

University of Texas MD Anderson Cancer Center
Department of Biostatistics

Multc Lean Desktop, Version 2.1
July 17, 2013
User's Guide

0. Overview

Multc is software for designing single-arm clinical trials based on comparing multiple events or outcomes. *Multc Lean Desktop* is a simplified version for Microsoft Windows that considers only two events: response and toxicity. The method is based on the following paper:

Peter F. Thall, Richard M. Simon, and Elihu H. Estey,
Bayesian sequential monitoring designs for single-arm clinical trials with multiple outcomes,
Statistics in Medicine, vol 14, 357-379 (1995).

Multc also allows a modification to the method in the paper; instead of modeling the standard treatment event rate as being a random variable with a beta distribution, it may also be modeled as a constant rate.

The software may be downloaded from the web site:

<https://biostatistics.mdanderson.org/SoftwareDownload/>

This software occasionally sends usage statistics and crash reports to our biostatistics software support team to improve your experience using it.

License:

This program is distributed at no cost to the user. However, redistribution of this program is not permitted. Rather than giving this program to someone else, please have them obtain their own copy directly from MD Anderson Cancer Center. This allows us to keep a record of who is using the software and allows us to notify all users when program enhancements become available.

NO WARRANTY

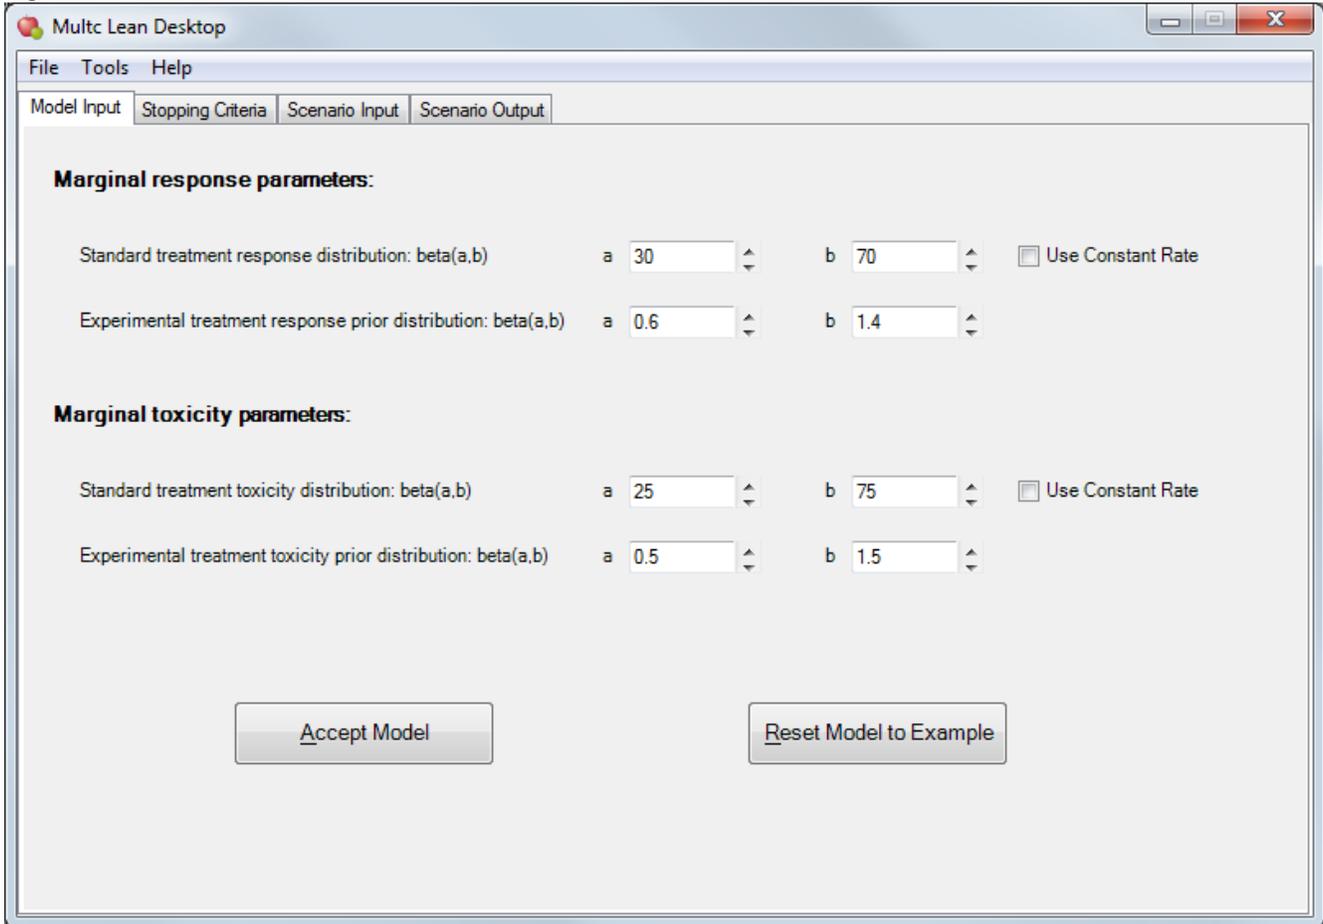
We provide absolutely no warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. The entire risk as to the quality and performance of the program is with the user. Should this program prove defective, the user assumes the cost of all necessary servicing, repair, or correction.

