

University of Texas M. D. Anderson Cancer Center
Department of Biostatistics

Adaptive Randomization Version 5.2

User's Guide

0. Overview

Adaptive Randomization is software for simulating clinical trials using outcome-based adaptive randomization to assign treatments. The program is designed for the Microsoft Windows XP operating system.

The software may be downloaded from the web site:
<http://biostatistics.mdanderson.org/SoftwareDownload/>

We recommend users refer to the technical report *Power and Bias in Adaptively Randomized Clinical Trials* by J. Kyle Wathen and John D. Cook, for more background on adaptive randomization.

http://www.mdanderson.org/pdf/biostats_utmdabtr_002_06.pdf

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 M. D. 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.

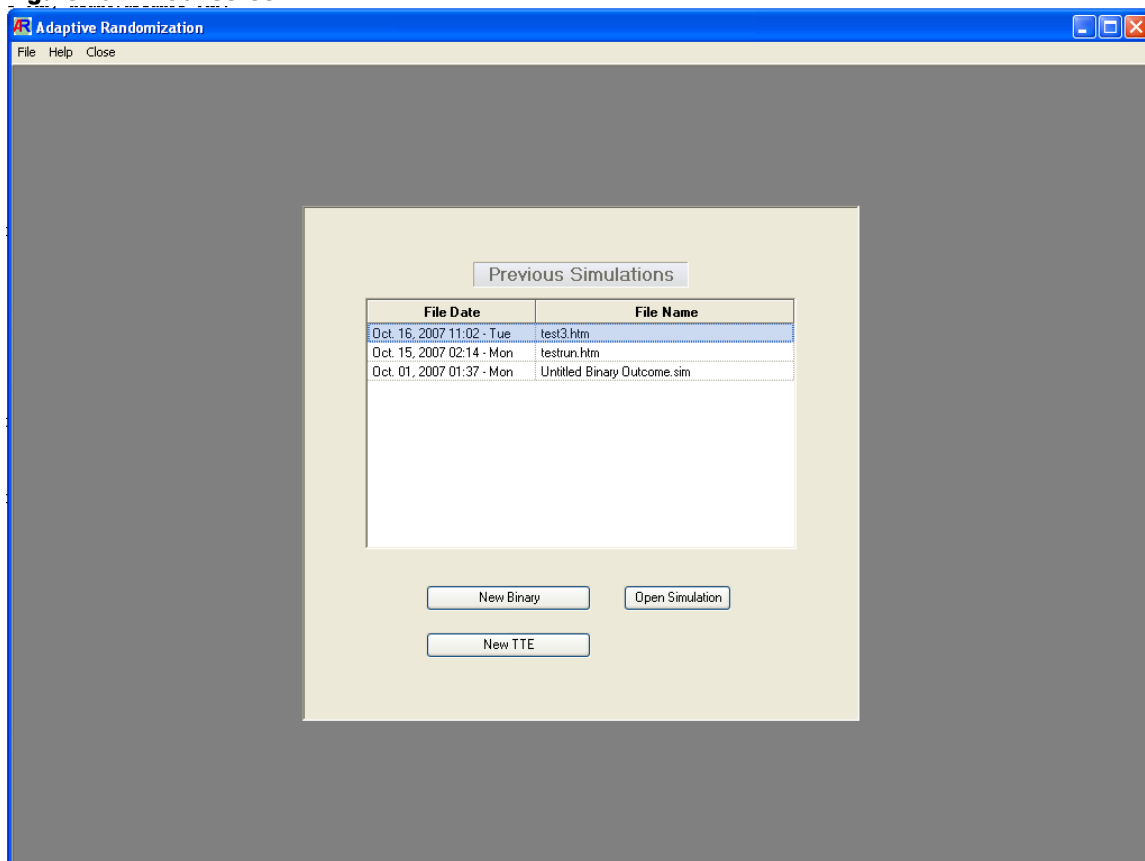
In no event shall The University of Texas or any of its component institutions, including M. D. Anderson Cancer Center, be liable for damages, including any lost profits, lost monies, or other special, incidental or consequential damages arising out of the use of or inability to use the program (including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

1. Introduction

The *Adaptive Randomization* program enables the simulations of clinical trials using a randomization method that depends on the latest available outcomes from the trial. The main idea is to unbalance the randomization probabilities to favor the treatment that has superior interim results. The simulated trial may use either a binary outcome model or a time-to-event model to determine which is the better treatment. As each patient enters a trial, the program considers all currently available data and compares the posterior distribution of a key model parameter (e.g., the response rate or the average time to toxicity) of each arm. The posterior probability of an arm's parameter being the best among all arms is used to set the relative chance of assigning the new patient to that arm.

