

# EffTox Users Guide and Tutorial (version 2.9)

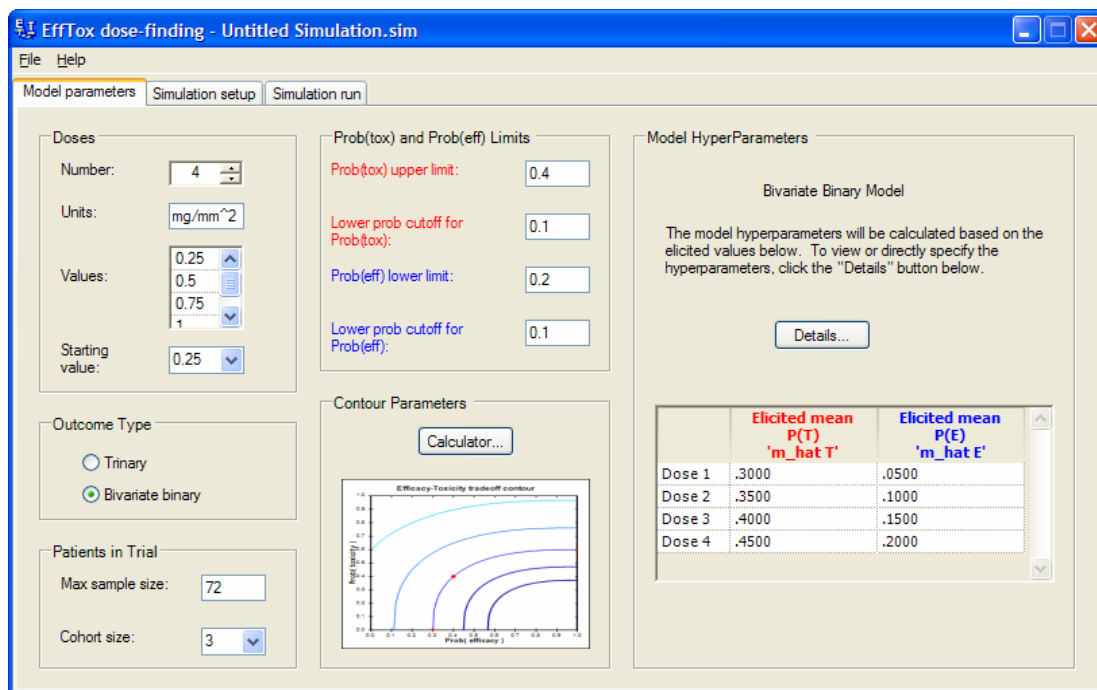
## Introduction

This program is a (beta) implementation of the dose-finding method described in "Dose-Finding Based on Efficacy-Toxicity Trade-Offs" by Peter F. Thall and John D. Cook, which appeared in *Biometrics* in September, 2004. One must read the *Biometrics* paper in order to fully understand how to use this software. The notation in the software matches that of the paper.

This document explains how each of the components of the software are used and gives some guidance for how to select the design parameters.

## Starting a new simulation

EffTox opens to the **Model parameters** tab for a new simulation:



Please note that the values that are filled in when the program opens are *examples* rather than *defaults*. Indeed, there can hardly be default values for clinical trial designs since individual trials differ so greatly in their details.

First, fill out the **Doses** group:

The 'Doses' window contains the following settings:

- Number:** 4
- Units:** mg/mm<sup>2</sup>
- Values:** 0.25, 0.5, 0.75, 1
- Starting value:** 0.25

Here **Number** is the number of doses in the trial.

**Units** are for display purposes; they are not used in any calculations.

The **Values** are numeric values of the doses, such as milligrams per square millimeter in this example. These must be non-negative numbers and must be strictly increasing.

The **Starting value** is the value of the dose to be given to the first cohort. This can be any dose value.

Dose-finding trials have traditionally started at the lowest dose for fear of starting with a toxic dose. However, since EffTox incorporates efficacy outcomes as well as toxicity, one could as easily argue for starting at the *highest* dose for fear of starting with an ineffective dose. The best starting dose depends on the particular trial, but it is often reasonable to start in the middle of the dose range.

