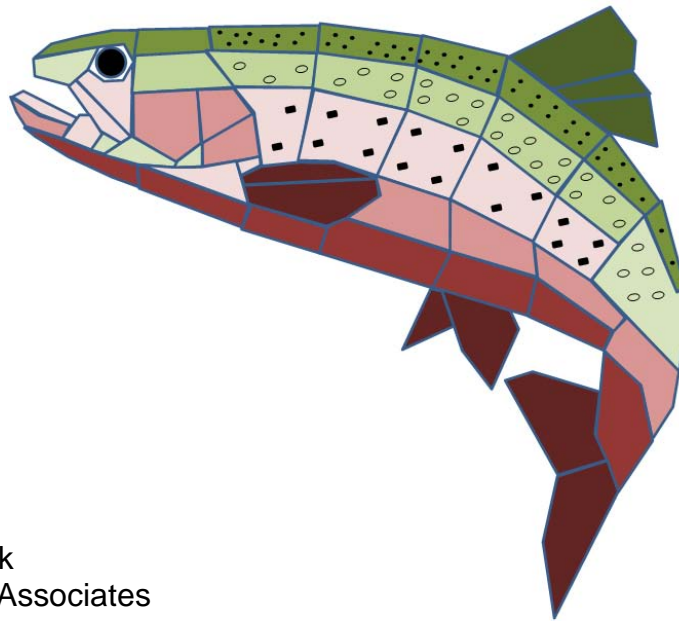


# ***inSTREAM*: The Individual-based Stream Trout Research and Environmental Assessment Model**

**Version 5.0**

## **Limiting Factors Tool Guide**



### **Prepared by:**

Steven F. Railsback  
Lang, Railsback & Associates  
Arcata, CA  
[www.LangRailsback.com](http://www.LangRailsback.com)

Bret C. Harvey  
USDA Forest Service, Pacific Southwest Research Station  
Redwood Sciences Laboratory  
Arcata, CA

Colin Sheppard  
Arcata, CA

Last updated: February 3, 2012

Copyright 2011, 2012 Lang, Railsback & Associates

## Table of Contents

<b>1. Objectives and Overview of the Limiting Factors Tool .....</b>	<b>3</b>
1.1. What are “limiting factors”? .....	3
1.2. The approach: sensitivity analysis with parameter uncertainty .....	3
1.3. Interactions among factors.....	4
1.4. When to use the LFT .....	4
1.5. LFT parameters .....	4
<b>2. Limiting Factor Experiments.....</b>	<b>5</b>
2.1. Base flow .....	5
2.2. Food availability .....	6
2.3. Winter water temperature.....	6
2.4. Summer water temperature .....	6
2.5. Spawning gravel availability .....	7
2.6. Velocity shelter availability .....	7
2.7. Hiding cover availability .....	7
2.8. Piscivory risk.....	8
2.9. Redd scouring.....	9
<b>3. Uncertainty Consideration .....</b>	<b>10</b>
<b>4. Results Analysis and Interpretation .....</b>	<b>12</b>
4.1. Analysis methods and interpretation of results .....	12
4.2. Results formats.....	12
<b>5. Parameters Controlling the Limiting Factors Tool .....</b>	<b>13</b>
<b>6. Using the Tool.....</b>	<b>17</b>
6.1. Running an LFT analysis .....	17
6.1. Error checking.....	20
<b>7. Computational Considerations .....</b>	<b>21</b>
<b>LFT Parameter Index.....</b>	<b>22</b>

## 1. Objectives and Overview of the Limiting Factors Tool

The Limiting Factors Tool (LFT) is a software add-on to the *inSTREAM* individual-based trout model. It automates simulation experiments that address an important habitat management question: which manageable habitat factors have high potential for enhancing fish populations?

**The LFT was originally developed for the *inSALMO* individual-based salmon model, and this document is adapted—largely verbatim—from the *inSALMO* LFT documentation.**

(That documentation can be cited as: Railsback, S. F., B. C. Harvey, and J. L. White. 2011. *inSALMO* version 1.0: Model improvements and demonstration application to Chinook salmon spawning, incubation, and rearing in Clear Creek, California. Report prepared for U.S. Bureau of Reclamation by Lang, Railsback & Associates, Arcata, CA.) The original *inSALMO* LFT was funded by the U. S. Bureau of Reclamation, Mid-Pacific Regional Office, under contract R09PS20027 with Lang, Railsback & Associates.

### 1.1. What are “limiting factors”?

“Limiting factors” is a common but not completely accurate term for environmental variables or processes that strongly affect a fish population. A traditional notion of limiting factors is that at any one time there is one factor that “limits” a population from growing; typically, food availability, predation pressure, reproductive output, and habitat space are considered potential limiting factors. Improving the limiting factor is expected to increase the population while the “non-limiting” factors are expected to have little effect.

In contrast, individual-based models such as *inSTREAM*, in which the individual fish have adaptive behaviors to trade off growth and predation risk, predict that multiple factors can have strong effects on population growth at any one time, and there is not necessarily a threshold above (or below) which a factor such as food availability has no effect (Railsback, S. F., and B. C. Harvey. 2011. Importance of fish behaviour in modelling conservation problems: food limitation as an example. *Journal of Fish Biology* 79:1648–1662. doi:10.1111/j.1095-8649.2011.03050.x). A more useful way to think about limiting factors, at least in the context of *inSTREAM*, is: how strongly do various environmental factors affect the population, and which factors offer high potential for population enhancement?

The *inSTREAM* LFT was designed to evaluate a number of habitat factors considered especially important for management of stream trout populations. In some cases, other factors may seem important to evaluate; when this occurs, experiments similar to those performed automatically by the LFT can usually be conducted using *inSTREAM*'s Experiment Manager (see the model's software documentation) and parameters or inputs that control those factors. However, *inSTREAM* and the LFT can only evaluate factors and their effects if those effects are explicitly represented in *inSTREAM*. Some potentially important limiting factors such as siltation of spawning gravel or angler harvest are not represented in *inSTREAM* and hence cannot be evaluated with the LFT.

