3, \theta_4)$, subjects on dose j have independently distributed responses

$$X_{ij} \sim N(\mu(d_j), \sigma^2).$$

We shall assume $\sigma = 9$.

Combining the likelihood of these responses with the prior for $(\theta_1, \theta_2, \theta_3, \theta_4)$ gives the posterior distribution for the dose response curve for use in designing Phase III.

Phase III responses

Suppose it is decided to test dose d_j against control in a Phase III trial.

Here, $2n_3$ subjects are randomised equally between dose zero and dose d_j .

Responses are distributed as

$$Y_{i0} \sim N(\mu(d_0), \sigma^2) \quad \text{on dose zero,}$$

$$Y_{ij} \sim N(\mu(d_j), \sigma^2) \quad \text{on dose } j.$$

We test $H_{0j}: \mu(d_j) - \mu(d_0) \leq 0$ against $\mu(d_j) - \mu(d_0) > 0$.

If H_{0j} is rejected at a significance level below α , efficacy of dose j is established.

We shall use $\alpha = 0.0005$ in our example, rather than consider two separate Phase III trials each with a target significance level of 0.025.

Gain function and sampling costs

We suppose a positive outcome in Phase III leads to approval of the new drug and a financial gain g .

Running the Phase IIb trial incurs a sampling cost of c_2 per subject.

Running Phase III incurs a cost of c_3 per subject.

In our example, we shall take

$$c_2 = 1,$$

$$c_3 = 1,$$

$$g = 12,000.$$

A cost of 1 unit may be \$10,000 to \$50,000, depending on the condition being investigated — so g represents a multi-million dollar return.

Risk of failure for safety

We suppose there is a probability $\gamma(d_j)$ that dose j will eventually fail on safety grounds.

This could occur in Phase III or later on in post-marketing surveillance.

We assume $\gamma(d)$ is a known, increasing function of d .

The function $\gamma(d)$ is specified before Phase IIb and patient follow-up in Phase IIb is not long enough to learn more about the safety profile.

In our example, we shall take $\gamma(d)$ to be quadratic with $\gamma(d_J) = 0.2$.

Thus, when Phase III has a positive outcome, we calculate the expected gain by discounting the gain function by a factor $1 - \gamma(d_j)$.

Optimising the Phase IIb / Phase III design

Before Phase IIb

We choose the Phase IIb sample size, n_2 .

At the end of Phase IIb

We decide whether to proceed to run Phase III and, if so, select

The dose to test in Phase III d_j ,

The Phase III sample size n_3 .

We wish to optimise:

The choice of n_2 ,

The rule for deciding whether to proceed to Phase III,

The rule for choosing d_j and n_3 .

Optimisation algorithm

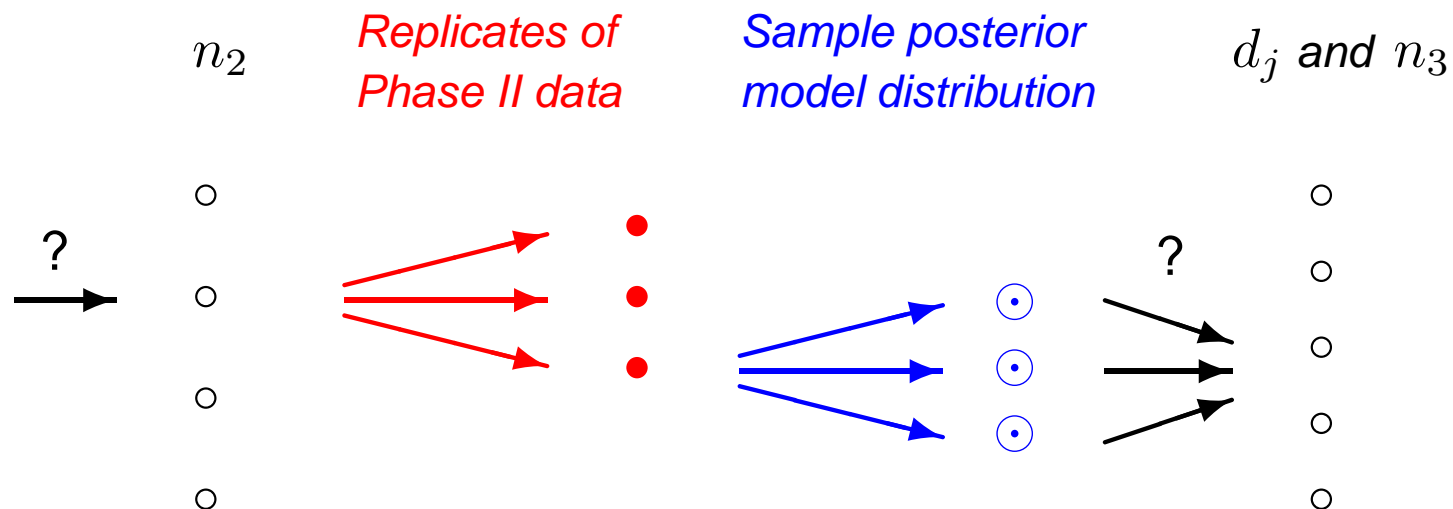
For a particular n_2 :