In no event shall The University of Texas or any of its component institutions including MD Anderson Cancer Center be liable for damages, including any lost profits, lost monies, or other special, incidental or consequential damages arising out of the use of or inability to use the program (including but not limited to loss of data or its analysis being rendered inaccurate or losses sustained by third parties).

1. Model Input

The program begins by displaying the first of four tab pages, labeled “Model Input” (Figure 1). The user defines a model by entering parameters into the text boxes on the tab page. The upper numbers are for comparing the response rate of an experimental treatment versus a standard treatment, while the lower numbers are for comparing toxicity between the treatments.

Figure 1:



1.1 Marginal response parameters

The probability of response for the standard treatment is modeled either by a random variable with a beta distribution $beta(a,b)$ or by a constant rate c ; the choice to use a constant instead of a random variable with a beta distribution is made by checking the “Use Constant Rate” checkbox directly to the right of the standard treatment response rate model specification. The $beta(a,b)$ parameters a and b are often based on historical data. When they are, typically the first parameter a is given by the number of responses observed with the standard treatment, while the second parameter b is the number of non-responses. These values are sometimes reduced proportionally while maintaining the same beta distribution mean in order to discount the data by increasing the variance of the distribution. These values may also be obtained by other methods, such as solving for them from a desired mean and variance or pair of quantile values using the Parameter Solver tool available from the Tools drop-down menu. The values of a and b must be positive (greater than 0) and less than or equal to 1,000. The value of the constant rate c must be between 0 and 1, inclusive.

The prior probability for the experimental treatment response is modeled by a beta distribution $beta(a,b)$. Since the prior distribution should not be too informative, these a and b parameters should not be given large values. Often a and b are chosen such that their sum $a + b$ is 2 and the mean of the beta distribution, $a / (a + b)$, is the same as either the mean of the beta distribution, or the constant rate, of the standard treatment response rate model. The values of a and b must be positive (greater than 0) and less than or equal to 100.

1.2 Marginal toxicity parameters

The probability of toxicity is modeled in a way that is similar to that of the probability of response. If the standard treatment toxicity rate is chosen to be modeled by a $beta(a,b)$ distribution, historical data for toxicity with the standard treatment is often used to determine the a and b parameters. In this case, the a parameter is given by the number of observed toxicities with the standard treatment, while b is given by the number of non-toxicities. As with the standard treatment response rate, the values of the a and b parameters of the beta distribution are sometimes reduced proportionally to discount the data, or are obtained by other methods. The standard treatment toxicity rate may also be modeled by a constant rate; this is chosen by checking the “Use Constant Rate” checkbox directly to the right of the standard treatment toxicity rate model specification. The values of a and b must be positive (greater than 0) and less than or equal to 1,000. The value of the constant rate c must be between 0 and 1, inclusive.

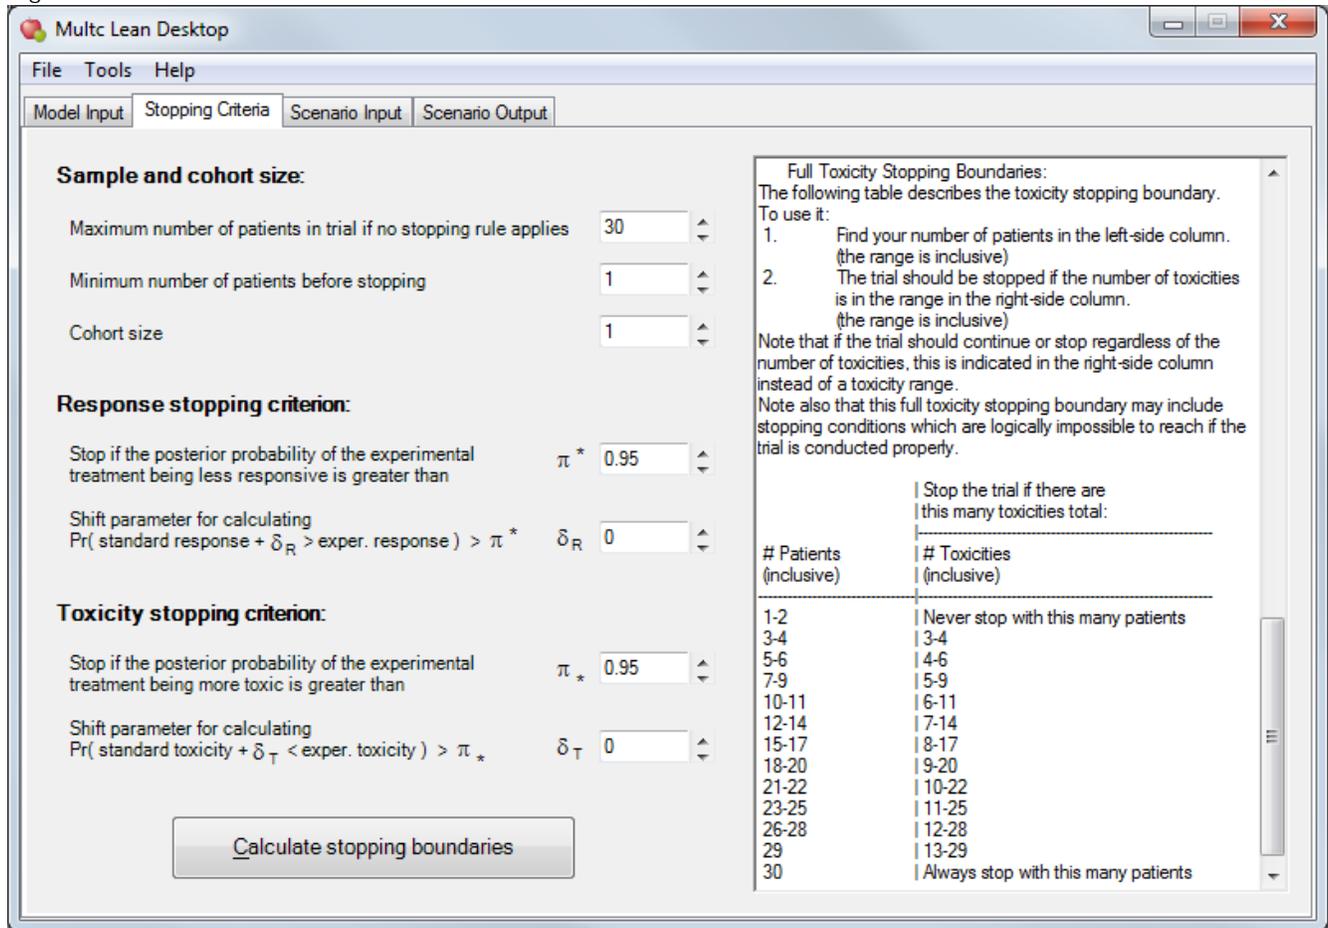
As with response, the prior probability for the experimental treatment toxicity is modeled by a beta distribution $beta(a,b)$. The values of the a and b parameters of the toxicity prior distribution are often obtained in the same way as those for the response prior distribution, but such that the beta distribution mean is the same as either the mean of the beta distribution, or the constant rate, of the standard treatment toxicity rate model. The values of a and b must be positive (greater than 0) and less than or equal to 100.

