

TRIM.FaTE USER'S GUIDE

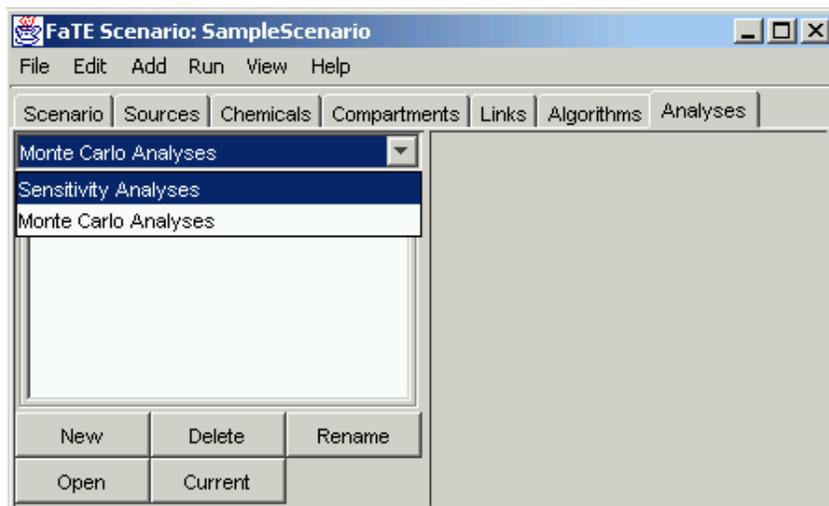
MODULE 15: TRIM.FaTE SENSITIVITY AND MONTE CARLO ANALYSES

1. INTRODUCTION

TRIM.FaTE includes features that allow the user to perform sensitivity analyses or Monte Carlo variability and uncertainty analyses on an existing TRIM.FaTE scenario. For these analyses, the user defines a scenario and all of its associated properties for the “base run.” Then, the values for one or more properties are varied, the scenario is run again, and the results for the new run are compared to the base run results or to each other. The results can provide quantitative sensitivity or variability and/or uncertainty information about that scenario set-up.

Within the TRIM.FaTE framework, sensitivity and Monte Carlo variability and uncertainty analyses are accessed via the Analyses tab of the Scenario window shown in Figure 1. The view attached to this tab contains a scroll-down menu located in the top left corner that allows the user to select the type of analysis (the menu is open in the figure below). Below this menu is a window that lists the specific sensitivity or Monte Carlo analyses that have been created. “Default” sensitivity and Monte Carlo analyses are automatically created for a new scenario and will be labeled with the same name as the current scenario (note that these analyses

Figure 1
Analyses View in Scenario Window



are saved within the same TRIM.FaTE project file with the scenario). The pane on the right side of the Analyses view is not currently used for TRIM.FaTE analyses. Settings and specifications for the analyses (e.g., parameters to be varied, output values to be assessed, output directory to be used) are set via separate, analysis-specific windows that the user accesses by opening a specific analysis. Detailed descriptions of how sensitivity and Monte Carlo analyses are set up and executed within TRIM.FaTE are provided in Sections 2 and 3, respectively.

Steady-state vs. Dynamic Mode

Sensitivity and Monte Carlo analyses can be carried out within TRIM.FaTE by running simulations in either the steady-state or the dynamic mode. For sensitivity and Monte Carlo analyses on TRIM.FaTE scenarios, the steady-state mode offers the advantage of a greatly reduced run-time relative to the dynamic mode. Run-time can be an important consideration because multiple model runs are required for these analyses. However, if the user uses the steady-state mode to investigate the sensitivity or uncertainty/variability associated with a dynamic simulation, he or she must determine the applicability of those results to the dynamic model runs of interest.

The steady-state mode for a scenario is specified via the scenario property *simulateSteadyState*. Refer to Module 13, Operation of TRIM.FaTE in Steady-state Mode, for details regarding the steady-state mode.

2. SENSITIVITY ANALYSES

2.1 Conceptual Background

In a sensitivity analysis, one or more input values are changed, or “perturbed,” and the results are compared to a base run (i.e., the results obtained for an unperturbed scenario). The **sensitivity** is the ratio of the input change to the change in a given output. Two additional measures of sensitivity can be calculated for each input/output pair: elasticity and a sensitivity score. The **elasticity** is the percent change in a model output value resulting from a 1 percent change in an input parameter, with all other parameters held constant. A positive value of elasticity results when an increase in an input value results in an increased output value or vice-versa (i.e., a decrease in an input value results in a decreased output value). A negative value of elasticity indicates that an increase in input value results in an output decrease or vice-versa. For example, if a change of -10 percent in the input parameter “algae growth rate” results in a +11 percent change in mercury concentrations in fish, then the elasticity is -1.1. The **sensitivity score** weights the elasticity by the coefficient of variation (CV) of the input parameter and provides a measure of the variability in the output value resulting from the variability of the input parameter. The CV is the standard deviation of the input parameter expressed as a fraction of the parameter’s mean (i.e., [std. deviation]/[mean]). See Chapter 6 in Volume I of the TRIM.FaTE TSD (EPA 2002) for additional information on concepts related to sensitivity analysis.

In a run or simulation conducted as part of a TRIM.FaTE sensitivity analysis, an object’s property value is varied by a given amount, the TRIM.FaTE scenario is executed, and the results in mass, moles, and/or concentration (as selected by the user) are compared to the results of the

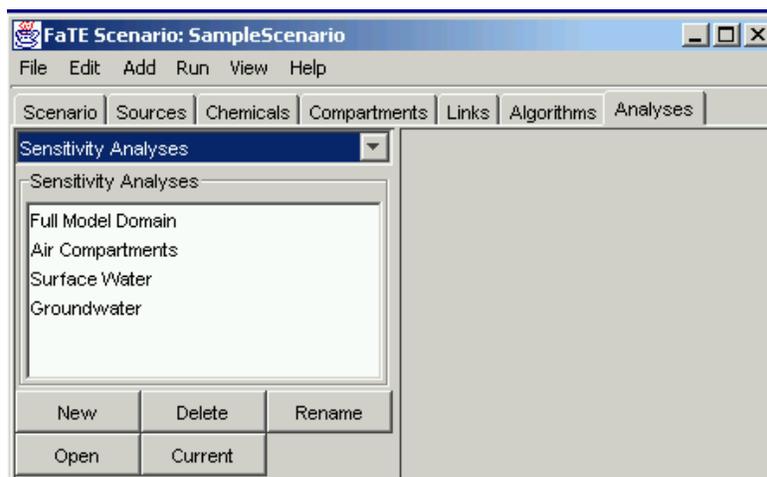
base run. For a given sensitivity analysis run, a single property can be changed (e.g., the soil porosity of one specific surface soil compartment), or a property value common to multiple objects of the same type can be changed simultaneously (e.g., the soil porosity of *all* surface soil compartments included in the scenario). TRIM.FaTE saves the complete results file for each run conducted (i.e., the masses/moles/concentrations for all compartments included in the scenario are saved) as part of the analysis. Then, based on user specifications, TRIM.FaTE generates one sensitivity output file for each chemical and form of results (i.e., mass/moles/concentration) that summarizes the sensitivity calculations for all properties perturbed during the runs of the sensitivity analysis. The file contains the specifications of the sensitivity run (e.g., names, initial values, and perturbed values for properties included in the analysis) and the sensitivity, elasticity, and a sensitivity score for each input/output combination designated by the user. If additional chemical/form of results combinations are selected by the user for sensitivity analysis, a separate sensitivity summary file can be created for each combination.

