bCRM – Defining and Running Simulations

Tessella Project Number 3760

System User Guide

Tessella Support Services plc

e-mail: post@tessella.com **www.tessella.com**

Head Office: Abingdon

3 Vineyard Chambers, Abingdon, Oxfordshire OX14 3PX, England
Tel: +44 (0)1235 555511
Fax: +44 (0)1235 553301

Document Control

This document is under document control. All members of the distribution list will receive updated copies whenever alterations are made.

Contributors

Person	Role	Company	Contribution
Geraint Lloyd	Author	Tessella	
Tom Parke	Reviewer	Tessella	

Document Approval

Person	Role	Company

Distribution

Person	Role	Company
Notes Projects Database 3760		Tessella

Revision History

Issue	Author	Date	Description
V1.R1.M0	Geraint Lloyd	11-Mar-2005	1st version written for release v1.0.0
V1.R2.M0	Robert Nelson	27-Jul-2005	Version for release v1.0.1
V1.R3.M0	Robert Nelson	03-Oct-2005	Version for release v1.1.2
V1.R4.M0	Robert Nelson	11-Oct-2005	Version for release v1.1.3

References

Ref.	Document	Date	Details and Version
[Goodman]	"Some practical improvements in the continual reassessment method for phase I studies", S.N.Goodman, M.L.Zahurak, S.Piantadosi	Statistics in Medicine 14:1149-1161, 1995	Paper describing a practical application of the CRM to phase I studies
[Braun]	"The bivariate continual reassessment method: extending the CRM to phase I trials of two competing outcomes", T.M.Braun	Controlled Clinical Trials 23:240-256, 2002	Original paper describing the bivariate continual reassessment method

Table of Contents

1 Introduction

1.1 Purpose of this document

This document describes how to use the bCRM application graphical user interface. It is intended for all end users of the system. Users are encouraged to read the 1995 [Goodman] paper and the original bCRM paper [Braun] before using the bCRM application.

1.2 Scope of this document

This guide only describes using the bCRM graphical interface and the bCRM facilities accessible via that interface for running a trial; it does not cover all the statistical details of the model and method that it uses for dose allocation and simulation.

1.3 Context of this Issue

This issue documents version 1.1.3 of the system.

1.4 Definition of Terms

bCRM	refers to the computer programs that implement the bivariate continual reassessment method dose-finding algorithm, simulate and run phase 1 clinical trials using the bivariate continual reassessment method design and provide charts and graphs of the results.
A Simulation Series	is comprised of a 'Design' for a trial using the bCRM algorithm, a set of one or more 'Scenarios' – possible probabilities of toxicities and efficacies – and one or more 'Variants' where some of the Design parameters can be varied to explore the effect of the variation on the operating characteristics of a Trial. A Simulation Series creates a number of Runs = (number of Scenarios) * (number of Variants).
A Design	is a set of parameters that define a trial design to be carried out using the bCRM algorithm. This consists of: the specific drug dose strengths that can be used during the trial; the 'alpha' value of the modelling logistic curve along with a set of estimated toxicity and efficacy probabilities which are used to define the set of curves to model the response; the start dose and maximum dose level increment allowed during the trial; and the trial size details – the number of subjects assigned in each cohort used in the trial, the maximum number of subjects and the trial termination conditions.
A Scenario	is a set of fixed values of the probability of toxicity and efficacy for each of the dose strengths to be used in the trial. These define the 'true' drug response during a simulation.
A Variant	is a small modification to a Design – for example a different cohort size, study sample size or value of alpha.
A Run	A run is a set of all the parameters required to simulate a trial, it is the combination of a Design and a Scenario. A Simulation Series creates one or more runs. Each run is organised as a separate folder, all the parameters for a run are held in a single file called 'init.bcrm' and all the outputs for the simulations of the run are held in the folder in '.csv' files. (Plain text files with one record per line, individual values separated by commas – this format of file is readily imported into many other programs such as Excel, Access, SAS and Oracle). bCRM can have at most one run directory open at any one time and it is the files in this directory that provide the source data for the parameter values and visualisation.

Simulation	A simulation is the result of probabilistically generating a single clinical trial using bCRM. bCRM creates data for subjects and their responses by randomly selecting from the relevant probability distributions of the scenario defined in the bCRM input file. A particular run may be simulated a number of times, indeed to accurately analyse the characteristics of a particular Design it must be simulated usually ~10,000 times over a range of scenarios.
Trial	A trial is a real clinical trial, where the user enters subjects' data and responses. The bCRM Algorithm is run to determine the doses to allocate and to analyse the results. This facility allows bCRM's adaptive allocation to be used on small (single centre) trials.
MTD	The maximum tolerable dose. This is the strongest of the doses in the trial that has a toxicity below the target toxicity.
MED	The minimum efficacious dose. This is the weakest of the doses in the trial that has an efficacy above the target efficacy.

2 Quick Guide

2.1 Starting bCRM

To start bCRM, double-click the bCRM icon that is installed on your computer desktop.

After bCRM has started, the main window will be displayed.

2.2 Main Window

The main window is displayed after bCRM has started and is present until bCRM exits. It is used to open and create designs and to access other parts of the application.

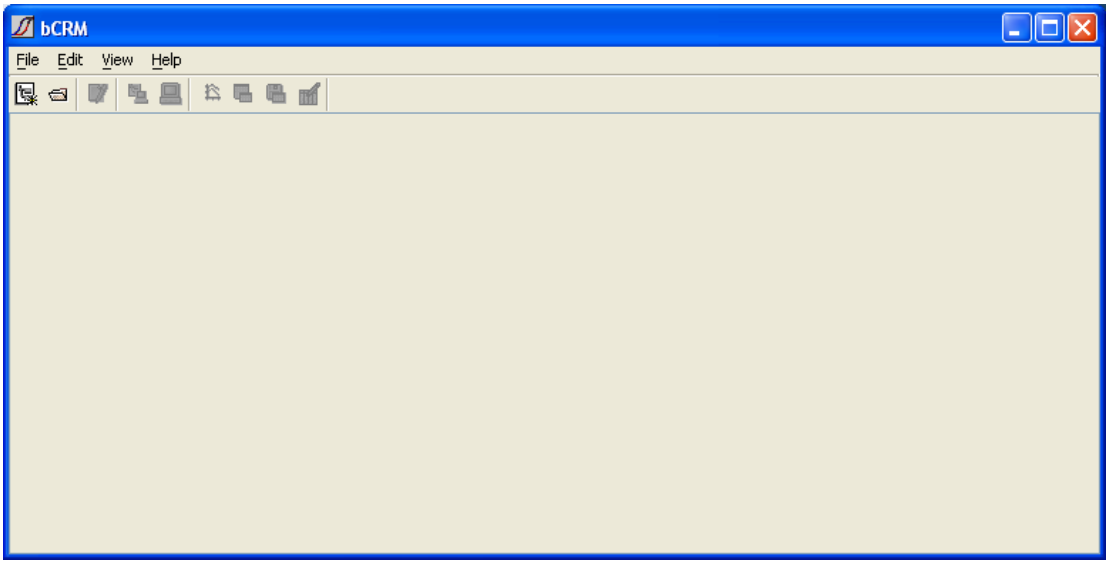

Initially the only enabled buttons are:

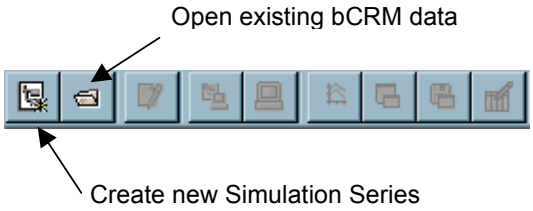

These commands are also available from the **File** menu.

2.3 Exiting bCRM

To exit bCRM, select **Exit** from the **File** menu of the main window.

2.4 Displaying Existing Data

To display existing bCRM Data, select **Open** from the **File** menu, click the tool button or press the `ctrl-o` key. If there are currently any open parameter wizard, run manager, Desktop bCRM or visualisation windows, you will be notified that they will be closed before changing the current simulation series and asked if you wish to proceed.

Use the file browser to select the directory to open and then click the **Open** button. If you change your mind about opening a particular directory then click the **Cancel** button. As well as folders, the file browser displays the bCRM input ('init.bcrm') and output files ('*.csv') in each directory to help you locate the correct one.

If the bCRM input files contain any invalid parameter values, then you will be notified via a message box. These values must be corrected in bCRM before bCRM can be run with this design.

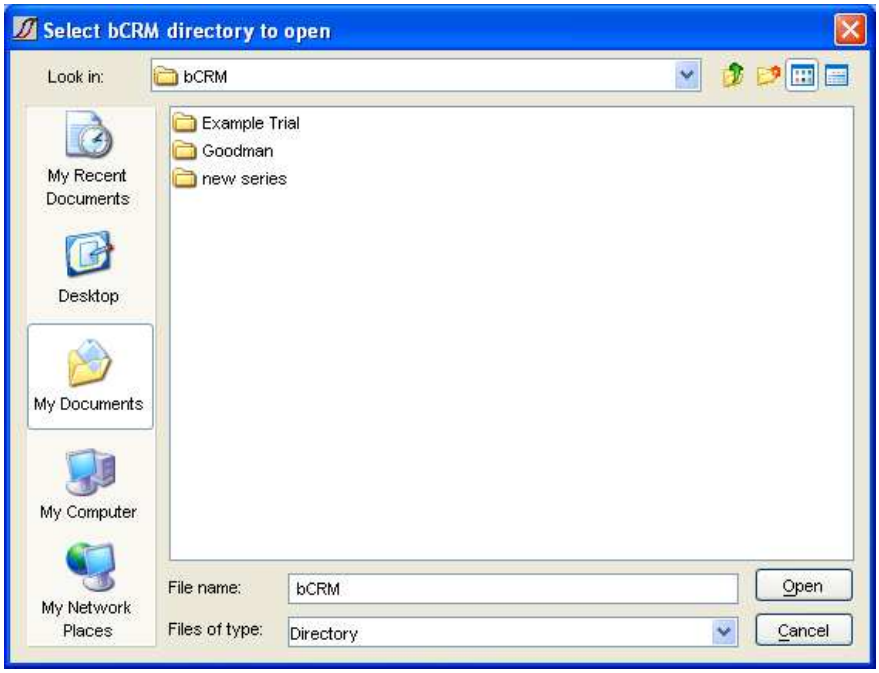

In the folder selected for opening bCRM will display all its folders that contain bCRM data:

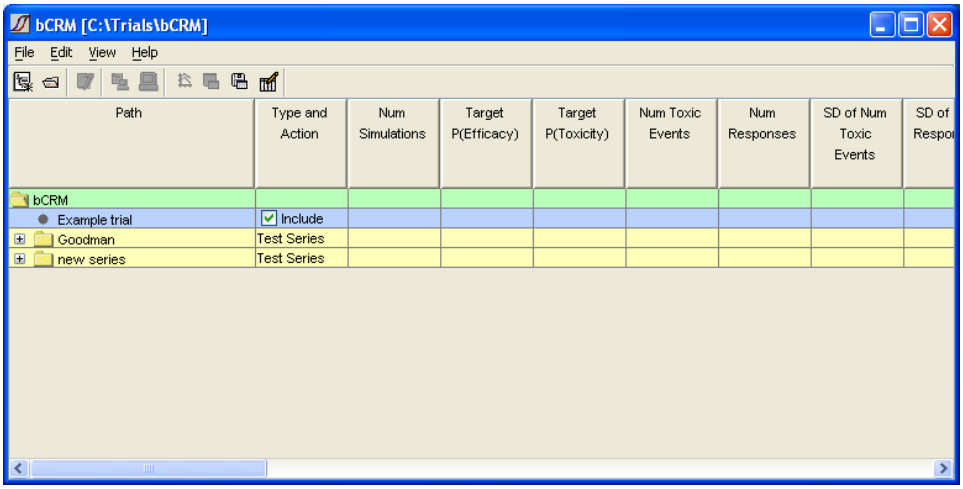

2.5 The Toolbar

Once a simulation series has been created, or a pre-existing one has been opened, the application's main panel becomes more interesting:

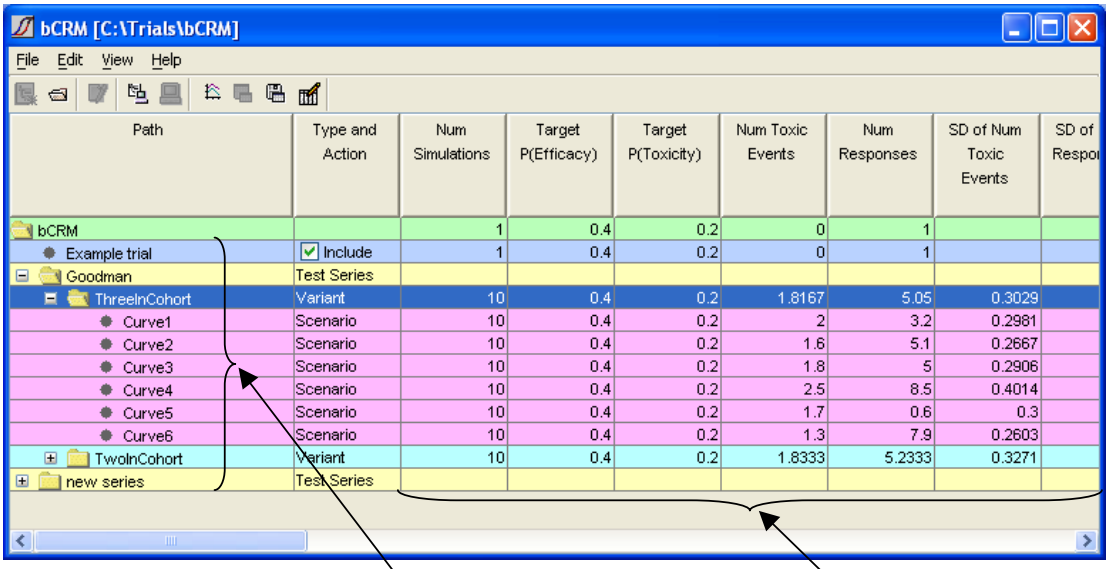

Directory containing the set of Simulation Series & Trials

Simulation Series and Trials

Summary data from simulations or actual trial

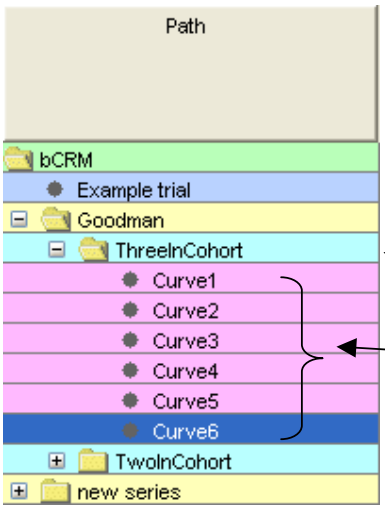

A Trial being run by bCRM

A Simulation Series

The design variants of the simulation series

The scenarios of design variant "ThreeInCohort", all variants are simulated over the same scenarios.

The currently selected series

Levels can be expanded or hidden by clicking on the '+' or '-' boxes.

Select a simulation series or trial by clicking on its name – the line becomes highlighted. More of the bCRM Toolbar buttons become enabled:

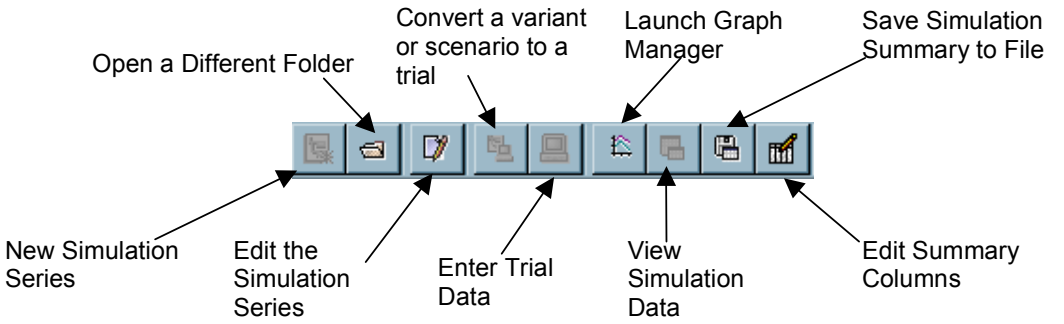

All of these operations are also available from the application's menus.

Right clicking on a Trial or Simulation Series selects it and opens a pop-up menu:

This gives short cuts to commands described elsewhere – it also allows the simulation series (the simulation series directory) to be renamed or deleted.

3 Design and Simulation of bCRM Trials

The main mechanism for designing and simulating trials in bCRM is the 'Simulation Series'.

A Simulation Series consists of:

1. a 'Design' (which provides the 'Default Parameters' for all of the different simulations),
2. a set of 'Scenarios' – probability curves describing the probability of toxicity and efficacy for all the drug dose strengths which may be used during a trial
3. and a set of 'Variants' – variations on the underlying design in order to evaluate options in the trial design.

To create a new simulation series, select either **New Simulation Series** from the **File** menu, or click the New Simulation Series button ().

A file browser will appear in which you should select the directory in which to save the new series:

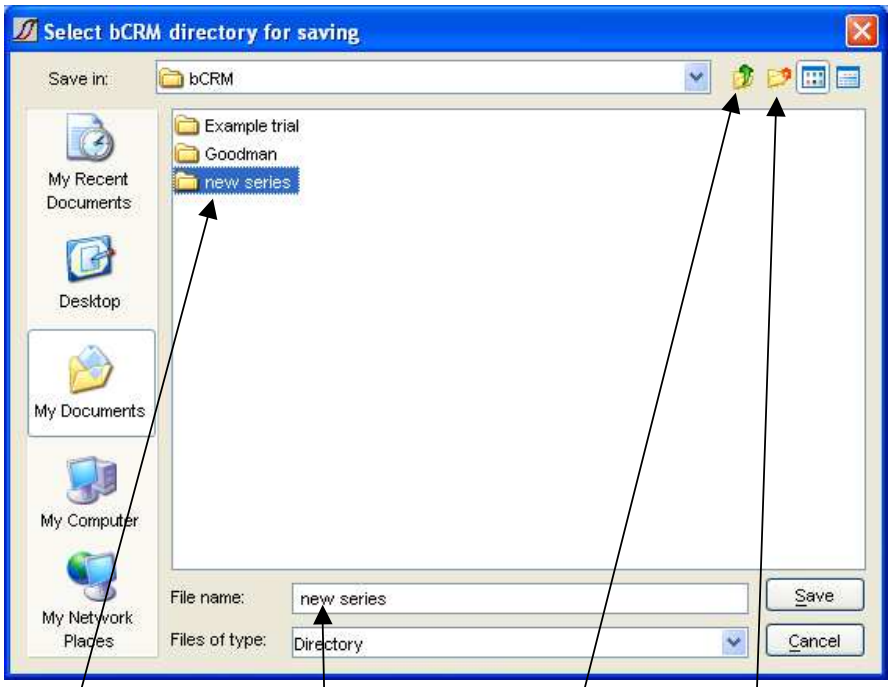

You should navigate to where you wish the new simulation series to be stored:

- double click on folders to move into a folder
- click on the 'parent folder' button to move up to the parent folder
- once in the folder that you want to act as the parent folder for the simulation series, enter the name of the new simulation series – this will be the folder for holding the simulation series data.
- click on the 'new folder' button if you wish to create a new parent folder – this will be useful for organising a number of related bCRM simulation series together:

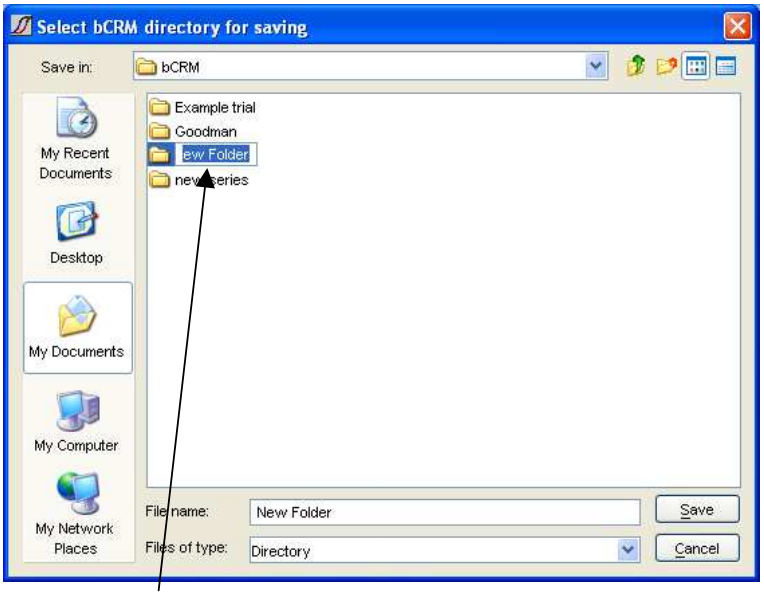

- if you select 'New Folder' it will be called 'New Folder' but the name will be selected and you can type a more meaningful one.
- you can then double click on that folder and enter the name of a new directory for the simulation series.