One way to select doses would be as follows. Ask the investigator “What would you guess the best dose is?” Or perhaps “If there were only one patient in this trial, how much of the agent would you give that person?” Use the answer as the starting dose. Then add doses above and below.

Version 2.9 has a symmetric version of the no-skip rule. The previous rule (version 2.7 and earlier) was not to skip over an untried dose when escalating doses. The current rule is simply not to skip over untried doses, *i.e.*, whether escalating or de-escalating.

Next select the **Outcome Type**.

The 'Outcome Type' window shows the following selection:

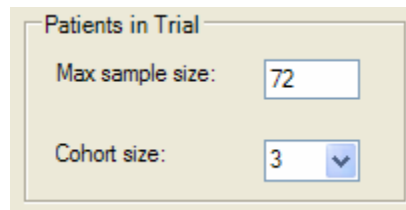
- Trinary
- Bivariate binary

This refers to the model to be used. The trinary model has three possible outcomes: efficacy, toxicity, or neutral. The bivariate binary model has four possible outcomes: efficacy without toxicity, efficacy with toxicity, failure without toxicity, and failure with toxicity. In particular, the bivariate binary model accommodates the possibility of toxicity and efficacy both occurring whereas the trinary model assumes that efficacy and toxicity cannot both occur in the same patient.

When in doubt, use the bivariate binary model. Sometimes investigators can be certain that efficacy and toxicity cannot both occur in the same patient, but then they do. Also, one may choose to distinguish toxicity with and without response in the statistical model even if clinically the distinction does not matter.

*Other parts of this dialog depend on the number of doses and outcome type. Filling in these sections first makes the software easier to use.*

Next, fill out the **Patients in Trial** group.



The image shows a dialog box titled "Patients in Trial". It contains two input fields. The first is labeled "Max sample size:" and has a text box with the number "72" entered. The second is labeled "Cohort size:" and has a dropdown menu with "3" selected and a downward arrow.

Here **Max sample size** is the maximum number of patients in the trial, the number of patients that will be treated unless an early stopping rule kicks in.

**Cohort size** is the number of patients treated at a given dose before a new dose decision is made. Setting this value to 1 effectively eliminates cohorts.

The cohort size should evenly divide the maximum sample size.

Traditionally, dose-finding methods have often used cohorts of size three. Feel free to try other cohort sizes, especially smaller cohorts. *Large cohorts are dangerous* since stopping rules cannot stop a trial until an entire cohort has been treated. Therefore the cohort size is the number of patients that could potentially be treated with an ineffective and toxic treatment.

Next fill out the **Probability Limits** group.

Prob(tox) and Prob(eff) Limits	
Prob(tox) upper limit:	<input type="text" value="0.4"/>
Lower prob cutoff for Prob(tox):	<input type="text" value="0.1"/>
Prob(eff) lower limit:	<input type="text" value="0.2"/>
Lower prob cutoff for Prob(eff):	<input type="text" value="0.1"/>

The probability limits are used to specify dose acceptability and stopping rules. At each step, the method chooses between acceptable doses. If no dose is acceptable, the trial stops.

For a given dose, toxicity is acceptable if the posterior probability of the probability of toxicity being below the **Prob(tox) upper limit** is at least **Lower prob cutoff for Prob(tox)**. Similarly, efficacy is acceptable if the posterior probability of the probability of efficacy being above the **Prob(eff) lower limit** is at least **Lower prob cutoff for Prob(eff)**. These probabilities of probabilities can be confusing. A little notation will help. A dose  $x$  is acceptable if

$$\Pr[ \pi_T(x, \theta) < \text{Prob(tox) upper limit} \mid \text{data} ] \geq \text{Lower prob cutoff for Prob(tox)}$$

and

$$\Pr[ \pi_E(x, \theta) > \text{Prob(eff) lower limit} \mid \text{data} ] \geq \text{Lower prob cutoff for Prob(eff)}.$$

In the illustration, the acceptability criteria are

$$\Pr[ \pi_T(x, \theta) < 0.4 \mid \text{data} ] \geq 0.1$$

and

$$\Pr[ \pi_E(x, \theta) > 0.2 \mid \text{data} ] \geq 0.1.$$

The cutoff probabilities are small numbers, typically 0.1 or smaller. The smaller these values are, the less evidence of desirability needed to continue, and the less often the trial will stop.