When the program starts up for the first time, the user is given the choice of creating a new binary event trial, or a time-to-event trial (Figure 1a).

Figure 1a: Initial screen



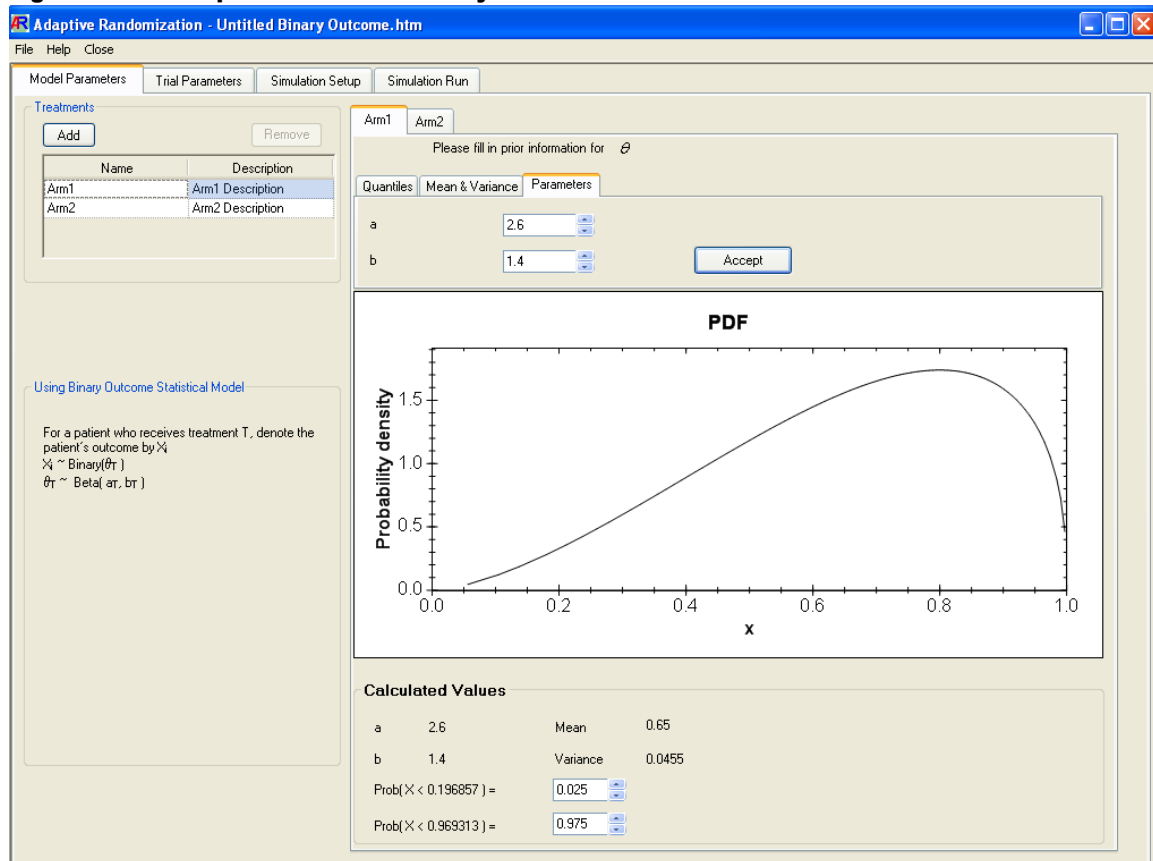
After setting up a simulation, the user can save all parameters to a file with the Save As command under the File menu. The “Open Simulation” button or the File/Open menu command allows reloading of the saved simulations.

2. Model Parameters

After a model type has been selected, the program displays the first tab page, labeled “Model Parameters” (Figure 2a). The upper left section of the page contains a list of the treatment arms, while the lower left section displays a short description of the statistical model.

The model parameter θ represents the key variable for deciding which treatment to assign during a trial, and which treatment to select as the best at the end of a trial. In a binary-event trial, the parameter θ typically represents a response rate or a toxicity rate (a value between 0 and 1.) In a time-to-event trial, θ typically represents the average or median time until response or toxicity. For both model types, the program determines the relative merit of the treatment arms by comparing their posterior distributions of θ . These posterior distributions are computed by the Bayesian method, where the prior distribution of θ for each arm is modified by the likelihood of the most recently available data.

Figure 2a: Model parameters for a binary event trial



The prior distribution of θ must be specified for each treatment arm before a simulation can be run. The program allows specification of these distributions for both model types in a consistent way, via the large tab control on the right side. The three options (tabs) available for specifying each distribution are:

- (i) Quantiles: The distribution is specified by two cumulative probability values.
- (ii) Mean & Variance: The distribution is specified by its desired mean and variance.
- (iii) Parameters: The distribution is specified by its two standard parameters.