2.2 Setting Up and Running a Sensitivity Analysis in TRIM.FaTE

2.2.1 Accessing the Sensitivity Analysis Feature

The TRIM.FaTE sensitivity analysis functionality is accessed via the Analyses tab of the Scenario window by selecting “Sensitivity Analyses” from the drop-down box in the upper left corner. The window below the drop-down box (i.e., the “Sensitivity Analyses” pane) lists the sensitivity analyses that are in the current scenario. By default, an “empty” sensitivity analysis is created with every scenario and is labeled with the name of the scenario. To add a new sensitivity analysis, click the “New” button and enter a name when prompted. The other buttons below the analyses window can be used to delete, rename, or open analyses. In the example shown in Figure 2, four sensitivity analyses have been created and are named according to the type of medium in which the properties are varied.

Figure 2
TRIM.FaTE Sensitivity Analysis Window



To view specific information about an analysis, the user can bring up the Sensitivity Analysis Window. This window is opened by selecting a sensitivity analysis name and clicking the “Open” button or by double-clicking the analysis name in the Sensitivity Analyses listing. Alternatively, the user can access this window by clicking on the “Run” pull-down menu in the scenario window, selecting “Sensitivity Analysis,” and then selecting the appropriate analysis in the dialog box that appears. The Sensitivity Analysis Window contains tabs that access three views:

- (1) The **Properties** view, used to specify the variation of the values included in the sensitivity analysis;
- (2) The **Runs** view, which conveys information about the status of the sensitivity run; and
- (3) The **Analysis** view, used to specify how the results of the analysis will be calculated.

The features associated with each of these views is described in more detail in the sections that follow.

2.2.2 Selecting Properties to be Varied

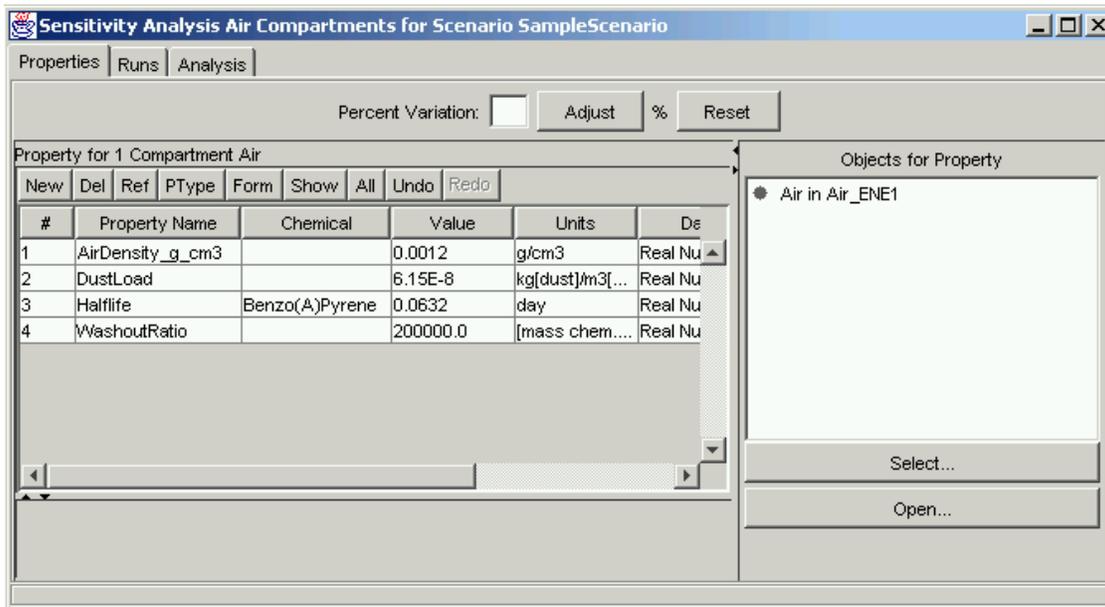
There is no default set of properties for a TRIM.FaTE sensitivity analysis. To designate the properties that will be varied, the user must add properties to the sensitivity analysis from the scenario window. The following steps describe this process.

- (1) Designate a “Current” sensitivity analysis by selecting one of the analyses from the Sensitivity Analyses list below the drop-down menu in the Analyses view of the Scenario Window (see Figure 2) and then clicking the “Current” button below the pane. After specifying an analysis as “current,” all properties added will be included in this analysis until the user designates a different analysis as current.
- (2) Click the tab in the main Scenario window (e.g., Chemical Tab, Compartments Tab) that contains the object property to be varied.
- (3) Select the object(s) containing the property or properties to be varied. This can be done manually, by finding the object and clicking on it to highlight the object name, or with assistance from the TRIM.FaTE “Select” tool by clicking the “Select” button and using the dialog window that appears (this feature can help the user find and select multiple, similar objects). **Note that objects in the scenario with properties to be varied must be selected from the Outdoor Environment pane on the left, not the library pane that appears in the center.**
- (4) Click the “Properties” button to bring the properties for the highlighted object(s) up in the Property Editor.
- (5) Select one or more properties by highlighting them in the Property Editor.

- (6) From the Add pull-down menu at the top of the scenario window, select “Add Selected Properties to Sensitivity Analysis [Analysis Name].” It is important to note that the selected properties will be added for ALL of the selected objects. When the property or properties have been added to the sensitivity analysis, a message will appear at the bottom of the TRIM window stating “### properties added to sensitivity analysis for ### objects” (each “###” will be replaced with the appropriate number depending on what the user has selected).

To view the properties that have been added to an analysis, open the Sensitivity Analysis window and select the Properties tab. An example of the Properties view in the Sensitivity Analysis window for a sample analysis entitled “Air Compartments” is presented in Figure 3. In this example, four properties of the air compartment “Air in Air_ENE1” have been added to the current analysis.

Figure 3
Properties View, Sensitivity Analysis Window



To select properties to be varied in a separate sensitivity analysis for the same scenario, designate a different sensitivity analysis as “Current” as described in Step 1. Repeat Steps 2 through 6 to add the desired properties to that sensitivity analysis. To remove a property from a sensitivity analysis after it has been added, highlight the property in the table in the Properties view and click the “Del” button just above the table.

2.2.3 Varying Property Values for the Analysis

Once the appropriate properties have been added to the analysis, the new values to be included in the sensitivity analysis runs can be set via the Properties view. As shown in Figure 3, each property is listed on a separate row of the table that appears on the left side of this view, with values for the following data fields in the columns:¹

- # (contains a number that TRIM.FaTE assigns to the property for identification);
- Property name;
- Chemical to which the property pertains (if applicable);
- Value (i.e., the adjusted value to be used in the sensitivity analysis);
- Units of the value;
- Data type (only properties defined as real numbers should be added to a sensitivity analysis);
- Reference; and
- Old value (i.e., the original value for the property as defined in the base case).

The objects to which the currently selected properties apply (e.g., compartments, links, chemicals) will appear in the pane on the right under the label “Objects for [object name(s)],” where the name(s) in the brackets will be filled in depending on which properties are currently highlighted in the table.