### 1.2. The approach: sensitivity analysis with parameter uncertainty

The LFT addresses its objective by implementing a sensitivity analysis approach. It includes built-in experiments for habitat variables that (a) likely have—or are often believed to have—strong effects on simulated trout populations, and (b) potentially could be changed via management actions. Each LFT experiment runs the model repeatedly using a wide range of

values for one factor. The tool then computes and compares the degree to which the model's key outputs—the number and biomass of adult trout—responds to each factor.

Results of the limiting factors experiments depend on the values used for *inSTREAM*'s parameters, and these parameters are uncertain. The LFT therefore explicitly addresses how robust its results are with respect to parameter uncertainty. Each experiment is executed multiple times using combinations of several parameters which are particularly important and uncertain, and the results analysis considers how consistent the importance of factors is among these multiple executions (Section 3).

### **1.3. Interactions among factors**

Managers using the LFT are likely to often be concerned about interactions among factors: would the LFT produce different results and lead to different conclusions if different conditions were assumed in the model inputs? Would it produce different results, for example, if the number of spawners was low instead of high, or for years with high vs. low winter flows?

Designing the LFT to automatically evaluate such interactions would make it substantially more complex and difficult to use and interpret. Therefore, the way to address potential interactions is the simplest: by just executing the LFT for several sets of conditions (low vs. high spawner abundance; low vs. high runoff years) and looking for differences in results.

### **1.4. When to use the LFT**

For its primary purpose of analyzing limiting factors to support management decisions, the LFT should of course be used after *inSTREAM* has been fully applied and calibrated to the study site, with the best available input and parameter values. (Even then, LFT results might identify processes and parameters that deserve additional scrutiny because of their importance.)

However, the LFT can be very useful in earlier stages of applying *inSTREAM*. Its sensitivity experiments provide information that can be useful, for example, in calibrating the model. The LFT can be thought of as a model exploration tool potentially useful throughout the model application process.

### **1.5. LFT parameters**

The experiments executed by the LFT are controlled by a number of parameters that belong not to *inSTREAM* but just to the LFT. The parameters are defined throughout Section 2, summarized with example values in Section 5, and indexed.

## 2. Limiting Factor Experiments

The following experiments are built into the LFT. Each experiment executes *inSTREAM* for each of a specified number of scenarios. Each scenario uses standard inputs except for the one factor being analyzed (and the parameters varied to evaluate parameter uncertainty effects; Section 3). The scenarios vary that factor over a wide range, from below to above the values represented in the standard input. The number of scenarios is set by the LFT parameter *numScenarios*, which must have a value of 2 or higher. Section 7 discusses computational issues that are important in selecting the number of scenarios.

To make results easy to interpret, the LFT analyzes only two outputs from *inSTREAM*: abundance and total biomass of adult trout. The specific measures of abundance and biomass are set by two LFT parameters. The parameter *resultsAgeThreshold* defines the age (years, with age 1 attained on the individual's first January 1) of simulated trout included in the LFT results: trout with age equal or exceeding *resultsAgeThreshold* are included in results. The value of *resultsAgeThreshold* can be set to reflect the age at which trout at a study site typically spawn or (if angling is an important objective) become catchable.

The second parameter defining results is *resultsCensusDay* (day of year, in mm/dd format): LFT results are "censused" from the model on this day of each simulated year, and then averaged over the model run. Typically, *resultsCensusDay* is the day of the last date simulated and 364 days after the model starts. Model runs are typically started on (for example) October 1 and end on September 30 of a later year, so *resultsCensusDay* would be set to "9/30". This way, LFT results do not include the first simulation year (when results are especially controlled by the initial population conditions) and do include the last simulated date.

As an example, if *resultsAgeThreshold* is set to 1 and *resultsCensusDay* is set to 10/01, then the LFT will measure the number and total biomass (g) of all age 1 and older trout on each October 1<sup>st</sup> in the simulations. The results analyzed for each scenario will be the mean, over all simulated years, of these October 1 values.

The first simulation day is never included in results analyzed by the LFT. If simulations run from 10/1/2000 to 10/1/2005 and *resultsCensusDay* is set to "10/1", then results from October first of 2001 through 2005 will be analyzed by the LFT.

It is very important to understand that the LFT applies the same experiments and parameters to all habitat reaches in an *inSTREAM* application. If the application includes multiple reaches differing widely in factors such as flow and temperature, the results may not be as meaningful.

The LFT also applies its changes to all trout species represented in the *inSTREAM* application.

### 2.1. Base flow

This experiment examines flow as a limiting factor. The base flow experiment is designed around the assumption that flows at the study site are controlled by reservoir releases and, possibly, in part by tributary inflow. The experiment simulates the effect of changing the reservoir releases that make up base flow, without changing the tributary inflows and flow variability in the flow input data.

Each scenario adds a constant to the daily flow input, for all simulated days. The constant flow change varies among the scenarios over a range defined by parameters *baseFlowRangeLow* (which is typically negative) and *baseFlowRangeHigh* (both with units of m<sup>3</sup>/s). The value of *baseFlowRangeLow* is typically negative, producing flows below the standard values. If the sum of *baseFlowRangeLow* and the standard flow is less than zero (so the flow would be negative) for any day, then an error message is created and the LFT will not proceed.

## **2.2. Food availability**

Production of invertebrate food for trout is rarely managed directly, but there are management actions that can indirectly affect food availability. Examples include watershed restoration efforts to control fine sediment inputs and supplementation of nutrient availability.

The food availability experiment varies both kinds of food represented in *inSTREAM*, assuming that changes in productivity would be reflected in both drift and stationary search food. The standard values of reach habitat variables *habDriftConc* and *habSearchProd* are both multiplied by a ratio that varies among scenarios within a range defined by LFT parameters *foodAvailabilityRatioLow* and *foodAvailabilityRatioHigh*. Food availability can vary widely (over orders of magnitude) among measured values, in part because it is naturally variable and difficult to measure precisely. However, useful values of *inSTREAM*'s food availability parameters vary over narrower ranges because they represent long-term and spatial averages. Useful ranges for LFT experiments are *foodAvailabilityRatioLow* = 0.5 and *foodAvailabilityRatioHigh* = 2.0.