After entering values into the appropriate fields for each arm, the user should click on the “Calculate” button to set the distribution for that arm.

2.1 Binary Model

In a binary-event trial, each patient’s outcome is assigned one of two possible values (e.g., response or non-response.) The probability of response (or toxicity) for a patient who receives treatment k is given by the unknown response rate of that treatment, designated as θ_k .

At the beginning of the trial, the uncertainty in the value of the response rate θ_k of each treatment arm is modeled by a prior beta distribution. That is, $\theta_k \sim \text{beta}(a_k, b_k)$, with positive parameters a_k and b_k . These parameters can be interpreted as the amount of prior knowledge about each arm. For example, if we are monitoring the response rate of treatment k , then the parameter a_k corresponds to the number of prior responses while b_k represents the number of prior non-responses. It is suggested that these values be set from historical data, suitably down-weighted, if available.

If the trial data for treatment k consists of p_k positive responses and n_k negative responses, then the posterior distribution is updated to

$$\theta_k | \text{data} \sim \text{beta}(a_k + p_k, b_k + n_k)$$

since the beta prior distribution is conjugate for the binomial likelihood.

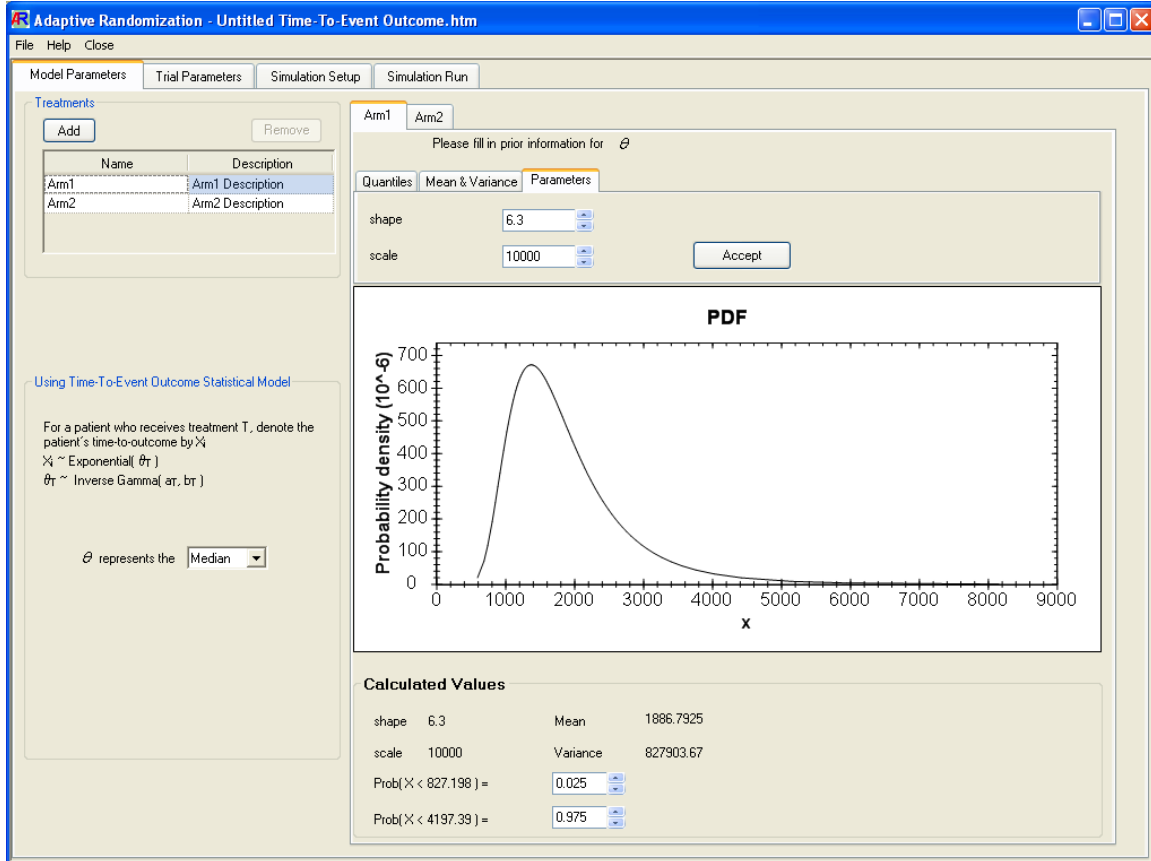
2.2 Time-To-Event Model

In the time-to-event trial, each patient’s outcome is a positive real number representing the elapsed time from enrollment until an event (e.g., a positive response or an observed toxicity) or last follow-up. These random time values for each treatment arm k are modeled by an exponential likelihood, which is described by a single parameter θ_k . The user may choose whether this parameter represents the mean or the median of the exponential distribution.