The precise new value for a property can be set directly or by specifying a percentage change from the base value. To set the value directly, the user can edit the value displayed in the “Value” column for that property in the Sensitivity Analysis window. Alternatively, the value can be changed by a percentage of the original value by selecting the property, entering a percent variation in the text box at the top of the Properties view, and clicking the “Adjust” button. A positive percent variation will increase the value (a positive sign (“+”) is not required). Negative numbers can be entered using the negative sign (“-”) to denote a decrease in the value. Note that the actual percent should be entered in this window (e.g., to increase the values by 10 percent, enter “10,” not “0.10”). A different percent variation can be used for different properties, or the same percent variation can be implemented for multiple properties by selecting multiple properties (using the shift or control key) and entering a value for variation.²

When adjusting property values by a given percent for a sensitivity run, the values are always adjusted from the original base case value. For example, if property values are initially varied by 10 percent, and the variation is then changed to 20 percent, the new values will vary by 20 percent of the original base case values, not by 20 percent of the 10 percent varied values. If desired, the user can reset one (or all) Property values to their original base case values by selecting the property (or all properties) and clicking the “Reset” button.

¹ Note that not all of these columns are visible in the window presented in Figure 3. The horizontal scroll bar at the bottom of the window can be used to view the hidden columns.

² Note: if properties from multiple objects are selected in the left pane, all of the object names may not appear in the “Objects for Property” window (all property values can still be varied simultaneously for the selected objects).

2.2.4 Executing the Sensitivity Analysis Runs

Information about the status and location of the simulation runs that are being performed is summarized on the Runs view within the Sensitivity Analysis window (see Figure 4). In the table that appears on this view, information about individual runs is presented on separate rows, including:

- the name of the run in the left column of the row;³
- the status of the run in the middle column (i.e., whether or not the run has been completed, and whether or not the run was successful); and
- the name of the subdirectory where the run results will be saved in the right column.

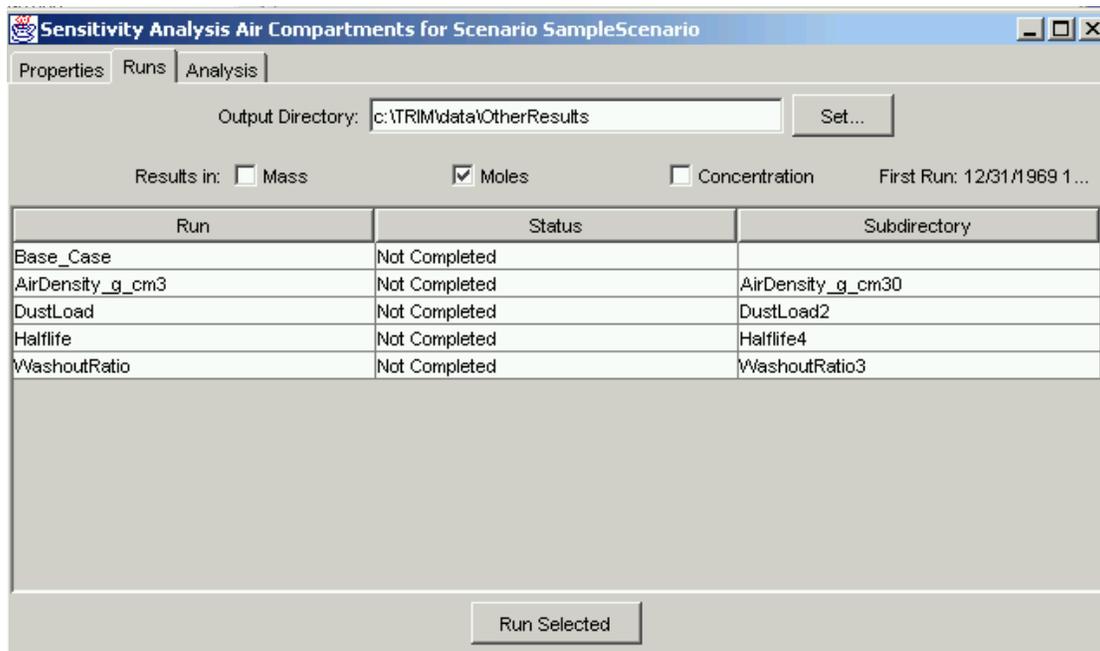
Note that the table on the Runs view is simply used to present information to the user (i.e., the user cannot modify this table directly to change run information).

Two important options must be set by the user via the Runs view:

- (1) The **output directory** where the results files for the simulations will be stored should be specified in the window at the top of this view. This directory can be entered manually by the user, by clicking within the directory window and typing in the location. Alternatively, the user can click the “Set” button to bring up a file browser window that allows the user to find and select a directory. In the example in Figure 4, the directory “c:\TRIM\data\OtherResults” has been set as the output directory.
- (2) The **form of results** to be used for presenting the sensitivity analysis (i.e., mass, moles and/or concentration) by checking one (or more) of the three check boxes below the output directory window. In the example in Figure 4, moles has been selected as the only form of results.

³TRIM.FaTE assigns a subdirectory name for the run results that matches the name of the property being varied in the analysis with the addition of a numeral. The numeral corresponds to the order in which the property was selected for the analysis, starting with zero and skipping one (i.e., 0, 2, 3, 4...). This number is used to differentiate the subdirectory with these run results from the subdirectory with results corresponding to a different run involving the same property names.

Figure 4
Runs View, Sensitivity Analysis Window



Once these options have been set, the model simulations associated with the sensitivity analysis can be run. Note that a separate simulation is completed for the base case and each of the “runs” (i.e., one for each of the properties that is selected to be varied). To execute these simulations, select one or more run in the table on the Runs view and click the “Run Selected” button at the bottom of the view. This will start the selected runs as a batch. While the simulations are processing, a small window with a “Stop” button will appear that can be used to stop the batch of simulations at any point during the run (it may take a minute or two for the simulation to stop after the button is clicked).

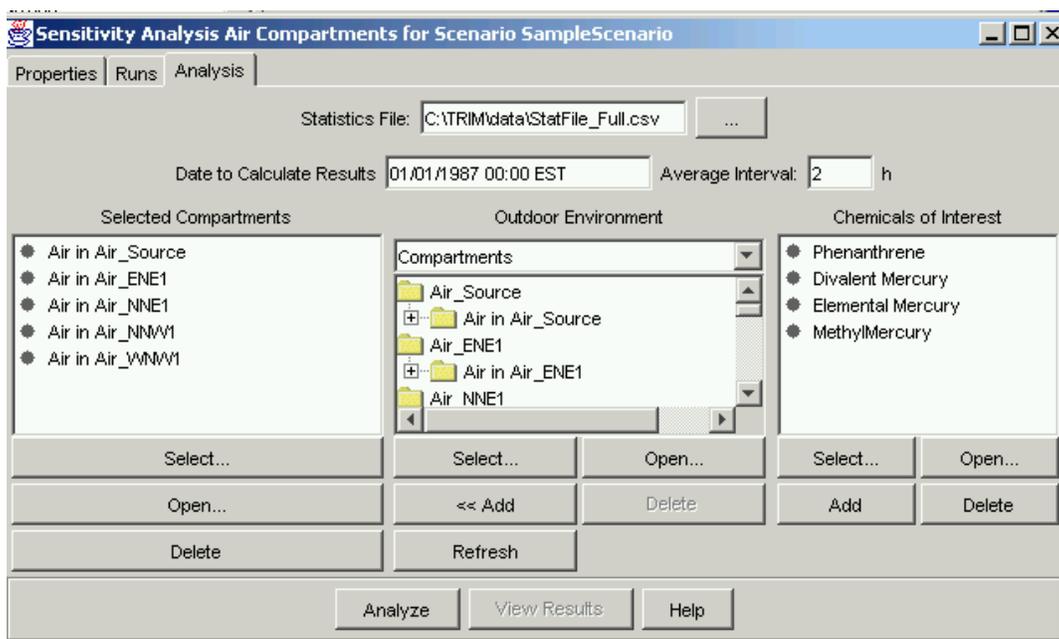
After the simulation for a given run has been completed, the status of that run in the table on the Runs view will change to “Successful,” “Warnings,” or “Failed.” An output subdirectory will be created within the main output directory (specified above) for each run as listed in the “Subdirectory” column. For each run included in the analysis (i.e., the base case run plus each of the runs for the properties that were perturbed), a complete set of output files will be saved in the corresponding subdirectory. These simulation results are used by TRIM.FaTE to calculate quantitative sensitivity results, as described in the following section.