## **2.3. Winter water temperature**

Winter water temperature is defined as daily mean water temperature, for days in a winter period. The beginning and ending dates of the “winter” period are defined by the LFT parameters *winterTemperatureStartDay* and *winterTemperatureEndDay*. For example, in California winter can be defined as November-April, the months when temperatures are generally lowest.

The experiment adds a constant to each day's daily temperature input, for days within the winter period. This constant varies among the scenarios over a range set by the parameters *winterTemperatureRangeLow* (which is typically negative) and *winterTemperatureRangeHigh* (°C). Any negative temperatures are set to zero.

## **2.4. Summer water temperature**

The summer water temperature experiment is identical to the winter temperature experiment, except that temperatures are manipulated for days in warm months. The “summer” period is defined by parameters *summerTemperatureStartDay* and *summerTemperatureEndDay*. (The summer period can overlap with the “winter” period used in the winter water temperature experiment. Users can select the summer and winter periods to include or exclude specific life stages such as egg incubation and fry rearing; and to affect fall-spawning vs. spring-spawning trout differently.)

The range of the constant (°C) added to each summer day's temperature input is controlled by the LFT parameters *summerTemperatureRangeLow* (typically negative) and *summerTemperatureRangeHigh*. Any negative temperatures are set to zero.

## 2.5. Spawning gravel availability

This experiment examines the effect of spawning gravel availability by varying (a) how many cells have spawning gravel, and (b) the amount of cell area containing gravel. The parameters *gravelAvailabilityRangeLow* and *gravelAvailabilityRangeHigh* define the limits of a variable (referred to here as *G*) for the amount of gravel simulated, relative to the amount in the standard input for the modeled sites. The spawning gravel experiment manipulates the variable *cellFracSpawn*, the fraction of cell area that is spawning gravel, for each habitat cell.

For values of *G* less than 1.0, the experiment makes these changes to each cell:

- The standard value of *cellFracSpawn* is multiplied by *G*. Hence, the area of gravel in cells containing gravel is reduced.
- If the standard value of *cellFracSpawn* is greater than zero, its value is set to zero if a random Bernoulli trial is true. The probability of this trial being true is  $(1-G)$ . Therefore, the number of cells containing gravel is reduced.

For values of *G* equal to 1.0, no changes to spawning gravel are made. When *G* is greater than 1.0, these changes are made to each cell:

- The value of *cellFracSpawn* is multiplied by *G*, but limited to a maximum of 1.0. The area of gravel in cells containing gravel is therefore increased.
- If *cellFracSpawn* is zero, its value is set to 0.5 if a random Bernoulli trial is true. The probability of this trial being true is  $(G-1.0)$ . Therefore, the number of cells containing gravel is increased.

The parameters *gravelAvailabilityRangeLow* and *gravelAvailabilityRangeHigh* are limited to values that produce meaningful values of *G*. When *G* is 0.0, there will be no spawning gravel, and when *G* is 2.0 all cells will have gravel. Therefore, the value of *gravelAvailabilityRangeLow* must be between 0.0 and 1.0, and *gravelAvailabilityRangeHigh* must be between 1.0 and 2.0.

## 2.6. Velocity shelter availability

The velocity shelter experiment examines the value of cover such as large rocks that reduce the swimming speed of drift-feeding fish. It manipulates the values of the cell variable *cellFracShelter*. The experiment works exactly as the spawning gravel experiment does, using parameters *shelterAvailabilityRangeLow* and *shelterAvailabilityRangeHigh* to control the range of velocity shelter availability.

## 2.7. Hiding cover availability

The hiding cover experiment evaluates the potential benefit of cover such as submerged trees and rocks with crevices, which provide places to hide from predators. It manipulates the cell variable *cellDistToHide*, a characteristic distance (m) between fish feeding in the cell and the nearest hiding cover. The risk from both fish and terrestrial predators (birds, mammals) in *inSTREAM* is sensitive to small values of *cellDistToHide* but not to large values: hiding cover is assumed to reduce predation risk only if it is close.

The method for evaluating hiding cover as a limiting factor is similar to that for spawning gravel. The LFT parameters *hidingCoverRangeLow* and *hidingCoverRangeHigh* define the limits of a variable *H* for the simulated availability of close hiding cover in each scenario, relative to the standard input data. “Close” hiding cover is defined by the LFT parameter *hidingCoverThreshold*: hiding cover in a cell is assumed to be valuable to trout if *cellDistToHide* is less than *hidingCoverThreshold*. The value of *hidingCoverThreshold* can, for example, be set to a value of 2.0 m.

For values of *H* less than 1.0, the experiment makes these changes to each cell:

- The value of *cellDistToHide* is divided by *H*. Hence, the distance to hiding cover is increased. If *H* is zero, *cellDistToHide* is set to 99 m, a distance at which hiding cover has no benefit.

For values of *H* equal to 1.0, no changes to *cellDistToHide* are made. When *H* is greater than 1.0, these changes are made to each cell:

- The value of *cellDistToHide* is divided by *H*. The distance to hiding cover is therefore reduced in all cells.
- If *cellDistToHide* is greater than *hidingCoverThreshold*, its value is set to  $(0.5 \times \text{hidingCoverThreshold})$  if a random Bernoulli trial is true. The probability of this trial being true is  $(H-1.0)$ . Therefore, the number of cells providing close hiding cover is increased.

The parameters *hidingCoverRangeLow* and *hidingCoverRangeHigh* are limited to produce meaningful values of *H*. When *H* is 0.0, distances to hiding cover will be high for all cells. When *H* is 2.0 all cells will have close hiding cover. Therefore, the value of *hidingCoverRangeLow* must be between 0.0 and 1.0, and *hidingCoverRangeHigh* must be between 1.0 and 2.0.