At the beginning of the trial, the uncertainty in the value of θ_k is modeled by a prior inverse gamma distribution: $\theta_k \sim \text{inverse gamma}(a_k, b_k)$, with positive “shape” parameter a_k and positive “scale” parameter b_k . The mean value of this distribution is $b_k/(a_k - 1)$, so that $(a_k - 1)$ may be interpreted as the prior number of events, while b_k may be interpreted as the sum of the prior time values.

If historical data are available for treatment arm k , then it is suggested that the parameters a_k and b_k of the prior distribution be set to reflect these data. If θ_k represents the mean time to event, set a_k to the number of prior events + 1 and set b_k to the sum of prior time values. If θ_k represents the median TTE, set a_k as before but set b_k to the sum of prior time values times $\ln(2)$. To down-weight the historical data, carry out the above steps for a subset of the data.

Figure 2b: Model parameters for a TTE trial



Suppose that the trial produces data for treatment arm k consisting of n_k observations and total time sum S_k . If θ_k represents the mean TTE then the posterior distribution is updated to

$$\theta_k | \text{data} \sim \text{inverse gamma}(a_k + n_k, b_k + S_k)$$

while if θ_k represents the median TTE then it is updated to

$$\theta_k | \text{data} \sim \text{inverse gamma}(a_k + n_k, b_k + \log(2) S_k).$$

3. Trial Parameters

After specifying a prior distribution for each treatment arm, the user should next set the parameters for conducting a trial. These parameters are organized into four groups under the second tab page “Trial Parameters” (Figure 3a).

3.1 Accrual

In the Accrual group box, the user can limit the trial length either by the maximum number of patients to enroll, by the maximum trial duration (in months), or both. The user also has the option of specifying a minimum number of patients to enroll before a trial can stop.

Figure 3a: Trial parameters

The screenshot shows the 'Trial Parameters' tab in a software application titled 'Adaptive Randomization - Untitled Time-To-Event Outcome.htm'. The interface is divided into several sections:

- Accrual:** Contains a sub-section 'Please select at least one' with two options: 'Max patient accrual' (checked, value 100) and 'Max trial length, (in months)' (unchecked, empty). Below it is an 'Optional' section with 'Min number of patients' (unchecked, empty).
- Stopping Rule:** Contains three probability settings: 'Suspend accrual to an arm if the probability that it is the best is less than' (P_L = .025), 'Stop the trial and select an arm as superior if the probability that it is the best is greater than' (P_U = .975), and 'After all patients are enrolled, select a treatment as superior if the probability that it is the best is greater than' (P_U^{*} = .975). There is also an option to 'Suspend accrual to an arm if Pr(θ > θ_{min} | Data) < P_L^{*}' with θ_{min} = 500 and P_L^{*} = .050.
- Randomization:** Contains 'Tuning parameter' (1), 'Initial number of patients to randomize fairly' (10), 'Equal number patients per arm at the end of the initial randomization phase' (checked), and 'Minimum randomization probability' (.10).
- Miscellaneous:** Contains 'Should θ be maximized or minimized?' (Maximize) and 'After all patients are enrolled, how many months will patients be followed before the final analysis?' (.00).

3.2 Randomization

Each time a new patient enters a trial, the randomization probability for each treatment arm is re-evaluated using the recorded outcomes of all patients enrolled so far. From these data, the program computes the posterior probability for each arm being the best among all arms

$$\pi_k = \text{Prob}(\text{ Arm } k \text{ is better than all other arms }).$$

If the trial has m arms, then the probability of the k -th arm being assigned next is

$$\frac{\pi_k^\lambda}{\sum_{i=1}^m \pi_i^\lambda},$$

where the non-negative power λ is the tuning parameter. As this parameter is increased, the currently best arm becomes increasingly favored over other arms. Typical values for λ are 2, 1, 0.5 or 0, which should be chosen with consideration for the patient accrual rate in relation to the observation time. A zero value for λ would mean equal randomization among all arms, and would be appropriate for a trial with very fast accrual and very long observation time. For a trial with relatively slow accrual in relation to the observation time, a positive λ value would be appropriate.

The user has the option of requiring equal randomization ($\lambda = 0$) for a certain number of patients at the beginning of the trial. The program will not compute the probability π_k for adaptive randomization until this specified number of patients has been enrolled. A conservative approach is to set this parameter to a significant fraction of the maximum sample size, allowing more information to be gained about the treatment arms before any adapting is attempted. While this approach tends to increase the average sample size, it also tends to reduce the variability in the trial results.