When all of the model parameters have been defined, the user should click on the “Accept Model” button to validate the parameters and move on to the next tab page. To reset the model and the stopping criteria to match the first example in this document, the user can click the “Reset Model to Example” button.

2. Stopping Criteria

The second tab page, labeled “Stopping Criteria”, (Figure 2a) is for specifying when a trial should be stopped. In general, there are three ways a trial can be stopped: (1) when the desired total number of patients has been reached, (2) when the response rate of the experimental treatment does not compare favorably to that of the standard treatment, and (3) when the toxicity rate of the experimental treatment is too high compared to that of the standard treatment.

Figure 2a:



2.1 Sample and cohort size

The user must specify the maximum number of patients to enroll if the trial does not stop early due to low response or high toxicity. This value must be between 3 and 1,000, inclusive.

It is possible to set a minimum number of patients to enroll before evaluating the stopping rule. This minimum must be less than or equal to the maximum number of patients and a multiple of the cohort size, or less than the cohort size. This value must be between 1, the default, and 1,000, inclusive.

The cohort size defaults to 1, and may be any integer factor of the maximum number of patients. In other words, the maximum number of patients must be a multiple of the cohort size. The program computes only stopping boundaries that are some multiple of this cohort size. This value must be between 1, the default, and 1,000, inclusive.

2.2 Response stopping criterion

The response stopping rule says that the trial should be stopped if the posterior probability of the experimental treatment being less responsive than the standard treatment is greater than some value π^* . Typically, π^* is set between 0.90 and 0.99, with 0.95 being the default value. This value must be between 0 and 1 inclusive. A value of 1 means the trial will never be stopped for response. A value of 0 means the trial will be stopped for response in all circumstances, in other words it makes the response requirement for continuing the trial impossible to satisfy.

By default, the probability of being less responsive is the result of comparing directly the posterior beta distribution of the experimental treatment response rate to the standard treatment response rate, which is specified on the “Model Input” tab page to either be modeled as a random variable with a beta distribution or to be modeled a constant. An alternative is to shift the standard treatment response rate by some amount before comparison. That is, the trial should stop when

$$\Pr(\text{standard response} + \delta_R > \text{experimental response}) > \pi^*.$$

By setting the parameter δ_R to a positive value (typically ≤ 0.1), the user requires the experimental treatment to have a response rate that improves by at least δ_R over the standard response rate. In some cases, the user may set δ_R to a negative value to allow the experimental treatment to be slightly less responsive than the standard treatment, hoping that it also leads to a significantly lower toxicity rate. Setting δ_R to 0, the default, is equivalent to not using it. The value of δ_R must be greater than -1 and less than 1.

2.3 Toxicity stopping criterion

The toxicity stopping rule says that the trial should be stopped if the posterior probability of the experimental treatment being more toxic than the standard treatment is greater than some value π_* . Typically, π_* is set between 0.90 and 0.99, with 0.95 being the default value. This value must be between 0 and 1 inclusive. A value of 1 means the trial will never be stopped for toxicity. A value of 0 means the trial will be stopped for toxicity in all circumstances, in other words it makes the toxicity requirement for continuing the trial impossible to satisfy.

The user may specify a shift parameter δ_T before comparing the posterior beta distribution of the experimental treatment toxicity rate to the standard treatment toxicity rate, which is specified on the “Model Input” tab page to either be modeled as a random variable with a beta distribution or to be modeled as a constant rate, so that the trial will stop when

$$\Pr(\text{standard toxicity} + \delta_T < \text{experimental toxicity}) > \pi_*.$$

Setting the shift parameter δ_T to a negative value requires the toxicity rate of the experimental treatment to be lower than the standard treatment. Alternatively, the value of δ_T may be positive, implying that the experimental treatment is allowed to have a slightly higher toxicity rate, hoping that it also means a higher response rate. Setting δ_T to 0, the default, is equivalent to not using it. The value of δ_T must be greater than -1 and less than 1.

