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

Modelling and Simulation in

Pharmaceutical Drug Development

PSI Workshop, Chesham

October 2012

The Phase II / Phase III problem

Phases of drug development occur sequentially.

There has been a lot of work on “optimising” phases individually — but not much on designing 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.

We shall need to estimate these parameters after Phase IIb and make decisions for Phase III design based on these estimates.

In a Bayesian approach we shall need a prior for the parameters and methods for dealing with their posterior distribution given Phase IIb data.

A full problem formulation will require further assumptions:

Risk of safety problems *vs* dose,

Costs of Phase IIb and Phase III trials,

The benefit from a successful outcome in Phase III.

DIA (formerly PhRMA) Working Group

I shall describe work on the joint design of Phase II and Phase III trials by the Adaptive Progams stream of the Adaptive Design Scientific Working Group

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 the Phase IIb / Phase III process

1. Phase IIb trial

   Compare several doses against control

   Decide whether to proceed to Phase III

   If so, choose a dose and specify the Phase III sample size

2. Phase III trial

   Run two Phase III trials comparing the selected dose against control

   If both trials provide significant evidence of a treatment effect, we have a success!

Design questions:

Sample size for Phase IIb,

Decision making after Phase IIb: Stop/go, dose, Phase III sample size.

Outline of the optimal design process

To maximise expected net gain in a Bayes decision theoretic approach, we shall

Specify a model for dose response in Phase IIb and Phase III

Specify a prior distribution for parameters in the dose response model

Design the Phase IIb trial (doses, sample size, etc)

Run the Phase IIb trial

Find the posterior distribution of model parameters given Phase IIb data

Design the Phase III trials (stop/go, dose, sample size)

Run two Phase III trials

Analyse the Phase III data and see if the process has a successful outcome.

Here, the colour coding distinguishes between:

Model specification, Simulation, Design optimisation.

How can we optimise the design in this complex problem?

Optimising the overall design is a complicated problem.

For instance, the best way to design the Phase IIb trial depends on how we use the results of Phase IIb in designing Phase III.

The first step is to be able to work forwards through the whole process, using some (non-optimal) decision rules where necessary.

We shall consider how to:

1. Choose the prior distribution for dose response model parameters,
2. Simulate Phase IIb trial data,
3. Find the posterior distribution of model parameters given Phase IIb data,
4. Optimise the Phase III design given this posterior distribution,
5. Optimise the Phase IIb design.

The Emax dose response model

We shall assume a 4 parameter Emax dose-response model.

In this model, the mean response at dose d is

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

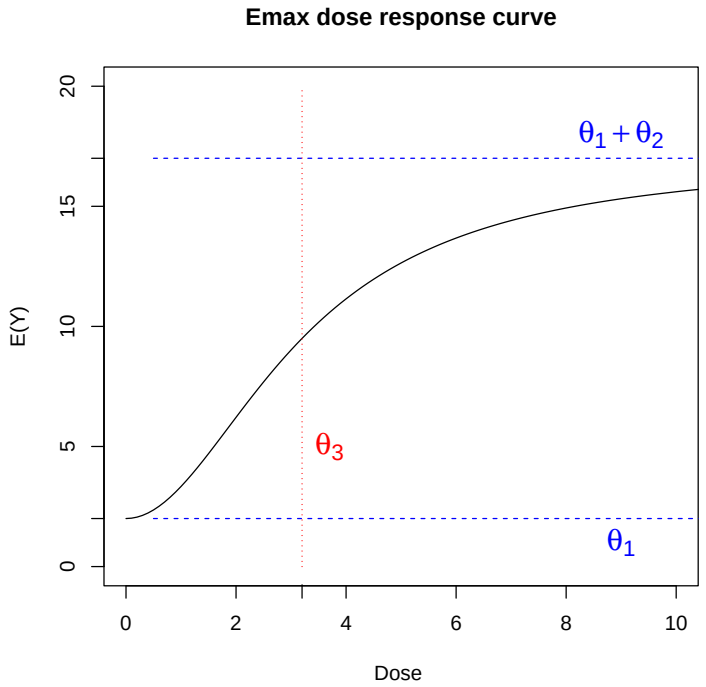

θ_1: Mean response at dose zero (placebo effect)

θ_2: Increase in mean response from dose zero to a very high dose

θ_3: ED50, the dose achieving half this maximum increase

θ_4: Governs the steepness of the dose response curve

1. Specifying the prior distribution of dose response parameters

In our example, we have 7 active doses and the control. Suppose units are defined so that the control dose is $d_0 = 0$ and the active doses are $d_j = j$, $j = 1, \ldots, 7$.

In reality, the values 0, 1, ..., 7 might represent doses of, say, 0, 50, ..., 350 *mg*.

So, the Emax model gives mean responses

$$\mu_j \;=\; \theta_1 \;+\; \theta_2 \;\frac{j^{\,\theta_4}}{\theta_3^{\theta_4} + j^{\,\theta_4}}\,, \quad j = 0, \ldots, 7.$$

Suppose that, in the prior, we assume the four parameters are independent and

$$\begin{aligned}
\theta_1 &\sim N(a_1,\, b_1^2), \\
\theta_2 &\sim N(a_2,\, b_2^2), \\
\theta_3 &\sim N^+(a_3,\, b_3^2), \\
\theta_4 &\sim N^+(a_4,\, b_4^2).
\end{aligned}$$

Here, N^+ denotes a normal distribution restricted to values greater than 0.01.