The minimum randomization probability is a safety parameter for preventing excessive adapting away from a treatment arm. In some extreme cases (possibly due to early bad luck), the randomization probability for an arm may be reduced too quickly, which effectively prevents that arm from accruing some desired minimum number of patients.

3.3 Stopping Rules

The user may specify one or more criteria for stopping the trial early and declaring a winner before the maximum number of patients has been accrued.

(i) Early loser selection parameter P_L :

During the trial, if the probability π_k for treatment arm k being the best is less than P_L , then that arm is declared a loser and is suspended. It is possible for that arm to return to active status later in the trial if the results for other arms grow sufficiently worse. The value of P_L should be set low, typically less than 0.10.

(ii) Early winner selection parameter P_U :

During the trial, if the probability π_k for arm k is greater than P_U then that arm is declared the winner and the trial is stopped early. The value of P_U should be set high, typically 0.90 and above. For a two-arm trial, the user should set $P_U = 1 - P_L$ so that rules (i) and (ii) are equivalent. Otherwise, only one of the two rules would really be in effect. When there are more than two arms, the values of P_U and P_L are not constrained in this way.

(iii) Final winner selection parameter P_U^* :

After all patients have been evaluated, if the probability π_k for arm k is greater than P_U^* , then that arm is declared the winner of the trial. If no treatment arm can meet this criterion then the program does not make a default final selection. Therefore, the value of P_U^* is often set smaller than P_U (typically from 0.7 to 0.9) to increase the chance of having a final winner.

(iv) Futility parameters θ_{\min} and P_L^* :

During the trial, if the probability of θ being larger than θ_{\min} is less than P_L^* for any treatment arm, then that treatment is declared futile and will not accrue more patients. This rule is enforced only when the model parameter θ is being maximized. The appropriate value for θ_{\min} depends on the model type (i.e., θ_{\min} is a rate under the binary-event model and a time value under the TTE model.) The value of P_L^* should be set low, typically 0.10 or less. Since this rule does not compare one treatment arm to another, a suspended futile arm cannot be re-activated later.

3.4 Miscellaneous

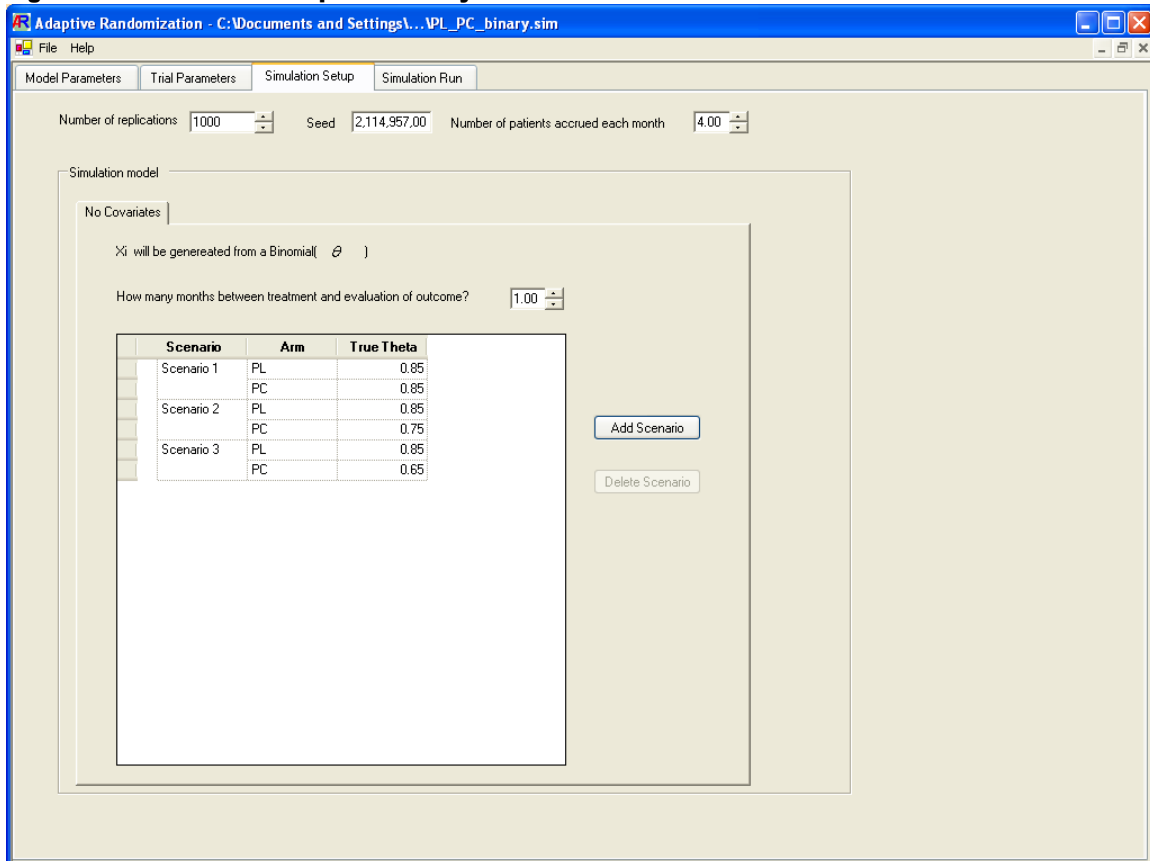
The user has the option of specifying whether the model parameter θ should be maximized or minimized. For example, if a binary-event trial is investigating the effectiveness of two treatments, then the response rate θ should be maximized. However, if θ represents a toxicity rate then it should be minimized. In a time-to-event trial, the average time to a positive response should be minimized, while the average time to toxicity (or survival time) should be maximized.