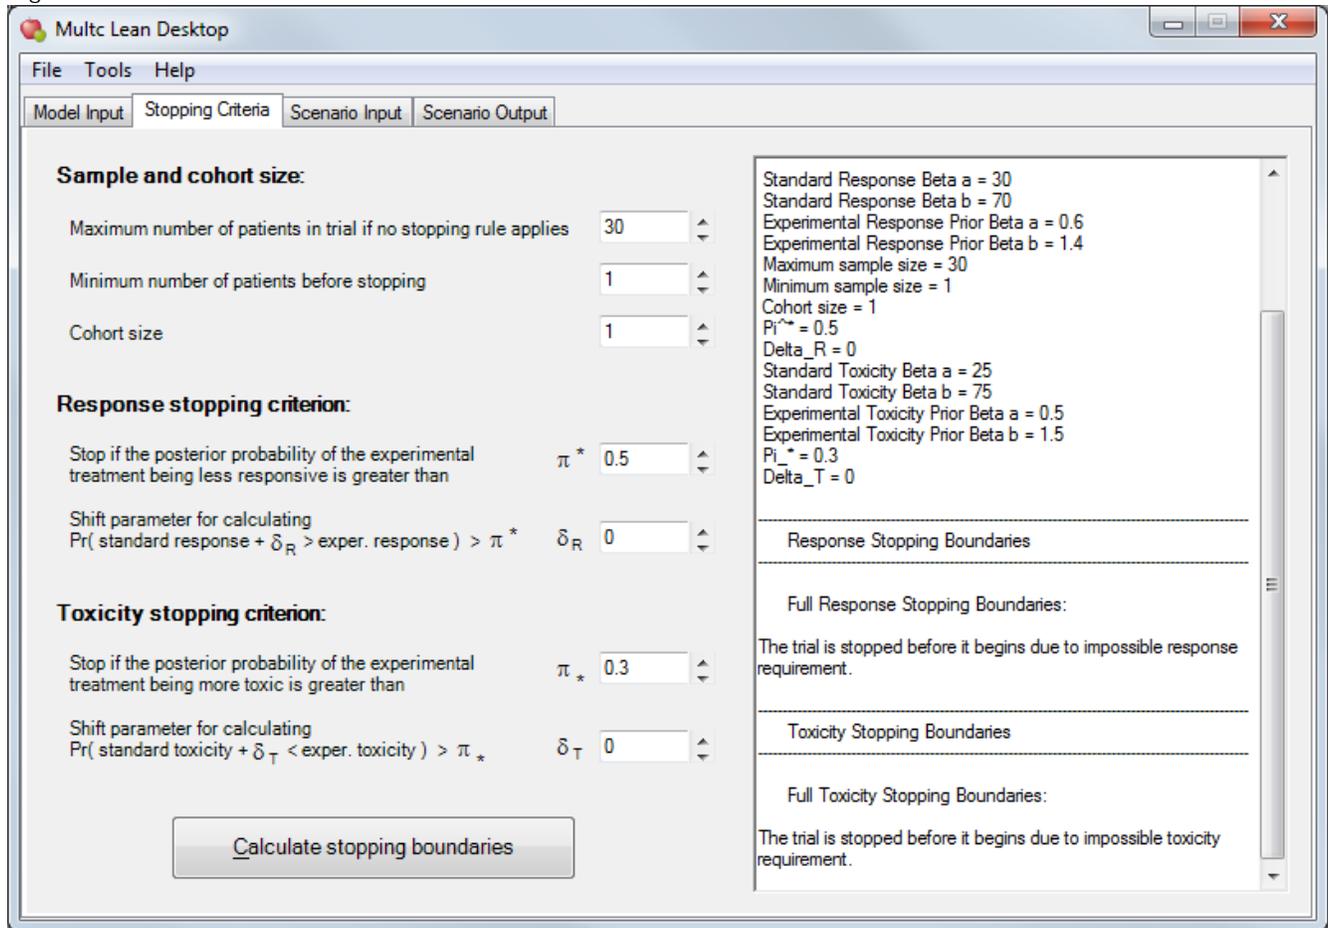
The program does not permit the sign combination $\delta_R < 0$ and $\delta_T > 0$, because the experimental treatment is allowed to be worse on either the response or toxicity criterion, but not both.

2.4 Calculations

When all parameters have been set, the user should click on the “Calculate stopping boundaries” button to continue. The program validates all model and stopping parameters before calculating the response and toxicity stopping boundaries. For detailed information on the statistical theory behind the calculations, please refer to the document *Multic Statistical Tutorial*, available under the Help menu.

In some extreme cases, a requirement on the experimental treatment may be so stringent that the stopping rule would kick in before any patients have been enrolled. In the example shown in Figure 2b, we set $\pi^* = 0.5$ and $\pi_* = 0.3$, making it impossible to satisfy either criterion before the trial starts. This is denoted by the statement that the trial is stopped before it begins, instead of a boundary.

Figure 2b:



2.5 Alternative output style for stopping boundaries

With version 2.0 of Multc Lean Desktop, the stopping boundaries are displayed as tables giving the stopping criteria for all numbers of patients in multiples of the cohort size up to the maximum sample size. Older versions of Multc Lean Desktop displayed only the stopping criteria which are logically achievable during the conduct of a trial, skipping the criteria for numbers of patients at which it is not possible to stop the trial if it is conducted properly and it has not stopped previously. To have this older style of output displayed in addition to the newer style, please see Appendix A.

