

The "Simulation view" is mainly used when you want to run a simulation on a model. This view is very similar to the "Node table view". The only difference is that in the "Simulation view" you have access to some additional parameters used during the simulations. These are found at the upper region of the window.

Simulation sets

DynRisk lets you run simulations at any level of a model. You can in fact run simulations on arbitrary sets of nodes.

To create a simulation set, make sure that all the nodes you want to include in the set, are represented in a folder, either by their "ego" node or by an alias. Since you may create as many aliases you like of a single ego node, you can define as many simulation sets you like as well. Simply create a folder for each set, and create aliases for each of the nodes in the sets inside their respective folders.

A simulation on a simulation set, can be divided into two different "phases":

- Main simulation
- Crit. simulation

However, a particular simulation may be set up such that only one of the phases actually occurs.

The main simulation phase occurs if the following is satisfied:

- 1) The "# Main simul. runs" field contains a positive integer.
- 2) At least one node has one or more of its "Sim.", "Crit." and "Sens." attributes turned "on".

The crit. simulation phase occurs if the following is satisfied:

- 1) The "# Crit. simul. runs" field contains a positive integer.
- 2) At least one node has one or more of its "Crit." and "Sens." attributes turned "on".
- 3) You have selected a "Goal node" for the simulation set.

Although it is possible to run a simulation consisting only of the crit. simulation phase, this is not recommended. The reason for this, is that the crit. simulation phase ideally needs some statistics from the main simulation phase to achieve acceptable numerical precision.

It is perfectly OK to run simulations consisting of only the main simulation phase, though. In fact you will probably do this more often than you run

complete two-phase sessions. To quickly exclude the crit. simulation phase, simply unselect the "Goal node". This is done by clicking the "Goal node" field while none of the nodes are selected. Similarly, to quickly include the crit. simulation phase again, simply reselect the "Goal node".

During a simulation, the output values of those nodes in the simulation set which have their "Sim." attribute switched "on", are stored in a result file. Furthermore, if you include the crit. simulation phase, DynRisk will estimate criticality and sensitivity for the nodes in the simulation set according to their "Crit." and "Sens." attributes. Thus, by choosing a particular simulation set, you focus your attention on this set, at least for the moment.

Note that if you switch off the "Sim.", "Crit." and "Sens." attributes of a node in a simulation set, this node is no longer a member of this simulation set, even though it is represented in the folder. It is sometimes useful to be able to temporarily leave out nodes of the simulation set this way without having to move the nodes out of the folder.

Note that if a node is selected as "Goal node", its "Sim." attribute is automatically switched "on", and locked in this position. Furthermore, the "Crit." and "Sens." attributes are automatically switched "off", and locked as well. Note also that events can never be included in simulation sets.

In order to simulate the nodes in the simulation set, we need to calculate all direct and indirect predecessors to these nodes as well. When you run a simulation on a set of nodes, DynRisk starts out with the chosen simulation set, and then traces all the paths of edges from the nodes in the set, and backwards through the model. Each node encountered in this process either belongs to the simulation set itself, or is a direct or indirect predecessor of a node in the set. When all the edge paths are traced, DynRisk has a complete list of all nodes it needs to calculate during the simulations of the chosen simulation set.

If a simulation set contains all "end" nodes in a model, the resulting list of traced nodes, may in fact contain every node in the model. In other cases, a simulation set can be calculated by including just a small subset of the entire node set. Anyway, DynRisk takes care of the model tracing for you, so you can rest assured that the simulations will be done using a minimal but yet sufficiently large set of nodes.

To understand the purpose of a simulation set, it is important to understand the difference between the simulation set and the set of traced nodes. The simulation set implicitly determines the usually much larger set of traced nodes. By selecting a simulation set, instead of specifying the set of traced nodes, we save a lot of work. Moreover, if we later restructure our model, the old simulation sets will still work just fine, while the set of traced nodes

may change a lot. In fact if we delete a node from the model which used to belong to a simulation set, DynRisk will automatically remove it from this set for us.

The only time we may want to update the simulation set, is when we decide to include yet another node in the set, or when we want to exclude a node from the set. Such decisions are very easy to make just by examining the simulation set itself.