Users should interpret results of this experiment carefully, as it does not consider the possibility that cover provided to protect trout could also shelter piscivorous fish that prey on them, at sites where such piscivores are present.

## 2.8. Piscivory risk

This experiment varies the *inSTREAM* parameter (*mortFishAqPredMin*) that controls the risk of trout being eaten by fish. This parameter is a daily survival probability. It is generally easier to think about predation as a risk (probability of being eaten per day); survival equals one minus risk. The experiment varies the fish predation risk (one minus *mortFishAqPredMin*) over a range.

The range of variation in risk, as a fraction of its standard value, is defined by the LFT parameters *piscivoryRiskRatioLow* and *piscivoryRiskRatioHigh*; typical values for these parameters are 0.9 and 1.1. (These risks are *daily* probabilities of dying; survival over a month, for example, varies sharply with seemingly small changes in daily risk.)

The risk for each scenario is set by multiplying  $(1.0 - \text{mortFishAqPredMin})$  by the risk ratio. A typical value of *mortFishAqPredMin* is 0.92, corresponding to a daily risk of 0.08. For a range of




0.9 to 1.1, this risk would range among scenarios from 0.072 to 0.088. The values of *mortFishAqPredMin* then range 0.928 to 0.912. The resulting probability of surviving 30 days then ranges 0.11 to 0.063.

(Keep in mind that *mortFishAqPredMin* is the daily survival of very small trout in the riskiest possible habitat; the actual survival of individual trout is adjusted upwards in accordance with factors such as fish length, depth, and distance to hiding cover.)

Unlike the other experiments, there is no clear and direct relation between the variable representing piscivory as a “limiting factor” and any management variable that could be measured or controlled in the field. While we expect that the value of *mortFishAqPredMin* should decrease when the density of predatory fish such as bass and pikeminnow increases, the relation between piscivore density and simulated predation risk is not clear. The lack of a clear relation is partly because it is not clear how the density of predators affects the individual risk for trout; doubling the number of predators may not double the risk because, e.g., predators interfere with each other. The ability of model fish to compensate for changes in risk by changing their choice of habitat also makes the relation between piscivore density and trout survival complex.

## **2.9. Redd scouring**

This experiment evaluates the importance of scouring mortality of redds during high flows, by varying a parameter that controls how vulnerable a redd is to being scoured. The redd scour in *inSTREAM* represents the probability of each redd being scoured each time the flow peaks during incubation; this probability is a function of parameters relating reach-scale bed shear stress to flow and a parameter (*mortReddScourDepth*) for how deeply buried redd egg pockets are. The following figure (from a site with *inSTREAM* redd scour parameters *habShearParamA* = 0.010 and *habShearParamB* = 0.49) illustrates that, especially at flows where scouring is just initiated, the extent of scouring is strongly affected by *mortReddScourDepth*: scouring is more extensive and occurs at lower flows when *mortReddScourDepth* is lower.



The redd scour experiment simply varies the value of *mortReddScourDepth*, with scour being more common at low values. The standard value of *mortReddScourDepth* is multiplied by a factor bounded by the parameters *reddScourRatioLow* and *reddScourRatioHigh*.

### 3. Uncertainty Consideration

This section describes how the LFT considers parameter uncertainty in *inSTREAM*. Uncertainties in *inSTREAM* are discussed in Section 15 of the *inSTREAM* version 4.2 documentation (Railsback, S. F., B. C. Harvey, S. K. Jackson, and R. H. Lamberson. 2009. *InSTREAM: the individual-based stream trout research and environmental assessment model*. PSW-GTR-218, USDA Forest Service, Pacific Southwest Research Station, Albany, California.). Those analyses showed that primary results (e.g., predicted fish abundance) can be quite sensitive to some parameters that have relatively uncertain values, but *relative* results (e.g., the ranks of alternative management scenarios; the relative strength of potential limiting factors) are typically much less affected by parameter uncertainty. Demonstrating this robustness is important, though, for developing confidence in model results. Therefore, the LFT uses a limited form of parameter uncertainty analysis to evaluate how robust its results are.

Parameter uncertainty analysis usually involves running a model repeatedly using many combinations of values for all the uncertain parameters. This full approach would be computationally infeasible for the LFT because of how many parameters *inSTREAM* has and because the limiting factors evaluation by itself requires many model runs.

To make the uncertainty consideration feasible, the LFT allows the user to identify up to three habitat or fish parameters that are varied to represent parameter uncertainty. The user also specifies low and high values for these uncertainty parameters, and the number of values for each parameter. The LFT then executes each scenario of each experiment once for each value of each uncertainty parameter. Hence, if three uncertain parameters are used, and three values

specified for each, then each scenario is simulated 27 times. Computational considerations (Section 7) typically limit experiments to three values of each parameter, especially if two or three uncertainty parameters are used.

The number of uncertainty parameters is set by the LFT parameter *numUncertaintyParams*, which can have values between 0 and 3. The number of values per uncertainty parameter is set by the LFT parameter *numUncertaintyScenarios*; typically, values above 3 for this parameter are computationally infeasible.

The up-to-three uncertainty parameters are specified by LFT parameters *uncertaintyParam1*, *uncertaintyParam2*, and *uncertaintyParam3*. The values of these parameters are just the names of either fish or habitat parameters in *inSTREAM*. The value of *uncertaintyParam3* etc. can also be set to “none” when fewer than three uncertainty parameters are used. The low and high values of these uncertainty parameters are calculated by the LFT by multiplying their standard values by the LFT parameters *uncertaintyParam1LowFactor*, *uncertaintyParam1HighFactor*, etc. (The standard value is the parameter value specified by the user as model input.)

