

Configuring an Insight Optimisation Run

Configuring a Source Project Contents

Decision variables and objectives must be defined in the Source project as functions. Insight can then run the Source model many times, setting different values for the decision variable expressions and evaluating the corresponding objectives for every individual run. Therefore, in order for a Source parameter to set or evaluated for an optimisation problem, that parameter needs to be defined using the Function Manager in Source.

Two additional configuration requirements are:

1. Ensure that all Source project files contain only one scenario. It is possible to run Insight using a project file containing multiple scenarios. However, Insight will only run one scenario in each project file, which is the first scenario created by the user.
2. Ensure that auto-export of model results, simulation settings and other output files are turned off in Source. If you have these features enabled, file-write conflicts may occur causing the optimisation to fail.

Run Configuration

The **Run Configuration** menu contains general settings for Insight. It enables the user to load a project directory and time series data, configure the NSGA-II algorithm and choose the number and type of endpoints. The **Run Configuration** menu items are listed in the order that they are usually configured by the user.

Load Projects

Click **Browse**, or go to **File » Load RiverSystem project directory**, to choose the directory containing the Source project file(s). Project files can be located in sub-directories of the selected directory. **In Source (public version), you can only load a single project file;**

Once the project has finished loading (indicated by a green progress bar) you will be able to [Configure Expressions](#).

Load Time Series (Optional)

The **Load Time Series** feature is optional. It is designed to allow the user to easily optimise a Source scenario with a range of different time series inputs. This behaviour can also be achieved by loading a directory into Insight (via [Load Projects](#)) that contains multiple Source project files in sub-directories. However, if the only difference between the Source projects is the time-series data sources, then the **Load Time Series** option saves the user from needing to create multiple copies of the project file.

To use the Load Time Series option, the **Time Series Data Sources** in the eWater Source scenario should be configured as described below (these requirements only apply to time series that will be changed):

- Enable **Reload on Run**
- Check **Relative Path**
- Place time series files in the same directory (or sub-directories) of the project file

The "Master Directory" directory containing the project file and time series (possibly in sub-directories) should be loaded as normal via the [Load Projects](#) menu item.

In addition, the user should create a separate "Time Series Directory" containing the time series that will be changed for each run:

- Each set of files should be contained in a separate "Run Directory"
- The structure of each Run Directory should be identical to the Master Directory (minus the Source project file itself)
- The file names in each Run Directory should be the same as in the Master Directory
- The Data Source name must match the name of the file. For a filename "MyRainfallData.csv" the default Data Source name is "MyRainfallData_csv" NOTE the "." is replaced with an _ underscore.

Time series files that are not changed (i.e. are the same as in the Master Directory) do not have to be placed in the Run Directories

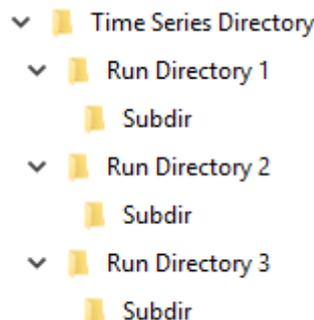


Figure 1. Example Time Series Directory structure.

Click **Browse** to select the directory containing the time series inputs.

Generations and Populations

The **Generations** and **Populations** fields set the number of generations and the population size of each generation for the NSGA-II algorithm, respectively. These parameters define how many simulation runs are performed by the optimiser. The required values will depend on the complexity of the optimisation project. If the number of runs is too small, the optimiser will not be given sufficient time to search for the multi-objective optimal set of results. The population size also sets the maximum number of points on the Pareto front or surface;

Advanced Options (Optional)

Advanced users can optionally configure additional NSGA-II algorithm parameters via the **Advanced Options** sub-menu.

Local Server Limit & Endpoints

The user can choose to run the optimisation locally or across server endpoints by enabling the respective radio buttons.