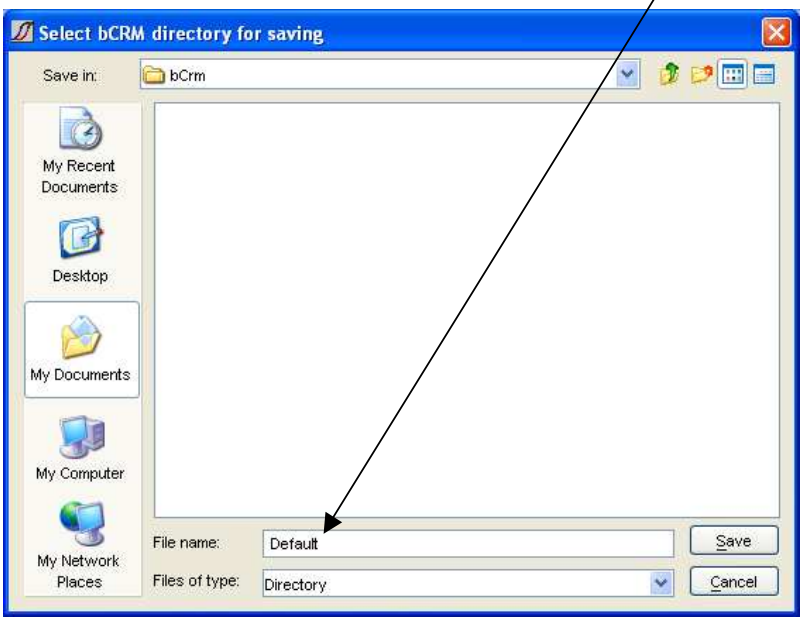

Click **Save** to save the new simulation series. The simulation series editor window will then open.

3.1 The Simulation Series Editor

- This panel allows you to select which aspect of the Simulation Series to define or edit. The step may be selected by clicking on it, or by using the **Previous** and **Next** buttons.
    - Series Definitions – enter the design of the trial.
    - Scenario Definitions – define the different toxicity and efficacy probability curves to simulate from in the simulations.
    - Variants – define variations on the main design of the trial for comparison.
    - Simulation – displays a screen for running the simulations.

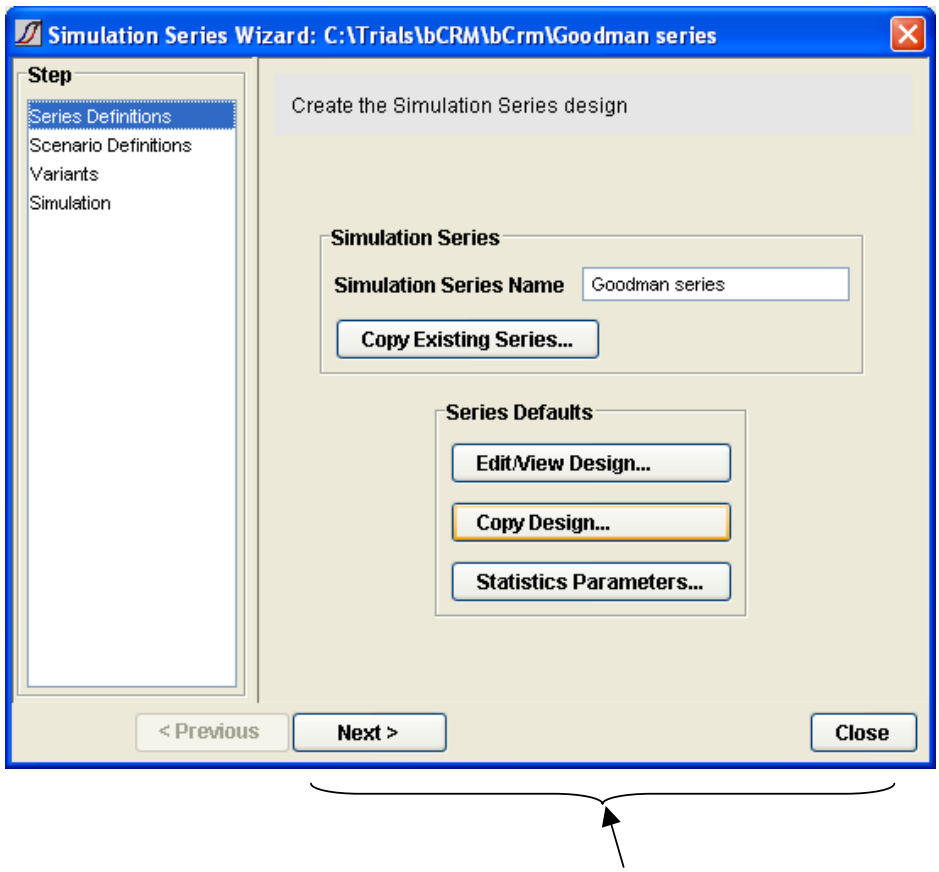

With **Series Definition** selected the rest of the screen looks like this. It allows the series name to be changed, the series design, scenarios and variants to be imported from an existing series (**Copy Existing Series**) or the Design to be setup from scratch (**Edit/View Design**) or imported from a particular run in another simulation series (**Copy Design**).

To design a simulation series from scratch, or modify a design that has been copied in, click on the **Edit/View Design** button.

3.2 Series Definitions - Editing The Design

We will now illustrate setting up a Trial design. This uses a further sequence of screens where the sequence of steps is listed in the left-hand panel and the controls of the current step are shown on the right.

3.2.1 Defining the Drugs and Target Toxicity and Efficacy

Enter the details of the drug that is being trialled:

- The name of the drug
- The units of the dose

Select whether the trial will monitor the toxicity of the drug, the efficacy of the drug, or both. (This document assumes both are being used. **If only one is used, then inappropriate items will be hidden or greyed out.**

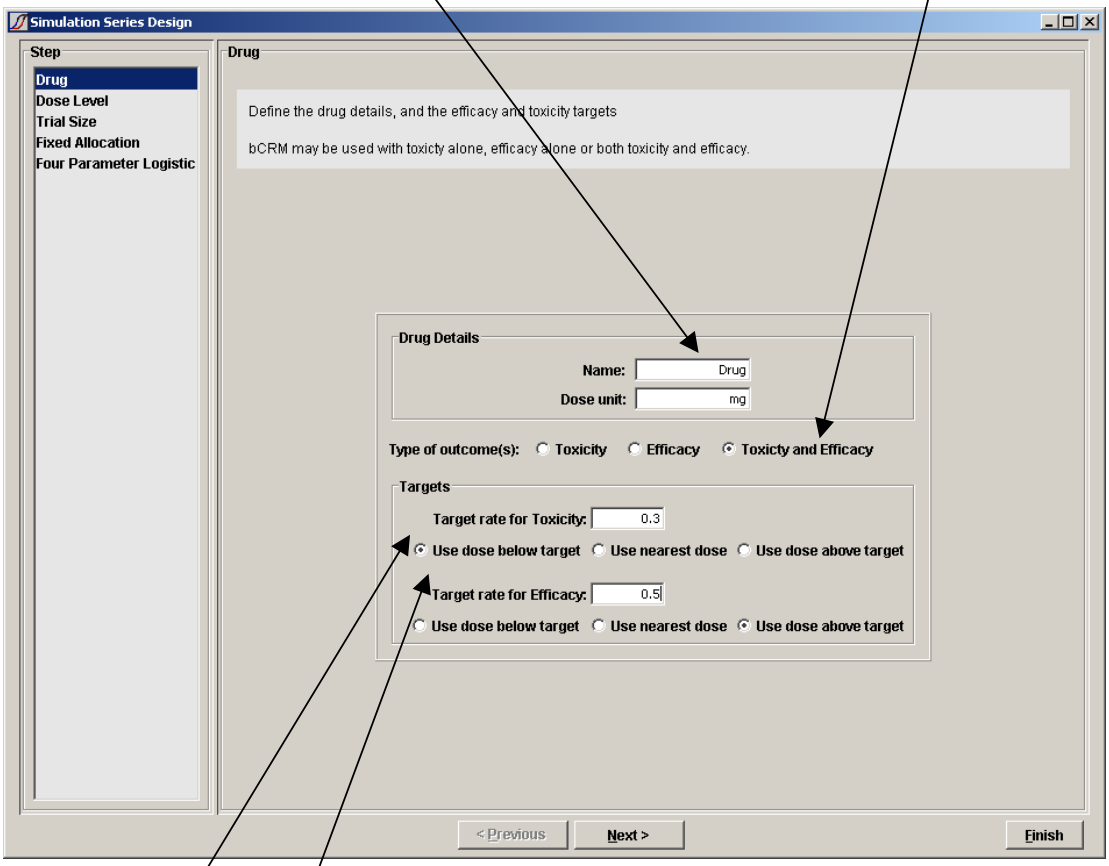

Enter the target toxicity and efficacy probabilities:

- The drug Maximum Tolerable Dose (MTD) will be the largest dose that has an (estimated) probability of toxicity that is below the target toxicity.
- The drug Minimum Efficacious Dose (MED) will be the smallest dose that has an (estimated) probability of efficacy that is above the target toxicity.
- The bCRM model is poor at determining the correct dose level for probabilities close to the value of 'x' = 0. This normally occurs at 5% or 95%. For this reason it is not recommended to run a trial with a target probability below 0.15 or above 0.85.

When the target level is not at an exact dose, the system can select either the closest dose, or the dose above or the dose below. The default is to use the dose below for toxicity and dose above for efficacy.

3.2.2 Entering the Dose Levels

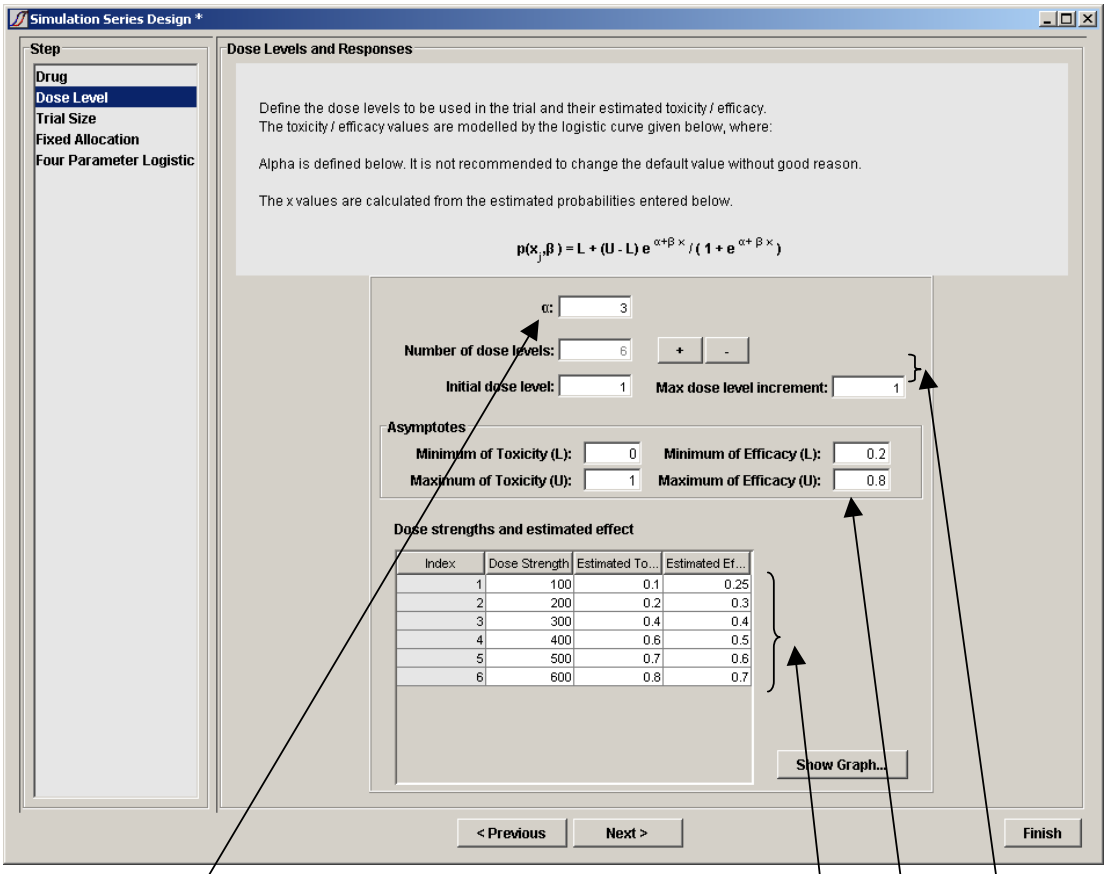

This screen is used to provide information on the drug doses that will be used in the trial, and on the set of logistic curves which bCRM uses to model the true drug responses. (If the series uses only toxicity or only efficacy then only those controls will be available)

- Enter the value of α used to define the set of logistic curves for modelling. α is usually either –3 or +3. The model generally works better with an α of +3 when the target probabilities are <0.5, and works better with an α of –3 when the target probabilities are >0.5
- During the trial only specific dose strengths will be given to subjects. The trial will not use any other dose strengths. You can specify how many different dose strengths will be used.
- Specify the index of the dose strength that will be given to the first cohort of subjects in the trial (1 to ensure the first cohort is given the weakest dose). Then enter how many increments to the dose strength are allowed between successive cohorts (enter 1 to ensure that the adaptive allocator will at most only use the next strongest dose above whatever has already been used).
- Specify the range of each curve. By default the efficacy and toxicity curves go from 0 to 1, but this may not be correct. For example if there is a know response rate on placebo, this can be entered here. Setting the correct asymptotes enables the logistic to fit more accurately and improves performance, but remember that the target value must be well clear of the asymptotes.
- For each dose to be used in the trial, specify the strength of the dose in the units given on the previous screen, and an estimated toxicity and efficacy for that dose level. Each of these values must increase with increasing dose index.

- The estimated toxicity and efficacy values do not need to be the same. Similarly the dose strengths do not need to be evenly spaced, they are not used by the model, but are included for reference.
- The toxicity and efficacy estimates determine the 'x' values that are used in the logistic curves that model the true drug responses. Each estimated probability is used to determine the 'x' value which gives that probability for a value of $\beta = 1$. The 'x' values from the estimated toxicity and efficacy probabilities are then averaged to give an average 'x' that is used in all future modelling

You can check the effect of the parameters entered and the range of logistic curves that can be used to model the true responses, by pressing **Show Graph…**

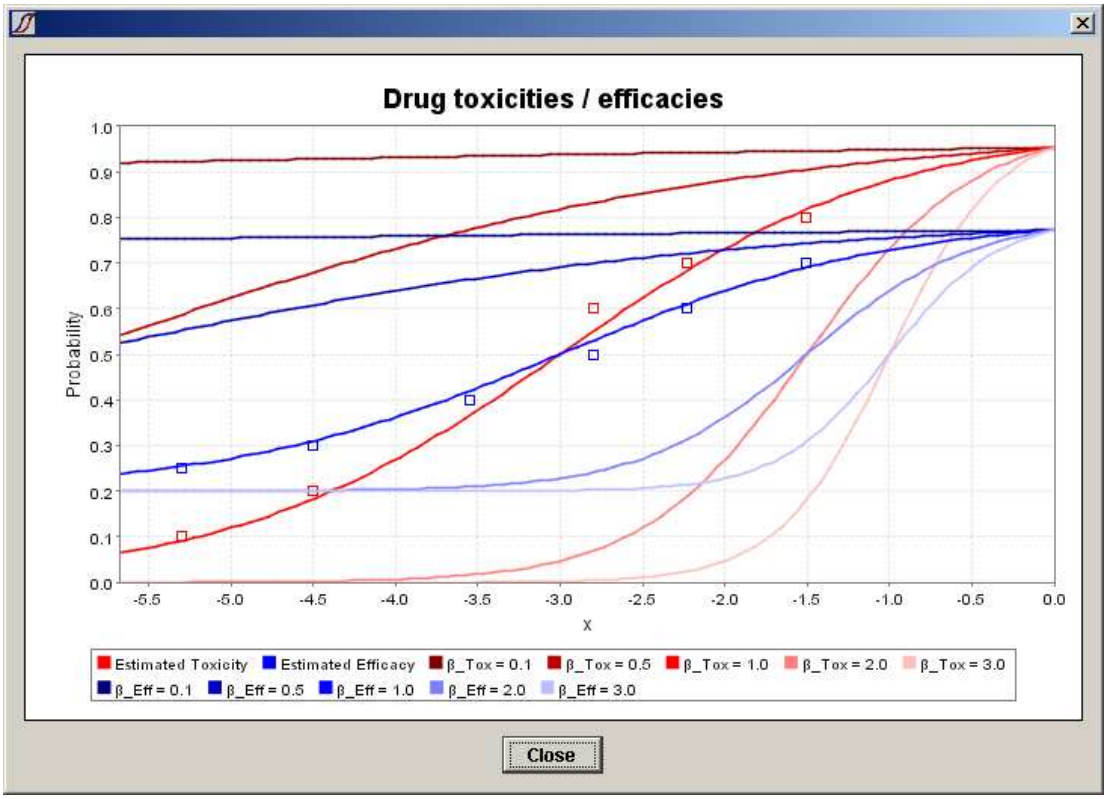

The estimated toxicity and efficacy probabilities will be plotted against the x values determined from them.

For each outcome a set of logistic curves will also be plotted for a range of values of β ranging from 0.1 to 3.0. The prior toxicity and response behaviour assumed by the model is for a uniform distribution of β ranging from 0 to 3. Thus you can see that the information entered on this page has defined a family of curves, distinguished solely by different values of β. The observed data will be modelled by selecting the best fit from this family. There will be some restrictions on how well the model fits the data as a result of this. In the example above:

- The curves can fit a toxicity of between 0.05 and 0.9 at all doses.
- But sharp changes in toxicity are only well modelled at the higher doses. For instance if probability of toxicity of 0.5 is observed at the second dose, the curve that fits that point exactly, implies a probability of toxicity of ~0.4 at the first dose.
- The curve that fits a low probability of toxicity at the higher doses, implies very low probability of toxicity at the lower doses.
- The curves will not fit very high toxicities at the top two (or more) doses without implying high toxicity at lower doses.