The LFT’s default uncertainty consideration uses three *inSTREAM* parameters that are particularly important and relatively uncertain. First is the fish parameter *fishRespParamA*, which relates respiration (metabolic energy cost) to fish weight. Respiration strongly affects the interactions among fish physiology, behavior, survival, and growth (in many individual-based models, not just *inSTREAM*). Second is a parameter relating maximum sustainable swimming speed to fish length, *fishMaxSwimParamA*. Maximum swimming speed is another key physiological process affecting both survival and growth. The third parameter is *mortFishTerrPredMin*, which controls the risk of predation by terrestrial animals. This form of predation is typically the dominant cause of mortality; risk is assumed to be highest for large fish in shallow and slow water, far from hiding cover. The value of *mortFishTerrPredMin* is the daily survival of a large trout in the riskiest habitat.

The uncertainty in *fishRespParamA* can be assumed  $\pm 20\%$ , so the default values of *uncertaintyParam1LowFactor* and *uncertaintyParam1HighFactor* are 0.8 and 1.2. For *fishMaxSwimParamA*, uncertainty can be assumed  $\pm 20\%$ , so the default values of *uncertaintyParam2LowFactor* and *uncertaintyParam2HighFactor* are 0.8 and 1.2. The uncertainty in *mortFishTerrPredMin* is assumed to be  $\pm \sim 20\%$  in survival for one month (*mortFishTerrPredMin*<sup>30</sup>), which is provided by default values of *uncertaintyParam3LowFactor* and *uncertaintyParam3HighFactor* of 0.993 and 1.007.

**When a survival probability such as *mortFishTerrPredMin* is used as one of the uncertainty variables, users must check carefully to make sure its highest value is never 1.0 or greater. Here, the value of *uncertaintyParam3HighFactor* must not be equal to or greater than  $1/mortFishTerrPredMin$ .**

Computational effort can be minimized by using only two values (the low and high end of the ranges) for these uncertainty parameters. The uncertainty formulation described above means that each experiment will be run for each of eight combinations of parameter values for *fishRespParamA*, *fishMaxSwimParamA*, and *mortFishTerrPredMin*. With nine potential limiting

factors and *numScenarios* set to 5 values, one complete limiting factors analysis therefore requires 360 model runs.

## 4. Results Analysis and Interpretation

### 4.1. Analysis methods and interpretation of results

The LFT includes an analysis spreadsheet that applies a standard linear regression approach to evaluate the relative strength of each factor's effect. The basic approach is to regress the model results—total number and biomass of adult trout—against the value of each factor. The independent (x) variable in the regression is not the actual value of the factor being analyzed (food concentration, base flow, etc.) but instead a scaled representation of the factor. The scaled values range from 0.0 for the lowest value of the factor simulated, to 1.0 for the factor's highest value. Other values are scaled between 0.0 and 1.0; hence, experiments with four scenarios for each factor have scaled factor values of 0.0, 0.33, 0.67, and 1.0. The purpose of this scaling is to make results comparable across all factors.

The slope of the regression indicates the strength of the factor's effect on model results: higher slope magnitudes (the absolute value of regression slope) indicate factors with stronger effects on trout production. The squared correlation coefficient  $R^2$  indicates how consistent or robust the factor's effect is over its range and among the uncertainty parameter values. Results of LFT experiments generally fall within three categories that can be interpreted as follows:

- Relatively high slope magnitude and high  $R^2$ : the factor has a strong and consistent effect on trout production.
- Relatively high slope magnitude but low  $R^2$ : the factor has a strong effect on trout production but the effect is nonlinear or interacts with the parameters varied to represent uncertainty (Section 3).
- Low slope magnitude and low  $R^2$ : the factor appears to have relatively little effect on trout production. However, it is possible instead that the factor has strong but highly nonlinear effects.

### 4.2. Results formats

The LFT's analysis spreadsheet presents results in two formats that are both essential to review. The first format is a table of the statistical results (slope magnitude and  $R^2$ ) for each factor and each of the two kinds of results (abundance and biomass of adult trout). This table is useful for quick and reproducible comparison of the different factors.

The second results format is the actual plots used to produce the regression statistics. Examination of these plots is important for understanding each factor's effect, especially to look for nonlinear responses. The plots are especially important for factors falling in the third category listed in Section 4.1: factors that are indicated by the statistics to have little effect but may in fact have strong but highly nonlinear effects. If, for example, the summer temperature experiment produces a low slope and  $R^2$ , does that mean summer temperature has little effect or that both decreasing and increasing the temperature from the standard input have strong but negative effects? This question is easily answered by looking at the results plot.

## 5. Parameters Controlling the Limiting Factors Tool

The LFT parameters were all described in preceding sections and are listed in Table 1. The table also includes example values. Values should be reconsidered for each study site, considering the following guidance.

The range over which management factors are varied in each experiment includes values both below and above the standard values specified by the input to *inSTREAM* for a study site. These ranges should approximate the range of feasible management actions, and not be broader than realistic. Using overly broad ranges of experimental factors increases the possibility of nonlinear responses that cause misinterpretation of limiting factors results, or results that are unhelpful because they represent conditions that cannot be attained by feasible management actions.

The range of variation need not be symmetric around the standard input. For factors such as spawning gravel availability that are much more likely to be increased instead of decreased by management actions, the range evaluated in the limiting factor experiments should be skewed toward higher values.

Table 1. Parameter definitions and example values for the LFT.

Parameter	Meaning	Example value
<i>numScenarios</i>	The number of scenarios per experiment (applies to all experiments)	5 (see Section 7)
<i>resultsAgeThreshold</i>	The age threshold for including trout in LFT results (years, with age 1 attained on a fish's first January 1)	1
<i>resultsCensusDay</i>	Day on which annual results are censused from the model (mm/dd)	10/01
<i>winterTemperatureStartDay</i>	The first day (MM/DD) of the "winter" period for winter temperature experiments	11/1
<i>winterTemperatureEndDay</i>	The last day (MM/DD) of the "winter" period for winter temperature experiments	3/31
<i>winterTemperatureRangeLow</i>	The constant change in temperature (°C) for the lowest-temperature scenario in the winter temperature experiment.	-4.0
<i>winterTemperatureRangeHigh</i>	The constant change in temperature (°C) for the highest-temperature scenario in the winter temperature experiment.	4.0
<i>summerTemperatureStartDay</i>	The first day (MM/DD) of the "summer" period for summer temperature experiments	4/1