- If **Run Locally** is selected, then the user can choose the number of local processes via the **Local server limit** field.
- Server endpoints allow you to run the optimisation across several machines simultaneously. Load the endpoints settings file by clicking **Load Endpoints** and point Insight to the appropriate file;

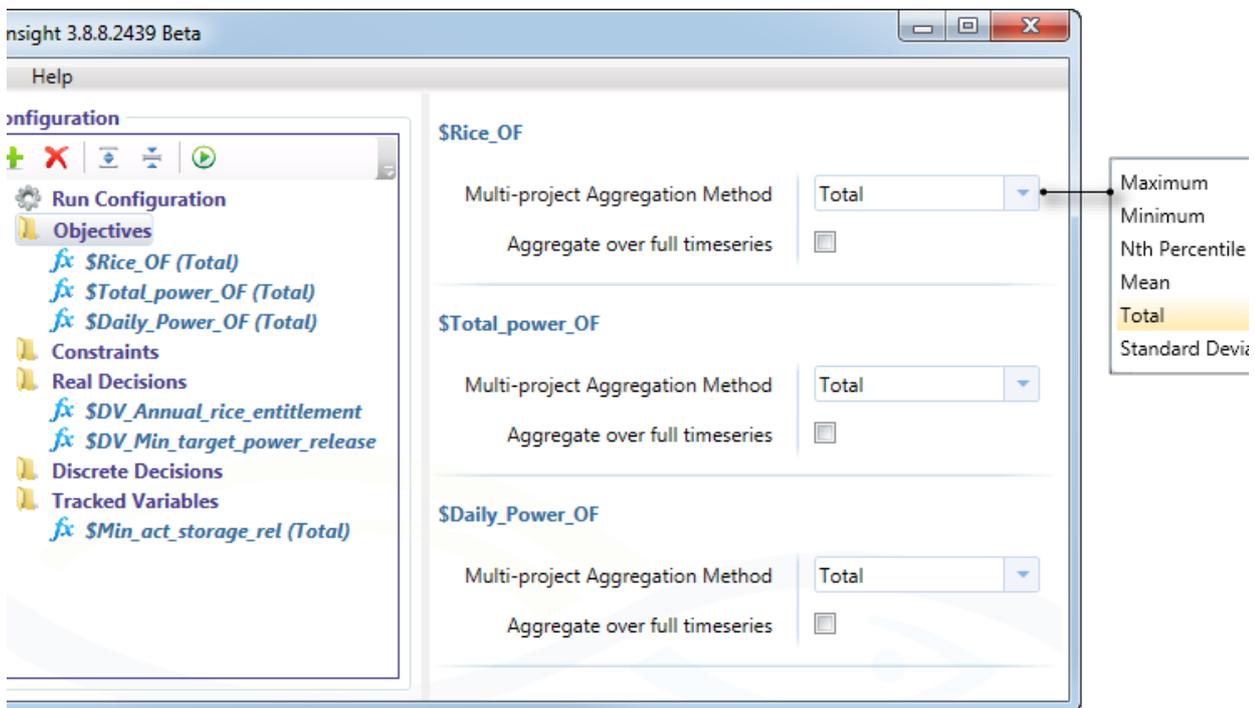
Run & Show Results

Once the configuration is complete (including [Configuring Expressions](#)) the **Run** and **Show Results** buttons run the optimisation and visualise results, respectively. See [Running an optimisation and viewing results](#).

Configuring Expressions

- Click **New Expressions** (on the green +) in the Insight toolbar and choose the component you wish to add;
- The resulting dialog contains a list of all pre-defined functions in the model. Choose the one you wish to add and click **OK**;
- Continue this process until you have added all components. Figure 2 shows an example of objectives, real decisions and tracked variables that have been added to the Insight project. Each component has additional parameters that must be defined:
 - For discrete decisions, specify the allowed values of the function;
 - For objectives and tracked variables, specify **Multi-project Aggregation Method** drop-down menu allows you to choose how you want to process the objective function values passed from individual scenarios. By default Insight uses the final value of the specified function as the value of the objective function for each scenario. The **Aggregate over full timeseries** checkbox allows you to use the values of the objective function determined at each time-step, rather than just the last value (via the selected statistical function).
 - For real decisions, specify the allowed minimum and maximum values.

Figure 2. Insight, Objectives



Note: The difference between objectives and tracked variables are that the former are optimised, whereas the latter are not. Tracked variables are a set of secondary objectives that may affect decision-making, but do not get optimised within Insight.

Saving The Configuration

Save your settings using **File » Save Settings**. This creates a settings file, which can be re-loaded into Insight.