Some people find it easier to think in terms of what makes a dose *unacceptable*. In the example given in the illustration above, a dose is *unacceptable* if either

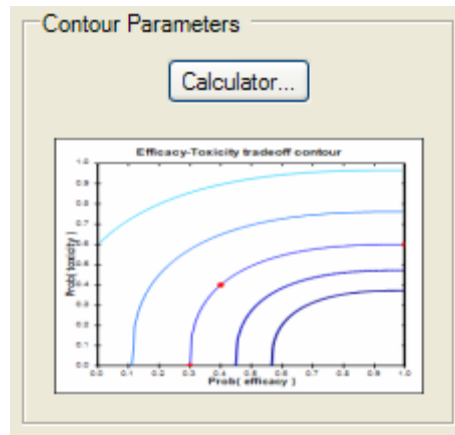
$$\Pr[ \pi_T(x, \theta) > 0.4 \mid \text{data} ] \geq 0.9$$

or

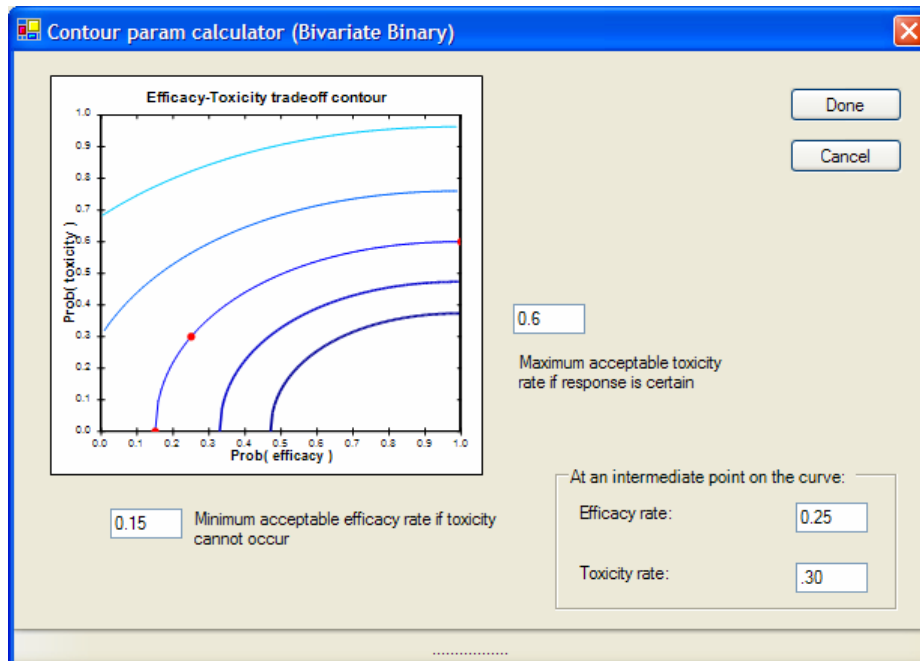
$$\Pr[ \pi_E(x, \theta) < 0.2 \mid \text{data} ] \geq 0.9.$$

You can hold your mouse over the text labels for these inputs to get a popup tip reminding you of the definitions.

Next examine the **Contour Parameters** group. Click the **Calculator** button to access the Contour Parameters dialog.



The **Calculator** button brings up a dialog which allows you to specify three equally desirable points that define the primary contour curve.



The screen shot above shows the contour parameter calculator. Three equally desirable  $(P(\text{eff}), P(\text{tox}))$  pairs are given:  $(0.15, 0)$ ,  $(0.25, 0.30)$ , and  $(1.0, 0.6)$ . These points are plotted with red dots. The contours become progressively darker blue as the contours increase in utility.