Important Note: Depending on the complexity of the scenario and, likely, whether the model is run in dynamic or steady state mode, including a large number of runs in a single batch may cause model execution problems. Running multiple batches of runs (each containing a subset of the runs comprising a sensitivity analysis) may reduce the potential for this problem.

2.2.5 Performing Sensitivity Calculations

After the simulations for the base run and the perturbed properties have been completed, the sensitivity calculations to be completed using the simulation results can be specified via the Analysis view within the Sensitivity Analysis window (see Figure 5). To obtain sensitivity results, the user selects compartments and chemicals for which to evaluate the sensitivity of the properties that were varied from their base values. The steps involved in specifying and executing the calculations are described below.

Figure 5
Analysis View, Sensitivity Analysis Window



- (1) Using the small window at the top of the Analysis view, enter the name of the TRIM.FaTE Statistics file to be used for the calculations. The file name can be entered manually, or the user can click the “...” button to bring up a file browser and select a file. This Statistics file contains the coefficient of variation (CV) needed for each property included in the analysis. See Module 3 (Library and Scenario Data Files), Section 2.2, for details regarding the content and formatting of the statistics file. Note that the user is responsible for ensuring that the appropriate values for property CVs are used in this analysis.
- (2) For analyses involving results from dynamic simulations, specify a date and time from the simulations for which the sensitivity results will be calculated. This date and time should be entered in the window below the file name.⁴ If the simulation

⁴ Steps 2 and 3 are not required if the sensitivity simulations are run in steady-state mode. If the simulation is being run in steady-state mode, TRIM.FaTE will automatically shade the date and averaging interval windows and the user will not be able to access them.

results did not include that particular date and time, TRIM.FaTE will use the last time point at which results were output. For example, if results for a run are calculated every four hours (e.g., hours 00:00, 4:00:00, 8:00:00, ... during a day), the results for 4:00:00 would be used if the user enters any time ranging from 4:00:00 to 7:59:59. If the sensitivity calculations are to be performed on averaged results, TRIM.FaTE will use the results for the time interval inclusive of this date and time.

- (3) For analyses involving results from dynamic simulations, specify the Average Interval for the simulation results (i.e., the time interval in hours within which the model simulation outputs will be averaged) in the text box to the right of the date window.⁴ This interval must be specified as either a whole-number multiple of the scenario output time interval or “monthly” or “annual.” For example, if simulation results are output every 3 hours (per user specification of the scenario properties for the simulation), the Average Interval could be specified as 3 hours, 6 hours, 12 hours, any other time period into which 3 hours is evenly divided, or as “monthly” or “annual.” For the sensitivity calculations to be performed on the simulation results as they are output from TRIM.FaTE (i.e., without subsequent averaging), enter the value (in hours) of the simulation output time step.⁵
- (4) In the “Outdoor Environment” pane in the center of the view, select the compartment(s) for which the output results will be analyzed. Multiple compartments can be selected simultaneously using the shift or control keys, and the “Select” button can be used to find and select groups of compartments. Click the “<<Add” button to add the selected compartments to the analysis. The selected compartments should appear in the “Selected Compartments” pane on the left side of the window. The “Delete” button below the pane can be used to remove compartments that have been added to the analysis.
- (5) Click the “Add” button below the “Chemicals of Interest” pane on the right-hand side of the window to bring up a “Select Chemicals” dialog box. From the list that appears, select the chemical(s) for which the output results will be analyzed. The selected chemicals should then appear in the “Chemicals of Interest” pane. The “Delete” button below the pane can be used to remove chemicals that have been added to the analysis.
- (6) Click the “Analyze” button at the bottom of the view to generate the results files. If the simulation is large (i.e., many compartments and links), or a large number of simulations are to be completed (i.e., many properties were perturbed), the analysis step may take some time.

With the completion of the analysis step, TRIM.FaTE creates a sensitivity summary file for each combination of chemical (per user specification in step 5) and form of results (per user

⁵As discussed in Module 12, the output time step is the product of the *simulationStepsPerOutputStep* and the *simulationTimeStep* properties (which are specified in the scenario window of the scenario).

specification in the Runs view; see Section 2.2.4). The names, location, and format of these files are described in Section 2.2.6.

To perform additional sets of sensitivity calculations from the same set of sensitivity runs (e.g., to generate sensitivity analysis results for other combinations of compartments and chemicals, or to evaluate the sensitivity at a different date/time by entering a different date/time in Step 2), Steps 2 through 6 can be repeated. Note, however, that TRIM.FaTE will *not* change the name of the output file(s). That is, existing summary results files will be overwritten if Steps 2 through 6 are repeated for the same TRIM.FaTE sensitivity analysis. Therefore, if additional calculations are to be performed to produce new summary files, the user should first rename the previous summary files in order to preserve them.

For sensitivity calculations on results from dynamic TRIM.FaTE simulations, Steps 2 and 3 specify which results (i.e., from which time point or, for averaged results, which time interval) to use in the sensitivity calculations. The “Average Interval” input provides a way for the user to use time-averaged results rather than an instantaneous point in time. For example, if the Average Interval is specified as “annual” and the user enters a date/time that falls within the 20th year of the simulation, then the annual average of the simulation results for the 20th year will be used to calculate that set of sensitivity results. The user should note that the “sensitivity score” value that is calculated by TRIM.FaTE could vary depending on the averaging interval used and the selected time.

2.2.6 Sensitivity Analysis Results

Summary Files

TRIM.FaTE generates summary results for the sensitivity analysis based on the user specifications listed in the previous section. These summary results will be saved to a series of text files in a subdirectory of the output directory specified in the Runs view of the Sensitivity Analysis window (see Figure 4). The subdirectory will be named “Sensitivity Analysis” and will include an additional subdirectory with the name of the analysis. For example, using the output directory specified in Figure 4, divalent mercury concentration results for the sensitivity analysis named “Air Compartments” will be saved in a summary file named *Sensitivity_Divalent_Mercury_conc.txt* in the following folder:

c:\TRIM\data\OtherResults\SensitivityAnalysis\Air Compartments\

One summary file will be created for each combination of chemical and form of results (i.e., mass, moles and/or concentration). Each row within a given summary file contains sensitivity results for a property that was perturbed and a selected compartment for which the output was of interest (i.e., the compartments listed in the “Selected Compartments” pane). For example, if four properties were varied in the sensitivity analysis, and three compartments were selected on the Analysis view, the sensitivity results file for a selected chemical and form of results (e.g., moles of divalent mercury) would contain 12 rows of data (see Figure 6). Each row in the summary results file will include, in separate columns, values for the sensitivity parameters generated by TRIM.FaTE as presented in Table 1. Key measures of sensitivity are highlighted in boldface in Table 1. Note that the calculation of the sensitivity score relies on the

coefficient of variation provided by the user for each perturbed property in the TRIM.FaTE statistics file.