3. Scenario Input

The third tab page is for setting up and running scenarios. For each scenario, the user specifies some “true” probabilities of response and toxicity for the experimental treatment, then runs the scenario to see how likely it is that the trial will stop at various sample sizes. The scenarios all assume the same underlying model input parameters and stopping criteria that have been entered in the first two tab pages. There is also an option to simulate the trial to estimate the average trial duration.

The user has a choice of specifying either dependent response and toxicity (Figure 3a), or independent response and toxicity (Figure 3b). In the latter case, the program automatically calculates the equivalent dependent probabilities.

3.1 Dependent response and toxicity

The dependency between response and toxicity is expressed in terms of four probability values that must each be between 0 and 1 inclusive and that must sum up to one:

- (1) probability of response and toxicity,
- (2) probability of response and no toxicity,
- (3) probability of no response and toxicity, and
- (4) probability of no response and no toxicity.

After entering these four values, the user should click on “Add Scenario” to add the set of parameters to the table on the right. To remove a scenario from the table, the user should click on the left-most column to select the row and then click on “Delete Scenario”. The values in the table also may be edited directly.

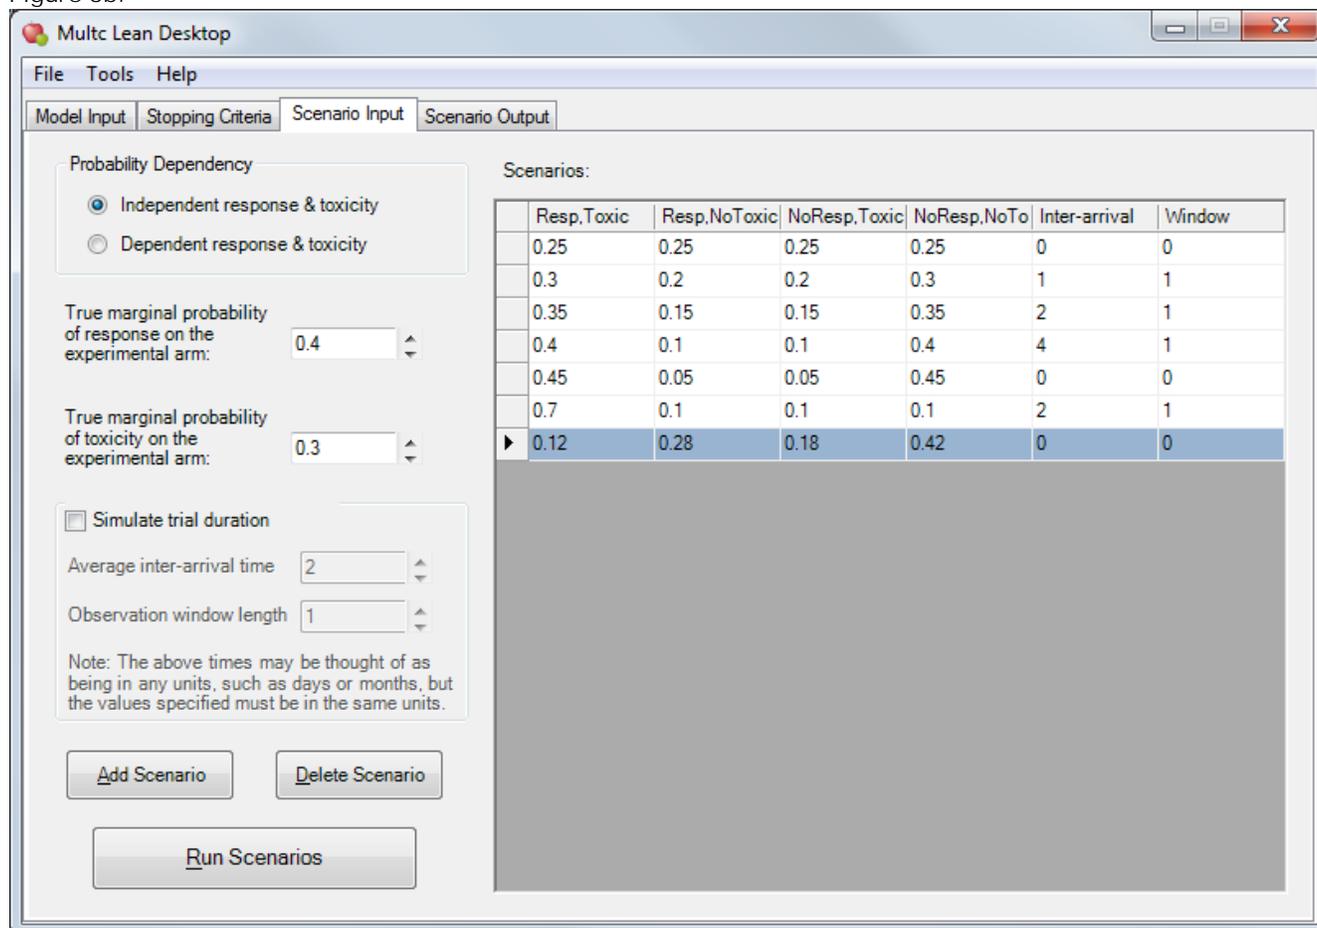
Figure 3a:

	Resp,Toxic	Resp,NoToxic	NoResp,Toxic	NoResp,NoTo	Inter-arrival	Window
	0.25	0.25	0.25	0.25	0	0
	0.3	0.2	0.2	0.3	1	1
	0.35	0.15	0.15	0.35	2	1
	0.4	0.1	0.1	0.4	4	1
	0.45	0.05	0.05	0.45	0	0
▶	0.7	0.1	0.1	0.1	2	1

3.2 Independent response and toxicity

When the probabilities of response and toxicity for the experimental treatment are assumed to be independent, the user needs to enter only two parameters for each scenario. These are the true marginal probability of response and the true marginal probability of toxicity. These values must be between 0 and 1 inclusive. The program automatically calculates the equivalent set of four dependent probabilities when the “Add Scenario” button is clicked. The example in Figure 3b shows that the two true marginal probabilities 0.4 and 0.3 have been converted to the probability set (0.12, 0.28, 0.18, 0.42).