Specifying the prior distribution of dose response parameters

The plot shows a sample of 20 Emax dose response curves from the prior with

$$a_1 = 2,\ b_1 = 1,\ \ a_2 = 4,\ b_2 = 3,\ \ a_3 = 4,\ b_3 = 2,\ \ a_4 = 4,\ b_4 = 0.5.$$

Note that the plot is of the increase in mean response over dose zero.

The curves with a negative treatment effect are coloured blue.

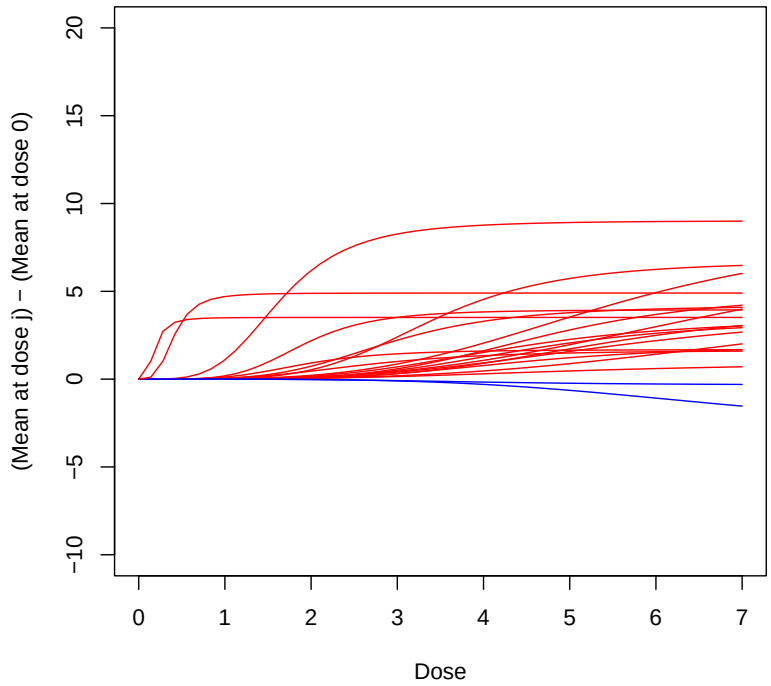

Specifying the prior distribution of dose response parameters

The investigators tell you that a treatment effect of 5 would be clinically meaningful — but they hope to see a higher effect.

There may be a placebo effect of 5 or 10 at dose zero.

A treatment effect (vs placebo) of 5 or 10 should be attainable.

A dose of 6 or 7 may be necessary to achieve this.

Opinion varies as to how steep the dose response curve might be.

Historical results for similar compounds at this stage indicate

$$P\{\text{No positive treatment effect}\} \;=\; 0.4.$$

The R routine prior_sampler.R generates a sample of 20 dose response curves from a specified prior distribution.

See run_prior_sampler.R for an example of how to call this routine.

Find values of $a_1, b_1, \ldots, a_4, b_4$ that match the above expectations.

(1) Running R: Initial set up

You can run R in a particular directory on your computer.

It is advisable to create this directory first and copy into it the files provided on the memory stick.

After you start R, click on File, then on Change dir . . . .

Navigate to find the directory you want to be in and double click on its name.

As a check, you can type

```
> getwd()
```

to find the current directory.

If you are working in a directory containing copies of the files from the memory stick, these will appear in the list of options when you try to load a file or open a script.

(1) Running R: prior_sampler

To use the R routine prior_sampler.R, you need to load this into your workspace.

To do this, click on File, then on Source R Code. Navigate to the directory containing prior_sampler.R, then double click on this to load the code.

If you then type

```
> ls()
```

You should see the functions “emaxmodel” and “prior_sampler” listed as being present in the workspace.

You can run commands in run_prior_sampler.R by opening this as a script.

Click on File, then on Open Script . . . . Navigate to run_prior_sampler.R, then double click on this to open a script.

If you highlight commands in the script and type Control R, these commands will run in the R Console.

You can edit the script file to modify the commands or add new ones.

(1) Running R: prior_sampler

The output from the commands in run_prior_sampler.R is a plot of 20 dose response curves generated from the specified prior.

Note that the plot is of the increase in mean response over dose zero, so all curves start at zero.

The command

```
> x11()
```

in run_prior_sampler.R creates a new graphics window each time you run the script.

Graphics windows are overlaid, so you have to move the top ones aside to see the old ones below.

In modifying the prior so that the sample of dose response curves matches the specification, you need to know what each of the parameters $\theta_1, \ldots, \theta_4$ controls.

Look back to Slide 7 for an illustration of each parameter's role.

(1) Notes on using Scripts

Having an open script file is a useful way to develop or modify a set of R commands.

If all you want to do is run the commands, you can simply load the file (click on File, then on Source R Code, etc.) and the full set of commands will be implemented.

With an open script, you can choose to run a few commands at a time, edit some lines and run them again, and so forth.

You can save changes to a script by clicking on File and Save while the script window is active.

I suggest that

(a) You keep the provided scripts unchanged, by saving using the Save as . . . option or by working with a copy of the original script.

(b) You close each script as you finish using it — it is easy enough to open it again if you need to.

2. Simulating Phase IIb responses

We assume a normal response distribution for patient response in Phase IIb — and the same response distribution in Phase III.

Given Emax model parameters $(\theta_1, \theta_2, \theta_3, \theta_4)$, we assume subjects on dose j have independent, normally distributed responses

$$X_{ij} \sim N(\mu_j, 9^2),$$

where μ_j is given by the Emax model formula.