Sensitivity analysis summary results files will be semi-colon delimited and can be opened in a spreadsheet program. Alternatively, they can be viewed within a TRIM.FaTE window as presented in Figure 6 by clicking the “View Results” button at the bottom of the Results view screen. A new TRIM.FaTE window will appear that contains an array of results for the current compartment(s) and chemical(s) (i.e., the compartment(s) and chemical(s) selected on the Analysis view). Note that, in the example shown, additional columns with more sensitivity results could be viewed by using the horizontal scroll bar at the bottom of the window (i.e., some columns included in Table 1 are not visible in Figure 6).

Table 1
Data Generated in Sensitivity Output File

Column Title	Description
Property	The name of the property that was varied.
Chemical	The chemical for the property that was varied, if appropriate.
ObjectType	The type of object in which the property was varied (e.g., compartment, link, chemical)
ObjectName	The name of the object for which the property was varied.
VE_Info	The volume element name (e.g., SW_lake), if the object named in the previous column is a volume element or compartment, <i>or</i> the name of the link if the object named in the previous column is a link (e.g., “from Soil - Surface, to Surface water” from SurfSoil_SE2 to SW_lake). If the object is a chemical, algorithm, or source, this column will be blank.
NumObjs	The number of objects for which the property was varied (e.g., if the property being varied is a property of three surface water bodies, this column would read “3”). The name of the first volume element is what appears in the VE_Info column.
CV	The coefficient of variation for this property (as provided in the statistics file).
Initial_Input	The initial input for the property value in the base case.
Perturbed_Input	The input value of the property after perturbation for the sensitivity analysis.
%deltaInput	The percent change in the input value between the base run and the perturbed run.
Initial_Output	The initial output value (i.e., base case) in the compartment of interest.
Perturbed_Output	The perturbed output value in the compartment of interest.
%deltaOutput	The percent change in the output value in the compartment of interest.

Column Title	Description
Sensitivity	The sensitivity is the ratio of the change in the output value (relative to the base case) to the perturbation of the input value.
Elasticity	The elasticity is a measure of the sensitivity relative to the initial values that were used and is calculated by the formula $(\Delta \text{Output} / \text{InitialOutput}) / (\Delta \text{Input} / \text{InitialInput})$.
Score	The sensitivity score is the elasticity weighted by a normalized measure of the variability of the parameter and is calculated by the formula $\text{Elasticity} * \text{CV}$ (where the CV is expressed as a fraction, not a percentage).
Absolute Score	The absolute score is the absolute value of the Sensitivity Score.

Figure 6
Sensitivity Analysis Summary File (Opened in “View Results” Window)

	Output_Compartment	Property	Chemical	ObjectType	ObjectName	CV	Initial_Input	Perturbed_Input	%deltaInput	Initial_Outg
1	Air in Air_ENE1	AirDensity_g_cm3		Compartment	Air	5.000000E-2	1.200000E-3	1.320000E-3	1.000000E1	1.204000
2	Air in Air_NNE1	AirDensity_g_cm3		Compartment	Air	5.000000E-2	1.200000E-3	1.320000E-3	1.000000E1	1.695000
3	Air in Air_NNW1	AirDensity_g_cm3		Compartment	Air	5.000000E-2	1.200000E-3	1.320000E-3	1.000000E1	2.209000
4	Air in Air_ENE1	DustDensity		Compartment	Air	3.000000E-1	1.400000E3	1.540000E3	1.000000E1	1.204000
5	Air in Air_NNE1	DustDensity		Compartment	Air	3.000000E-1	1.400000E3	1.540000E3	1.000000E1	1.695000
6	Air in Air_NNW1	DustDensity		Compartment	Air	3.000000E-1	1.400000E3	1.540000E3	1.000000E1	2.209000
7	Air in Air_ENE1	DustLoad		Compartment	Air	3.000000E-1	6.150000E-8	6.765000E-8	1.000000E1	1.204000
8	Air in Air_NNE1	DustLoad		Compartment	Air	3.000000E-1	6.150000E-8	6.765000E-8	1.000000E1	1.695000
9	Air in Air_NNW1	DustLoad		Compartment	Air	3.000000E-1	6.150000E-8	6.765000E-8	1.000000E1	2.209000
10	Air in Air_ENE1	WashoutRatio		Compartment	Air	3.000000E-1	2.000000E5	2.200000E5	1.000000E1	1.204000
11	Air in Air_NNE1	WashoutRatio		Compartment	Air	3.000000E-1	2.000000E5	2.200000E5	1.000000E1	1.695000
12	Air in Air_NNW1	WashoutRatio		Compartment	Air	3.000000E-1	2.000000E5	2.200000E5	1.000000E1	2.209000
Total	n/a	n/a	n/a	n/a	n/a	2.8499999999...	604200.003600...	664620.003960...	120.0	0.02C
Units	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/

Individual Sensitivity Run Files

The run results from which the sensitivity measures are calculated consist of model results files for each unique combination of modeled chemical, form of results (i.e., moles, mass, or concentration), and perturbed property value. The format of these files is identical to the format of the outputs from a deterministic simulation. These results files are saved in subdirectories of the analysis results folder that are named using the property names included in the analysis. For example, using the output directory specified in Figure 4, the results for the runs in which the *WashoutRatio* property was perturbed will be saved in the following directory:

c:\TRIM\data\OtherResults\SensitivityAnalysis\Air Compartments\WashoutRatio3\

As described previously, the numeral “3” was attached to the folder name by TRIM.FaTE to differentiate these results from separate runs related to the same property name (e.g., if the property *WashoutRatio* for a different compartment was added to the analysis).

2.3 Additional Information for the User

This module does not include guidance on how a sensitivity analysis should be implemented for TRIM.FaTE scenarios (i.e., how to select properties and assign new values). However, based on the results of TRIM.FaTE test cases to date, the following information is presented for the user.

- To obtain useful or relevant sensitivity results, the new value for a parameter should be set within the range of realistic values for that parameter. Units should be considered with respect to the realistic range. For example, if surface water temperature is included as one of the properties to be varied in an analysis, it may be unreasonable to assume a 10 percent variation because the surface water temperature property is set on the Kelvin scale. In these units, if the base case water temperature is set as a constant equal to 290 K, a 10 percent variation below the base case yields 261 K (i.e., 5° Fahrenheit), which represents a sub-freezing water temperature.

3. MONTE CARLO VARIABILITY AND UNCERTAINTY ANALYSES

3.1 Conceptual Background

In a Monte Carlo analysis, statistical distributions are assigned to model input parameters to represent the variability and uncertainty associated with that parameter.⁶ The distributions are sampled to provide values for each parameter, and a model run is executed. This process is repeated many times to generate a set of results. The distribution of the results that are obtained should provide information about the range and distribution of possible model outcomes.