Figure 3b:



3.3 Simulate trial duration

For any scenario, the probabilities of stopping at the various boundaries can be calculated directly without simulation since the model is simple enough. However, the program does use Monte Carlo simulation to estimate the average trial duration if desired. In this case, two additional input time parameters are required: the average inter-arrival time and the observation window length. These are displayed in the last two columns of the scenario table. These time parameters have no explicit units such as days or months, so the user is free to enter them using whatever units are most convenient. However, the values specified must be in the same units for the simulation results to be meaningful, for example both in terms of days or both in terms of months. The average trial duration reported will then be in terms of the same units, for example days or months.

In order to request that the average trial duration be determined by simulation, the user checks the 'Simulate trial duration' check box. This enables the time parameter specification boxes so that these parameters can be entered. The user then specifies both time parameters, which must be positive (greater than 0) and less than or equal to 10,000. When the "Add Scenario" button is pushed, if the 'Simulate trial duration' check box is checked, these values will be copied into the last two columns of the scenario table. The average trial duration will be determined for scenarios for which the values in the last two columns of the table, the time parameters, are both greater than 0. If the user unchecks the 'Simulate trial duration' checkbox (or leaves it unchecked), the time parameter specification boxes will be disabled, and when the 'Add Scenario' button is pushed values of 0 will be copied into the last two columns of the scenario table, indicating that the average trial duration will not be determined by simulation for that scenario.

The values of these time parameters may be edited directly in the scenario table. If so, both of these parameters must be greater than or equal to 0 and less than or equal to 10,000. If either of these parameter values in the table is 0, the average trial duration will not be determined by simulation.

When both of these time parameters are set to positive values in the table, the average trial duration simulation option is turned on automatically for that scenario. The program simulates 10,000 trials for each scenario, and uses a random seed based on the current time. The average trial duration which is determined will be displayed along with other information in that scenario's output.

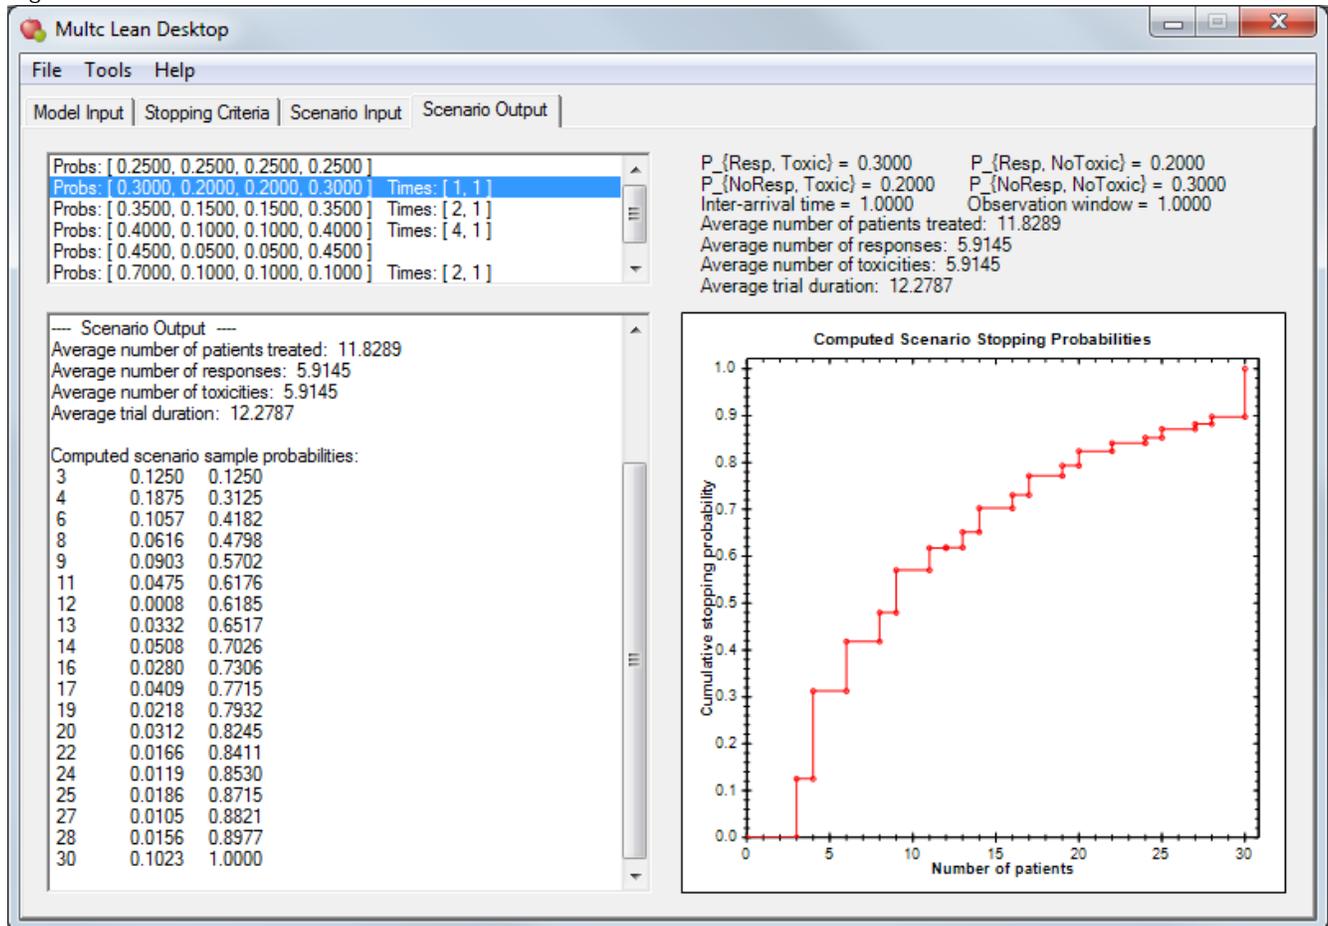
The program models the time interval between two consecutive patient arrivals by an exponential distribution. For each scenario, the mean of this exponential distribution is set equal to the average inter-arrival time parameter.