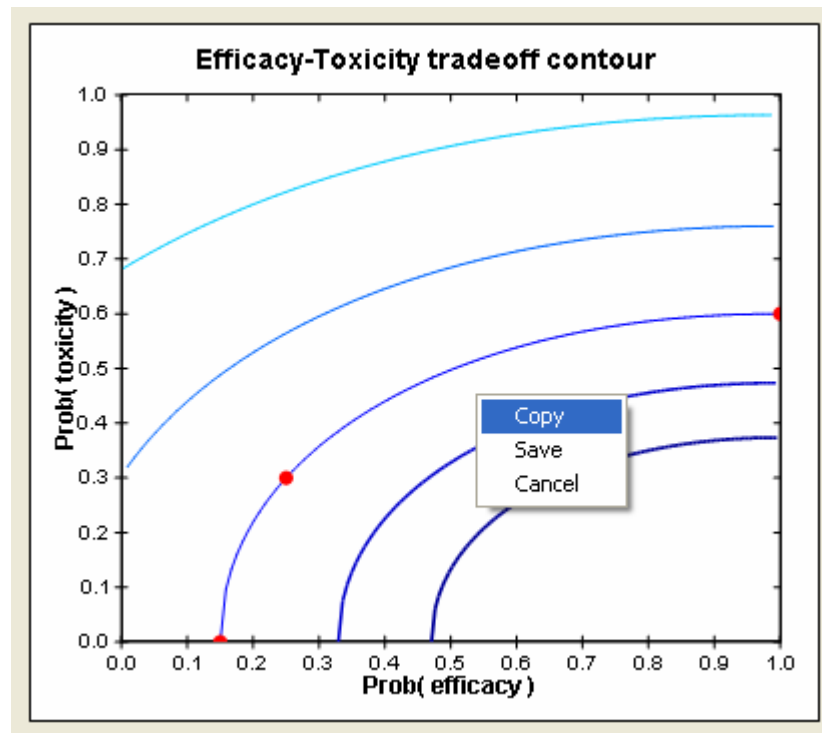
The size of this dialog box can be increased by grabbing the bottom-right corner with the mouse and pulling down and to the right.

The coordinates of the three elicited points can be changed either by directly editing the values in the edit boxes, or by placing the cursor in a box and using the up and down arrow keys to increment or decrement the values repeatedly by 0.01. The curves are recalculated and redrawn continuously as the points are changed.

The contour calculator behavior changes according to the outcome type because the domain of possible efficacy and toxicity probabilities depends on the model: the bivariate domain is a full square, but the trivariate domain is a triangle because the sum of the probabilities of toxicity and response must sum to no more than one in this model. Therefore the model must be selected before using this calculator.

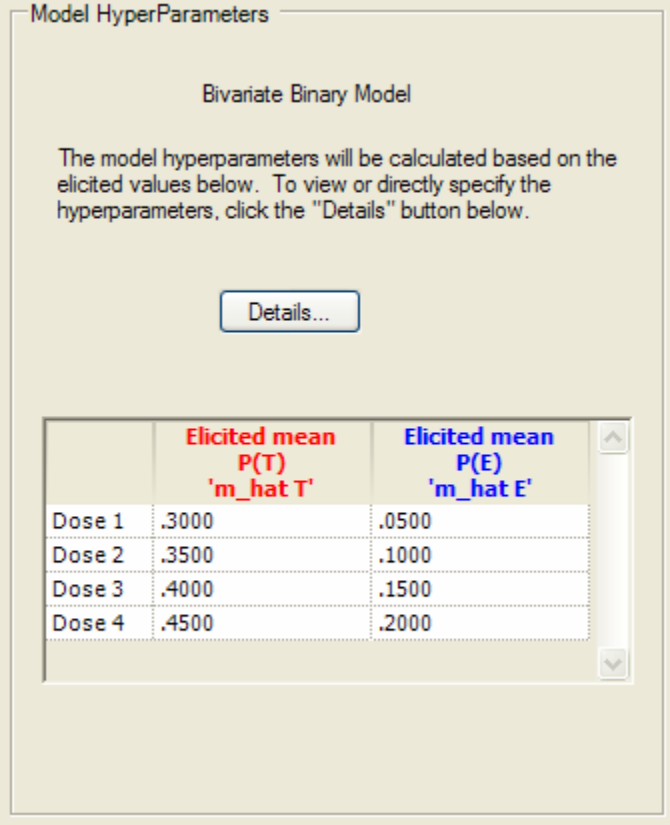
The contour parameters are subject to constraints that make the resulting contour clinically meaningful. In particular, the curve must be strictly increasing since desirability must increase with probability of efficacy. At times the up or down arrow keys are disabled in a particular box because allowing a change in the requested direction might violate a constraint. If this happens and there are other coordinates you want to change, modify the other coordinates and then come back to the coordinate that was against a constraint.