Parameter	Meaning	Example value
<i>summerTemperatureEndDay</i>	The last day (MM/DD) of the “summer” period for summer temperature experiments	10/31
<i>summerTemperatureRangeLow</i>	The constant change in temperature (°C) for the lowest-temperature scenario in the summer temperature experiment.	-4.0
<i>summerTemperatureRangeHigh</i>	The constant change in temperature (°C) for the highest-temperature scenario in the summer temperature experiment.	4.0
<i>baseFlowRangeLow</i>	The constant change in flow (m <sup>3</sup> /s) for the lowest-flow scenario in the base flow experiment.	-2.0
<i>baseFlowRangeHigh</i>	The constant change in flow (m <sup>3</sup> /s) for the highest-flow scenario in the base flow experiment.	4.0
<i>gravelAvailabilityRangeLow</i>	The lowest value of relative gravel availability in the spawning gravel experiment. The value must be between 0.0 and 1.0	0.5
<i>gravelAvailabilityRangeHigh</i>	The highest value of relative gravel availability in the spawning gravel experiment. The value must be between 1.0 and 2.0	1.5
<i>shelterAvailabilityRangeLow</i>	The lowest value of relative velocity shelter availability in the velocity shelter experiment. The value must be between 0.0 and 1.0	0.5
<i>shelterAvailabilityRangeHigh</i>	The highest value of relative velocity shelter availability in the velocity shelter experiment. The value must be between 1.0 and 2.0	1.5
<i>hidingCoverRangeLow</i>	The lowest value of hiding cover availability in the hiding cover experiment. The value must be between 0.0 and 1.0	0.5
<i>hidingCoverRangeHigh</i>	The highest value of hiding cover availability in the hiding cover experiment. The value must be between 1.0 and 2.0	1.5

Parameter	Meaning	Example value
<i>hidingCoverThreshold</i>	In the hiding cover experiment, the cell distance to hiding cover (m) defining cover close enough to provide substantial benefit	2.0
<i>piscivoryRiskRatioLow</i>	The change in piscivory risk (fraction of the standard value) for the lowest-risk scenario in the piscivory risk experiment.	0.9
<i>piscivoryRiskRatioHigh</i>	The change in piscivory risk (fraction of the standard value) for the lowest-risk scenario in the piscivory risk experiment.	1.1
<i>foodAvailabilityRatioLow</i>	The fraction of standard food availability for the lowest-food scenario in the food availability experiment.	0.5
<i>foodAvailabilityRatioHigh</i>	The fraction of standard food availability for the highest-food scenario in the food availability experiment.	2.0
<i>reddScourRatioLow</i>	The lowest factor by which redd depth is multiplied in the redd scour experiment	0.5
<i>reddScourRatioHigh</i>	The highest factor by which redd depth is multiplied in the redd scour experiment	1.5
<i>numUncertaintyParameters</i>	The number of parameters varied to represent parameter uncertainty (Section 3).	3
<i>numUncertaintyScenarios</i>	The number of values used per uncertainty parameter.	2
<i>uncertaintyParam1</i>	The name of the first uncertainty parameter (a text string with the name of a fish or habitat parameter; set to "none" if no uncertainty parameters are used).	<i>fishRespParamA</i>
<i>uncertaintyParam1LowFactor</i>	The ratio multiplied by the standard value of uncertainty parameter 1 to set the low end of its range.	0.8
<i>uncertaintyParam1HighFactor</i>	The ratio multiplied by the standard value of uncertainty parameter 1 to set the high end of its range.	1.2

Parameter	Meaning	Example value
<i>uncertaintyParam2</i>	The name of the second uncertainty parameter (“none” if fewer than two uncertainty parameters are used).	<i>fishMaxSwimParamA</i>
<i>uncertaintyParam2LowFactor</i>	The ratio multiplied by the standard value of uncertainty parameter 2 to set the low end of its range.	0.8
<i>uncertaintyParam2HighFactor</i>	The ratio multiplied by the standard value of uncertainty parameter 2 to set the high end of its range.	1.2
<i>uncertaintyParam3</i>	The name of the third uncertainty parameter.	<i>mortFishTerrPredMin</i>
<i>uncertaintyParam3LowFactor</i>	The ratio multiplied by the standard value of uncertainty parameter 3 to set the low end of its range.	0.993
<i>uncertaintyParam3HighFactor</i>	The ratio multiplied by the standard value of uncertainty parameter 3 to set the high end of its range.	1.007



## 6. Using the Tool

The LFT exists as an option inside *inSTREAM*'s graphical user interface (GUI), which is documented separately.

### 6.1. Running an LFT analysis

When a model project is ready for use (all input files have been provided; parameter values are at least tentatively established; test runs have been completed), the user can hit the “Limiting Factors” button on the GUI interface. (It is generally smart to first save the project as a new project just for the LFT analysis.)

The “Limiting Factors” button takes the user to the GUI’s LFT view, which has two tabs labeled “Limiting Factors Tool Parameters” and “Execution Control/Output”. Normal use of the LFT follows these steps:

1. On the “Limiting Factors Tool Parameters” tab, edit the LFT parameter values to customize the analysis (Figure 1). Use the information in sections 2, 3, and 5 to select parameter values.
2. Click the “Start Limiting Factors Experiments” button. This creates separate subdirectories (under the project directory) for each experiment and builds the input and setup files (especially *Experiment.Setup*) for each. It also starts execution of *inSTREAM* for each experiment (explained in Section 7).
3. Switch to the “Execution Control/Output” tab (Figure 2) to observe execution of the experiments. Individual experiments can be terminated via a button on their tab. All runs can be terminated at once via a button on the “Limiting Factors Tool Parameters” tab.
4. When all the experiments have stopped (because they finished, you terminated them, or they stopped due to an error), a dialog pops up asking you whether you want to process and view results now or later. If you prefer, you can later re-open the project and LFT, and click on the “Process/View Results” button of the “Limiting Factors Tool Parameters” tab to open the results analysis spreadsheet.
5. (If not all experiments completely finished, you can still use the following steps to observe the partial results.)
6. When you do process/view the results, the LFT opens an Excel spreadsheet that contains code to import LFT results and report the statistical and graphical results. (These results are imported from a special output file named *LFT\_Output.rpt*.) **This Excel import process will very likely halt, with Excel displaying a “Security warning” saying that macros have been disabled. You must choose the option to allow macro content for the analysis to proceed.** (See Figure 3.)
7. After the spreadsheet has updated itself, you will be asked to save it under a new file name (Figure 4).
8. Examine the results displayed in the spreadsheet. It includes a table of statistical results (the “ResultsSummary” sheet) and a separate sheet displaying graphical results for each experiment. The table of statistical results includes a column labeled “n” that reports the number of model runs completed for each limiting factor experiment.