We shall consider Phase IIb designs in which 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, we have

$0.3\, n_2$ on dose zero,

$0.1\, n_2$ on each active dose $j = 1, \ldots, 7$.

Simulating Phase IIb responses

Given a Phase IIb design (e.g., sample size n_2) we wish to know properties such as

$$P\{\text{The Phase III trials are ultimately successful}\}$$

and expected total sample size.

We shall estimate such quantities by averaging over simulated Phase IIb data sets.

To create these data sets, we

- Simulate a vector $(\theta_1, \theta_2, \theta_3, \theta_4)$ from the prior,
- Calculate $\mu_j = E(X_{ij})$ at each dose j,
- Simulate responses X_{ij} for subjects on each dose,

producing a number of data sets from hypothetical realisations of the Phase IIb trial.

How should we run these simulations to serve our purposes most effectively?

Simulating Phase IIb responses

We shall create data sets from realisations of a Phase IIb trial by

Simulating $(\theta_1, \theta_2, \theta_3, \theta_4)$ from the prior,

Calculating the means μ_j,

Simulating responses X_{ij} for subjects on each dose.

Question 1.

We could generate one data set from each $(\theta_1, \theta_2, \theta_3, \theta_4)$ simulated from the prior — or a sample of several data sets for each $(\theta_1, \theta_2, \theta_3, \theta_4)$.

Which approach is better for estimating the quantities of interest?

Question 2.

Suppose we wish to compare properties of two Phase IIb designs with different sample sizes, n_2.

Are there useful ways to link the data sets generated for the two cases?

3. Generating a sample from the posterior distribution of $\boldsymbol{\theta = (\theta_1, \theta_2, \theta_3, \theta_4)}$ given Phase IIb data

Background information only: an implementation of this method is provided.

Combining the likelihood of the Phase IIb responses with the prior for θ gives the posterior distribution for θ after Phase IIb — which is not very tractable.

If we use Markov chain Monte Carlo simulation, we face the usual problems of

Uncertainty about how rapidly the Markov chain sampler converges,

Correlated samples.

For our choice of prior, the posterior distribution of θ has some useful properties:

In the posterior distribution, the conditional distribution of (θ_1, θ_2) given (θ_3, θ_4) is bivariate normal.

Thus, we can integrate out θ_1 and θ_2 to get an expression for the joint density of θ_3 and θ_4, up to a multiplicative constant.

Sampling the posterior distribution of $\boldsymbol{\theta = (\theta_1, \theta_2, \theta_3, \theta_4)}$

Once we have a formula for the joint posterior density of (θ_3, θ_4), we can:

- Create an envelope for the density of (θ_3, θ_4) on a two-dimensional grid,
- Use acceptance sampling to obtain exact, independent samples of (θ_3, θ_4),
- Combine each sample of (θ_3, θ_4) with a pair (θ_1, θ_2) from the known conditional bivariate normal distribution.

A problem and a solution:

- We find may our envelope for the density of θ_1 and θ_2 is not high enough in places and some acceptance probabilities are greater than one.
- We can correct for this by giving importance sampling weights (greater than 1) to these values.
- Typically, only a few samples need a weight greater than 1.

Sampling the posterior distribution of $\boldsymbol{\theta = (\theta_1, \theta_2, \theta_3, \theta_4)}$

Denote by $\pi_{\theta|X}(\theta|x)$ the posterior distribution of θ given Phase IIb data $X = x$.

The R routine posterior_sampler.R generates a sample from $\pi_{\theta|X}(\theta|x)$.

If a sample size S is specified, the output is a set of vectors θ^s, $s = 1, \ldots, S$.

Each value θ^s has a weight w_s (the importance weight, which is 1 in most cases).

Suppose we are interested in the conditional expectation of the function $f(\theta)$ given Phase IIb data $X = x$.

This is given by the integral

$$\int f(\theta)\, \pi_{\theta|X}(\theta|x)\, d\theta,$$

which we approximate by the sum

$$\frac{\sum_{s=1}^{S} w_s\, f(\theta^s)}{\sum_{s=1}^{S} w_s}.$$

Creating a set of data

The R routine ph2_sampler.R generates a sample of Phase IIb data for a specified parameter vector θ and sample size $n2$.

The output is a vector of observed mean responses at doses $0, 1, \ldots, 7$.

Use the commands in run_ph2_sampler.R to specify θ and n_2 and call the routine ph2_sampler.R — see next slide

The last two commands in run_ph2_sampler.R plot the mean responses against dose and superimpose the true dose response curve on this plot.

The R routine posterior_sampler.R generates a sample of vectors θ from the posterior distribution of θ given observed Phase IIb data.

Use the commands in run_posterior_sampler.R to sample from the posterior distribution of θ given the observed data generated by ph2_sampler.R.

Use the script run2_posterior_sampler.R to do more of the same — generating new Phase IIb data sets and sampling the posterior distributions.

(3) Running R: ph2_sampler

To load ph2_sampler.R, into your workspace, click on File, then on Source R Code.
Navigate to ph2_sampler.R, and double click on this to load the code.

If you then type

```
> ls()
```

You should see that the functions “emaxmodel” and “ph2_sampler” are present in the workspace.

Open the script run_ph2_sampler.R, by clicking on File, then on Open Script . . . .
Navigate to run_ph2_sampler.R, then double click on this to open the script.

Highlight commands in the script and type Control R to run these commands.

The `set.seed` command initialises the seed for the random number generator.