Simulate θ , the vector of dose response curve parameters, from the prior.

Simulate Phase IIb data, given θ .

Evaluate Phase III options given the posterior for θ and choose the best option.

Average over replicates to compute the expected net gain for this n_2 .



Compare $E(\text{Net gain})$ over possible choices of n_2 and choose the best n_2 .

Evaluating Phase III options

Given Phase II data $X = x$, denote the posterior distribution of the dose response curve parameters θ by

$$\pi(\theta|x).$$

Consider a Phase III trial with dose d_j and sample size n_3 .

The conditional expectation of the net gain is

$$\int \pi(\theta|x) [P_{\theta}\{\text{Positive Phase III}; d_j, n_3\} (1 - \gamma(d_j)) g - 2 n_3 c_3 - n_2 c_2] d\theta.$$

With a sample $\theta^1, \dots, \theta^S$ from $\pi(\theta|x)$, estimate this $E(\text{Net gain})$ by

$$\frac{1}{S} \sum_{s=1}^S P_{\theta^s}\{\text{Positive Phase III}; d_j, n_3\} (1 - \gamma(d_j)) g - 2 n_3 c_3 - n_2 c_2.$$

Results for a simple example

Consider a problem with 7 active dose levels $d_j = j, j = 1, \dots, 7$.

Following the earlier definition, the prior distribution for $\theta = (\theta_1, \theta_2, \theta_3, \theta_4)$ has

$$\theta_1 \sim N(5, 10^2), \quad \theta_2 \sim N(5, 10^2),$$

$$\theta_3 \sim N^+(7/2, 7^2), \quad \theta_4 \sim N^+(1, 1).$$

Phase IIb has $0.3 n_2$ subjects on dose zero and $0.1 n_2$ on each active dose.

The sampling cost is 1 unit for each Phase IIb and Phase III subject.

The financial gain for a positive Phase III trial is $g = 12,000$.

But dose d_j may fail on safety grounds with probability

$$\gamma(d_1) = 0.004, \quad \gamma(d_2) = 0.016, \quad \gamma(d_3) = 0.037, \quad \gamma(d_4) = 0.065,$$

$$\gamma(d_5) = 0.10, \quad \gamma(d_6) = 0.15, \quad \gamma(d_7) = 0.2.$$

Results for a simple example

We have optimised over Phase III sample sizes

$$n_3 \in \{100, 150, 200, 250, 300, 400, 500, 600, 800, 1000\}.$$

Comparing Phase IIb designs, we found:

n_2	$E(\text{Net gain})$	n_2	$E(\text{Net gain})$
25	4,515	200	4,770
50	4,595	250	4,780
75	4,660	300	4,795
100	4,695	350	4,795
125	4,715	400	4,795
150	4,740	450	4,780
175	4,755	500	4,755

So, we conclude the optimal choice is $n_2 = 300$.

Accuracy of comparisons

Comparisons of Phase IIb designs are based on:

500 replicates of Phase II data sets,

500 samples from posterior distribution of θ for each Phase II data set.

n_2	$E(\text{Net gain})$
200	4,770
250	4,780
300	4,795
350	4,795

Estimated values of $E(\text{Net gain})$ are subject to sampling error with



Standard errors of $E(\text{Net gain}) \approx 200$.

