

Main simul. runs

In the “# Main simul. runs” field you enter the number of runs in the “Main simulation” phase. Obviously, if you increase this number, the “Main simulation” phase will need more time to finish. However, your results will be more numerically precise. It is not possible to state a general rule for how many runs you need in any given situation. This depends heavily both on the model itself, and how precise results you need in your analysis.

If you have a large complex risk model, it may be a good idea to use a small number, e.g., 100-1000, during the model development. When your model is finished, however, you should consider increasing the number to at least 2000.

During the “Main simulation” phase DynRisk calculates all nodes in the traced set determined from your simulation set. Moreover, the output values of those nodes in the simulation set which have their “Sim.” attribute switched “on”, are stored in the result file. This is repeated for each simulation run. When the “Main simulation” phase is finished, DynRisk calculates some summary statistics for all nodes in the simulation set, and stores these in the result file.

Crit. simul. runs

In the “# Crit. simul. runs” field you enter the number of runs in the “Crit. simulation” phase. As with the “Main simulation” phase, if you increase this number, the “Crit. simulation” phase will need more time to finish. However, your results will be more numerically precise.

Both theoretical considerations and practical experiences indicate that you need at least as many runs in the “Crit. simulation” phase as in the “Main simulation” phase to obtain stable results.

Note that the time needed for the “Crit. simulation” phase is affected not only by the number of runs, but also by the number of “Crit.” and “Sens.” attributes being switched “on” in the simulation set. Moreover, if you calculate the criticality or sensitivity of a low level node with a high level node as goal node, this takes longer time than if you calculate the criticality or sensitivity of a high level node with the same goal node.

During the “Crit. simulation” phase DynRisk calculates once more all nodes in the traced set determined from your simulation set. However, in this phase no output values are stored in the result file. Instead DynRisk proceeds as follows:

Assume that in a particular run in the “Crit. simulation” phase, the output

values of the traced set of nodes are listed in the order of calculation:

$X(1), X(2), \dots, X(n)$

The values at the beginning of this sequence, are outputs from the nodes at the lowest level in the model, while the values at the end of the sequence, are outputs from the higher level nodes. Typically, (but not strictly necessary) the last node in the sequence, is the goal node. To simplify the present discussion we will assume that this is the case here.

Suppose that "Node A" is a node in the current simulation set whose "Crit." attribute is switched "on", and let $X(k)$ be the corresponding value in the sequence. Then DynRisk modifies the sequence temporarily first by replacing $X(k)$ by the mean value of "Node A" estimated during the "Main simulation" phase, and then by updating the sequence accordingly. In particular, DynRisk computes a modified value of the goal node. Thanks to the way the sequence is ordered, this can be done by recalculating only the subsequence:

$X(k+1), X(k+2), \dots, X(n)$

Similarly, if "Node A" has its "Sens." attribute is switched "on", and has an "S.fact." attribute equal to $S(k)$, then DynRisk modifies the sequence temporarily first by replacing $X(k)$ by $S(k)*X(k)$, and then by updating the sequence accordingly. In particular, DynRisk obtains another modified value of the goal node. Again thanks to the way the sequence is ordered, this is done by recalculating only the subsequence mentioned above.

Note, that if "Node A" is a low level node, the subsequence may still be very long. Hence, it typically takes longer time to calculate criticality and sensitivity of lower level nodes compared to higher level nodes.

During a single run in the "Crit. simulation" phase, DynRisk calculates modified values of the goal node as explained above, for each node whose "Crit." or "Sens." attributes are switched on. For nodes where both of these attributes are switched on, DynRisk needs to calculate two different modified goal node values. The modified goal node values are used to accumulate statistics, i.e., mean values and variances.

This is repeated for each simulation run. When the "Crit. simulation" phase is finished, DynRisk calculates the requested criticalities and sensitivities, and stores these in the result file.

Criticality

DynRisk provides two measures of the criticality of a node with respect to a goal node:

- Criticality mean
- Criticality variance

The criticality mean is calculated by using the following formula:

$$[M(G) - M(G | X = M(X))]/M(G)$$

where G is the goal node, X is the current node, M(G) and M(X) denotes the mean values of G and X respectively, and M(G | X = M(X)) denotes the conditional mean of G given that X is equal to its mean value. Thus, the criticality mean measures the effect of the uncertainty in X on the mean value of the goal node. M(G) and M(X) are estimated based on the results in the “Main simulation” phase, while M(G | X = M(X)) is estimated based on the modified goal node values obtained during the “Crit. simulation” phase.

The criticality variance is calculated by using the following formula:

$$[V(G) - V(G | X = M(X))]/V(G)$$

where G is the goal node, X is the current node, V(G) and V(X) denotes the variance values of G and X respectively, and V(G | X = M(X)) denotes the conditional variance of G given that X is equal to its mean value. Thus, the criticality variance measures the effect of the uncertainty in X on the variance of the goal node. V(G) and M(X) are estimated based on the results in the “Main simulation” phase, while V(G | X = M(X)) is estimated based on the modified goal node values obtained during the “Crit. simulation” phase.

Sensitivity

DynRisk provides two measures of the sensitivity of a node with respect to a goal node:

- Sensitivity mean
- Sensitivity variance

The sensitivity mean is calculated by using the following formula:

$$[M(G | S*X) - M(G)]/M(G)$$

where G is the goal node, X is the current node, S is the specified sensitivity factor, M(G) and M(X) denotes the mean values of G and X respectively, and M(G | S*X) denotes the conditional mean of G given that X is replaced by S*X in all calculations. Thus, the sensitivity mean measures the effect of a change in X on the mean value of the goal node. M(G) is estimated based on the results in the "Main simulation" phase, while M(G *X) is estimated based on the modified goal node values obtained during the "Crit. simulation" phase.

The sensitivity variance is calculated by using the following formula:

$$[V(G | S*X) - V(G)]/V(G)$$

where G is the goal node, X is the current node, S is the specified sensitivity factor, V(G) and V(X) denotes the variance values of G and X respectively, and V(G *X) denotes the conditional variance of G given that X is replaced by S*X in all calculations. Thus, the sensitivity variance measures the effect of a change in X on the variance value of the goal node. V(G) is estimated based on the results in the "Main simulation" phase, while V(G *X) is estimated based on the modified goal node values obtained during the "Crit. simulation" phase.