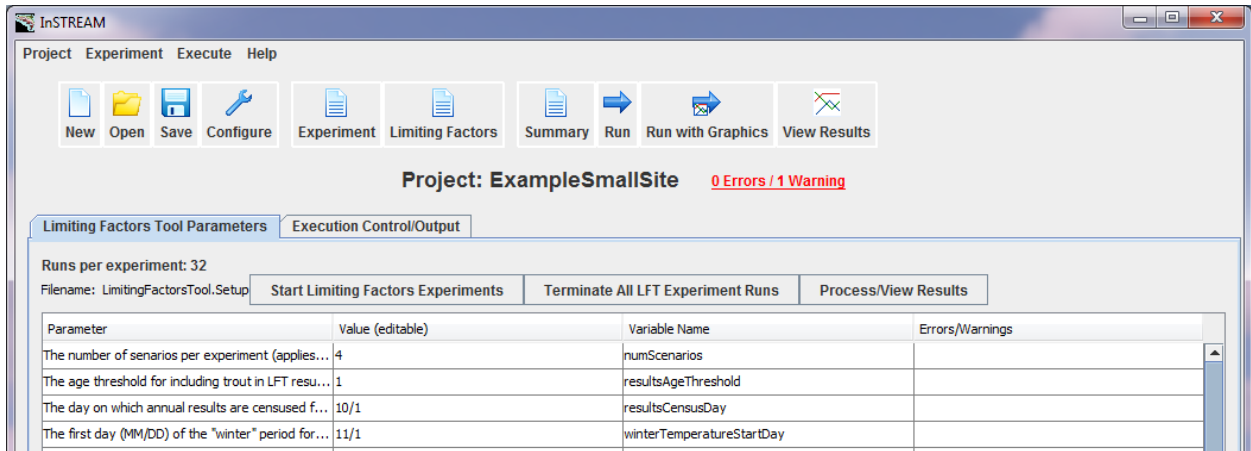


Figure 1. The *inSTREAM* GUI's LFT view and its "Limiting Factors Tool Parameters" tab.

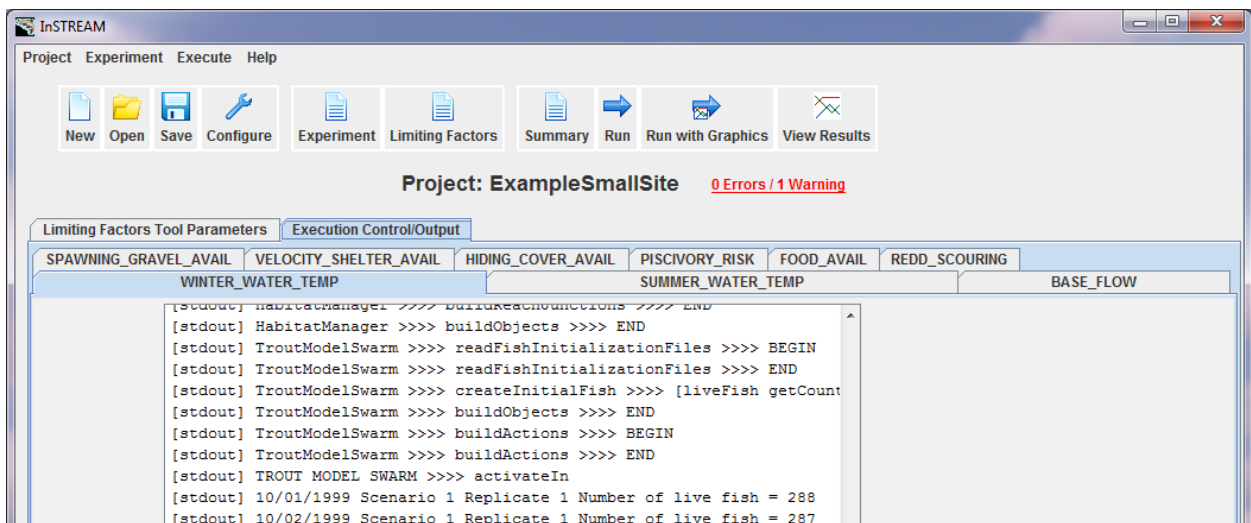


Figure 2. The LFT view and its "Execution Control/Output" tab. A separate tab is created for each LFT experiment, with an output window where *inSTREAM* output is displayed as the model runs.