However, coupling the simulations of Phase II data sets leads to

Standard errors of differences in $E(\text{Net gain}) \approx 10$.

Results for a simple example

Within replicates of Phase IIb data for $n_2 = 300$, the optimal choice of d_j and n_3 varies considerably:




The risk of safety problems guides the decision towards lower doses.

Sampling costs in Phase III argue for lower values of n_3 .

Results for a simple example

We can explore the decisions made in selecting a dose to go forward to Phase III.




A high dose is selected when the posterior samples of the dose response curve show modest treatment effects.

More promising results lead to lower doses being chosen (especially for higher n_2).

Results for a simple example

High posterior means for $E(Y)$ translate into high Phase III success probabilities.



For the highest doses, probability of Phase III success is offset by greater risk of safety problems.

Extending the methodology

Phase III options

Two Phase III trials.

Group sequential Phase III designs.

Allowing two or more active doses to be tested in Phase III.

Gain function and costs

Define the gain function to be the net present value based on:

patent life remaining after a successful Phase III,

true treatment effect (or estimated effect?) at selected dose.

Elicit a problem-specific gain function for two successful doses in Phase III.

Portfolio management: Choosing which of several candidate treatments (possibly for different indications) should go forward to a Phase III trial.

Extending the methodology

Additional model features

Learning about safety problems in Phase IIb.

Change of endpoint between Phase IIb and Phase III.

Phase IIb options

Different fixed patterns of dose allocation.

Adaptive dose-allocation.

Early stopping in Phase IIb.

Computational problems and possible solutions

Coupling We have used coupling of replicate data sets under different Phase IIb designs to increase the accuracy of comparisons *between* these designs.

Sampling the posterior distribution See Jane Temple's talk on how to sample directly from the posterior distribution of dose response curve parameters.

Multiple use of samples from the posterior model distribution

Rather than repeat simulations to sample the posterior distribution of θ for Phase IIb data sets which are similar due to coupling, values for a “central” case can be re-used with importance sampling weights to provide results for other cases.

Pre-computing for a reference set of cases

More complex Phase III designs (group sequential or multi-armed) can be evaluated up-front on a grid of parameter values, creating a look-up table for general cases.

This re-use of information for different interim states has a parallel with the dynamic programming (backwards induction) optimisation algorithm.

Results when Phase III has a group sequential design

Consider the previous example but now allow Phase III to be group sequential.

Members of the ρ -family of one-sided, error spending designs (Jennison & Turnbull, 2000, Ch. 7) are known to be highly efficient (Barber & Jennison, *Biometrika*, 2002).

We shall use this form of design with 5 groups and $\rho = 2$.

With the same values of n_3 for possible *maximum* Phase III sample sizes, we find:



n_2	$E(\text{Net gain})$	n_2	$E(\text{Net gain})$
25	5,240	200	5,280
50	5,270	250	5,260
75	5,300	300	5,255
100	5,300	350	5,235
125	5,290	400	5,215
150	5,290	450	5,190
175	5,280	500	5,155

Results when Phase III has a group sequential design

The Expected net gain is considerably higher (by about 500 units) than with a fixed sample size Phase III trial.

The optimal Phase IIb sample size is smaller (100 rather than 300).

The group sequential Phase III designs means it is less crucial to have an accurate estimate of the treatment effect on which to base the Phase III sample size.



Conclusions

A full treatment of the Phase IIb/ Phase III design process is possible, with joint optimisation of both stages under a Bayesian model.

The Bayesian approach allows propagation of uncertainty and provides a natural framework for decision making under uncertainty.

Simulations from the posterior distribution nested within replicates of Phase IIb data constitute a substantial computation task. However, there are several routes to improving computational efficiency and making this task feasible.

Rules for dose selection and Phase II design are often based on point estimates of the treatment effects at individual doses. Comparison with optimised decision rules will provide a check on such rules and insight into their strengths and weaknesses.

Rules can be assessed under a random distribution of θ values from the prior or under a specific θ (but optimisation is still based on the assumed prior for θ).