If we change the value of the Alpha parameter from 3 to –3 the family of curves changes to:

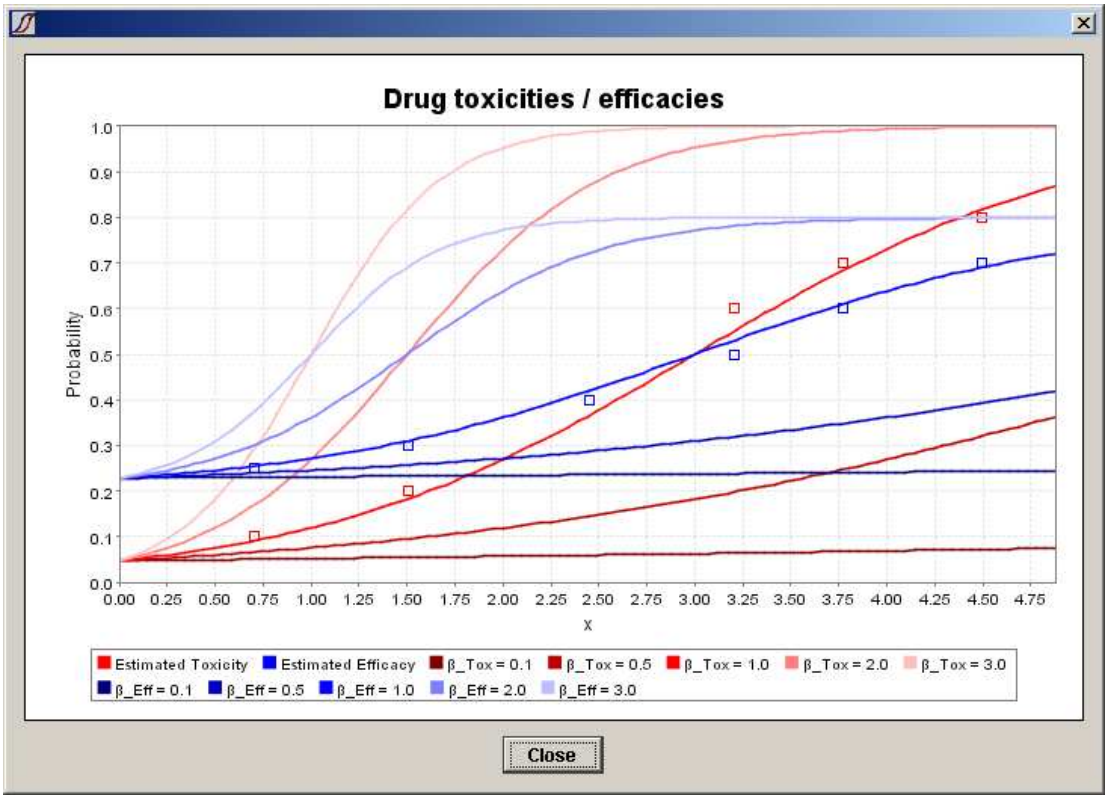

Notice this family of curves has very different strengths and weaknesses from the previous ones:

- There is no curve that fits the lowest dose having a high toxicity.
- The curves easily fit large increases in toxicity at the lower doses.
- The curves are very flexible in the levels of toxicity at the highest doses.

In each case efficacy behaves in a similar way to toxicity but, because we have changed the asymptotes, it is confined to a narrower range of probabilities.

The fit of the family of curves can also be 'tweaked' by changing the expected levels of toxicity for the doses:

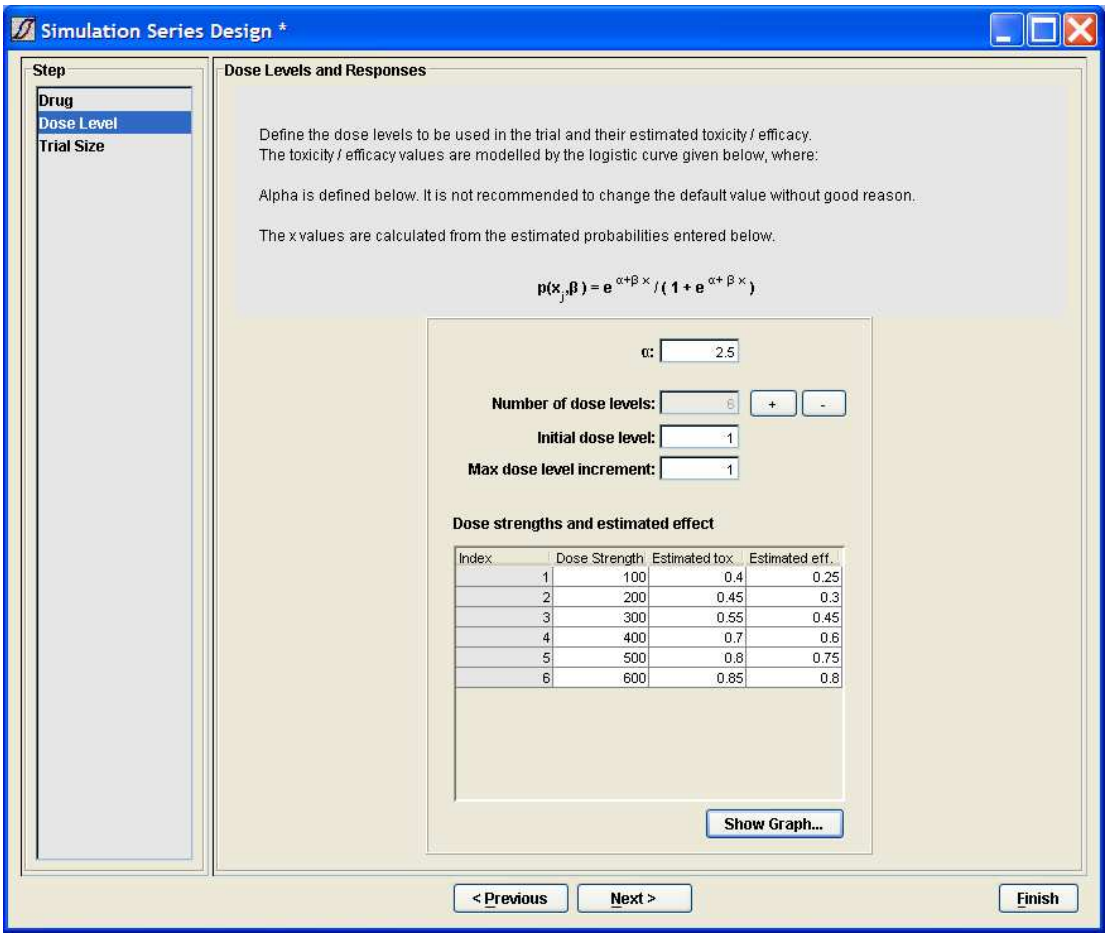

Here with an alpha of +3 we have increased the expected levels of toxicity, this bunches the doses around the more variable section of the curve family, but at the cost of having a prior expectation of toxicity that is greater than the maximum tolerated level.

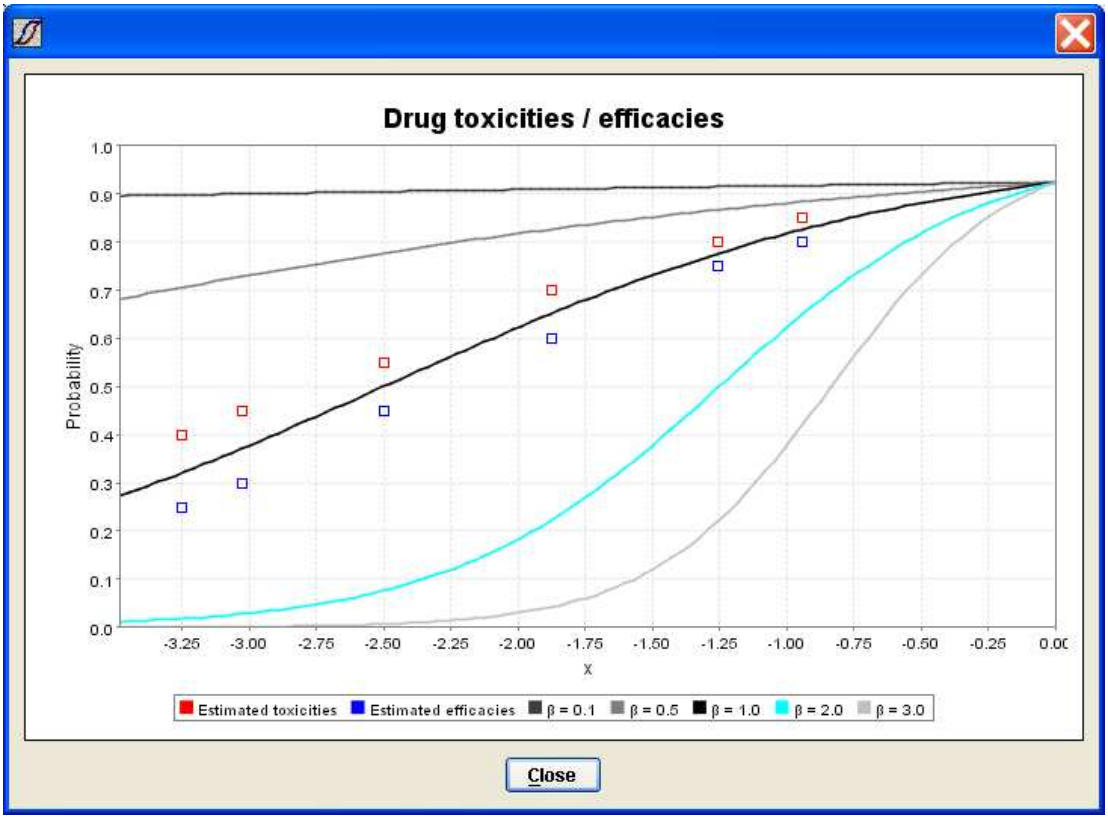

It is important to bare two things in mind:

1. There will be very little allocation to doses where the observed toxicity is high (relative to the acceptable level), this should mean that the characteristics of the family of curves at the higher toxicities is relatively unimportant.
2. The CRM is justified on its ability to determine the MTD whilst not exposing subjects to excessive toxicity. The important feature of the curves is where the MTD is placed for any given set of data – not the fit of the curve as a whole.

3.2.3 Specifying the trial size and termination conditions

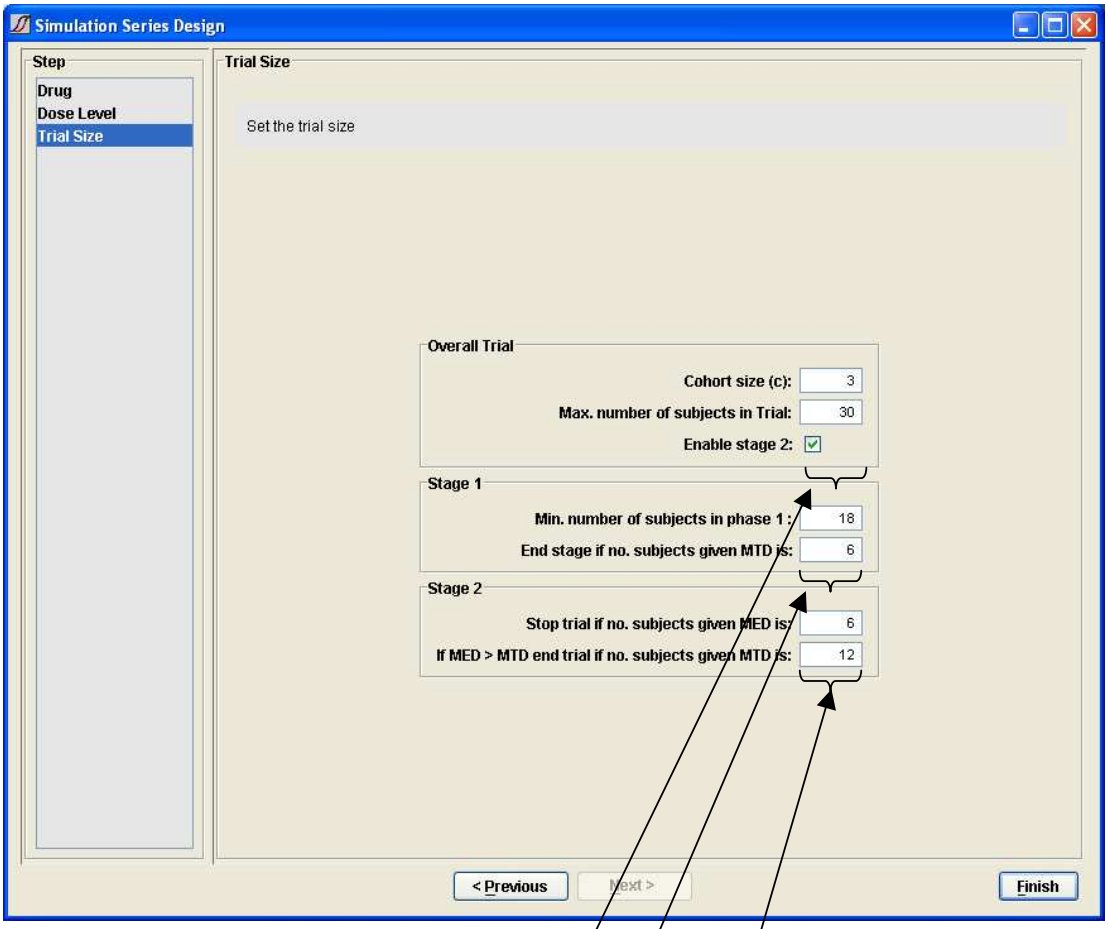

This final screen enables the setting of parameters to determine the size of the trial and its early termination conditions.

For the trial as a whole the following can be set:

- The number of subjects in each cohort
- The maximum number of subjects allowed in the trial. This must be a multiple of the cohort size
- Whether or not the trial should have a second stage (This is only possible if both toxicity and efficacy are being modelled). In the first stage, the allocator seeks to assign subjects to the MTD dose level to determine the MTD value most accurately. In the second stage the allocator endeavours to assign subjects to the MED value.

As well as a maximum number of subjects in a trial, you can specify the minimum number of subjects that should be allocated during the first stage. During this stage, the allocator will try to allocate subjects to the MTD dose level, or if it believes that all doses are toxic, will allocate to the lowest dose level. Once the minimum number of subjects has been allocated, the first stage can then terminate when either:

- the maximum number of subjects has been allocated, or
- a specified number of subjects have been allocated to the dose that the model currently believes is the MTD.

If the second stage of the trial has been enabled, then it terminates when either:

- The maximum number of subjects for the trial has been allocated.
- A specified number of subjects have been allocated to the dose that the model currently believes is the MED.
- If the model believes that the MED is a higher dose than the MTD, it will not assign subjects to the MED. Instead it will continue to allocate to the MTD, and the trial will terminate early if a specified number of subjects have been allocated to the current MTD. If you do not wish any additional subjects to be allocated to the MTD in these circumstances, set this value to the same number for ending stage 1 of the trial.

3.2.4 Specifying Extra Allocation Probabilities and Futility Rules

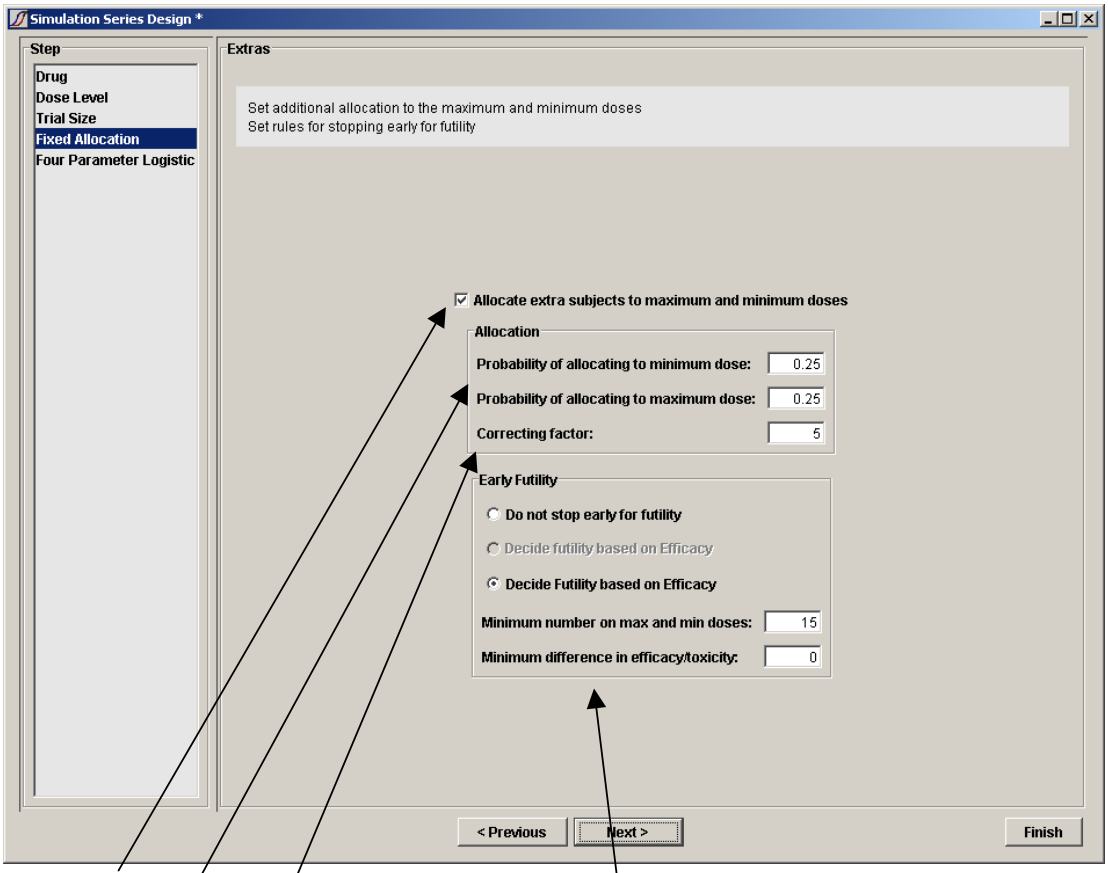

Decide whether to allocate additional subjects to the minimum and maximum doses. If the trial is being run to find the MTD, leave the checkbox unchecked to disable this feature.

If the trial is being run to measure efficacy, you may wish to allocate a certain fraction of subjects to the minimum and maximum doses.

The "correcting factor" controls a mechanism, which can adjust these probabilities during the trial so as to keep the numbers allocated closer to the target values. Probabilities are adjusted according to the formula

$$p = p_T^{1+\gamma(p_o - p_T)}$$

where: p is the allocation probability used

p_T is the target fraction entered by the user

p_o is the fraction already allocated to this dose

γ is the correcting factor

A value of zero disables this mechanism, while positive values give a greater degree of correction. The default is 2. To ensure some randomisation always occurs, p is constrained to lie between 0.1 and 0.5.

If extra subjects are being allocated to the extreme doses, you have the option of terminating early due to futility. Specify the minimum number of subjects which must have been allocated to **both** the minimum and the maximum doses and how much more effect the maximum dose must have compared to the minimum. If the minimum difference is zero, then once the minimum number of subjects have been reached, the trial will terminate unless it is 90% confident that the maximum dose has a higher effect than the minimum.