The last parameter on the tab page is the number of months to follow up after all patients have enrolled and before the final analysis. This time amount is included in the trial duration when the simulated trial runs to completion, but not when it is stopped early. This parameter is used to allow data to accrue after all patients have been enrolled and is especially useful when the outcome is not observed immediately. Both the binary event and the TTE trials use this parameter to estimate the average trial length. A similar time parameter that is available only in binary trials (the observation window) may be set on the next tab page.

4. Simulation Setup

The third tab page, “Simulation Setup”, allows the user to set some trial parameters and specify the scenarios for simulation (Figure 4a).

Figure 4a: Simulation setup for a binary event trial



The number of replications is the number of times each scenario will be simulated in order to generate the average results. For a quick evaluation, 100 replications is a good starting value. For a final verification, from 1,000 to 10,000 replications are recommended.

The seed parameter is an integer to initialize the program’s random number generator. A re-simulation using the same seed value (without changing any other parameter) should return exactly the same results.

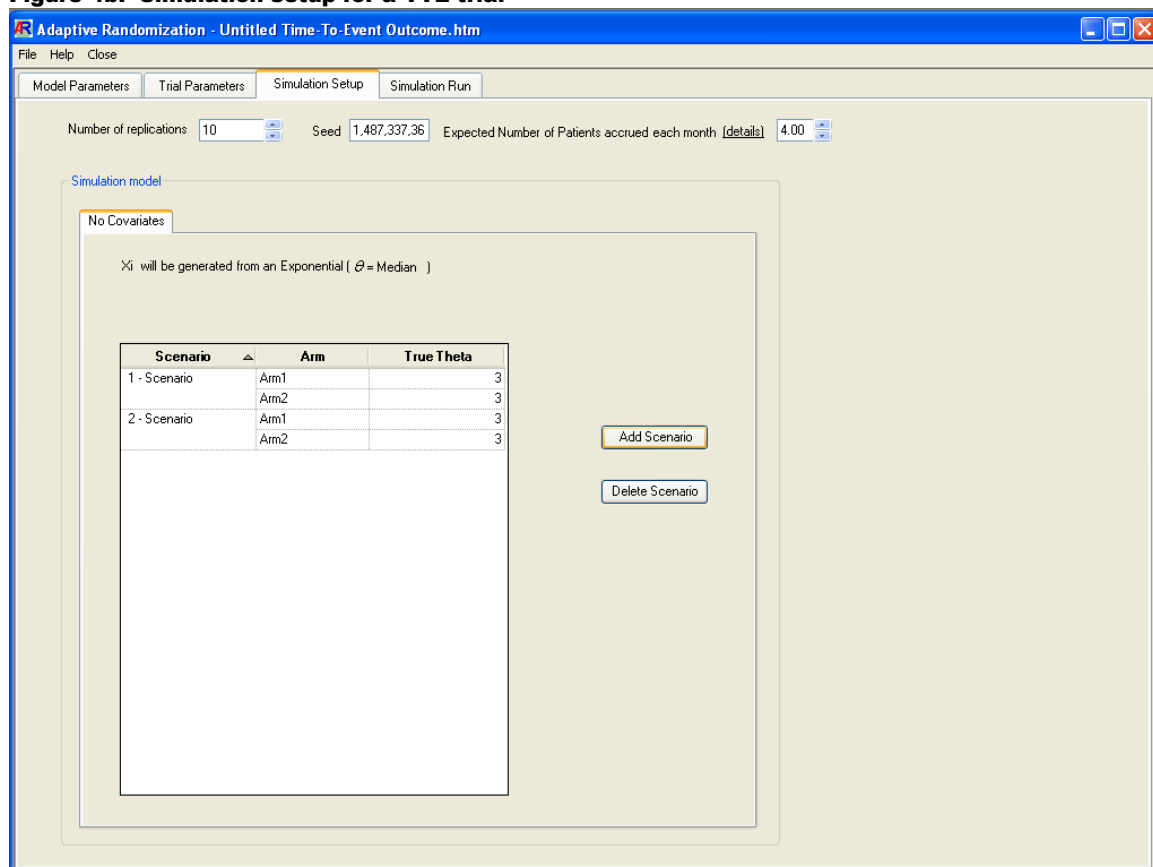
The average number of patients accrued each month is used to simulate the time interval between patient arrivals, which has a direct effect on the average trial length. Another important simulation parameter, available only to binary event trials, is the follow-up