You may right-click on the contour graphic to pull up a dialog that will let you save the image to disk or copy it to the Windows clipboard to paste the image into another program.



When you click the **Done** button, the parameters will be exported to the Model Parameters tab that launched this dialog.

Next, fill out the **Model HyperParameters** group. Enter elicited prior probabilities of toxicity and efficacy at each dose in the grid. Typically the model performs best if there is some spread between the elicited probabilities. As a rule of thumb, probabilities should differ by at least 0.05 and should be between 0.01 and 0.99.



Model HyperParameters

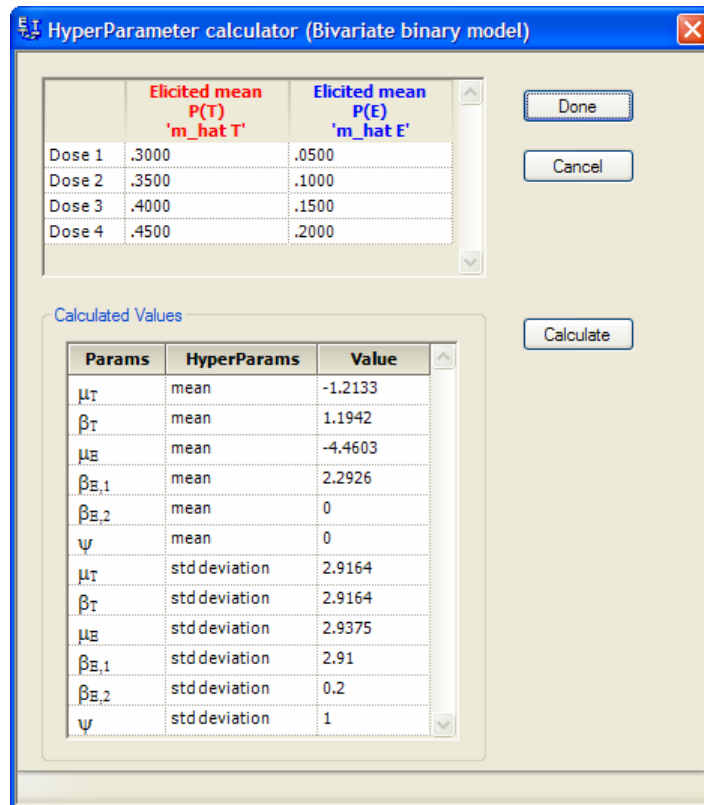
Bivariate Binary Model

The model hyperparameters will be calculated based on the elicited values below. To view or directly specify the hyperparameters, click the "Details" button below.

Details...

	Elicited mean P(T) 'm_hat T'	Elicited mean P(E) 'm_hat E'
Dose 1	.3000	.0500
Dose 2	.3500	.1000
Dose 3	.4000	.1500
Dose 4	.4500	.2000

When the first simulation is run, the software will compute the model hyperparameters that best satisfy the elicited values in a least squares sense. The optimization process for determining the hyperparameters also sees that the variance is sufficiently large and roughly evenly distributed among the model parameters. If you would like to see the hyperparameters or manually specify them, click on the **Details** button. This brings up the following dialog.