Stopping for futility is independent of the rules on the trial size page and the trial will stop when either set of rules tells it to.

3.2.5 Four Parameter Logistic

The final screen allows you to fit a four parameter logistic to the results. This is a summary calculated at the end and dose not affect dose allocation or decisions about when to stop. By fitting four parameters rather than one, this curve will usually fit the results more accurately but there may not be enough information to determine the parameters uniquely.

The four parameter logistic is only fitted to one outcome. If there are both toxicity and efficacy, it will be fitted to toxicity.

Specify the confidence interval. This controls the interval reported for the fitted curve and for the parameters. The default is 0.95.

Clicking on **Finish** saves the parameters.

3.3 Scenario Definitions – The Probability Curves to be simulated

The second step in the overall definition of the Simulation Series is to define some example toxicity and efficacy probability curves to simulate from to assess the design's performance.

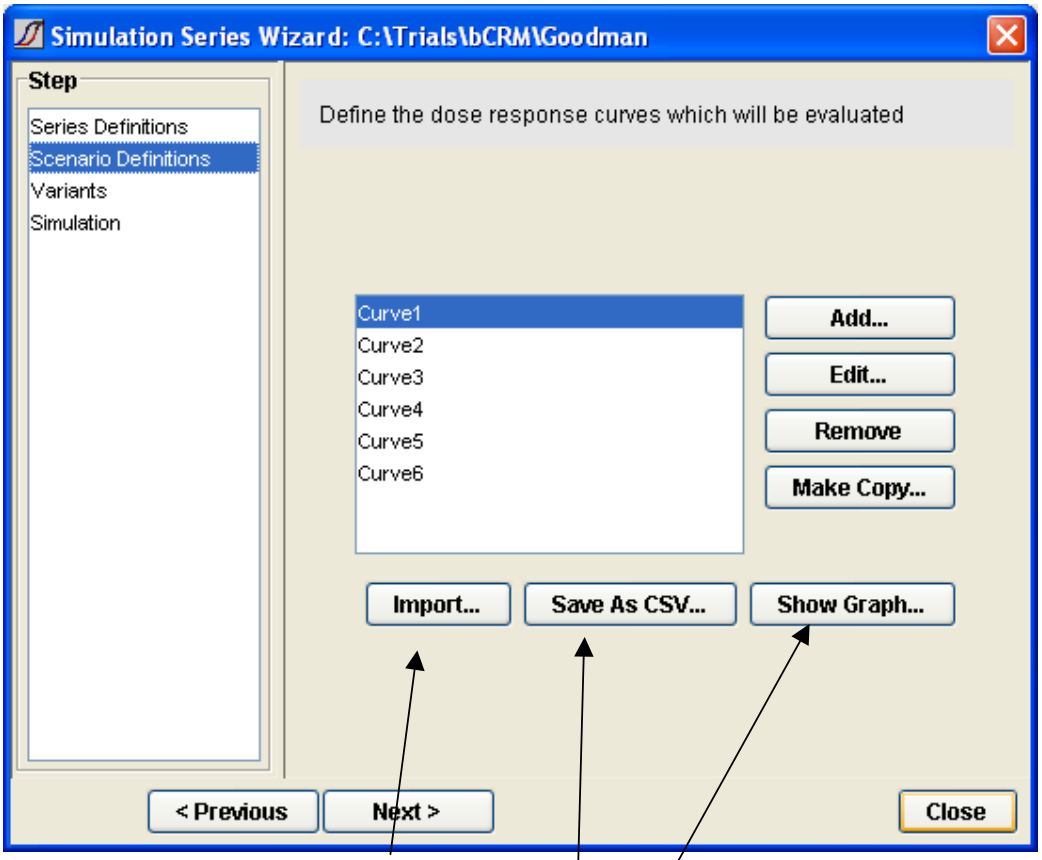

The main buttons allow curves to be imported from other simulation series, and for the scenario data to be saved as CSV files (e.g. for importation into Excel, S-Plus or other graphing tool). bCRM itself can also display the probability curves for a selected scenario.

The second set of buttons are self explanatory – they allow a new scenario to be added to the simulation series, or the scenarios already defined to be edited, removed or copied. Before using these last three, ensure the scenario you want the operation to apply to is selected in the list of scenarios.

Defining Toxicity and Efficacy Probability Curves

Pressing **Add** or **Edit** displays a screen for defining the toxicity and efficacy probability curves.

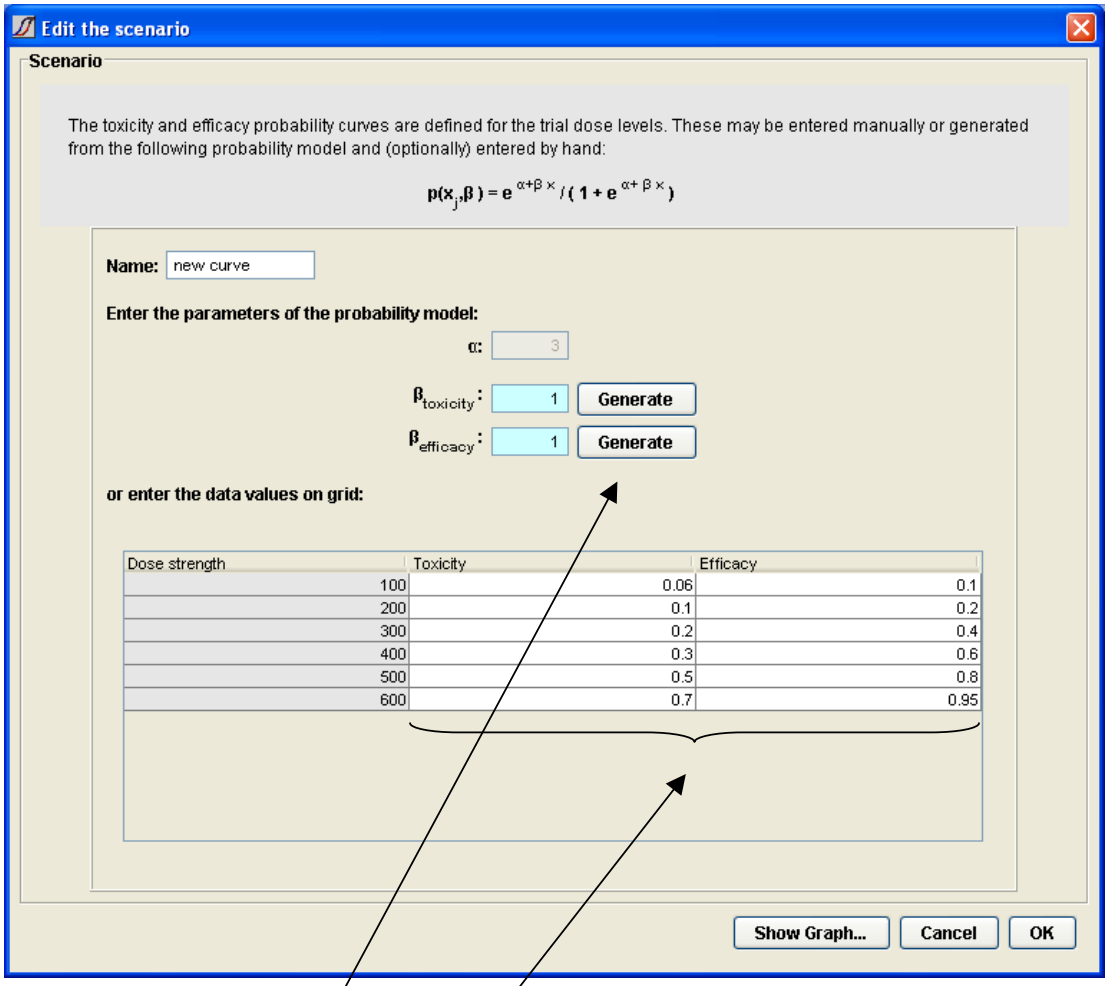

There are two methods for specifying the toxicity and efficacy probabilities to be simulated.

1. The user may enter or modify by hand any of the existing probability values displayed in the table
2. The user may automatically generate the probability values corresponding to a logistic curve for any specific value of β. Clicking on **Generate** causes the probability values to be entered in the table.

The two methods may be combined. Initial probability curves may be generated automatically, and then individual values in the table modified by hand.

(If the series uses only toxicity or only efficacy then only those controls will be available)

Pressing **Show Graph…** will display the probability values entered against the dose strengths:

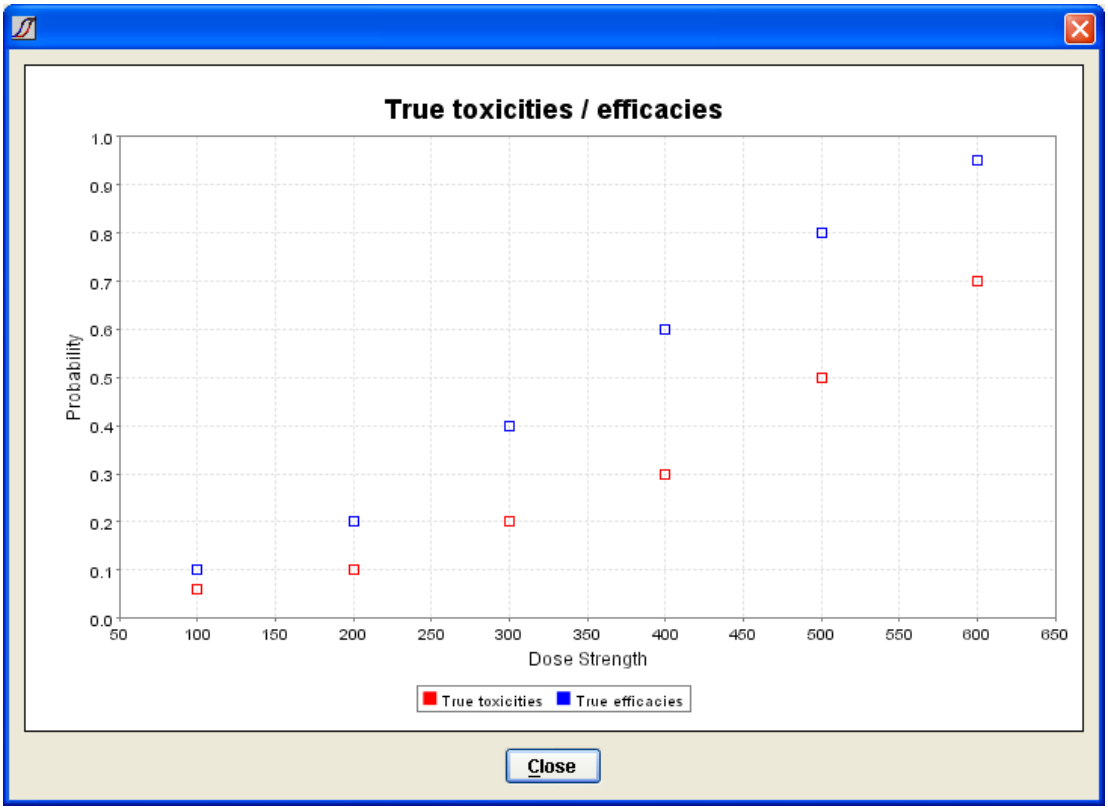

3.4 Defining Variants

Variants are variations on the default parameters set in the 'Series Definitions' or trial design. In this illustration two versions of the trial design with cohort sizes of 3 and 2 are specified. At least one variant is required – but this can simply be a default with no changes.

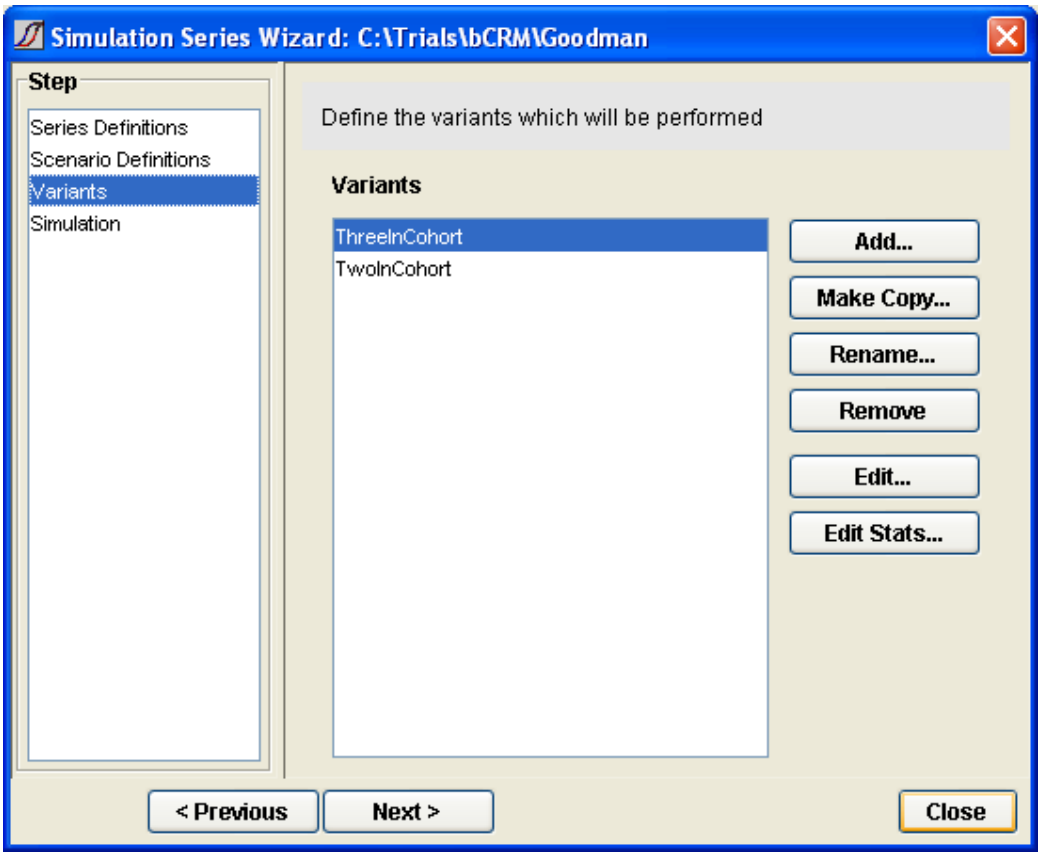

3.4.1 Editing Variants

Editing the variants presents the same screens as used in defining the series, all pre-populated with the design parameter values entered in the series definition. The type of outcome cannot be changed between variants. Otherwise there is no restriction on which parameters, or how many parameters can be changed. It is left to the user to judge what is sensible and what is not.

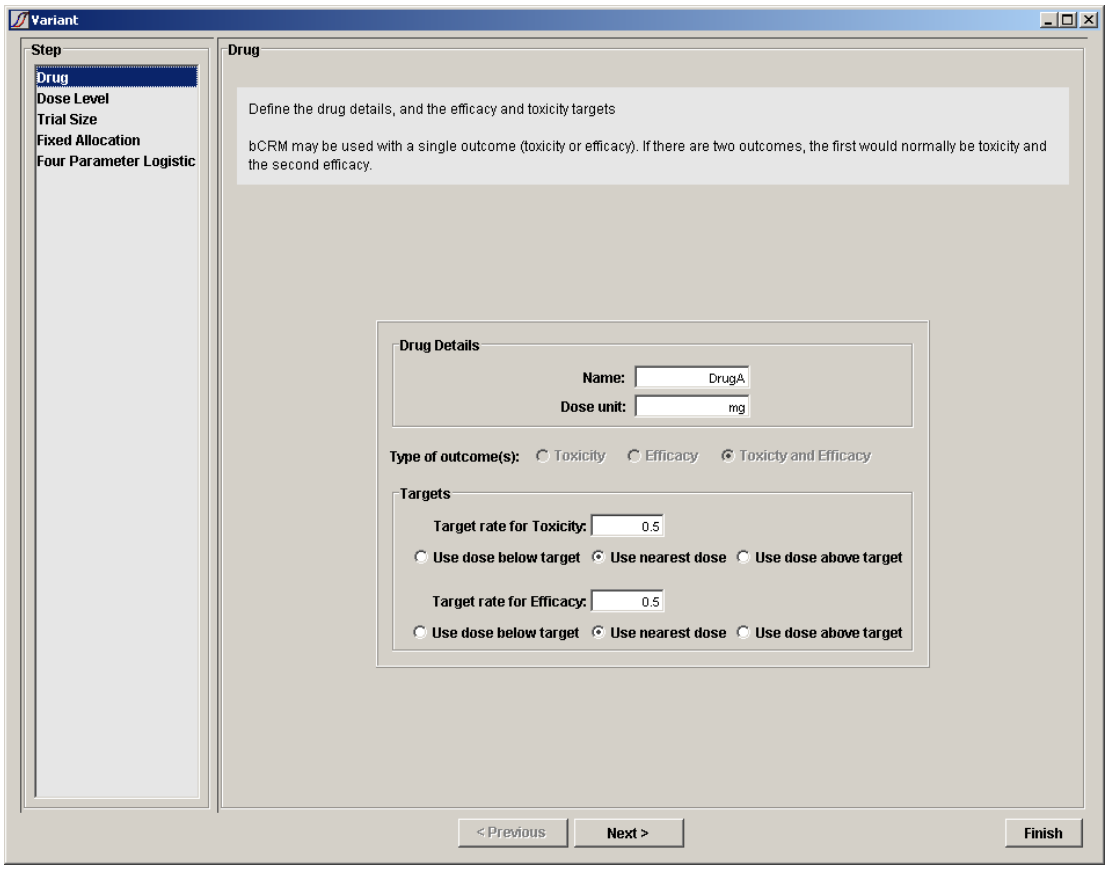

4 The Statistics Parameters

In both the series definition and the variant definition stage it is possible to change the 'statistics parameters'. These parameters are for statisticians who wish to investigate the effects of various choices in the underlying model.

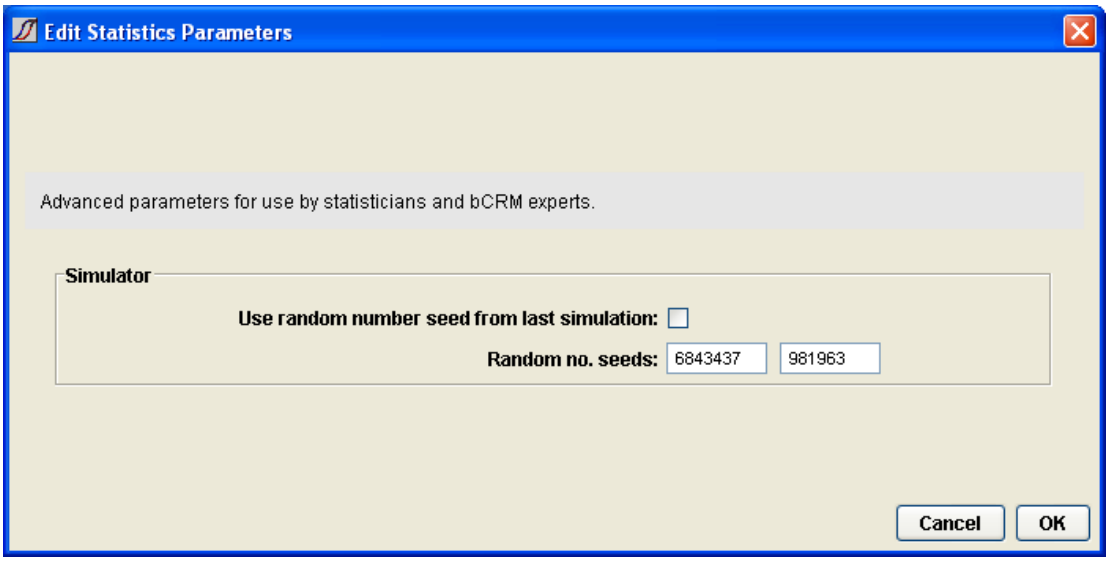

This allows the random number seed to be specified and whether or not to use the final seeds from the previous simulation. If this option is selected then all simulations are started with different random number seeds. If it is not selected then it is possible to reproduce the results from previous simulations.