time from treatment to evaluation of outcome (in months), which is also known as the observation window.

If the expected accrual rate is very fast compared to the observation window, then it is recommended that the randomization tuning parameter be set conservatively (e.g., 0.5 or less). Otherwise, the patient assignment may become unbalanced too quickly (possibly to the wrong treatment), resulting in some arms not accruing enough patients.

The table of scenarios is for specifying the set of “true” values for the model parameter θ for each treatment arm. In the binary event models, these values typically represent the true probability of response or toxicity for each treatment (Figure 4a). In time-to-event models, they represent the true value of the model parameter θ , which can be the mean or median of the exponential distribution describing the time values (Figure 4b). The program uses these scenario values for generating the patients’ response status or TTE values in the desired proportions.

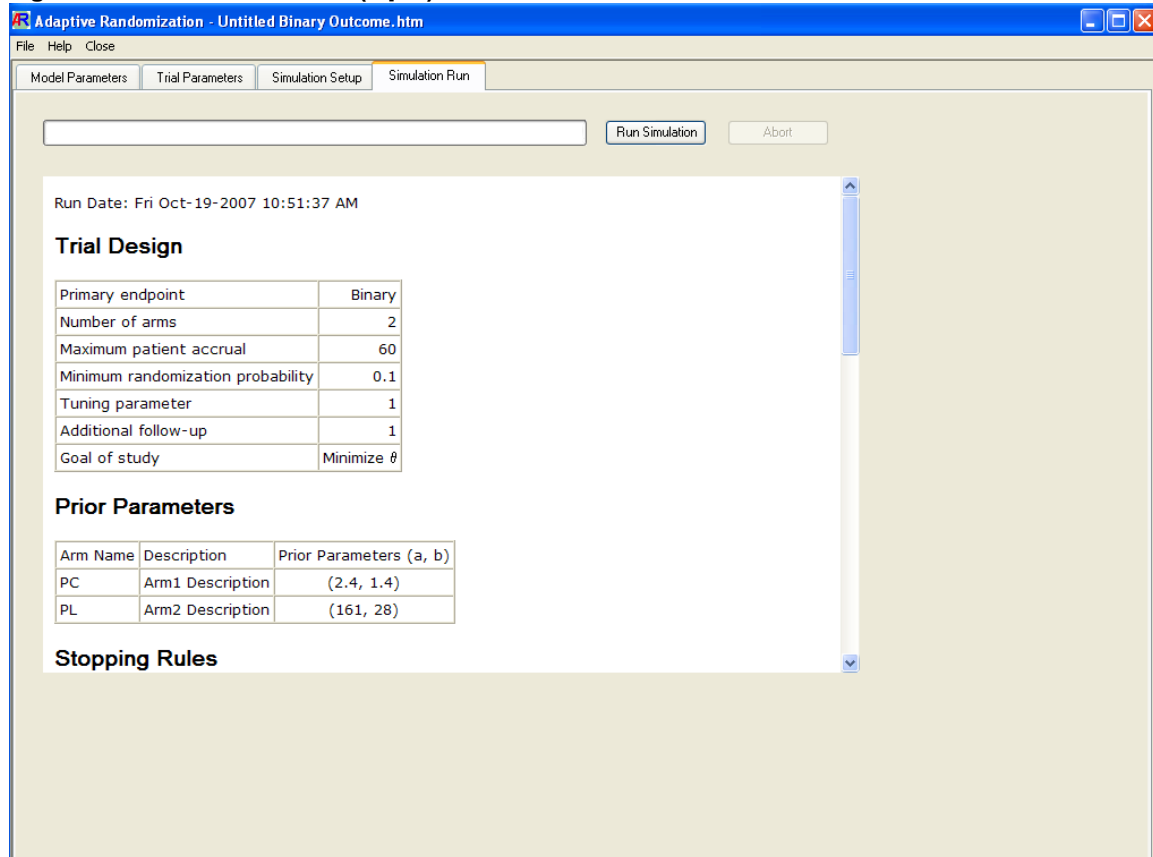
Figure 4b: Simulation setup for a TTE trial



5. Simulation Run

The last tab page, “Simulation Run”, allows the user to start the simulation and display the results for all scenarios entered on the “Simulation Setup” page.