If you re-run the command

```
> xbar=ph2_sampler(theta,n2)
```

by itself, the seed will have changed and you will get a new set of data.

(3) Running R: posterior_sampler

Having loaded ph2_sampler.R and run run_ph2_sampler.R to generate data xbar, load posterior_sampler.R, by clicking on File, then on Source R Code, etc.

Open run_posterior_sampler.R by clicking on File, then on Open Script . . . , etc.

Highlight commands in the script and type Control R to run these commands.

To run the whole script, press Control A (highlighting all the lines), then Control R.

You will be prompted to load the package MASS by double clicking on its name.

This package has a command for generating bivariate normal random variables.

You can also load this by clicking on Packages, then Load packages . . . , etc.

If the plots produced by run_posterior_sampler.R *are overlaid, move them around to see all the output.*

The script run2_posterior_sampler.R does not load MASS (it is already loaded).

You can run this repeatedly to see more examples.

Using the samples from the posterior distribution

For a particular n_2 we can:

Simulate $\boldsymbol{\theta}$, the vector of dose response curve parameters, from the prior.

Simulate Phase IIb data, given $\boldsymbol{\theta}$.

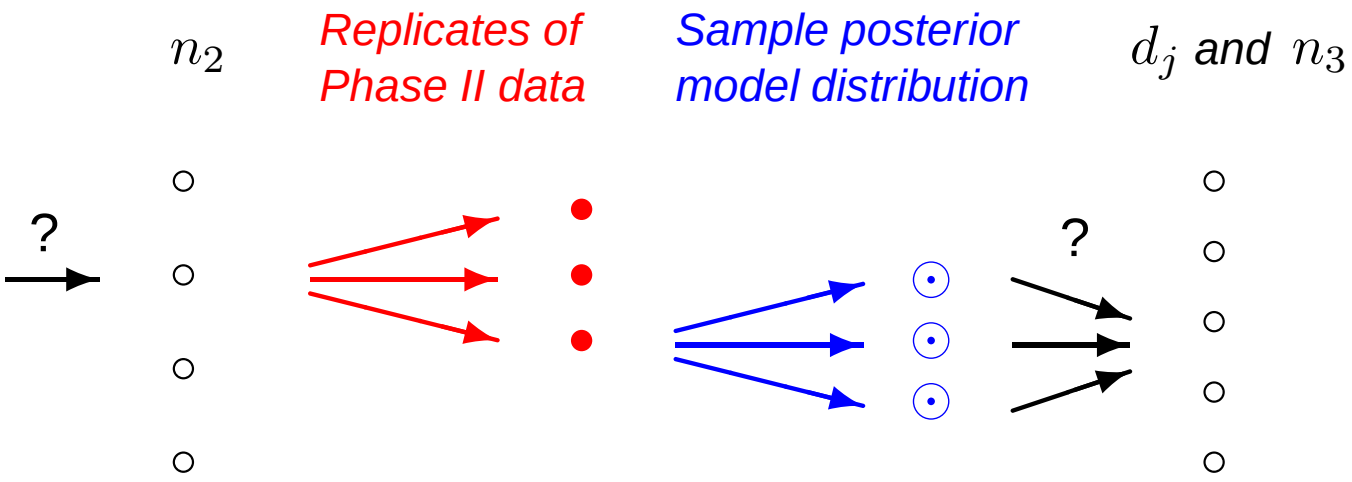

Next, we shall use the samples representing the posterior distribution after Phase IIb to evaluate the Phase III options and choose the best option.

Then, we can average over replicates of Phase IIb to compute the expected net gain for a given n_2 and compare results to choose the best n_2.

4. Optimising the Phase III design given Phase IIb data

The Phase III trials

Suppose it is decided to test dose j against control in Phase III.

We run two Phase III trials.

In each, $2n_3$ subjects are randomised equally between dose 0 and dose j.

Responses are distributed as

$$X_{i0} \sim N(\mu_0, \sigma^2) \quad \text{on dose zero,}$$

$$X_{ij} \sim N(\mu_j, \sigma^2) \quad \text{on dose } j.$$

In each trial, we test H_{0j}: $\mu_j - \mu_0 \leq 0$ against $\mu_j - \mu_0 > 0$.

(We assume the response distribution at each dose is exactly as in Phase IIb.)

If H_{0j} is rejected at a significance level below $\alpha = 0.025$ in both trials, efficacy of dose j is established.

Optimising the Phase III design

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

$$c_2 = 1,$$

$$c_3 = 1,$$

$$g = 12{,}000.$$

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

Optimising the Phase III design

Risk of failure for safety

Suppose the probability that dose d will eventually fail on safety grounds is $\gamma(d)$.

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 $\gamma(7) = 0.2$. Thus, the risk for dose j is

$$\gamma_j \;=\; (j/7)^2 \times 0.2.$$

When Phase III has a positive outcome, we calculate the expected gain by discounting the gain function by a factor $1 - \gamma_j$.

Optimising the Phase III design

At the end of Phase IIb

We must decide whether to proceed to run Phase III at all.

If we decide to run Phase III, we must select

The dose to test in Phase III j,

The sample size per arm, in each Phase III trial n_3.