5 Simulating Trials using bCRM

The implementation of the bCRM algorithm can be run on the simulation series from the last step in the Simulation Series editor. These all cause the Simulator screen to be displayed:

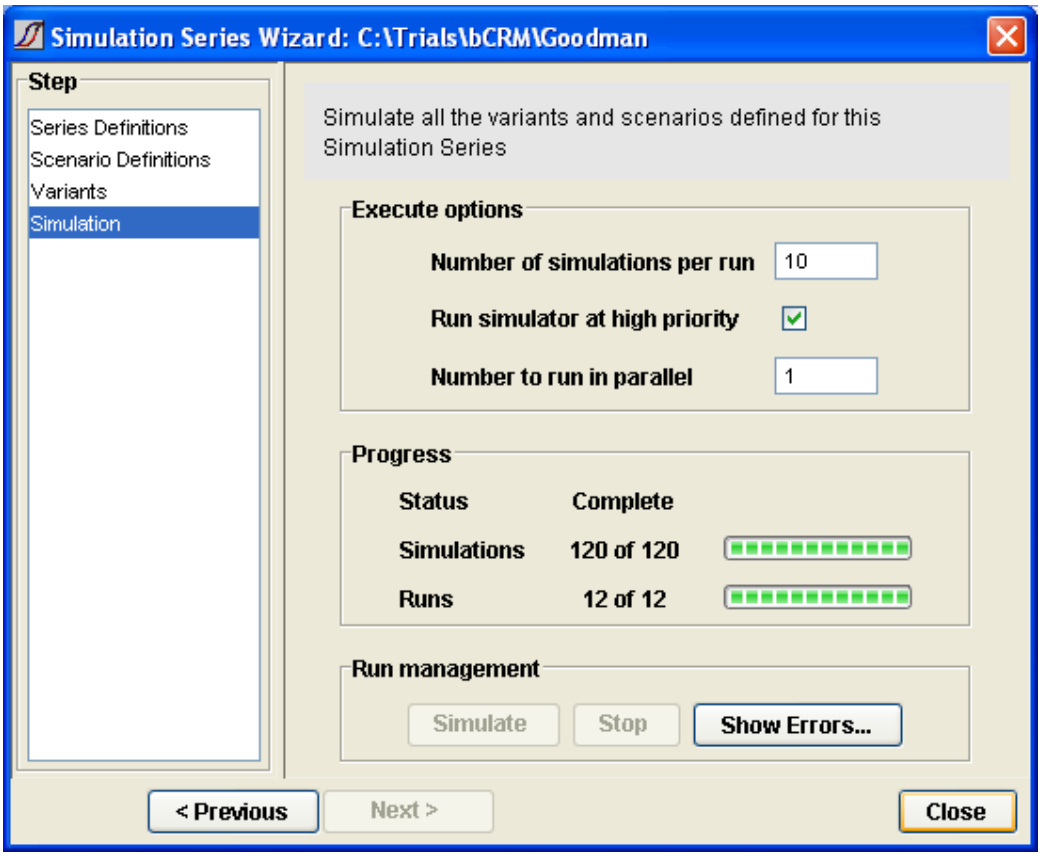

The 3 parameters to set control:

- the number of trials to simulate for each Run,
- whether to run them at high priority – this causes the simulations to run faster but makes it difficult to use your computer for anything else at the same time. If this option is left unchecked then the simulations are run in the background and your computer gives priority to whatever else the user is doing while they run.
- The 'number to run in parallel' should normally be left at 1, however on new PC's with CPU's with 'Hyper-Threaded' architectures, these CPUs can effectively run two programs simultaneously with no loss of performance to either one. On these PCs it is worth setting this parameter to '2'. To tell if your PC has this kind of CPU, right click on the Windows task bar (normally at the bottom of the screen). Select 'Task manager' from the pop-up menu and select the 'Performance' tab on the Windows Task Manager dialog. If you see two charts of 'CPU usage history' then you have a Hyper-threaded CPU, otherwise not.
- Once you have clicked 'Simulate' the progress bars will update as the simulations run.

6 Data

6.1 Summary Data

After running simulations the main window can show the summary data for the different scenarios and design variants simulated. The summary data falls into two groups – overall data and then data specific to the doses which the model has determined as the MTD and MED.

It is possible to edit the selection of columns displayed using the column editor by clicking on the 'column editor' button in the toolbar (), or selecting **Change Columns** from the **Run** menu. See below for a description of the column editor.

It is also possible to save the contents of all the columns to a 'comma separated values' (CSV) text file that can be easily imported into a number of other applications such as Excel, Access, SAS, S-Plus etc.. by clicking on the 'Save summary data to file' button on the toolbar (), or selecting **Export Table** from the **File** menu.

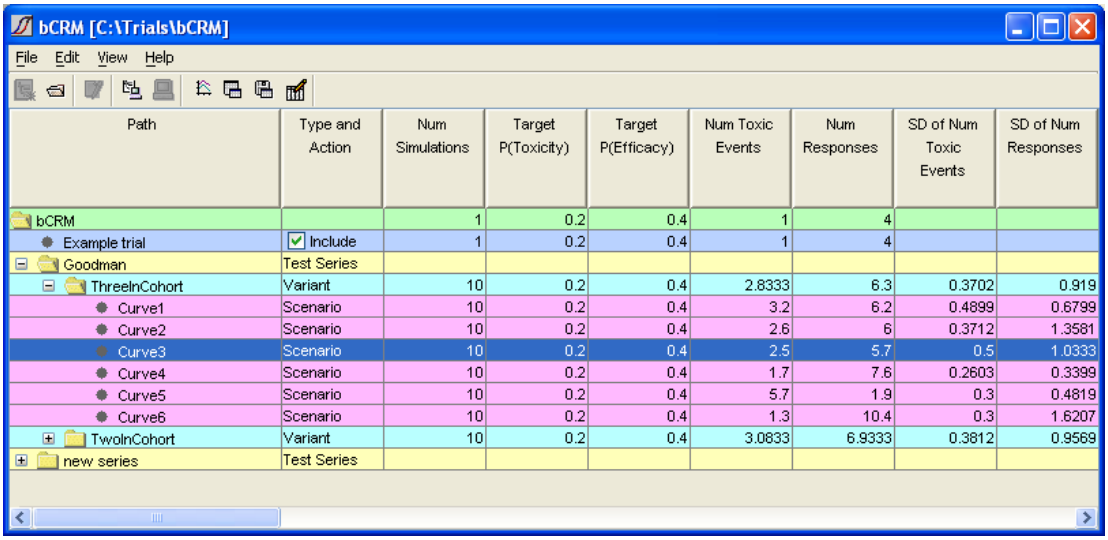

Path	Type and Action	Num Simulations	Target P(Toxicity)	Target P(Efficacy)	Num Toxic Events	Num Responses	SD of Num Toxic Events	SD of Num Responses
bCRM		1	0.2	0.4	1	4		
Example trial	Include	1	0.2	0.4	1	4		
Goodman	Test Series							
ThreeInCohort	Variant	10	0.2	0.4	2.8333	6.3	0.3702	0.919
Curve1	Scenario	10	0.2	0.4	3.2	6.2	0.4899	0.6799
Curve2	Scenario	10	0.2	0.4	2.6	6	0.3712	1.3581
Curve3	Scenario	10	0.2	0.4	2.5	5.7	0.5	1.0333
Curve4	Scenario	10	0.2	0.4	1.7	7.6	0.2603	0.3399
Curve5	Scenario	10	0.2	0.4	5.7	1.9	0.3	0.4819
Curve6	Scenario	10	0.2	0.4	1.3	10.4	0.3	1.6207
TwoInCohort	Variant	10	0.2	0.4	3.0833	6.9333	0.3812	0.9569
new series	Test Series							

These columns are:

Num Simulations	The number of trials simulated for each scenario, only 10 have been run in the example above – this is a wholly inadequate number to be able to draw reliable conclusions about the operating characteristics of the design.
Target P(Toxicity)	This is the maximum level of toxicity desired. The MTD will be the largest dose with a toxicity probability below this value.
Target P(Efficacy)	This is the minimum level of efficacy desired. The MED will be the smallest dose with an efficacy probability below this value.
Num Toxic Events	The mean number of subjects, per simulated trial, who were simulated to have a toxic event.
Num Responses	The mean number of subjects, per simulated trial, who were simulated to have a response (efficacious event).
SD of Num Toxic Events	The standard deviation of the mean of the result: "Num Toxic Events" over the set of simulated trials.
SD of Num Responses	The standard deviation of the mean of the result: "Num Responses" over the set of simulated trials.

Total Num Subjects	The number of subjects in each simulated trial.
Num Subjects Stage 1	The number of subjects in stage 1 of the trial.
Observed Avg. Toxicity	The mean, over the simulated trials, of the observed toxicity which for each simulation is 'Num Toxic Events' / 'Total Num Subjects'.
Avg. Toxic Exposure	The mean, over the simulated trials, of the average toxicity probability of the doses allocated to the subjects in the trial. This is, given the scenario being simulated, the average risk of experiencing toxicity during the trial.
SD of Avg Toxic Exposure	The SD of the mean of the result: "Avg Toxic Exposure" over the set of simulated trials.
Observed Avg. Response	The mean, over the simulated trials, of the observed response rate which for each simulation is 'Num Responses / 'Total Num Subjects'.
Avg. Response Exposure	The mean, over the simulated trials, of the average response probability of the doses allocated to the subjects in the trial. This is, given the scenario being simulated, the average likelihood of experiencing efficacy during the trial.
SD of Avg Response Exposure	The SD of the mean of the result: "Avg Response Exposure" over the set of simulated trials.
Num Simulations With Viable Drug	The number of simulations where the model results indicate that the drug is viable, that is: an MTD exists, an MED exists, and the MED <= MTD.
Scenario Viable	The number of simulations where the scenario (true) probabilities indicate that the drug is viable, that is: an MTD exists, an MED exists, and the MED <= MTD. This should either be 0 or Num Simulations.

For the final recommended MTD and MED doses the following summary data is presented (here the data for the MTD dose is shown). For the MED dose, the probabilities referenced are the probabilities of efficacy as opposed to toxicity for the MTD dose. Where a simulation fails to find an MTD dose or MED dose, then that simulation does not contribute to the average results.

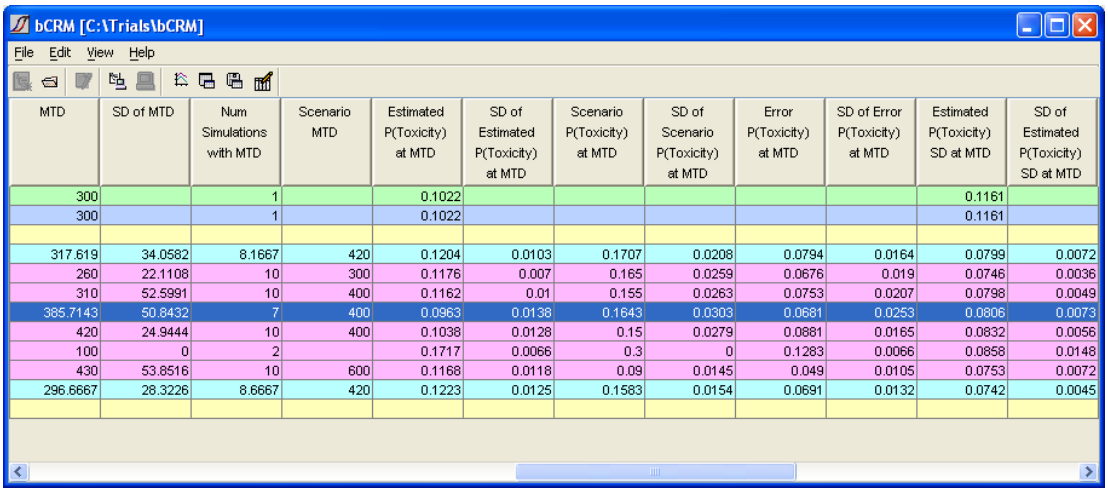

MTD	SD of MTD	Num Simulations with MTD	Scenario MTD	Estimated P(Toxicity) at MTD	SD of Estimated P(Toxicity) at MTD	Scenario P(Toxicity) at MTD	SD of Scenario P(Toxicity) at MTD	Error P(Toxicity) at MTD	SD of Error P(Toxicity) at MTD	Estimated P(Toxicity) SD at MTD	SD of Estimated P(Toxicity) SD at MTD
300		1		0.1022						0.1161	
300		1		0.1022						0.1161	
317.619	34.0582	8.1667	420	0.1204	0.0103	0.1707	0.0208	0.0794	0.0164	0.0799	0.0072
260	22.1108	10	300	0.1176	0.007	0.165	0.0259	0.0676	0.019	0.0746	0.0036
310	52.5991	10	400	0.1162	0.01	0.155	0.0263	0.0753	0.0207	0.0798	0.0049
385.7143	50.8432	7	400	0.0963	0.0138	0.1643	0.0303	0.0681	0.0253	0.0806	0.0073
420	24.9444	10	400	0.1038	0.0128	0.15	0.0279	0.0881	0.0165	0.0832	0.0056
100	0	2		0.1717	0.0066	0.3	0	0.1283	0.0066	0.0858	0.0148
430	53.8516	10	600	0.1168	0.0118	0.09	0.0145	0.049	0.0105	0.0753	0.0072
296.6667	28.3226	8.6667	420	0.1223	0.0125	0.1583	0.0154	0.0691	0.0132	0.0742	0.0045

MTD	The mean strength of the MTD dose selected over the simulated trials.
SD of MTD	The SD of the mean of the result: “MTD” over the set of simulated trials.
Num Simulations with MTD	The number of simulations where the model found an MTD. If the estimated probability of toxicity for all doses is greater than the target toxicity then no MTD will be found.
Scenario MTD	The MTD value from the scenario (true) toxicity curve.
Estimated P(Toxicity) at MTD	The mean of the estimated toxicity probabilities for the MTD selected over the simulated trials.
SD of Estimated P(Toxicity) at MTD	The SD of the mean of the result: “Estimated P(Toxicity) at MTD” over the set of simulated trials.
Scenario P(Toxicity) at MTD	The mean of the scenario (true) toxicity probabilities at the MTD selected over the set of simulated trials.
SD of Scenario P(Toxicity) at MTD	The SD of the mean of the result: “Scenario P(Toxicity) at MTD” over the set of simulated trials.
Error P(Toxicity) at MTD	The mean absolute error between the estimated toxicity and the scenario toxicity at the MTD selected over the set of simulated trials.
SD of Error P(Toxicity) at MTD	The SD of the mean of the result: “Error P(Toxicity) at MTD” over the set of simulated trials.
Estimated P(Toxicity) SD at MTD	The mean of the SD of the probability of toxicity at the MTD as estimated by the bCRM algorithm and model.
SD of Estimated P(Toxicity) SD at MTD	The SD of the result: “Estimated P(Toxicity) SD at MTD” over the set of simulated trials.

6.2 Modifying the Columns Displayed

For any of the data summary screens it is possible to configure the columns displayed and their order:

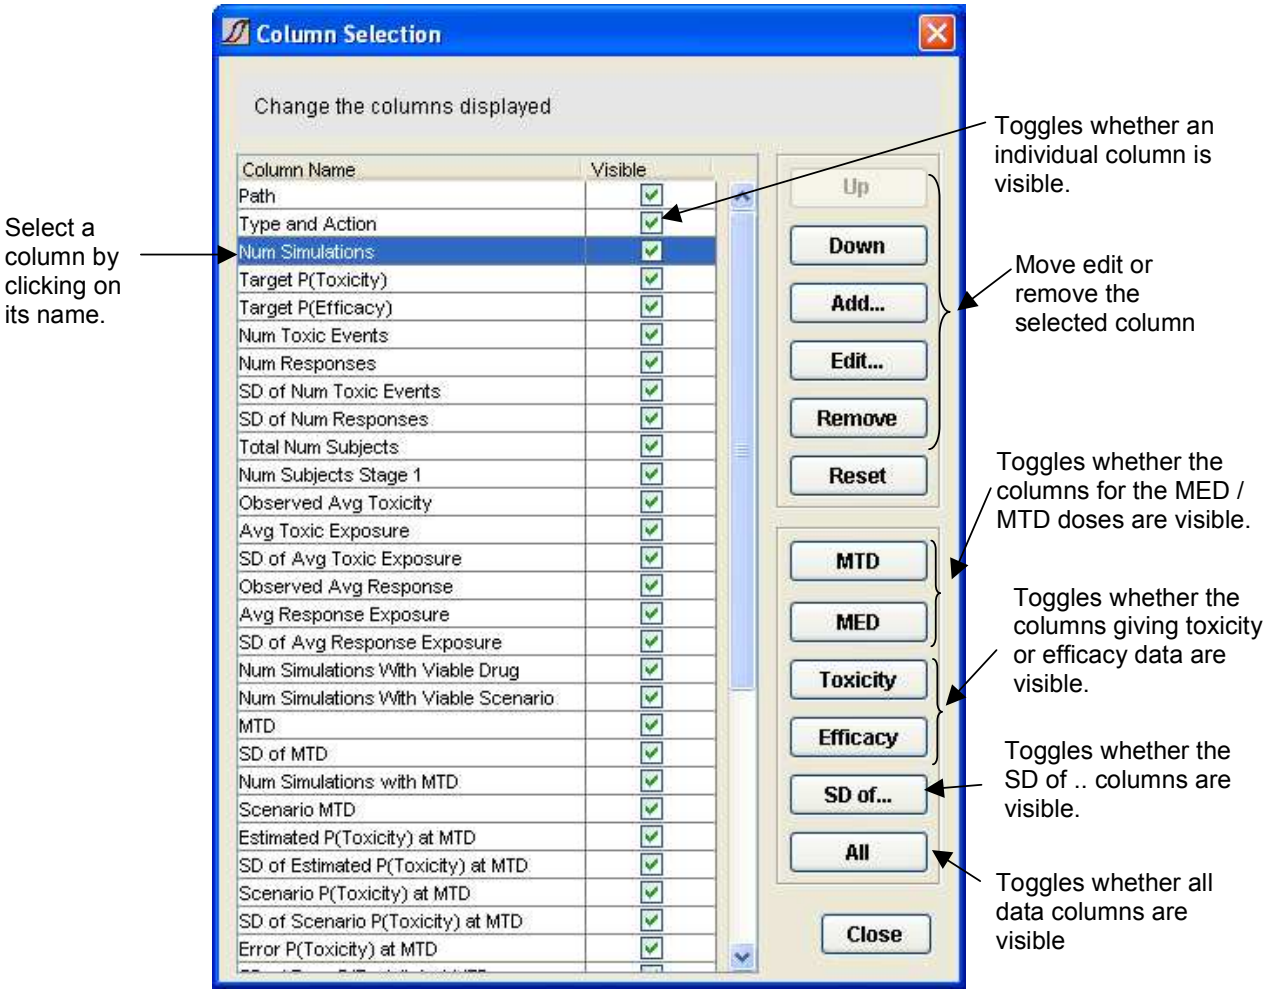

6.3 Editing or Adding a New Column

If you select to edit or add a new column the following dialog is displayed:

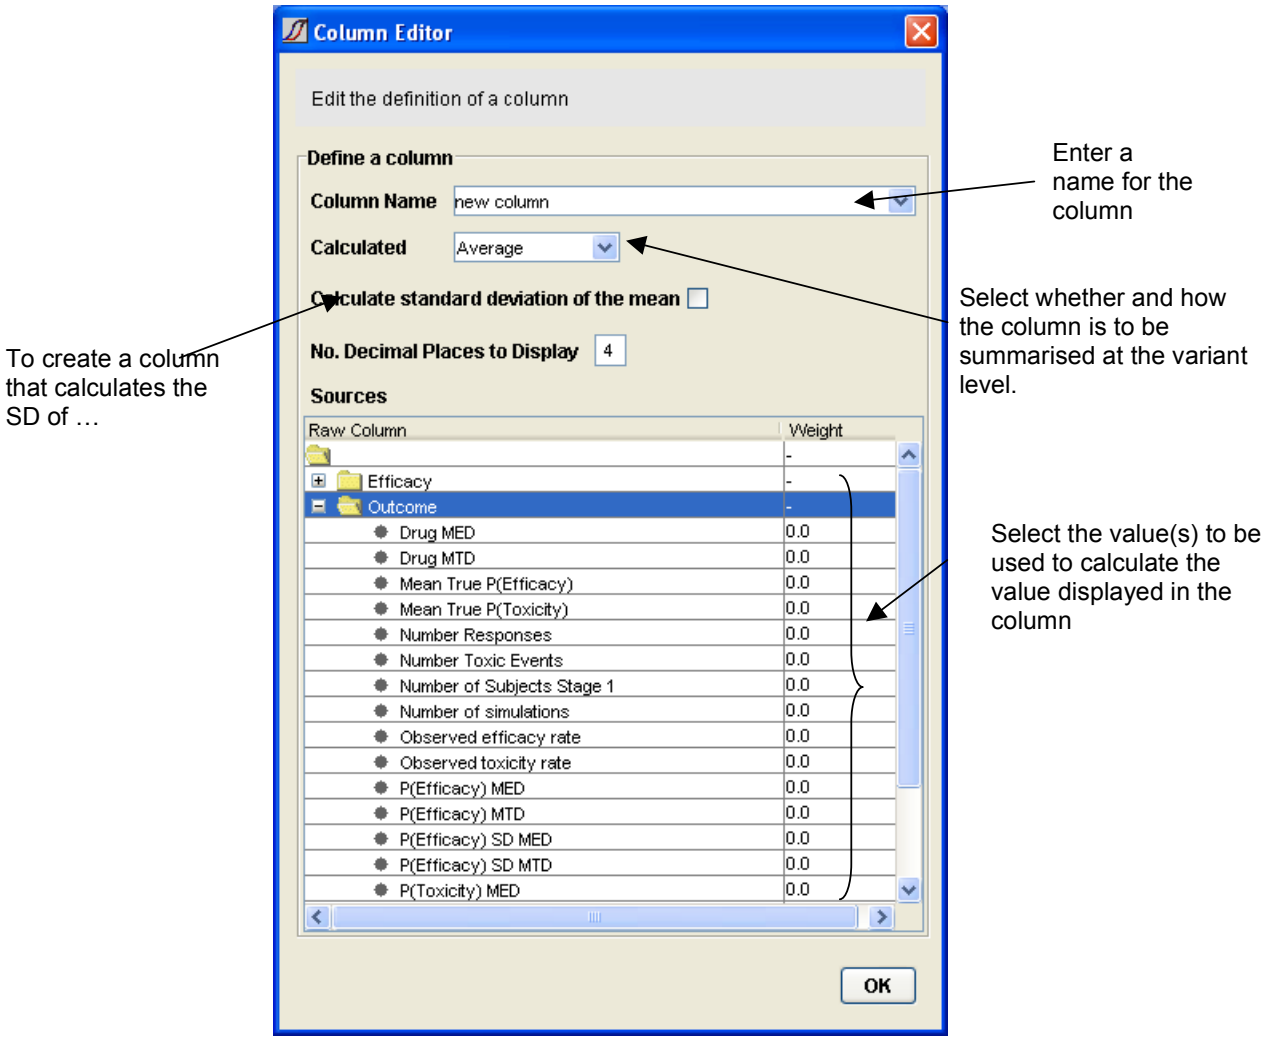

This dialog has a couple of features that warrant further explanation:

- The values to be displayed per run, can be defined as the weighted sum of a number of values. To display a single particular value simply set its weight to '1'. To scale it (e.g. by 100) enter the value to scale it by. To sum a set of values set all their weights to 1, to average a set of values set their weights to the appropriate decimal fraction for the number of values being averaged.
- The values can be summarised (or not) at the Design Variant level – either as average, ranking or not at all. For instance variants could be ranked by their mean number of toxic events.

6.4 Simulation Data

In addition to the summary data displayed in the bCRM main window at the individual run level, it is also possible to see simulation results data, safety data and response data. These can be accessed by right clicking on the Run's line in the main bCRM window and selecting '**Simulation**

Data …', '**Simulation Safety Data …'** or **Simulation Response Data …'** or by selecting the Run's line and accessing these commands from the **View** menu.

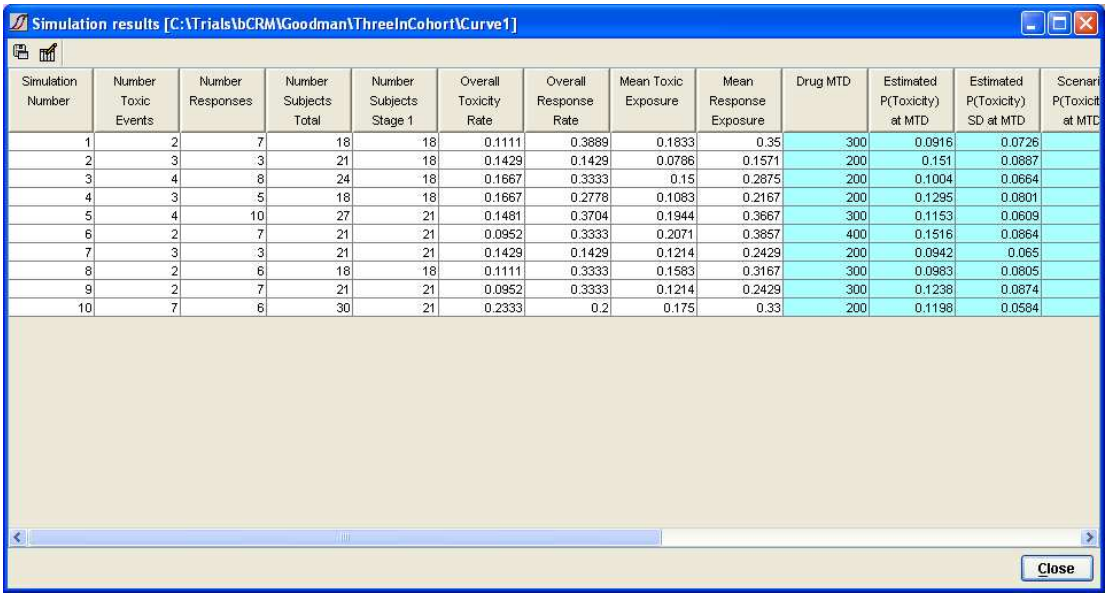

Simulation Number	Number Toxic Events	Number Responses	Number Subjects Total	Number Subjects Stage 1	Overall Toxicity Rate	Overall Response Rate	Mean Toxic Exposure	Mean Response Exposure	Drug MTD	Estimated P(Toxicity) at MTD	Estimated P(Toxicity) SD at MTD	Scenari P(Toxicit at MTD
1	2	7	18	18	0.1111	0.3889	0.1833	0.35	300	0.0916	0.0726	
2	3	3	21	18	0.1429	0.1429	0.0786	0.1571	200	0.151	0.0887	
3	4	8	24	18	0.1667	0.3333	0.15	0.2875	200	0.1004	0.0664	
4	3	5	18	18	0.1667	0.2778	0.1083	0.2167	200	0.1295	0.0801	
5	4	10	27	21	0.1481	0.3704	0.1944	0.3667	300	0.1153	0.0609	
6	2	7	21	21	0.0952	0.3333	0.2071	0.3857	400	0.1516	0.0864	
7	3	3	21	21	0.1429	0.1429	0.1214	0.2429	200	0.0942	0.065	
8	2	6	18	18	0.1111	0.3333	0.1583	0.3167	300	0.0983	0.0805	
9	2	7	21	21	0.0952	0.3333	0.1214	0.2429	300	0.1238	0.0874	
10	7	6	30	21	0.2333	0.2	0.175	0.33	200	0.1198	0.0584	

The columns are similar to those for the summary data, except they show the results per simulated trial, not the means. Toxicity and efficacy columns are only shown if they were used in this run. If the four parameter logistic was fitted, the final set of columns give the MLE estimate for the parameters together with confidence bounds.

6.5 Simulation Safety Data

(Only available if the run includes toxicity)

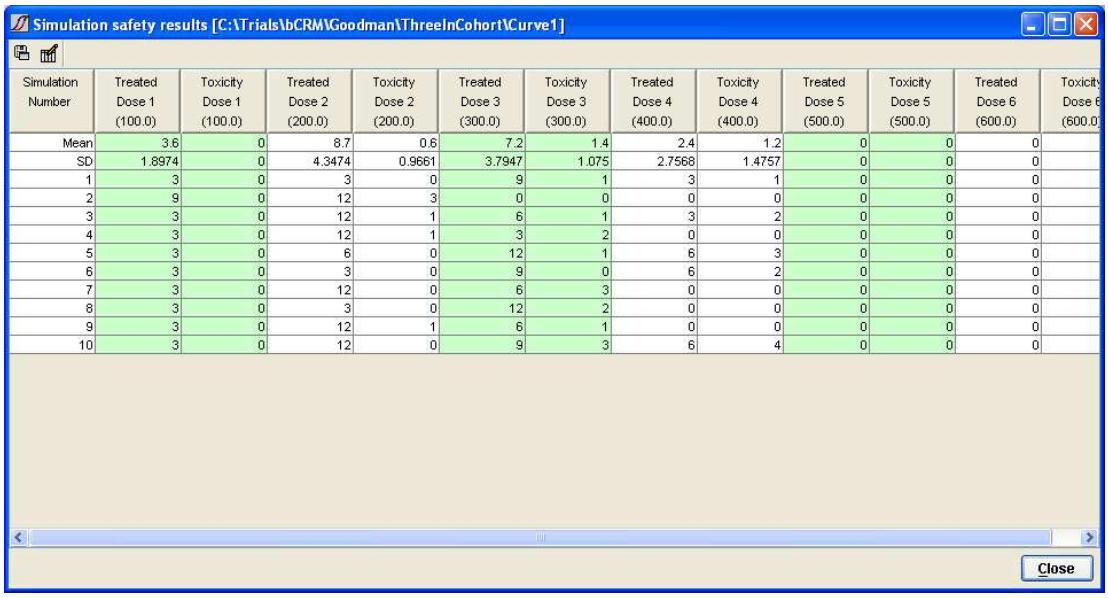

Simulation Number	Treated Dose 1 (100.0)	Toxicity Dose 1 (100.0)	Treated Dose 2 (200.0)	Toxicity Dose 2 (200.0)	Treated Dose 3 (300.0)	Toxicity Dose 3 (300.0)	Treated Dose 4 (400.0)	Toxicity Dose 4 (400.0)	Treated Dose 5 (500.0)	Toxicity Dose 5 (500.0)	Treated Dose 6 (600.0)	Toxicity Dose 6 (600.0
Mean	3.6	0	8.7	0.6	7.2	1.4	2.4	1.2	0	0	0	
SD	1.8974	0	4.3474	0.9661	3.7947	1.075	2.7568	1.4757	0	0	0	
1	3	0	3	0	9	1	3	1	0	0	0	
2	9	0	12	3	0	0	0	0	0	0	0	
3	3	0	12	1	6	1	3	2	0	0	0	
4	3	0	12	1	3	2	0	0	0	0	0	
5	3	0	6	0	12	1	6	3	0	0	0	
6	3	0	3	0	9	0	6	2	0	0	0	
7	3	0	12	0	6	3	0	0	0	0	0	
8	3	0	3	0	12	2	0	0	0	0	0	
9	3	0	12	1	6	1	0	0	0	0	0	
10	3	0	12	0	9	3	6	4	0	0	0	

The simulation safety table shows the number of subjects allocated to each of the trial doses, and the number of toxic events experienced by those subjects. The data shows the mean, SD and individual values for each simulated trial.

6.6 Simulation Response Data

(Only available if the run includes efficacy)

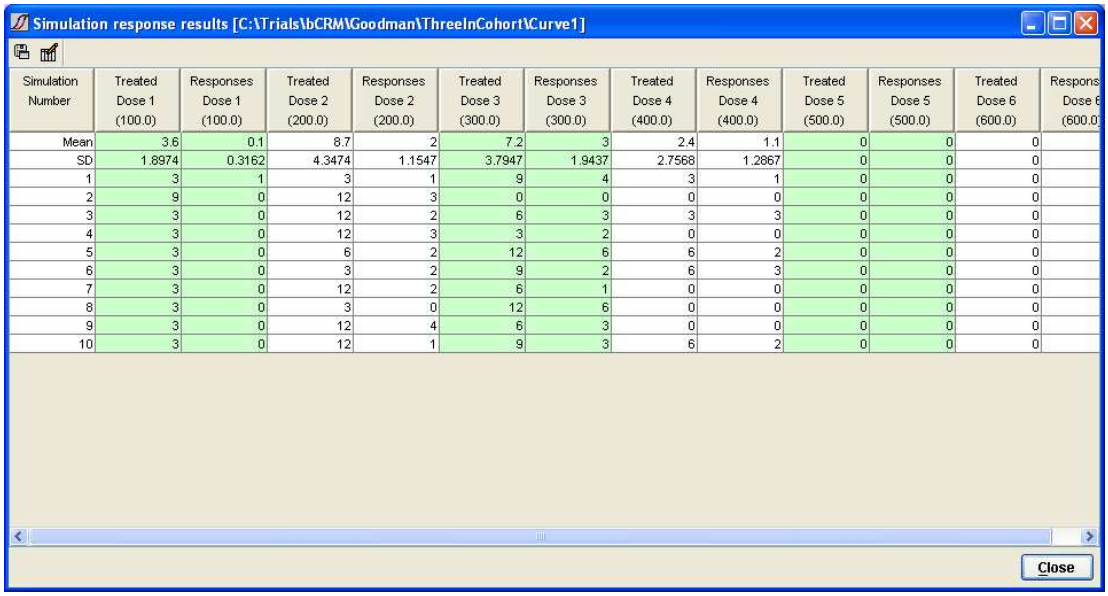

Simulation Number	Treated Dose 1 (100.0)	Responses Dose 1 (100.0)	Treated Dose 2 (200.0)	Responses Dose 2 (200.0)	Treated Dose 3 (300.0)	Responses Dose 3 (300.0)	Treated Dose 4 (400.0)	Responses Dose 4 (400.0)	Treated Dose 5 (500.0)	Responses Dose 5 (500.0)	Treated Dose 6 (600.0)	Respons Dose 6 (600.0
Mean	3.6	0.1	8.7	2	7.2	3	2.4	1.1	0	0	0	
SD	1.8974	0.3162	4.3474	1.1547	3.7947	1.9437	2.7568	1.2867	0	0	0	
1	3	1	3	1	9	4	3	1	0	0	0	
2	9	0	12	3	0	0	0	0	0	0	0	
3	3	0	12	2	6	3	3	3	0	0	0	
4	3	0	12	3	3	2	0	0	0	0	0	
5	3	0	6	2	12	6	6	2	0	0	0	
6	3	0	3	2	9	2	6	3	0	0	0	
7	3	0	12	2	6	1	0	0	0	0	0	
8	3	0	3	0	12	6	0	0	0	0	0	
9	3	0	12	4	6	3	0	0	0	0	0	
10	3	0	12	1	9	3	6	2	0	0	0	

The simulation response table shows the number of subjects allocated to each of the trial doses, and the number of response events experienced by those subjects. The data shows the mean, SD and individual values for each simulated trial.

7 Graphs and Charts

Once simulations have been run, it is possible to see charts and graphs of the results.

If you select a Simulation Series, Variant or individual Run -

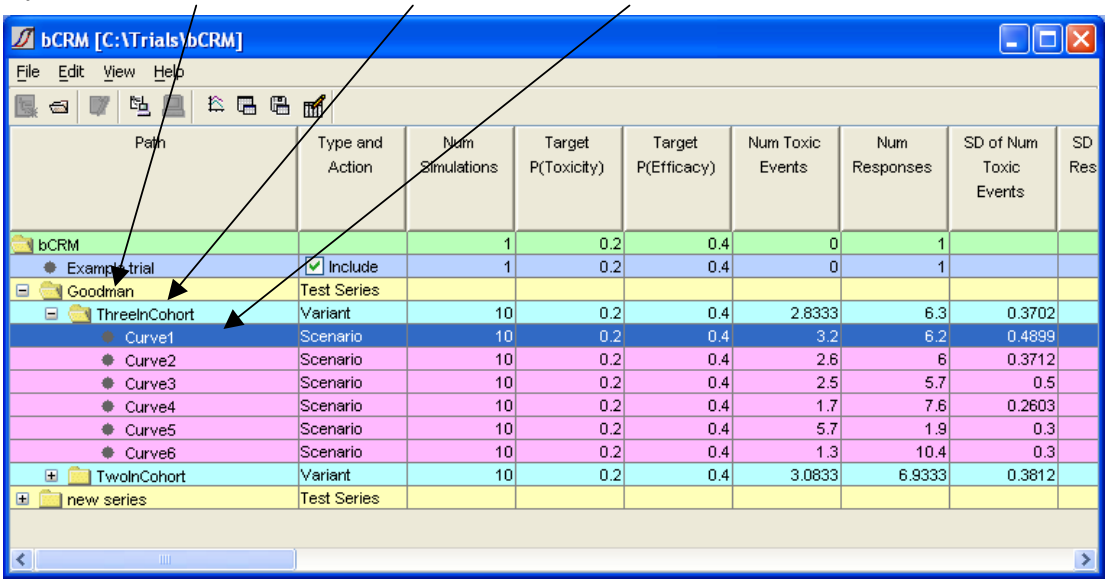

You can then view the graph manager by clicking on the graph manager button on the tool bar (), selecting **Graph Manager** from the **View** menu or right clicking on the Simulation Series, Variant or Run line in the bCRM main window and selecting **Graphs** from the pop-up menu.

7.1 Histograms for Series and Variants

If the Graph Manager is opened on a Simulation Series or a Variant, then you can select to display the data in any of the columns as a histogram.