The observation time for each patient depends on whether there is a response for that patient. For a non-response, the observation time is set equal to the observation window length. For a positive response, the observation time is modeled by a truncated exponential. The mean of this exponential is set so that the cumulative probability of the distribution is 95% at the observation window length.

4. Scenario Output

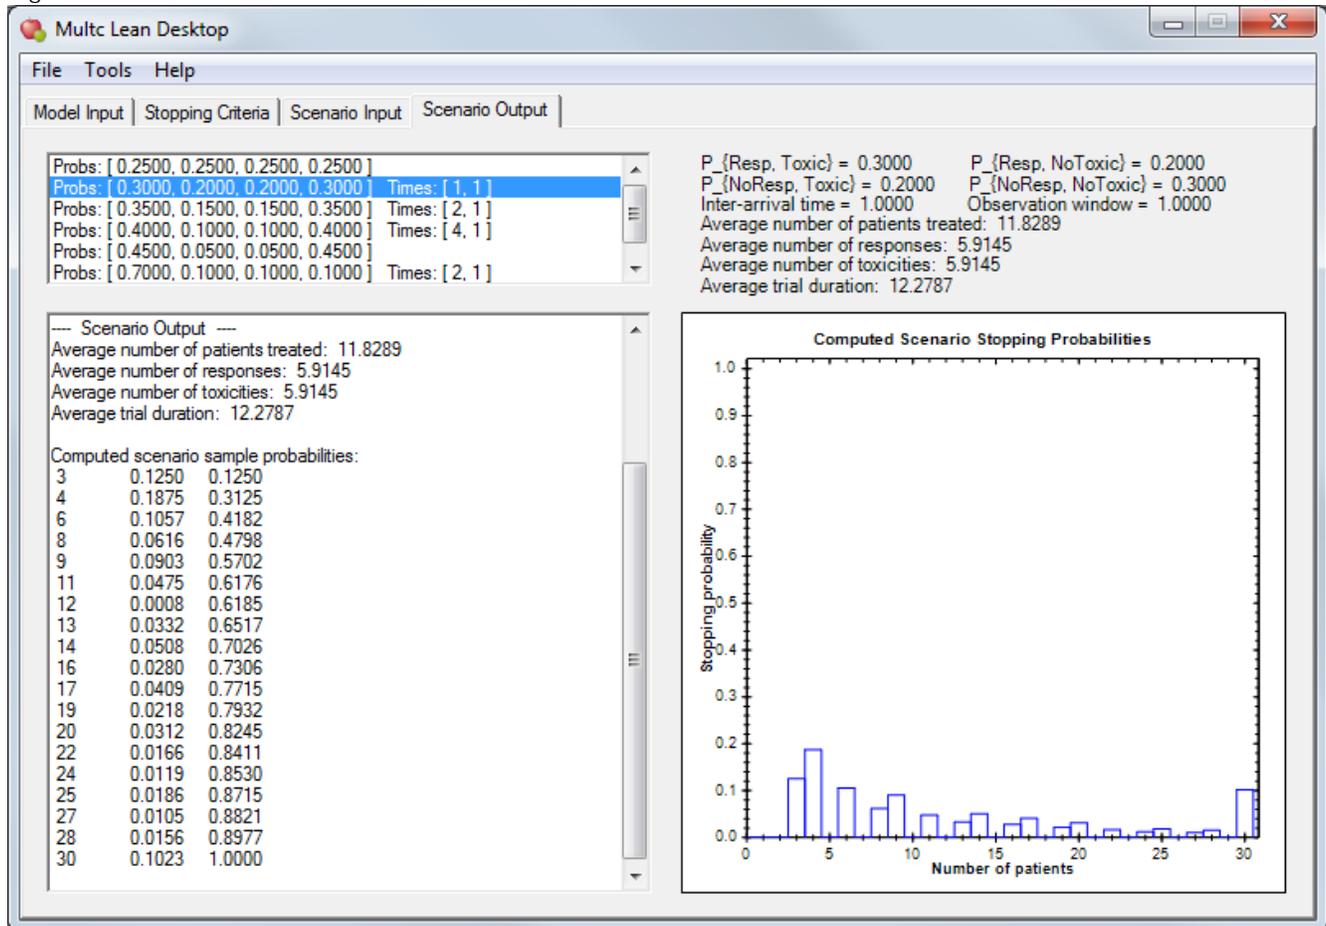
The last tab page (Figure 4a) displays the results from running the scenarios. The list box at the upper left section of the page displays all the scenarios in the same order as the table on the Scenario Input tab page. Each scenario is identifiable by its four probability parameters and the two optional time parameters. When a scenario in the list box is selected, the upper right section of the page shows the corresponding input and output parameters. In addition, the large scrollable text box at the lower left displays the stopping probabilities for that scenario in three columns. The first column is the sample size, the second column is the probability of stopping at that sample size, and the third column is the cumulative probability of stopping at or before that sample size. These values are computed analytically, based on the four input probability parameters. The average number of patients, responses and toxicities in a trial are computed analytically from the stopping probabilities.