We wish to make these decisions optimally

Overall, our optimality criterion is the expected net gain:

$$g \times P\{\text{Both Phase III trials significant at } \alpha = 0.025 \text{ and no safety problems}\}$$

$$-n_2\, c_2 - 4\, E(N_3)\, c_3.$$

Optimising the Phase III design given Phase IIb data

Evaluating Phase III options

If we decide to run Phase III trials with dose j and sample size n_3, the conditional expectation of the net gain, given Phase IIb data $X = x$, is

$$\int \left[P_\theta\{\text{Positive Phase III}; j, n_3\}\,(1-\gamma_j)\,g\; - 4\,n_3\,c_3 - n_2\,c_2\right]\,\pi_{\theta|X}(\theta|x)\,d\theta,$$

where $\pi_{\theta|X}(\theta|x)$ denotes the posterior density of θ given $X = x$.

We estimate this conditional expected net gain by

$$\frac{1}{S}\sum_{s=1}^{S} P_{\theta^s}\{\text{Positive Phase III}; j, n_3\}\,(1-\gamma_j)\,g\; - 4\,n_3\,c_3 - n_2\,c_2, \qquad (1)$$

where θ^s, $s = 1, \ldots, S$, is a sample from $\pi_{\theta|X}(\theta|x)$.

The optimal Phase III design is that which maximises (1) over j and n_3.

(4) Optimising the Phase III design given Phase IIb data

Evaluating Phase III options

Run create_theta_sample.R (which is similar to run_posterior_sampler.R) to generate a set of Phase IIb data, then simulate a data array of vectors $\theta^1, \ldots, \theta^{500}$ from the posterior distribution of θ.

Load the script create_theta_sample.R and run these commands. (If you have re-started R with an empty workspace, you will need to re-load ph2_sampler.R and posterior_sampler.R and the package MASS.)

Use Control A to highlight all the commands in create_theta_sample.R and Control R to run these commands.

The plots produced by the commands in create_theta_sample.R show

(i) The Phase IIb data, superimposed on the dose response curve from which the data were generated,

(ii) A contour plot of the posterior density of (θ_3, θ_4).

(4) Running R: create_theta_sample

The resulting 500×5 array theta_sample contains vectors $\theta^1, \ldots, \theta^{500}$ sampled from the posterior distribution of θ given the Phase IIb data set, with an importance weight (usually 1) associated with each vector.

Type

```
> theta_sample
```

to look at the values in the array theta_sample.

The final entry in each row is the importance sampling weight. Do any of these weights differ from 1?

You can give the command

```
> max(theta_sample[,5])
```

to find the largest value of the numbers in column 5 of theta_sample.

(4) Optimising the Phase III design given Phase IIb data

Evaluating Phase III options

The R script e_cond_net_gain.R evaluates terms that make up the conditional expected net gain

$$\frac{1}{S}\sum_{s=1}^{S} P_{\theta^s}\{\text{Positive Phase III}; j, n_3\}\,(1-\gamma_j)\,g \;-4\,n_3\,c_3 - n_2\,c_2$$

for a sample $\theta^1, \ldots, \theta^S$ and specified dose j and Phase III sample size n_3.

Starting with the posterior sample $\theta^1, \ldots, \theta^{500}$ produced by create_theta_sample.R, run the R script e_cond_net_gain.R to calculate the probability of a successful Phase III outcome for each combination of

i) ***Dose $j \in \{1, \ldots, 7\}$, and***

ii) ***Phase III sample size $n_3 \in \{50, 75, 100, 150, 200, 300, 400, 500\}$.***

Hence, determine whether it is worthwhile to conduct the Phase III trials and, if so, which dose and sample size should be chosen.

(4) Running R: e_cond_net_gain

Load the script e_cond_net_gain.R, and run these commands to analyse the posterior distribution represented by theta_sample.R.

Use Control A to highlight the whole script and Control R to run these commands.

The script produces a number of 7×8 arrays, with the rows representing the doses $j = 1, \ldots, 7$ and the columns the 8 options for n_3.

To see $P\{\text{Two successful Phase III trials}\}$ for each combination of j and n_3, averaged over the posterior sample of θ vectors, type

```
> epsuc
```

The probability of safety problems for each combination of j and n_3 is shown by

```
> mgamma
```

What is the meaning of the output from the following command?

```
> (1-mgamma)*epsuc
```

(4) Running R: e_cond_net_gain

Look at the array of sampling costs for each combination of j and n_3.

```
> ecost
```

Note that these values include the cost of n_2 observations in Phase IIb.

We can combine the variables epsuc, mgamma and ecost, to obtain the array of expected net gains

```
> (1-mgamma)*epsuc*g-ecost
```

Check this agrees with the array egain calculated by the script.

The location of the largest value in this array gives the optimal combination of dose j and Phase III sample size n_3.

How high does this largest value have to be for Phase III to be worthwhile?

Check you agree with the optimal dose and Phase III sample size reported by e_cond_net_gain.R.

5. Optimising the Phase IIb design

Once we know how to optimise the Phase III design for a given Phase IIb data set, we are ready to optimise Phase IIb.

Assume an Emax dose response model and a prior in which the Emax model parameters are independent with

$$\begin{aligned} \theta_1 &\sim N(5,\, 10^2), \\ \theta_2 &\sim N(5,\, 10^2), \\ \theta_3 &\sim N^+(3.5,\, 7^2), \\ \theta_4 &\sim N^+(1,\, 1^2), \end{aligned}$$

where N^+ denotes a normal distribution restricted to values greater than 0.01.

Let response distributions, gain and cost functions, and the risk of safety problems be as defined previously.

Given the Phase IIb data we can apply the methods we have just seen to find the optimal decisions regarding Phase III.

(5) Running R: e_net_gain

Evaluating Phase IIb options

The R script e_net_gain.R reads in data from the file posterior_samples to create a $6 \times 500 \times 500 \times 5$ data array.