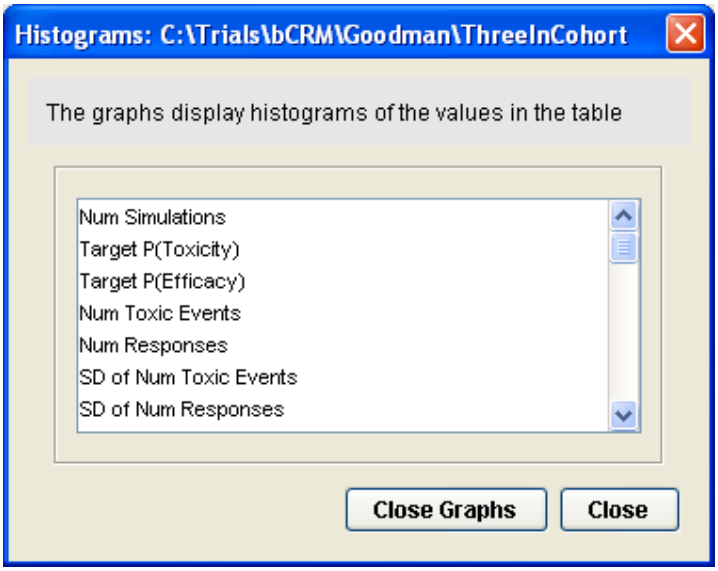

For a Simulation Series the histogram is plotted with a column per Design Variant in the series:

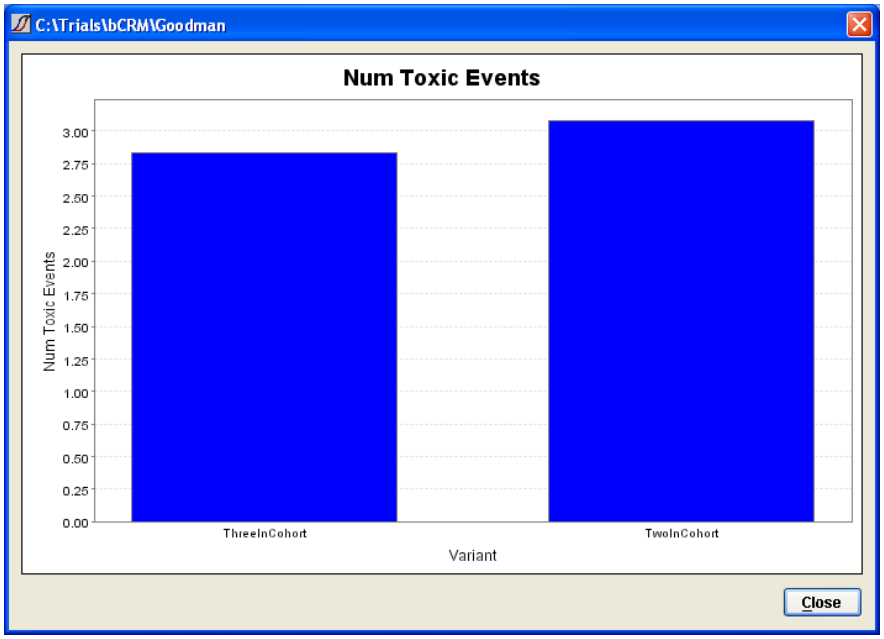

For an individual variant, the histogram is plotted with a column per Scenario:

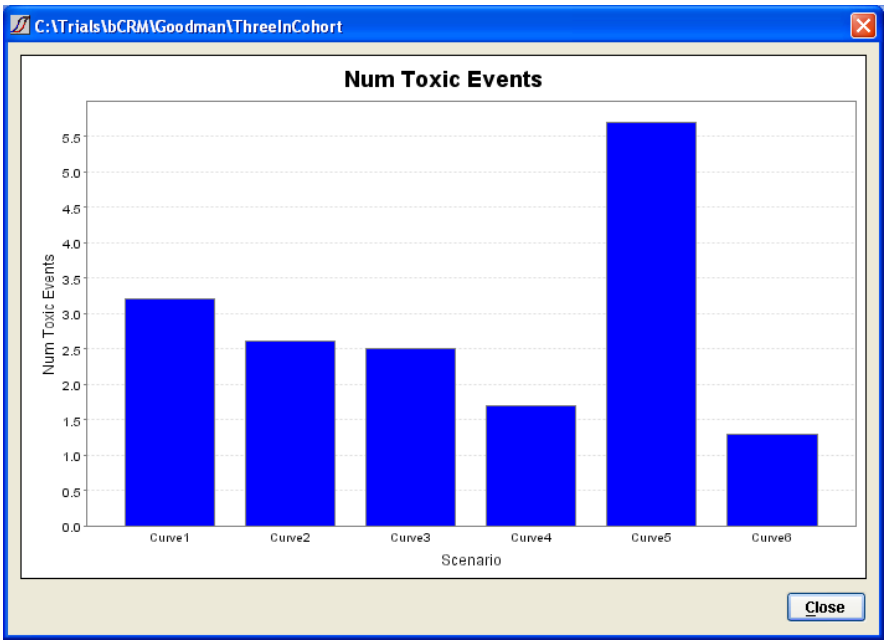

7.2 Charts for Runs

For an individual Run the charts available are quite different:

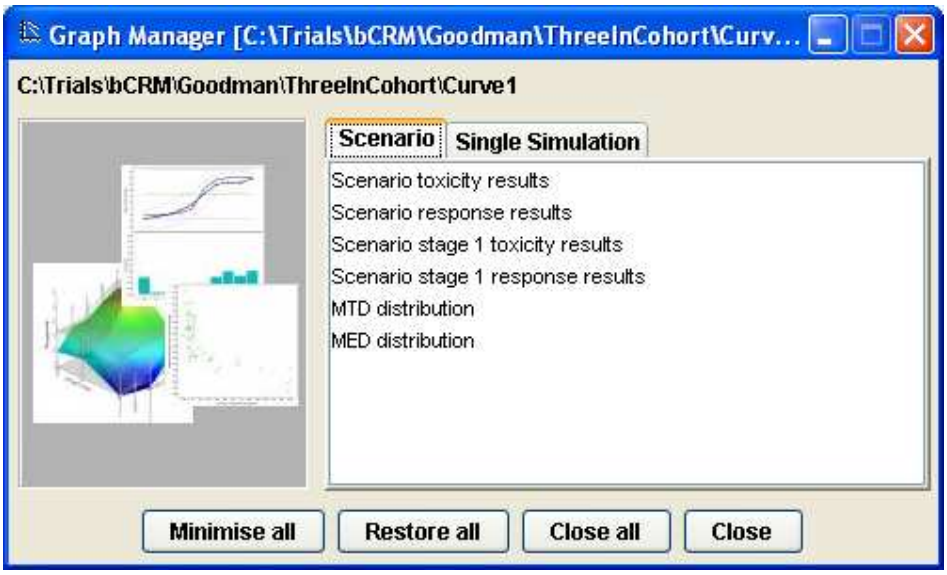

(If the run uses only toxicity or only efficacy then only those graphs will be available)

All charts have a standard set of buttons in the top left corner:

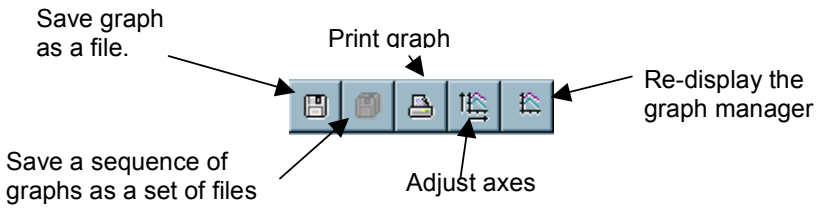

Graphs for individual simulations have a 'selector' control at the bottom:

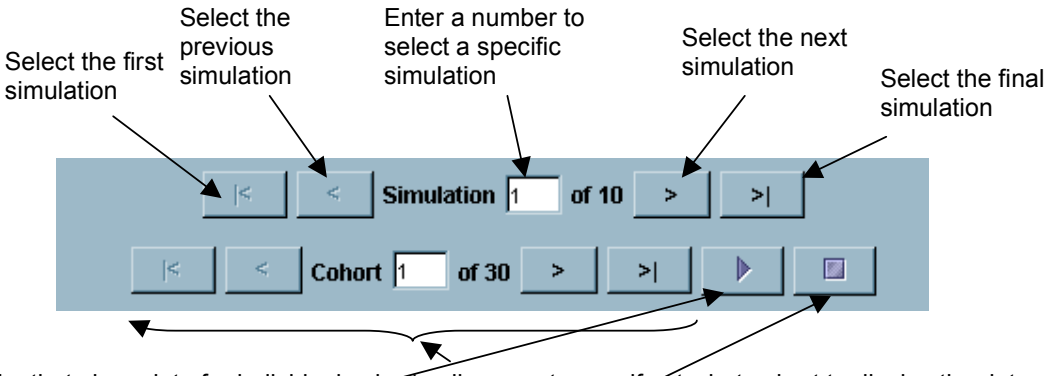

Graphs that show data for individual cohorts allow you to specify at what cohort to display the data, and to

- animate
- halt the animation of the graph.

7.2.1 Scenario toxicity results

The Scenario toxicity results display a histogram of the average number of subjects allocated to each dose, along with the average number that experienced toxicity at the dose. The mean MTD and MED dose strengths determined for each simulation are also marked as vertical lines.

Below the histogram are plotted the actual scenario toxicity probabilities (circles) along with the average estimated toxicity curve.

The variation in estimated toxicity curve across simulations is also used to plot an upper and lower bound (the dashed lines) for the estimated toxicity curve with a 95% confidence interval. Upper and lower bounds for difference confidence intervals can also be plotted by selected a different confidence level in the top right hand corner.

The grey line is the observed toxicity, i.e. the values in the histograms converted into toxicity rates.

The **Scenario stage1 toxicity results** displays the same information. For this plot, however, only the results from cohorts assigned in stage 1 of the trial are considered. A cohort is considered to be assigned in stage 1 of the trial if the allocator doesn't consider the stage 1 termination conditions to have been met. The MTD, MED and probability curve values plotted are those determined at the end of the last stage 1 cohort.

7.2.2 Scenario response results

The Scenario response results display a histogram of the average number of subjects allocated to each dose, along with the average number that experienced efficacy at the dose. The average MTD and MED dose strengths determined for each simulation are also marked as vertical lines.

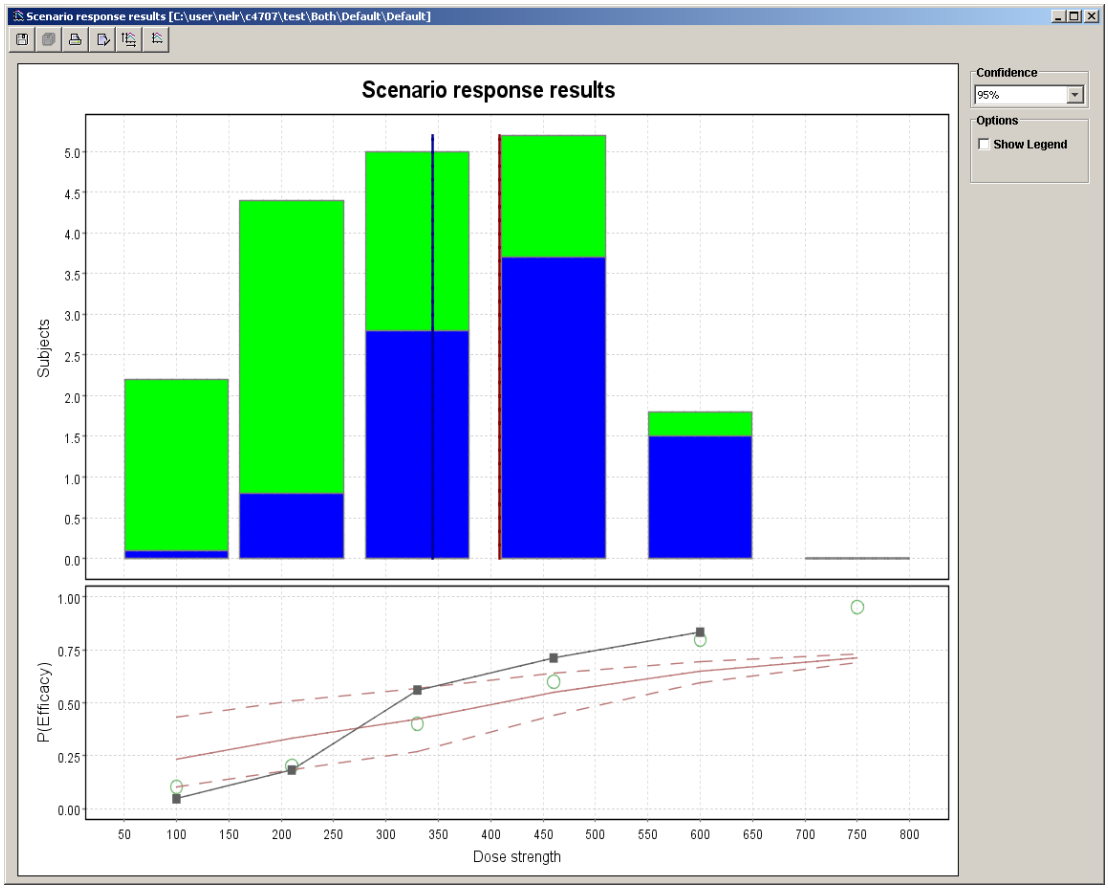

Below the histogram are plotted the actual scenario efficacy probabilities (circles) along with the average estimated efficacy curve.

The variation in estimated efficacy curve across simulations is also used to plot an upper and lower bound for the estimated efficacy curve with a 95% confidence interval. Upper and lower bounds for difference confidence intervals can also be plotted by selected a different confidence level in the top right hand corner.

The grey line is the observed efficacy, i.e. the values in the histograms converted into efficacy rates.

The **Scenario stage1 efficacy results** displays the same information but only taking data from stage 1 cohorts.

7.2.3 MTD distribution

The MTD distribution displays a histogram of the number of simulations for which a specified dose was determined to be the MTD.

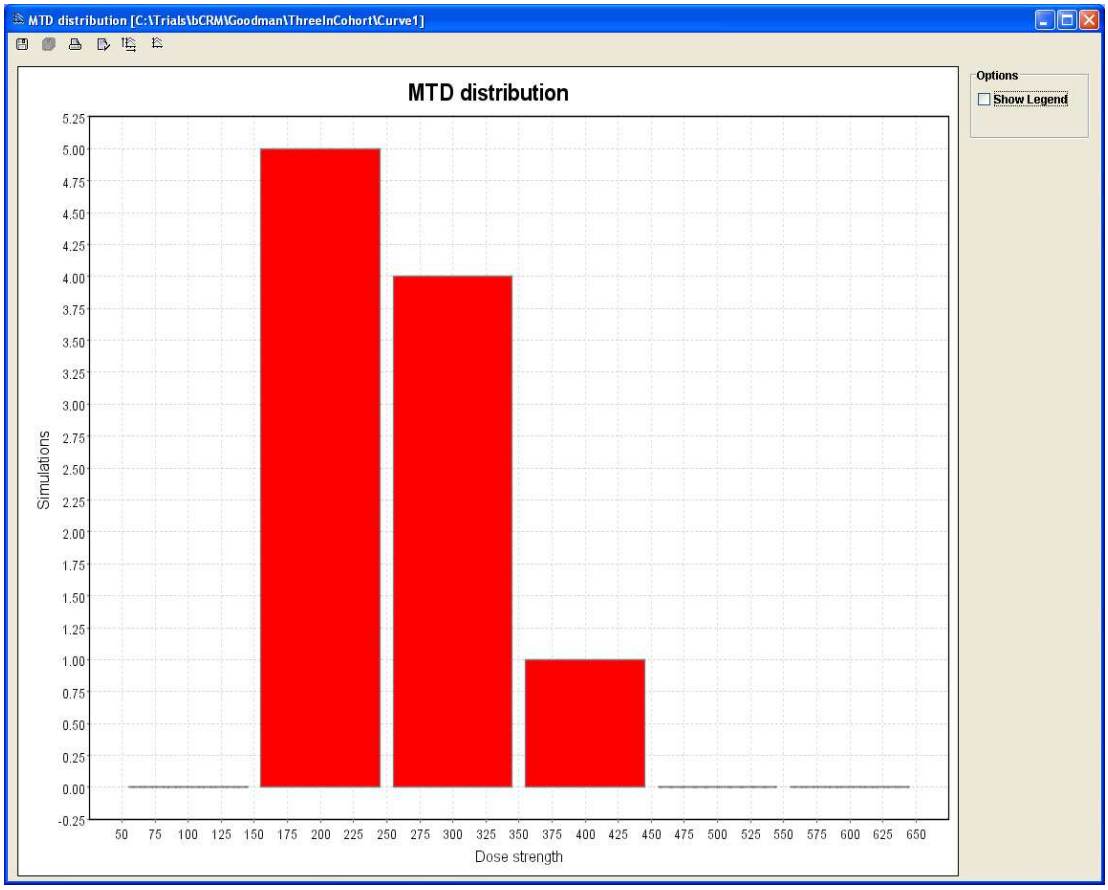

7.2.4 MED distribution

The MED distribution displays a histogram of the number of simulations for which a specified dose was determined to be the MED.