The TRIM.FaTE model includes a Monte Carlo analysis feature. This feature, accessed via the Analyses tab of the Scenario window, has some similarity to the sensitivity analysis feature described in Section 2 in that the user selects properties to be varied, the analysis entails repeated simulations, and the results of multiple simulations are used to generate analysis summaries. However, while properties are varied *one* at a time and according to *user specified* intervals in the sensitivity analysis, *multiple* properties are varied simultaneously and randomly according to sampling of their *distributions* in the Monte Carlo analysis. In TRIM.FaTE, the user can specify that values for chosen properties are selected using either simple random

⁶ The distribution of a parameter can reflect uncertainty (i.e., a lack of knowledge regarding the actual values of model input variables and of physical systems) and/or variability (i.e., the diversity or heterogeneity in a population or parameter). However, in practice, it is difficult to assess these concepts individually, and when a parameter distribution is developed, it is rarely separated into components of uncertainty and variability. This module explains the TRIM.FaTE uncertainty analysis feature in general terms and does not address these individual concepts explicitly. Refer to Volume I of the TRIM.FaTE Technical Support Document (EPA 2002) for additional information regarding uncertainty and variability and the conceptual implementation of Monte Carlo analyses within the TRIM.FaTE framework.

sampling or Latin hypercube sampling. For additional background information on the application of Monte Carlo analyses within TRIM.FaTE, refer to Chapter 6 in Volume I of the TRIM.FaTE TSD (EPA 2002).

Currently, the TRIM.FaTE Monte Carlo feature supports sampling from the following four distribution types for properties in Monte Carlo analyses: uniform, normal, lognormal, and triangular. For each property that is included in the Monte Carlo analysis, a type of distribution must be specified along with associated required information (i.e., minimum, maximum, coefficient of variation, and distribution type) in the TRIM.FaTE statistics data file (i.e., the same file described for the sensitivity analysis). The CV is the standard deviation of the input parameter expressed as a fraction of the parameter's mean (i.e., [std. deviation]/[mean]). For additional information on the statistics file required by TRIM.FaTE, refer to Section 2.2 of Module 3, Library and Scenario Data Files.

3.2 Setting up and Running a Monte Carlo Analysis in TRIM.FaTE

3.2.1 Accessing the Monte Carlo Analysis Feature

To access the Monte Carlo feature, go to the Analyses view within a TRIM.FaTE scenario and select “Monte Carlo Analyses” from the drop-down box in the upper left corner (see Figure 1). The window below the drop-down box (labeled the “Monte Carlo Analyses” pane) will list the Monte Carlo analyses that are in the current scenario. By default, an “empty” Monte Carlo analysis is created with every scenario and labeled with the scenario name. To add a new Monte Carlo analysis, click the “New” button and enter a name when prompted. The other buttons below the Analyses window can be used to delete, rename, or open analyses.

To view specific information about an analysis, the user can bring up the Monte Carlo Analysis Window. This window is opened by selecting a Monte Carlo analysis name and clicking the “Open” button or by double-clicking the analysis name in the Monte Carlo Analyses pane. (Alternatively, the user can access this window by clicking on the “Run” pull-down menu in the scenario window, selecting “Monte Carlo Analysis,” and then selecting the desired analysis in the dialog box that appears.) Like the Sensitivity Analysis Window, the Monte Carlo Analysis Window contains tabs that allow the user to access three views:

- (1) The **Properties** view, used to specify the properties to be varied and the general set-up of the analysis;
- (2) The **Runs** view, which conveys information about the status of the Monte Carlo run; and
- (3) The **Results** view, used to specify how the results of the analysis will be calculated.

Each of these views is described in more detail in the sections that follow.

3.2.2 Selecting Properties to be Varied

The process for selecting properties to be included in the Monte Carlo analysis (i.e., properties for which values will be varied according to a specified distribution) is nearly identical to the process used to add properties to the sensitivity analysis (see Section 2.2). The steps to be followed are described here.

- (1) Designate a “Current” Monte Carlo analysis by selecting one of the analyses from the “Monte Carlo Analyses” pane below the drop-down menu and then clicking the “Current” button below the pane. After specifying an analysis as “current,” all properties added will be included in this analysis until the user designates a different analysis as current.
- (2) Click on the tab in the main Scenario window (e.g., Chemical Tab, Compartments Tab) that contains the object property to be varied.
- (3) Select the object(s) containing the property or properties to be varied. This can be done manually, by finding the object and clicking on it to highlight the object name, or with assistance from the TRIM.FaTE “Select” tool by clicking the “Select” button and using the dialog window that appears (this can help the user find and select multiple, similar objects). **Note that objects in the scenario with properties to be varied must be selected from the Outdoor Environment pane on the left, *not* the library pane that appears in the center.**
- (4) Click the “Properties” button to bring the properties for the highlighted object(s) up in the Property Editor.
- (5) Select one or more properties by highlighting them in the Property Editor.
- (6) From the Add pull-down menu at the top of the scenario window, select “Add Selected Properties to Monte Carlo Analysis [Analysis Name].” It is important to note that the selected properties will be added for ALL of the selected objects. When the property or properties have been added to the Monte Carlo analysis, a message will appear at the bottom of the TRIM window stating “### properties added to Monte Carlo analysis for ## objects” (each “###” will be replaced with the appropriate number depending on what the user has selected).

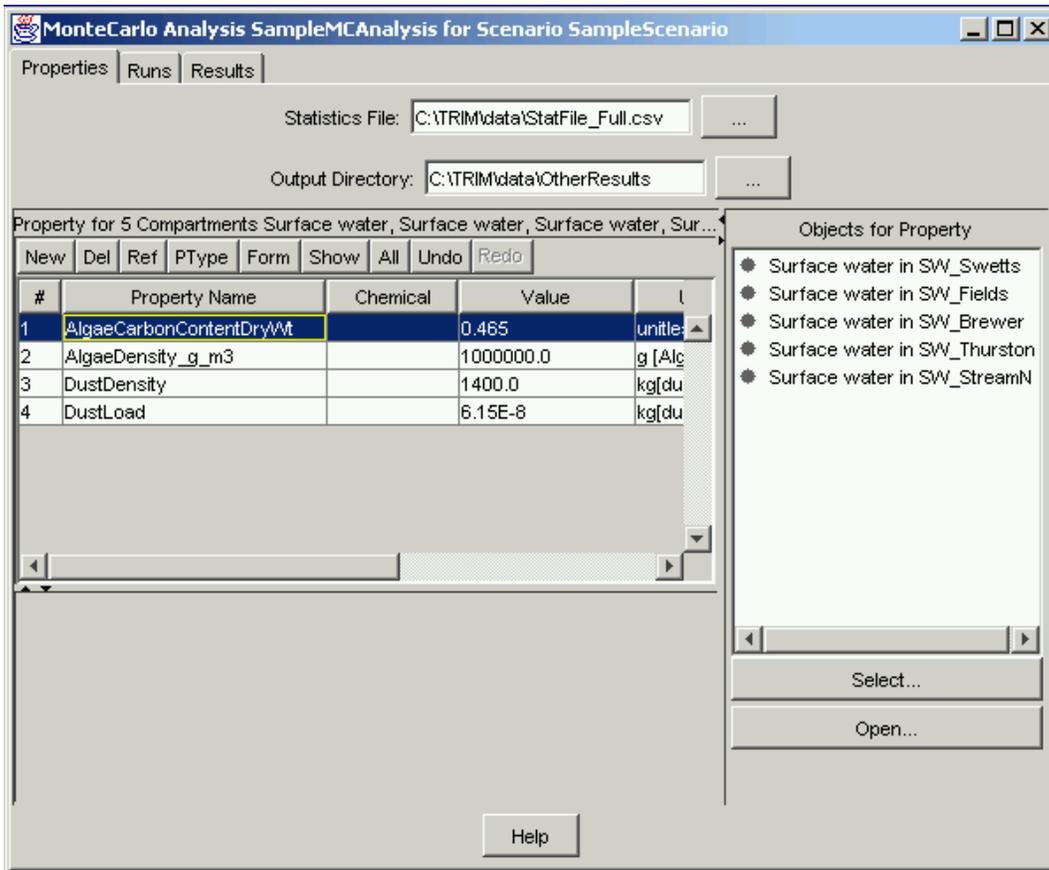
To view the properties that have been added to an analysis, click on the Properties tab within the Monte Carlo analysis window. The properties selected for inclusion in the Monte Carlo analysis will be listed in the left-hand window along with the original property values that were assigned for the simulation. The objects to which these properties apply will be shown in the right-hand window. An example of the Properties view within the Monte Carlo Analysis window for a sample analysis entitled “SampleMCAnalysis” is presented in Figure 7. In this example, four properties have been added to the current analysis. The highlighted property in the example (*AlgaeCarbonContentDryWt*) has been added for the five surface water compartments shown in the right-hand object window.