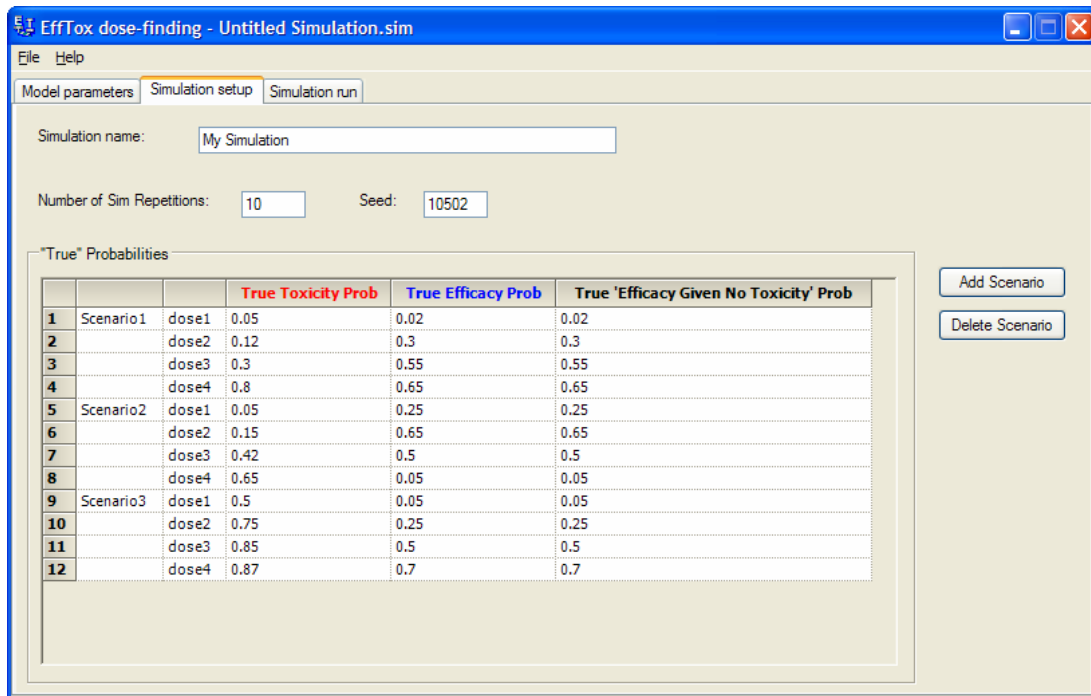
Press the **Calculate** button to see the hyperparameters derived from the elicited values. This calculation may take a little while to complete. Press **Done** to close this dialog and accept the displayed values.

The hyperparameters may be directly edited. If the hyperparameters do not match the values that would result from running the calculator using the elicited values as input, the software will use the last item edited. This means that if the elicited values were last edited, the hyperparameters will be recalculated at simulation time, and if the hyperparameters were the last edited values, those hyperparameter values will be used in the calculations. Both elicited values and hyperparameter values are saved to the input file. To exit the dialog without changing the hyperparameters or elicited values, press **Cancel**.

A useful diagnostic for the quality of the hyperparameters is the prior probabilities of toxicity and efficacy associated with each dose. These are printed out as part of the simulation results under the heading "Prior probabilities (hyperparameter diagnostic)".

## **Simulation setup**

Next, go to the **Simulation setup** tab.



The **Simulation name** is for your own documentation.

**Number of Sim Repetitions** is the number of times you would like each of the scenarios to be repeated. You may want to run a small number of repetitions (say 10) followed by a larger number (such as 100 or 1000). The small set will return quickly and may point out gross problems. If the small set looks feasible, then run a larger set to reduce the variance of the runs and obtain results that can be examined more closely.

**Seed** is the seed for the random number generator that is used to produce the simulated data. This can be any positive integer.

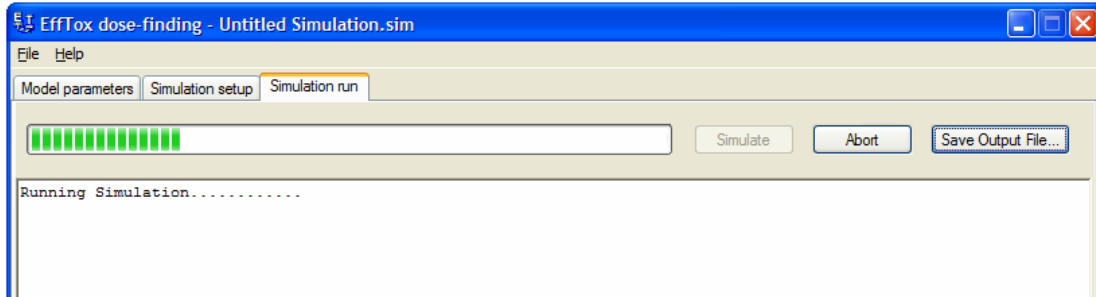
The **Add Scenario** button adds a row to the **True Probabilities** grid for each dose. For the bivariate binary model, there are columns for the probability of toxicity, the probability of efficacy, and the probability of efficacy given that toxicity did not occur. For simplicity, you may want to make the last two columns equal, unless you have a strong belief that efficacy and toxicity are correlated. The trinary model only has columns for probabilities of toxicity and efficacy since the probability of efficacy *given no toxicity* is built into the trinary model.