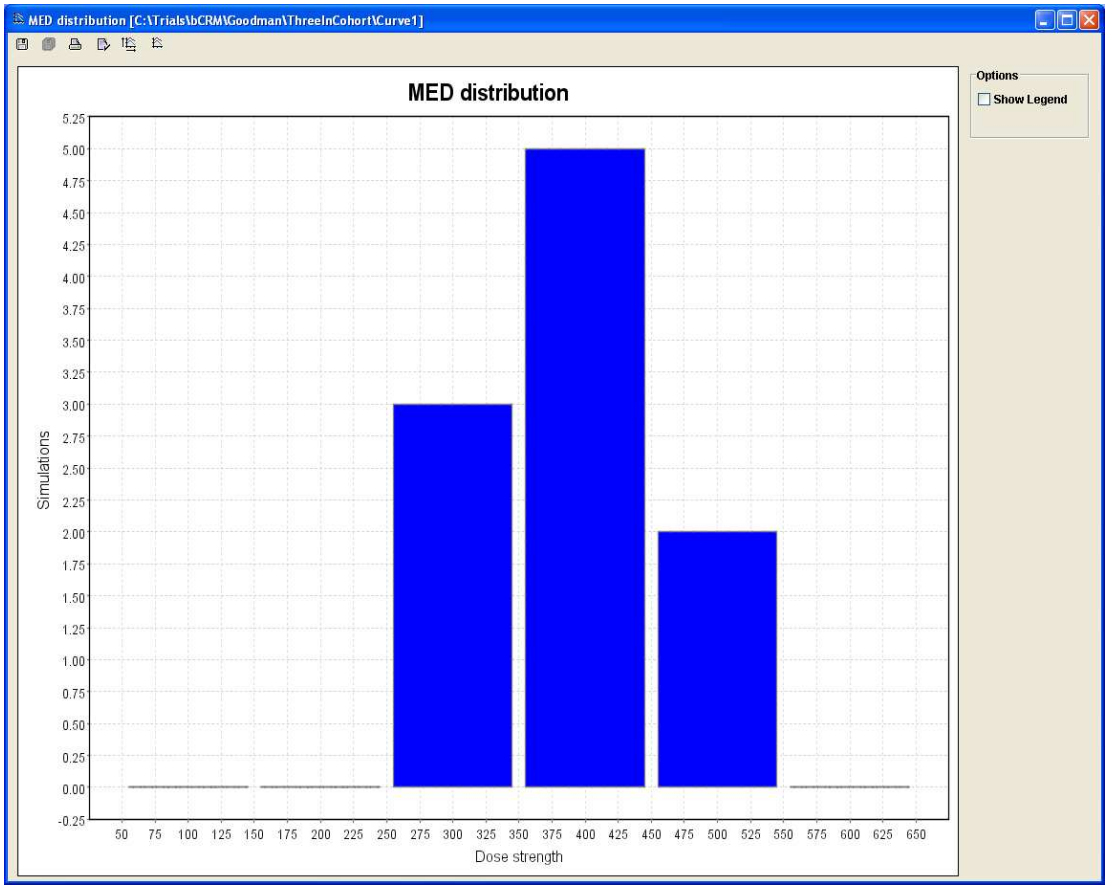

7.3 Charts for individual simulations

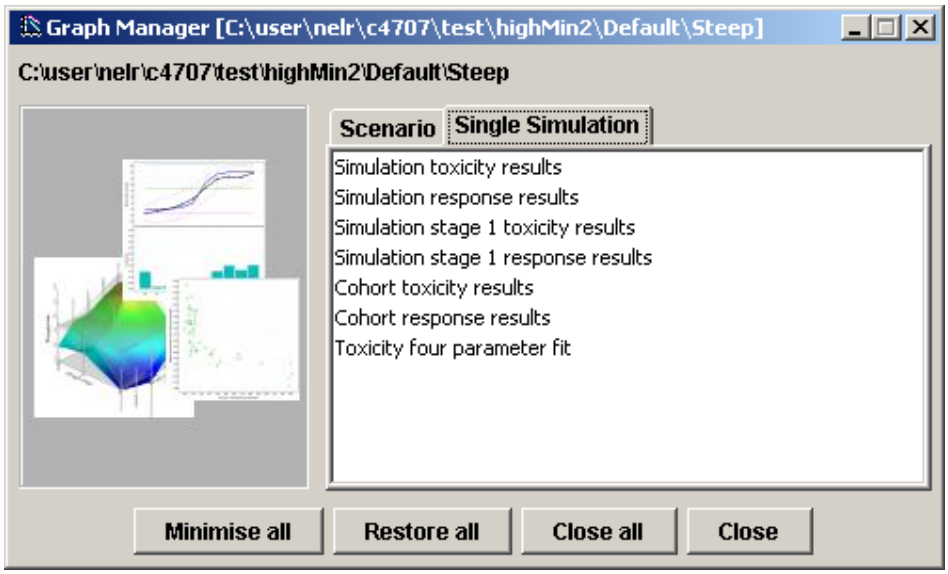

7.3.1 Simulation toxicity results

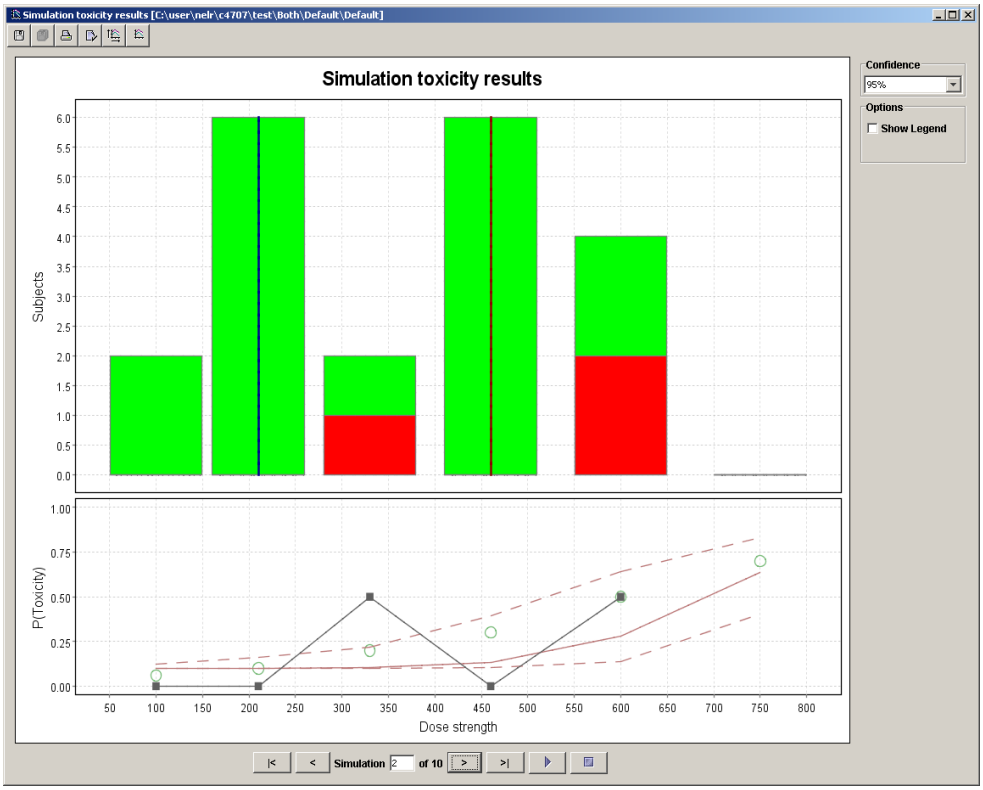

Displays the histogram plot of the number of patients allocated to and experiencing toxic events at each dose along with the estimated MTD and MED, for each simulation. It also displays the scenario and estimated toxicity curve from each simulation along with the estimated upper and lower bounds to the toxicity curve and the observed toxicity rates.

The Simulation stage 1 toxicity results are similar but only use data from cohorts assigned in stage 1 of the trial

7.3.2 Simulation response results

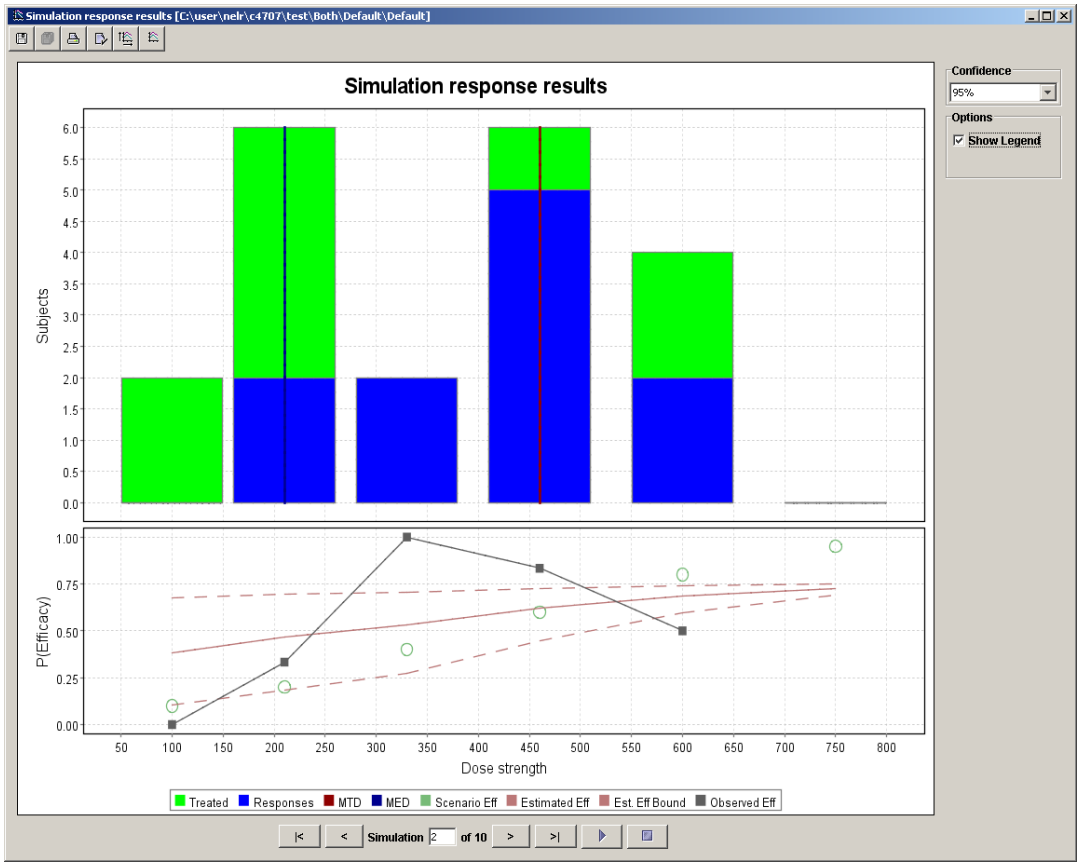

Displays the histogram plot of the number of patients allocated to and experiencing responses at each dose along with the estimated MTD and MED, for each simulation. It also displays the scenario and estimated efficacy curve from each simulation along with the estimated upper and lower bounds to the efficacy curve and the observed efficacy rates.

The Simulation stage 1efficacy results are similar but only use data from cohorts assigned in stage 1 of the trial.

7.3.3 Cohort toxicity results

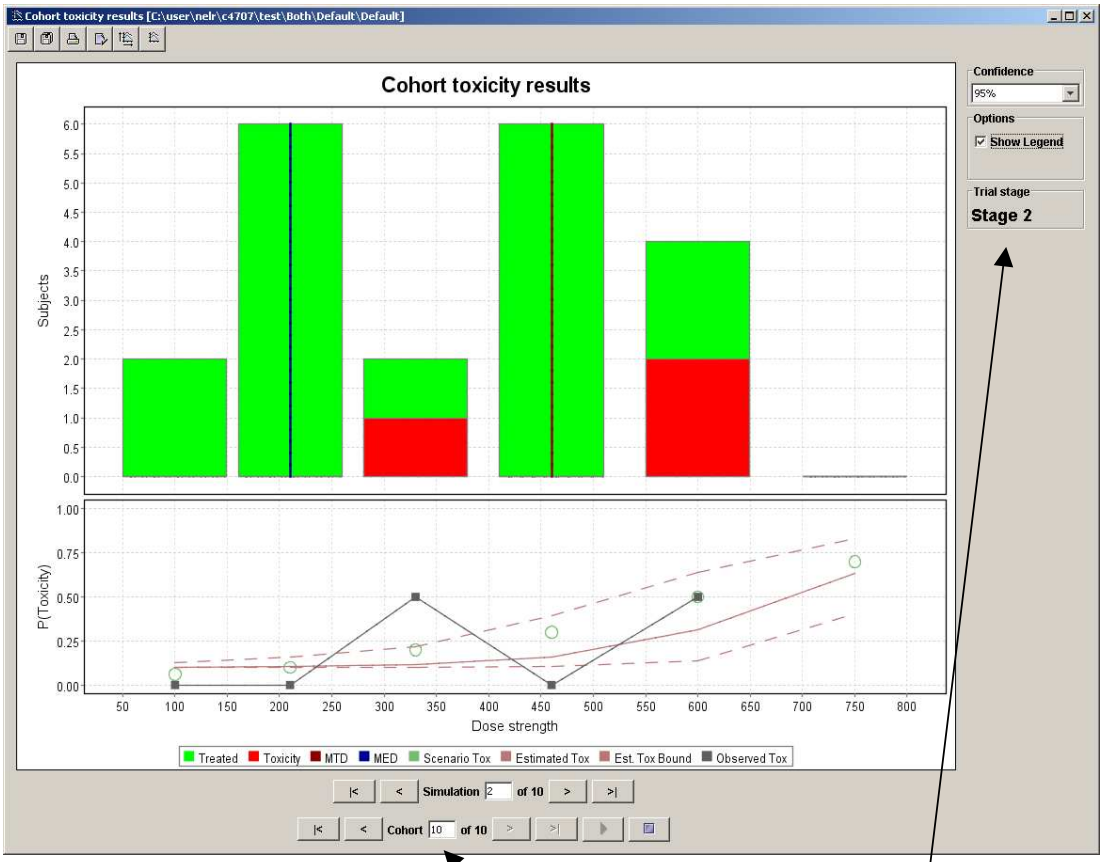

The cohort toxicity results are similar to the simulation toxicity results, however:

- The histograms display the doses allocated, and number of toxic event for subjects in all cohorts up to the currently selected index. Stepping through cohorts, or animating the results, enables you to see how the dose allocations and toxicities build up through the simulation.
- The observed and estimated toxicity curves, and the upper and lower bounds to the estimated toxicity are those produced by the model using the data from all cohorts up to the currently selected index. Stepping through cohorts enables you to see how the model converges to the final estimated curve at the end of the simulation.
- The plot displays which stage of the trial the currently selected cohort was assigned to, stage1 or stage2.

7.3.4 Cohort response results

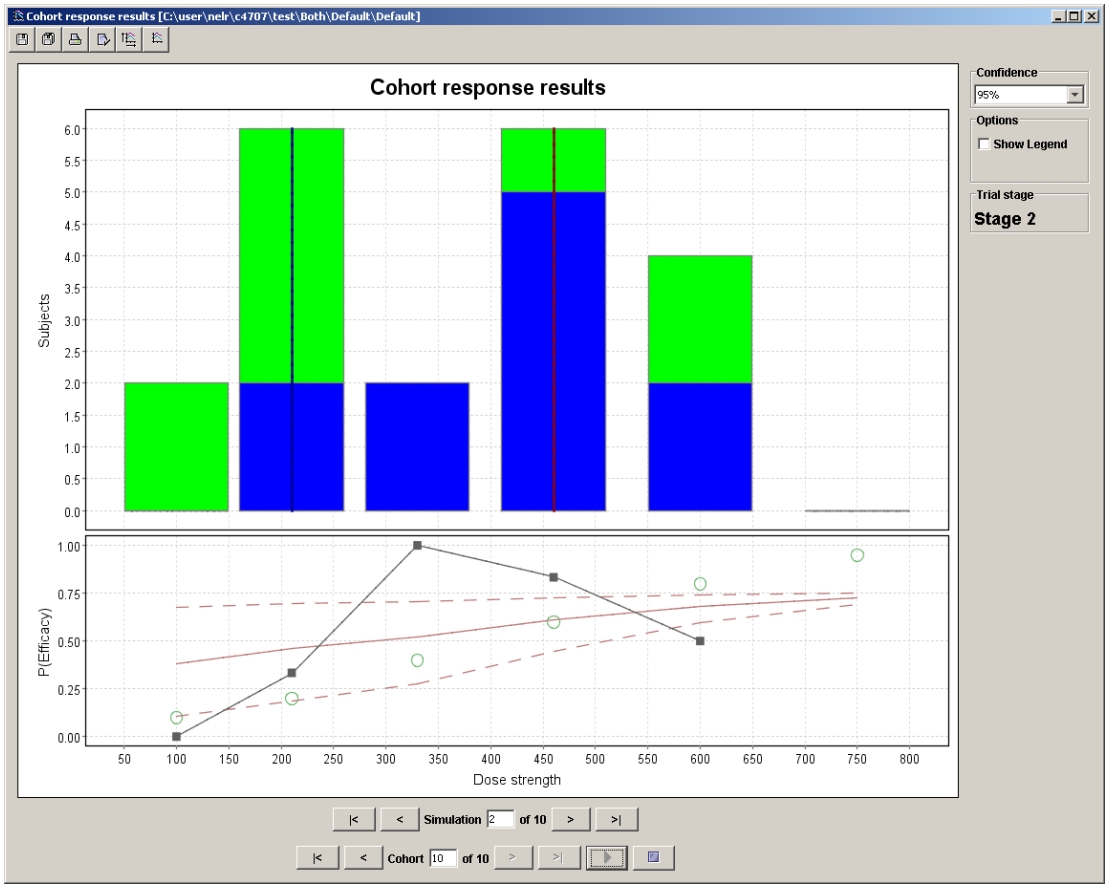

The cohort response results are similar to the cohort toxicity results. However the responses are displayed rather than the toxic events, and the efficacy curves plotted, not the toxicity curves.

7.3.5 Four Parameter Logistic

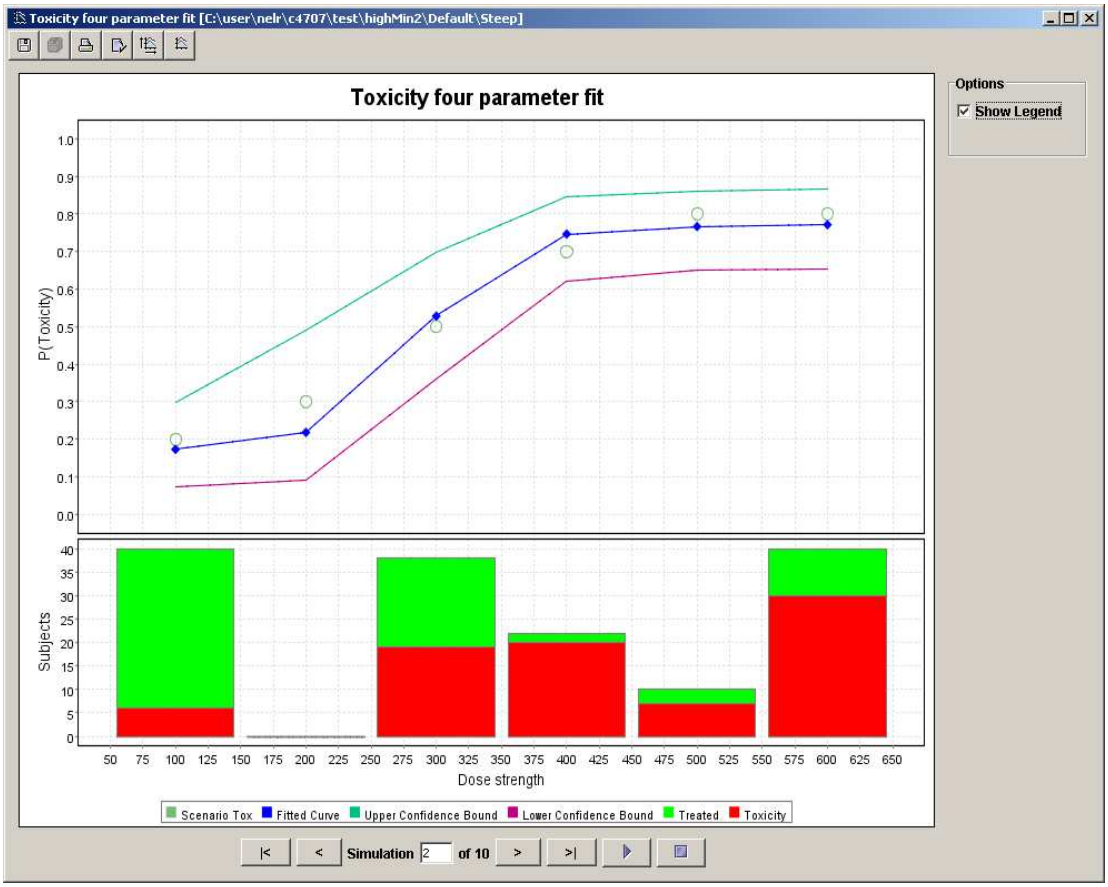

This shows the fit of the four parameter logistic to toxicity (or efficacy if there is only efficacy) together with the confidence bounds that were specified when editing the parameters.

In this case the confidence bounds clearly show that the drug clearly causes toxicity but that a wide range of curve could fit the data. Because no subjects were allocated to 200mg, the fit cannot distinguish between curves that reach a lower asymptote of 0.2 at 100mg and curves that reach zero at some lower dose. At the other extreme, the data cannot quite rule out a step function. Although it is unlikely, given this data, the response could be flat at 0.3 up to a dose of 275mg, then jump to a constant 0.65.

8 Instructions for evaluating the model implementation

8.1 Debug output:

In order to see in detail what the implementation is doing (and help determine that the implementations are indeed correct), diagnostic output from the model is available. This is written to standard out and is controlled by the parameter "verbose".

This parameter is not accessible through the user interface but can be manually edited in the ‘init.bcrm’ file – e.g. "verbose 1"

“verbose” takes a single integer value:

0 = The default. Simulation number and numbers of subjects only

1 = As 0 plus MTD, MED and the reason for stopping

2 = As 1 plus the values of Beta1 and Beta2 during allocation