To set up a separate Monte Carlo analysis, designate a different run as “Current” as described in Step 1, and repeat Steps 2 through 6 until all of the desired properties have been added to the run. Properties can be removed from an analysis by highlighting the unwanted property and clicking the “Del” button above the properties table.

3.2.3 Specifying Property Value Distributions and Output Directory

The Properties view is also used to specify where the model can find the statistical distribution data for the selected properties. These data are stored in the statistics file for the scenario (which can be the same file in which the coefficients of variation for the sensitivity analysis are stored). Refer to Section 2.2 of Module 3, Library and Scenario Data Files, for information on how to set up and format the statistics data file. The user should specify the location of the statistics file using the “Statistics File” window at the top of the Properties view (see Figure 7). The file name and directory can be entered directly by the user (by clicking in the window and typing the directory and name of the file), or the user can click the “...” button to bring up a file browser that allows the user to find and select the appropriate file.

Figure 7
Properties View, Monte Carlo Analysis Window

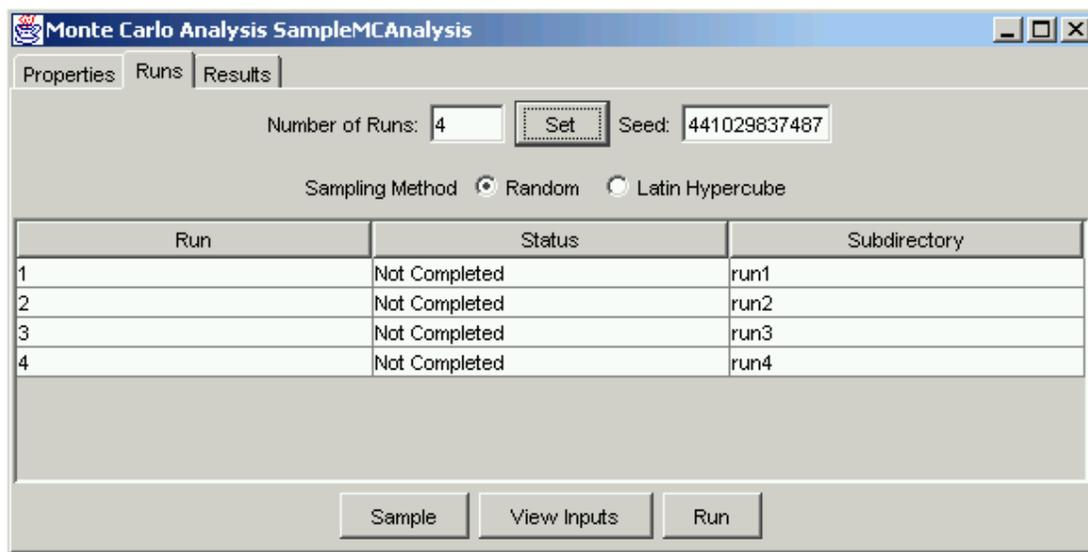


In addition, the directory where the Monte Carlo analysis output will be saved is also specified on the Properties view in the “Output Directory” window. Note that this procedure is slightly different from setting the output directory for the sensitivity analysis (when it is specified via the Runs view). In the example presented here, the folder “c:\TRIM\data\OtherResults” has been specified as the output directory.

3.2.4 Configuring and Executing the Monte Carlo Runs

Information about the status and location of the Monte Carlo analysis model runs that will be performed is summarized on the Runs view (see Figure 8). The user defines three analysis settings on this view.

Figure 8
Runs View, Monte Carlo Analysis Window



- (1) Set the **number of runs** (iterations) to be performed by entering a value in the “Number of Runs” text box at the top of the view. If simple random sampling is being used to sample property values from the distributions, the number of runs can be set to any positive integer. However, if the Latin Hypercube sampling option has been selected, the number of runs must be set to an even multiple of the number of properties that are included in the analysis. For example, if four properties will be varied, the number of runs must be set to a multiple of four (e.g., 12, 400, 1000). After the number of runs has been entered, click the “Set” button next to the window. The runs will appear in the table along with each run’s status and the sub-directory to which the run results will be saved.

- (2) Set the **numerical seed** to be used for initializing the random number generation in the “Seed” text box. This step is optional; if not specified, TRIM.FaTE will set the seed to a random number by default. Setting the seed to certain number may be useful when comparing separate uncertainty runs that vary similar parameters. Note that if a previous Monte Carlo analysis is to be replicated, the user must specify the seed used in the original analysis here.
- (3) Select the **sampling method** to be used (i.e., either Random or Latin Hypercube).

After following these steps, the user must click the “Sample” button at the bottom of the view to sample from the distributions for each property. If one of the Properties does not exist in the statistics file, TRIM.FaTE will notify the user via an error message, and the user must edit the statistics file and resave the file to the location specified on the Properties view. The property values assigned by TRIM.FaTE for the Monte Carlo analysis can be viewed by clicking the “View Inputs” button at the bottom of the view (the values will appear in a new TRIM.FaTE window).