These data were produced by simulating 500 Phase IIb data sets for each of 6 n_2 values. For each data set, a sample $(\theta^1, \ldots, \theta^{500})$ with related importance weights was simulated from the posterior distribution of θ.

The script e_net_gain.R goes on to evaluate the expected net gain

$$g \times P\{\text{Both Phase III trials significant at } \alpha = 0.025 \text{ and no safety problems}\}$$

$$-n_2\, c_2 - 4\, E(N_3)\, c_3$$

for Phase IIb designs using each possible choice of n_2.

Here, the decision whether to conduct Phase III trials and, if so, the choice of dose and sample size n_3 optimise this expected gain.

(5) Running R: e_net_gain

Evaluating Phase IIb options

The script e_net_gain.R averages the expected gain (optimised over Phase III designs) for each of the 500 Phase IIb data sets from a given n_2.

This allows us to compare the benefits of different Phase IIb sample sizes, n_2, and choose the optimal value from the set $\{50, 100, 200, 300, 400, 500\}$.

Apply the commands in e_net_gain.R to find the optimal choice of n_2 in this formulation of the overall design problem.

To do this, first load the routines in e_net_gain_routines.R and e_net_gain.R by clicking on File, then on Source R Code, etc.

Next, click on File and Source R Code and load run_e_net_gain.R to run this code.

There are a lot of calculations to carry out and you should see results appear gradually (one set every few minutes) for one value of n_2 at a time.

(5) Running R: e_net_gain

For each n_2, there is output giving estimates, and associated standard errors, of:

The expected gain,

The Probability of two successful Phase III trials,

$P\{\text{Two successful Phase III trials}\}$ discounted for the risk of safety failure,

The expected cost of Phase III sampling.

For each n_2, histograms are drawn to show the distribution of the optimal dose and the optimal Phase III sample size n_3 over the 500 Phase IIb data sets.

We find the optimal choice of n_2 in the overall design problem by comparing the expected gain under different choices of n_2. What is the optimal n_2?

The histograms of optimal dose and the optimal Phase III sample size for this optimal n_2 are drawn again at the end.

(5) Running R: e_net_gain

The final commands in e_net_gain.R draw plots showing the relationship between the posterior means of various quantities and the optimal dose for Phase III.

These quantities are

- The posterior mean of the treatment effect, $E(X)$,
- The posterior probability of two successful Phase III trials,
- The posterior mean of the expected total gain.

For each Phase IIb data set, the quantity is evaluated at the dose selected for Phase III.

Explore the output from e_net_gain.R to see how the success probability and Phase III sampling costs behave as n_2 increases.

Can you think of reasons why the final histograms and plots look the way they do?

6. Recap: 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 j,

The Phase III sample size n_3.

We aim to optimise:

The choice of n_2,

The rule for deciding whether to proceed to Phase III,

The rule for choosing j and n_3.

Optimisation algorithm

For a particular n_2:

Simulate $\boldsymbol{\theta}$, the vector of dose response curve parameters, from the prior.

Simulate Phase IIb data, given $\boldsymbol{\theta}$.

Evaluate Phase III options given the posterior for $\boldsymbol{\theta}$ and choose the best option.

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

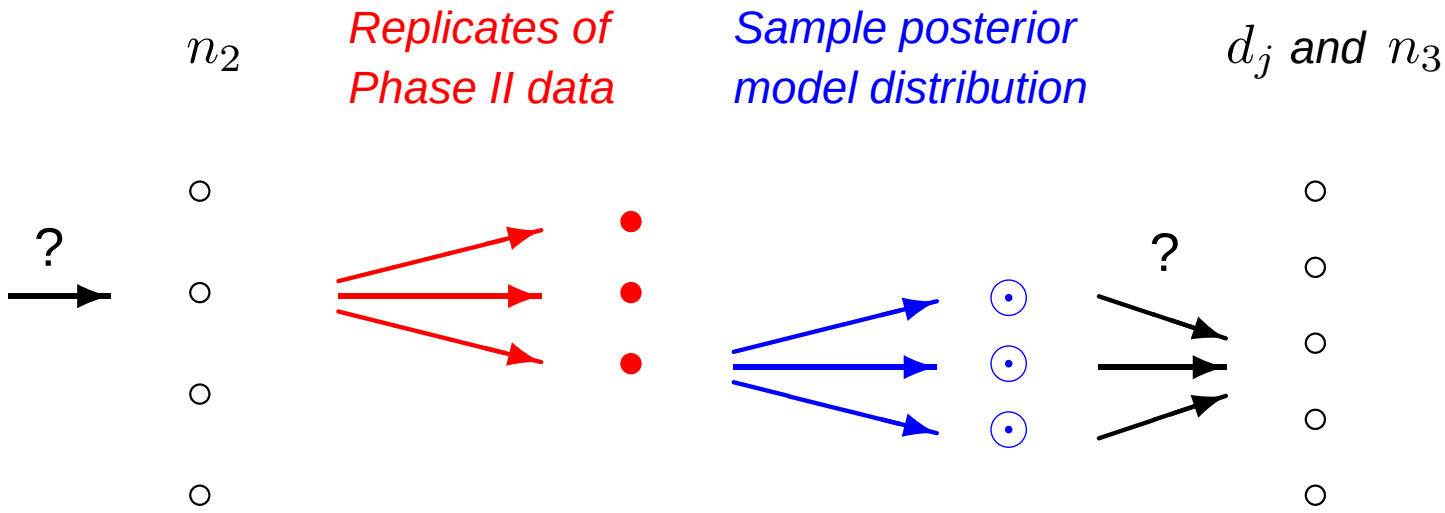

Compare $E(\text{Net gain})$ over possible choices of n_2 and choose the best n_2.