To delete a scenario, select one of the rows in that scenario and click the **Delete** button.

Finally, go to the **Simulation run** tab.

Clicking the **Simulate** button will start the simulations running. Depending on the number of simulation repetitions and the number of scenarios, it could take a long time for the results to be displayed. The progress bar near the top will provide an indicator that

the calculations are still running. The **Abort** button allows you to terminate the simulation process; partial results will then be displayed.

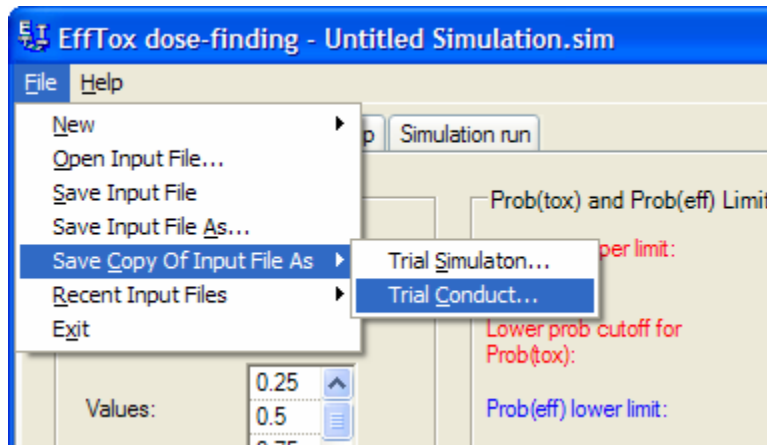


When the simulation is complete, the results will appear on this tab.

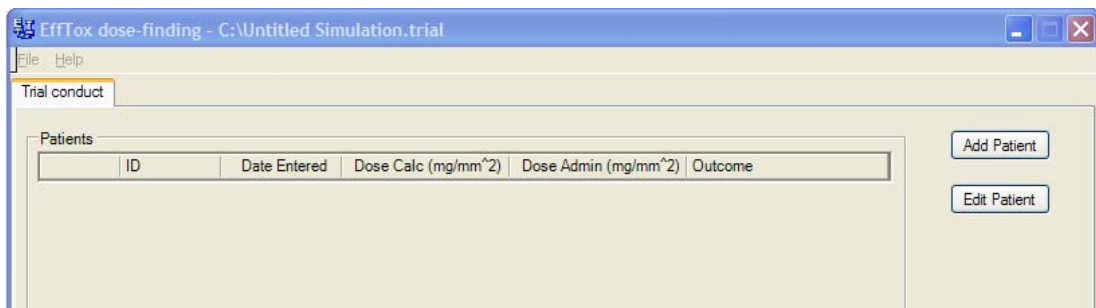
The **Save** button allows you to save the simulation output.

## ***Trial Conduct***

Once you are satisfied with the operating characteristics of a design from having studied the simulation results, you may save the simulation as a trial conduct file. Then you may use the software to run a hypothetical trial.

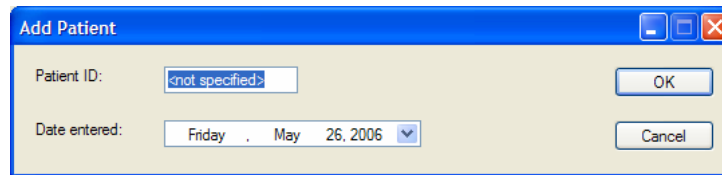


Opening a trial conduct file brings up the **Trial conduct** tab:



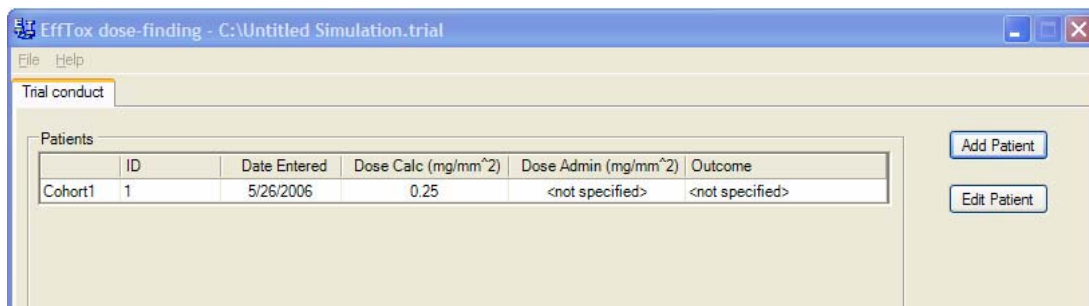
Patients are treated in a three-step process: enroll, assign dose, record outcomes.

Clicking the **Add Patient** button brings up this dialog to enroll a patient:



The **Patient ID** can be any alpha-numeric identifier. The **Patient ID** and **Date entered** are not used in the EffTox calculations but are there for convenience.

The patient then appears in a grid and the recommended dose is given in the **Dose Calc** column, 0.25 mg/mm<sup>2</sup> in this example.



	ID	Date Entered	Dose Calc (mg/mm <sup>2</sup> )	Dose Admin (mg/mm <sup>2</sup> )	Outcome
Cohort1	1	5/26/2006	0.25	<not specified>	<not specified>

The software does not assume that the recommended dose is administered. The recommended dose would nearly always be the administered dose, but one can experiment to see the effect of deviations from protocol. (If a patient is given a dose other than the recommended dose, either intentionally or unintentionally, the patient's data is still valuable and used in future dose decisions.)

To enter the administered dose for a patient, select that patient's row and click **Edit Patient**.

The screenshot shows the 'EffTox dose-finding' application window. It features a menu bar with 'File' and 'Help'. Below the menu is a 'Trial conduct' tab. The main area contains a 'Patients' section with a table of patient data. To the right of the table are two buttons: 'Add Patient' and 'Edit Patient'.

	ID	Date Entered	Dose Calc (mg/mm <sup>2</sup> )	Dose Admin (mg/mm <sup>2</sup> )	Outcome
Cohort1	1	5/26/2006	0.25	0.25	Toxicity
	2	5/26/2006	0.25	0.25	Neither
	3	5/27/2006	0.25	0.5	Both
Cohort2	4	5/27/2006	0.75	0.75	Both
	5	5/28/2006	0.75	0.75	Both
	6	5/28/2006	0.75	0.75	Both
Cohort3	7	6/4/2006	0.25	<not specified>	<not specified>


The 'Edit Patient' dialog box is shown. It includes fields for 'Patient ID' (set to 1) and 'Date entered' (set to Friday, May 26, 2006). There are 'OK' and 'Cancel' buttons. The 'Outcome' section has five radio button options: '(not determined yet)', 'Toxicity, without Efficacy', 'Neither Efficacy nor Toxicity', 'Both Efficacy and Toxicity', and 'Efficacy, without Toxicity'. The 'Dose values' section has a 'Calculated per protocol' field (0.25) and an 'Administered' dropdown menu. The dropdown menu is open, showing options: '<not specified>', '0.25', '0.5', '0.75', and '1'. The '0.25' option is currently selected.

Once the dose is selected, click OK.

To record a patient's outcomes, select that patient's row and again click the Edit Patient button.

### ***Troubleshooting***

The EffTox program has a number of validation rules built in to reduce the likelihood of error. The general method for indicating a violation of validation rule is a flashing icon, a red dot with a white exclamation point inside. Placing the mouse over the validation icon produces a tool tip which indicates the rule that has been violated. For example, the following show the result of entering 1.4 in a box expecting a probability.



In addition to the flashing icon, the software will generally not let you go on to another field until the error is corrected: If you're unable to go on to the next thing you want to do, look for a flashing red dot and mouse over it.