Figure 4a:



The graph at the bottom right plots the cumulative stopping probability as a function of the sample size. Another option is the density plot (Figure 4b), which can be toggled with the cumulative plot by first clicking on the graph with your mouse, then pressing the “C” or “D” key. The user may copy the plot to the clipboard or save it to a file (.PNG format) by right clicking on the plot and choosing “Copy” or “Save” from the menu.

Figure 4b:



5. Miscellaneous

The File menu has a “Save Scenario Output File ...” option that prompts for a file name to save all model and scenario input parameters and scenario results. Before saving the scenario output file, the scenario results must be up-to-date with respect to the inputs. In order to make them up-to-date, at least one scenario needs to be defined and run. If the scenario results are not up-to-date and the user attempts to save the scenario output file, the user will be prompted to run the scenarios. After the scenarios have been run and are up-to-date, when the user attempts to save the scenario output file he or she will be prompted for a file name. The file is in HTML format, allowing it to be read by a variety of software. The program also automatically saves the current model and scenario input parameters upon exit and restores them at the next session. The scenario results are not saved automatically, since they can be regenerated easily.

The Tools menu contains the option to run the Inequality Calculator program for computing the probability inequalities between two distributions. Another option is the Parameter Solver program for estimating the parameters of a distribution when given some input quantiles, or input mean and variance.

The Help menu has options to open the User’s Guide (this document), the Statistical Tutorial, and an Example Protocol Memo. Opening the User’s Guide or Statistical Tutorial requires that a PDF file viewer be installed on the user’s computer, such as the free Adobe Reader available from

<http://get.adobe.com/reader/> . Opening the Example Protocol Memo requires that Microsoft Word or another program which can open Microsoft Word documents be installed on the user's computer. The Example Protocol Memo is intended to help those who prepare statistical descriptions of protocols designed with Multc Lean Desktop. To edit the memo, first save it under a different file name in a convenient location.

The Help menu also contains a View License option to view the license under which Multc Lean Desktop is distributed, and an About Multc Lean Desktop option to see details about the software. The About Multc Lean Desktop option opens a window containing contact information for feedback, including a link which when clicked will automatically start an email message for the user using an email client installed on the computer. If no email client is installed on the computer, the user may alternatively send feedback via email to BiostatisticsSoftwareSupport@mdanderson.org . If so, we request that the email refer to Multc Lean Desktop and its version number.

Each numerical input field is equipped with an error indicator. When the symbol  is displayed, its adjacent field contains an invalid entry. By moving the mouse cursor over this symbol, a tip will be displayed as shown in Figure 5.

Figure 5:



Appendix A

Enabling Alternative Stopping Boundary Output

With version 2.0 of Multc Lean Desktop, the stopping boundaries are displayed as tables giving the stopping criteria for all numbers of patients in multiples of the cohort size up to the maximum sample size. We refer to this as the “full boundary” since it describes the complete boundary for all numbers of patients in complete cohorts.

Older versions of Multc Lean Desktop displayed only the stopping criteria which are logically achievable during the conduct of a trial, skipping the criteria for numbers of patients at which it is not possible to stop the trial if it is conducted properly and it has not stopped previously. We refer to this as the “potential boundary” since it describes exactly the conditions under which the trial can in fact potentially stop.

To enable the output of the potential boundaries in addition to the full boundaries, the configuration file for Multc Lean Desktop should be edited, according to the following procedure. Note that depending on the location of the Multc Lean Desktop installation folder, editing the configuration file may require administrative privileges on the computer.

First the file `MultcLeanDesktop.exe.config` should be located. It is placed in the Multc Lean Desktop installation folder when Multc Lean Desktop is installed. Depending on the Windows operating system of the computer on which it was installed, this folder is usually either:

`C:\Program Files (x86)\MD Anderson Cancer Center\Multc Lean Desktop\V2.1.0\`

Or

`C:\Program Files\MD Anderson Cancer Center\Multc Lean Desktop\V2.1.0\`

However, it may be different depending on the folder which was specified during installation.

Next, the file `MultcLeanDesktop.exe.config` in the installation folder should be edited. Note that saving the edited file may require administrative privileges on the computer. The section in the file, which specifies whether to output the potential boundaries or not, should be located. It is:

```
<setting name="boOutputPotentialStoppingBoundariesAlso" serializeAs="String">  
  <value>False</value>  
</setting>
```

The value in the middle line of this section, which is `False` above, should be changed to `True`:

```
<setting name="boOutputPotentialStoppingBoundariesAlso" serializeAs="String">  
  <value>True</value>  
</setting>
```

If Multc Lean Desktop is currently running, it should be closed and a new session should be started. After this is done, when the stopping boundaries are calculated the potential boundary will be displayed immediately following each full boundary.