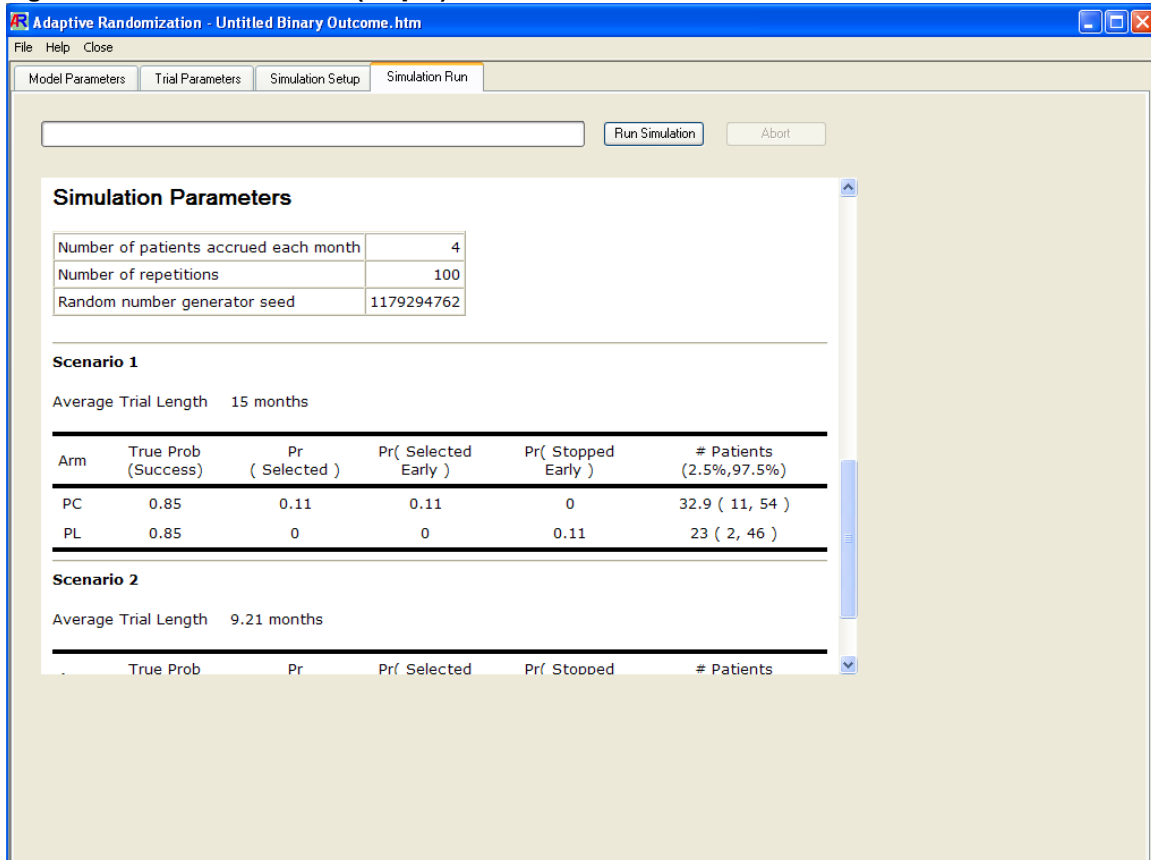
Figure 5a: Simulation results (input)



The top section of the output window contains the model and trial parameters that have been entered in the previous tab pages (Figure 5a). The bottom section displays the simulation results for each scenario. For each treatment arm, the results include: (Figure 5b)

- (i) The probability of being selected as the best treatment after the trial has ended.
- (ii) The probability of being selected early as the best (before all patients are enrolled).
- (iii) The probability of being dropped early for any reason (e.g., due to futility or another arm being superior).
- (iv) The average number of patients treated per trial, along with values for the credible interval (0.025, 0.975)

Figure 5b: Simulation results (output)



The entire content of the output window can be selected (by right-clicking and choosing “Select All”), copied and pasted. The format is HTML, allowing easy exchange with other software. The user may also save all input and output parameters to a file for later retrieval, by accessing the File menu.