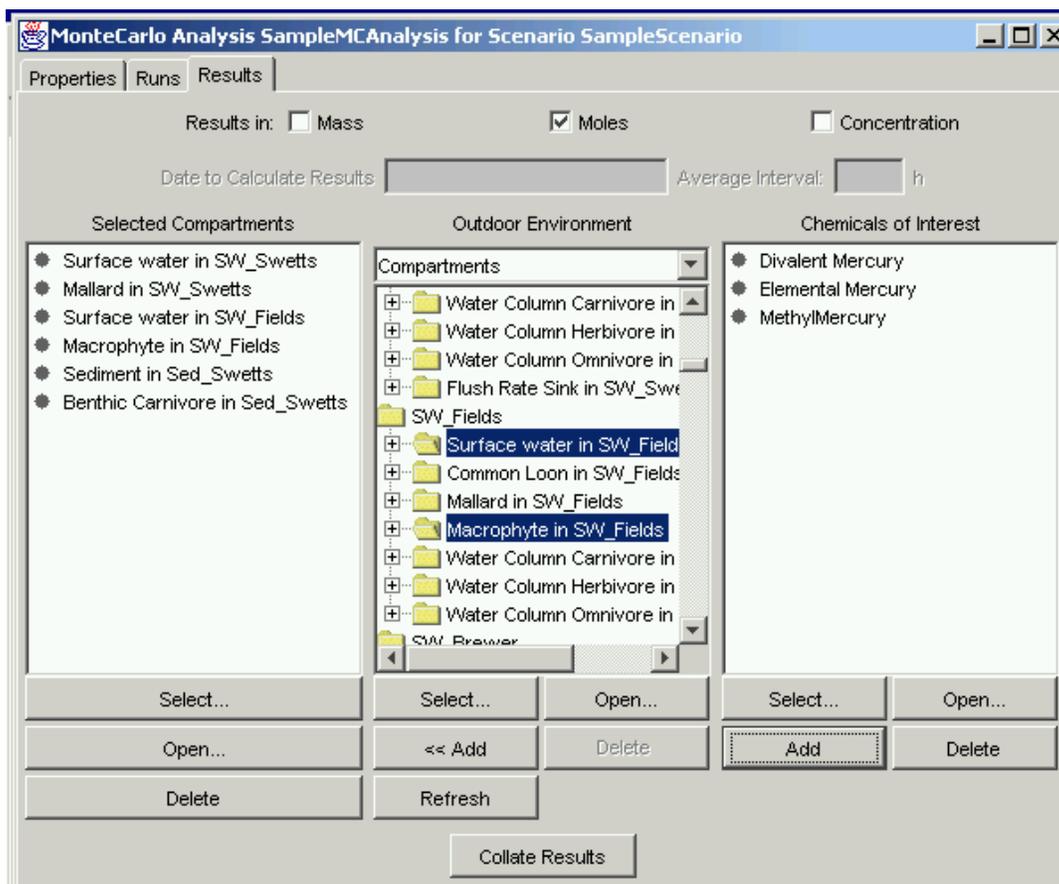
To start running the simulations for the Monte Carlo analysis, the user should select the runs to be performed by highlighting them in the table and clicking the “Run” button at the bottom of the view. A small window containing a “Stop” button will appear; this button can be clicked to abort the batch of runs.

Important Note: Depending on the complexity of the scenario and, likely, whether the model is run in dynamic or steady state mode, including a large number of runs in a single batch may cause model execution problems. Running multiple batches of runs (each containing a subset of the runs comprising a Monte Carlo analysis) may reduce the potential for this problem.

3.2.5 Collating Monte Carlo Analysis Results

After the simulations for the Monte Carlo analysis have been completed, the results can be collated (i.e., combined to a single summary file) via the Results view (see Figure 9). The following six steps should be completed to collate the results.

Figure 9
Results View, Monte Carlo Analysis Window



- (1) Select the units for which the results should be processed (i.e., mass, moles, or concentration, or any combination of these) by clicking the appropriate button(s) at the top of the view.
- (2) For analyses involving results from dynamic simulations, specify a date and time from the simulations for which the sensitivity results will be calculated. This date and time should be entered in the window below the check boxes for results units.⁷ If the simulation results did not include that particular date and time, TRIM.FaTE will use the last time point at which results were output. For example, if results for a run are calculated every four hours (e.g., hours 00:00, 4:00:00, 8:00:00, ... during a day), the results for 4:00:00 would be used if the user enters any time ranging from 4:00:00 to 7:59:59. If the Monte Carlo results are to be collated from averaged results, TRIM.FaTE will use the results for the time interval inclusive of this date and time

⁷ Steps 2 and 3 are not required if you are running the Monte Carlo simulations in steady-state mode. If the simulation is being run in steady-state mode, TRIM.FaTE will automatically shade the date and averaging interval windows and the user will not be able to access them.

- (3) For analyses involving results from dynamic simulations, specify the Average Interval for the simulation results (i.e., the time interval in hours within which the model simulation outputs will be averaged) in the text box to the right of the date window.⁷ This interval must be specified as either a whole-number multiple of the scenario output time interval or “monthly” or “annual” (see example in Step 3 of section 2.2.5, Performing Sensitivity Calculations). For the calculations to be performed on the simulation results as they are output from TRIM.FaTE (i.e., without subsequent averaging), enter the value (in hours) of the simulation output time step.⁸
- (4) In the “Outdoor Environment” pane in the middle of the view, specify the Compartments for which the results will be collated by highlighting them (multiple compartments can be selected simultaneously using the shift or control keys, and the “Select” button can be used to find and select groups of compartments). Click the “<<Add” button to add the selected compartments to the analysis. The selected compartments should appear in the “Selected Compartments” pane on the left side of the view.
- (5) To specify which chemicals to include in the analysis, click the “Add” button below the “Chemical of Interest” pane on the right-hand side of the view to bring up a “Select Chemicals” dialog box. Select the appropriate chemical(s) from the list that appears. The selected chemicals should appear in the “Chemicals of Interest” pane.
- (6) Click the “Collate Results” button at the bottom of the view to collate the results (i.e., bring all of the final values in the selected compartments for the selected chemicals together into a single summary file).

Once the results have been collated by TRIM.FaTE, there will be one results file for each combination of chemical and form of results. For example, if mass and concentration are both selected as the form of results and three chemicals are included in the simulations, six results files will be generated. Each results file that is created will contain an array of results, with results for each user-specified compartment saved in the columns of the array and the results for each run saved in the rows in the array.

To complete a Monte Carlo analysis for the same set of properties (i.e., to collate files at a different date/time by entering a different date/time in Step 2), Steps 2 through 6 can be repeated. Note, however, that TRIM.FaTE will *not* change the name of the output file(s). That is, existing collated results files will be overwritten if Steps 2 through 6 are repeated for the same TRIM.FaTE Monte Carlo analysis. Therefore, if additional calculations are to be performed to produce new collated results files, the user should first rename the previous results files in order to preserve them.

⁸As discussed in Module 12, the output time step is the product of the *simulationStepsPerOutputStep* and the *simulationTimeStep* properties (which are specified in the scenario view of the scenario window).

For collating Monte Carlo results from dynamic TRIM.FaTE simulations, Steps 2 and 3 specify which results (i.e., from which time point or, for averaged results, which time interval) to use in collating the Monte Carlo results. The “Average Interval” input provides a way for the user to use time-averaged results rather than an instantaneous point in time (see the analogous example in Section 2.2.6).

Each collated results file will be saved in a subdirectory of the output directory specified on the Properties view of the Monte Carlo analysis window (see Figure 6). The subdirectory will be named “MonteCarlo” and will include an additional subdirectory with the name of the analysis. The model results for individual runs will also be saved in folders within the scenario name subdirectory of the “MonteCarlo” subdirectory according to the names specified on the Runs view. For example, using the output directory specified in Figure 4, collated divalent mercury concentration results for the Monte Carlo analysis will be saved in a summary file named *SampleScenario_Divalent_Mercury_conc.txt* in the following folder:

```
c:\TRIM\data\OtherResults\MonteCarlo\SampleMCAnalysis\
```

The run results from which the collated results were generated (for each run, one file per unique combination of modeled chemical and form of output) are saved in the folder listed above, in subdirectories that correspond to the run name. For example, the first run in the above example will be saved in a subdirectory named Run1, or:

```
c:\TRIM\data\OtherResults\MonteCarlo\SampleMCAnalysis\Run1\
```

A collated Monte Carlo results file can be used to calculate statistics on the outputs of the Monte Carlo simulations (e.g., mean or median values in the compartments of interest; percentiles). TRIM.FaTE does not calculate such statistics internally; they must be calculated off-line by the user.

4. REFERENCES

U.S. EPA. 2002. TRIM.FaTE Technical Support Document. Volume I: Description of Module. EPA 453/R-02-011a. Office of Air Quality Planning and Standards. Research Triangle Park, NC.