Results for a worked example

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

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

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

$$\theta_3 \sim N^+(3.5,\, 7^2), \quad \theta_4 \sim N^+(1,\, 1).$$

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

$$\gamma_1 = 0.004, \quad \gamma_2 = 0.016, \quad \gamma_3 = 0.037, \quad \gamma_4 = 0.065,$$

$$\gamma_5 = 0.10, \quad \gamma_6 = 0.15, \quad \gamma_7 = 0.2.$$

Results for a simple example

We optimise over Phase III sample sizes

$$n_3 \in \{50, 75, 100, 125, 150, 200, 250, 300, 400, 500\}.$$

Comparing Phase IIb designs, we find:

n_2	$E(\text{Net gain})$	n_2	$E(\text{Net gain})$
25	4,375	200	4,630
50	4,450	250	4,635
75	4,520	**300**	**4,650**
100	4,555	350	4,645
125	4,575	400	4,645
150	4,600	450	4,630
175	4,615	500	4,605

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

Breakdown of the expected net gain

The $E(\text{Net gain})$ values are made up from:

n_2	$P(\text{Overall success}^*)$	$4\,E(N_3)$	$E(\text{Net gain})$
25	0.441	893	4,375
50	0.447	861	4,450
100	0.460	862	4,555
150	0.466	837	4,600
200	0.473	843	4,630
250	0.478	854	4,635
300	**0.483**	**850**	**4,650**
350	0.487	847	4,645
400	0.490	840	4,645
450	0.492	823	4,630
500	0.493	814	4,605

* Two successful Phase III trials and no safety problems.

Accuracy of comparisons

Comparisons of Phase IIb designs are based on:

500 replicates of Phase IIb data sets,

500 samples from posterior distribution of $\boldsymbol{\theta}$ for each Phase IIb data set.

n_2	$E(\text{Net gain})$
250	4,635
300	**4,650**
350	4,645
400	4,645

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

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

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

Standard errors of differences in $E(\textit{Net gain}) \approx$ *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:

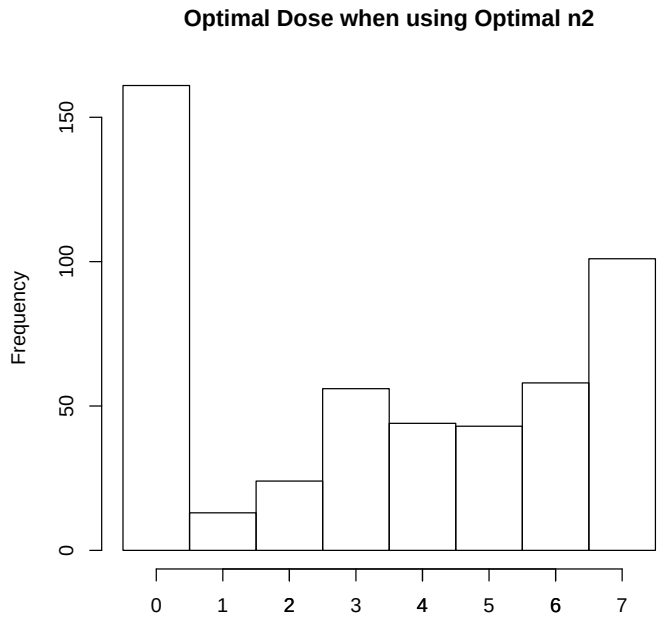

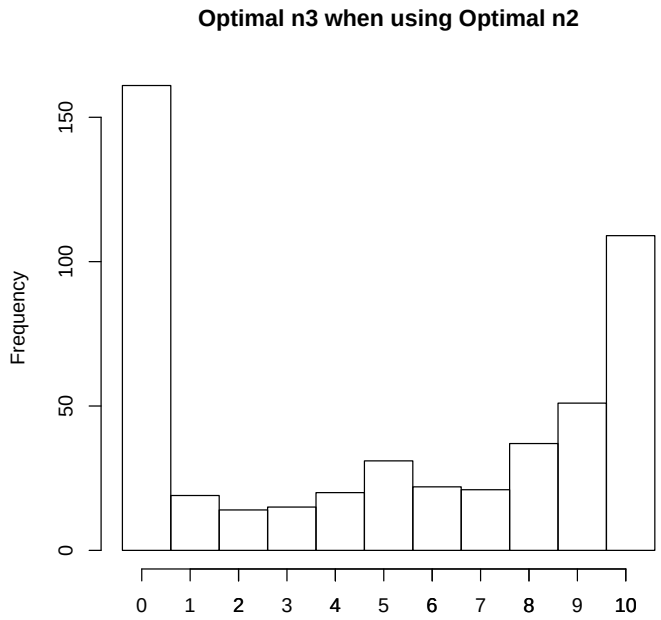

The risk of safety problems guides the decision towards lower doses.

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.

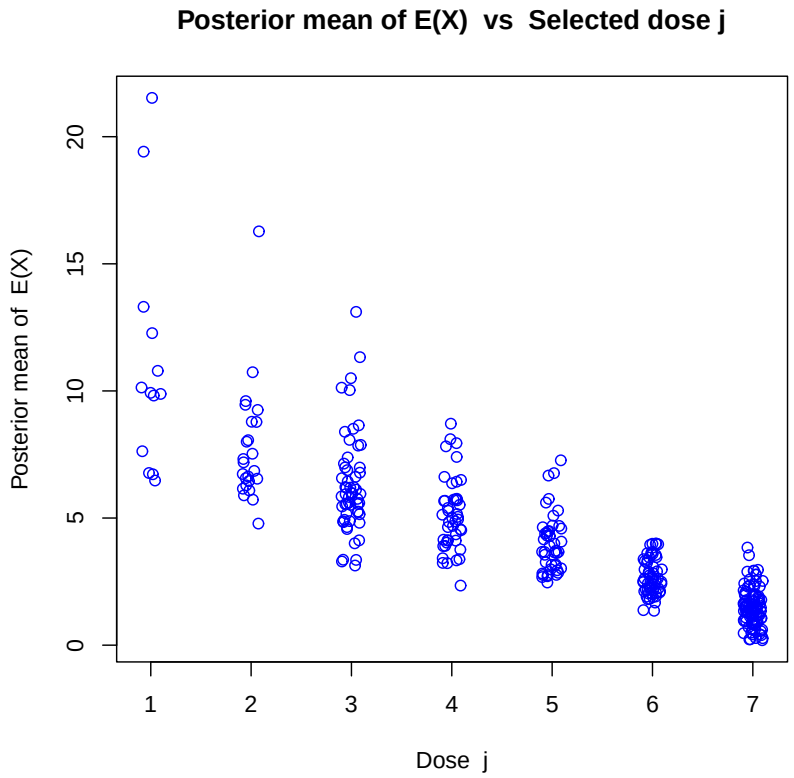

A high dose is selected when the posterior samples of the dose response curve show modest treatment effects.

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

Results for a simple example

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

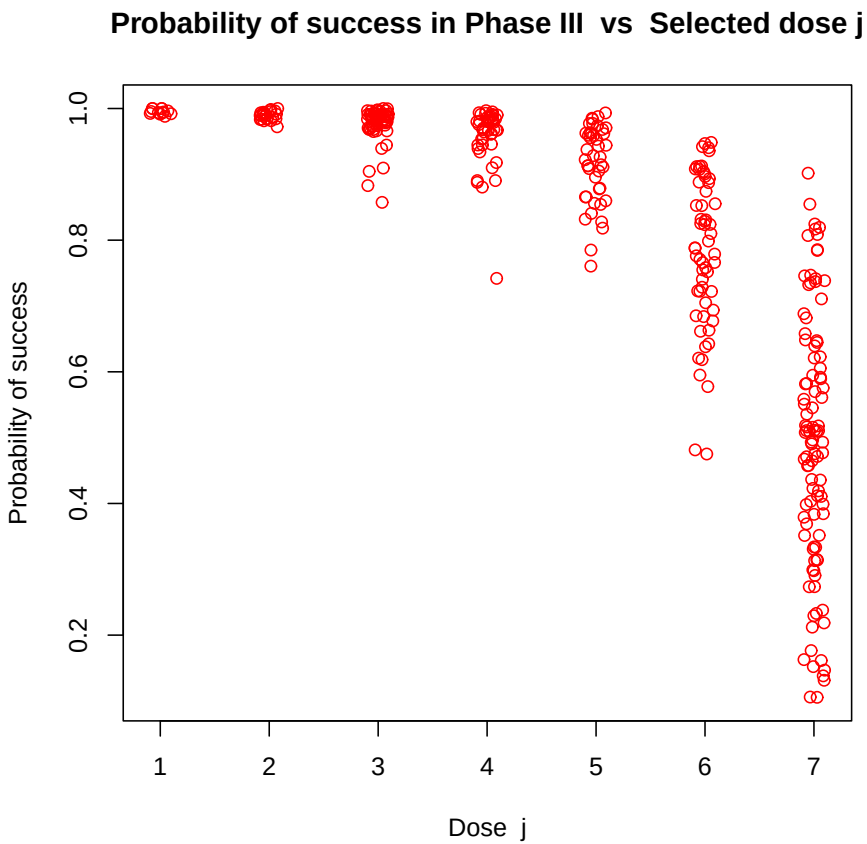

For the highest doses, probability of Phase III success is offset by greater risk of safety problems.

7. Extending the methodology

Consider the problem formulation and the methods we have used to solve the optimal design problem.

This problem might be regarded as about the simplest possible formulation that has sufficient ingredients to make a solution possible.

Question 1.

Which aspects of the problem would you like to modify in order to achieve a more realistic model for the Phase II – Phase III process?

Question 2.

Which of these modifications do you believe to be computationally feasible using the approach that we have followed?

(7) Implementing some extensions

The "net present value" function npvgain in e_net_gain_routines.R takes arguments

$n2$, the Phase IIb sample size,

$vn3$, the vector of possible Phase III sample sizes,

$effect$, the treatment effect at the dose under consideration.

One might wish to let this gain function depend on the time taken to reach a conclusion and the treatment effect size.

Explore this by modifying the function npvgain to reflect that:

Larger Phase III trials eat into patent lifetime and reduce the period in which a new drug earns income for the manufacturer,

The income a drug earns depends on how effective it is (over and above showing that it has at least some effect).

If you want some ideas, one alternative definition is available in alt_npvgain.R

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 computational task. However, there are several routes to improving computational efficiency and making this task feasible.

There are many directions in which to elaborate the problem we have studied. Some of these elaborations can be handled with a similar amount of computation — but others may be more challenging!