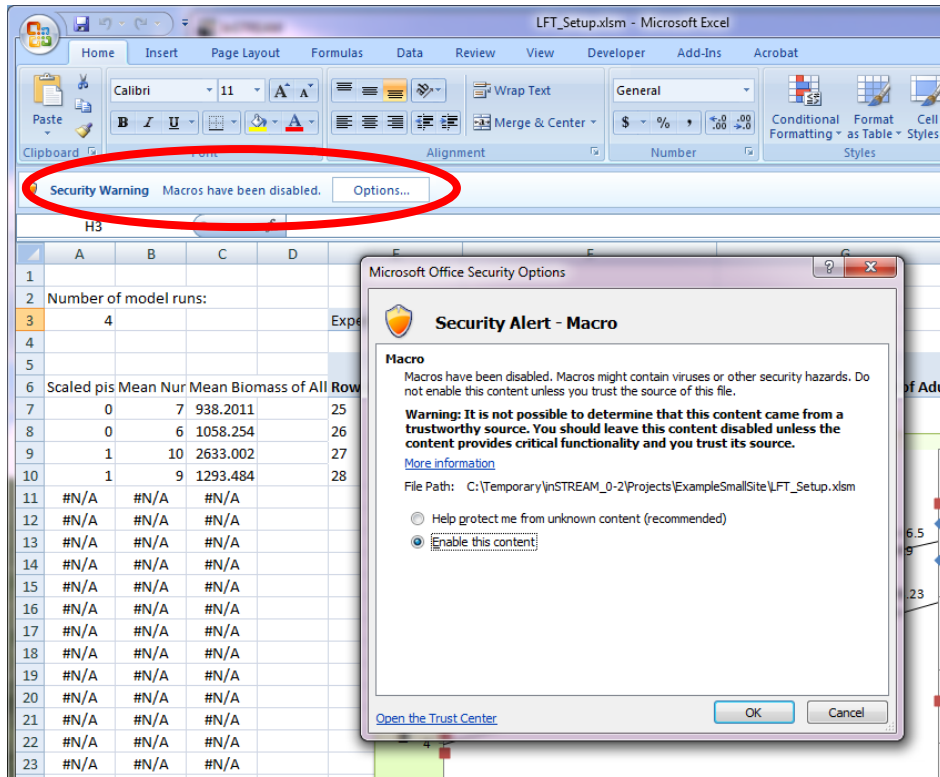


Figure 3. The "Security Warning" that can open when Excel tries to import LFT results. Click on "Options..." and then "Enable this content".

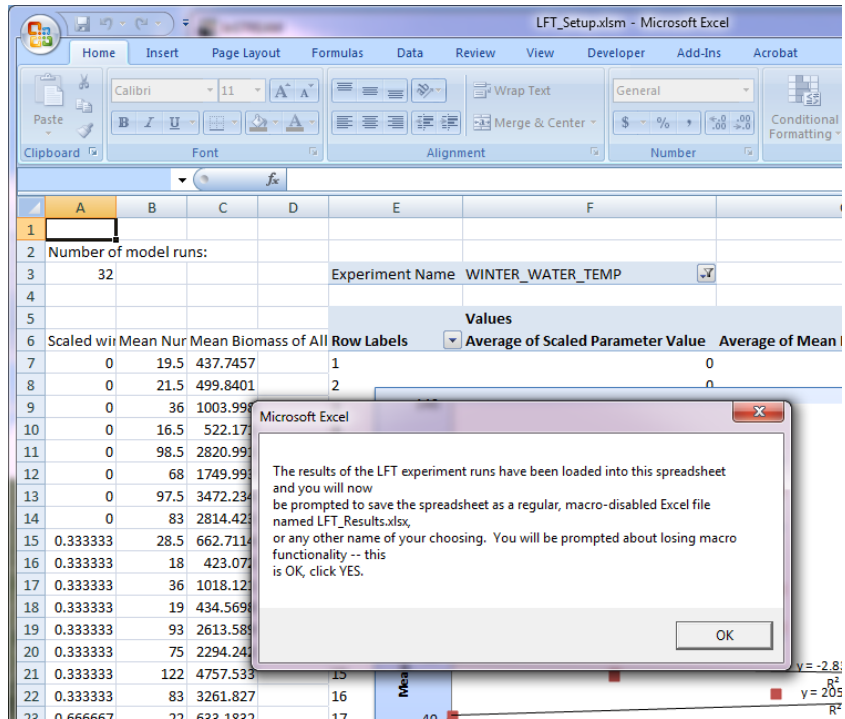


Figure 4. The LFT's analysis spreadsheet automatically imports results and asks you to save the file as a regular Excel file.

### 6.1. Error checking

The LFT does some error checking before executing. It will not run if there are any errors in the inSTREAM project (e.g., missing or incorrectly formatted parameter values), or if the LFT parameter *numUncertaintyParameters* is not within the range of 0-3. A warning is raised if the total number of model runs to be executed is greater than 500, but this warning lets the user continue if they want.

## 7. Computational Considerations

The limiting factors analyses are computationally demanding and the extent of the analysis is typically limited by computer considerations. The computational limitations are not just due to execution speed (which can be overcome with patience or a faster computer) but can also be due to memory availability.

When the user starts the limiting factor experiments, the LFT starts each experiment (as defined in Section 2) as a separate *inSTREAM* job: a set of model runs that are executed sequentially. All the experiment jobs are started at once and the operating system automatically assigns them to the available processors. The experiments therefore are executed more or less in parallel, with execution time depending on how many processors the computer has and how fast they are.

On a modern computer (e.g., Intel Core chip sets), the number of experiment jobs seems to affect only execution time: the jobs simply rotate to the available processor cores. (The LFT should not be used on old computers with Pentium-type hardware.) The *inSTREAM* code is memory-intensive, but execution does not appear to slow substantially when the available RAM is filled (the processor moves inactive memory use to the hard disk's swap file—page file, in Windows; however, a job may be killed by the operating system when the swap file is full).

Experience to date indicates that the most likely computation problem is that extremely large experiment jobs stop before completing because the model uses up all the available memory. The number of runs per LFT analysis is equal to  $9 \times \text{numScenarios} \times U^n$  where  $U$  is the number of values per uncertainty parameter (LFT parameter *numUncertaintyScenarios*) and  $n$  is the number (between 0 and 3) of uncertainty parameters used (LFT parameter *numUncertaintyParams*; Section 3). Memory use also depends very strongly on how big the *inSTREAM* simulation is: how many habitat cells, how many fish, and how many years are simulated.

(Large analyses can also fail because of a Windows limitation on how much memory can be allocated to each experiment and fragmentation of that memory. The symptom of this problem is that the model stops with an error in the function "malloc". If very large experiments stop due to inadequate RAM or other memory problems, the only solution for Windows users is to set up smaller experiments by using fewer scenarios, uncertainty parameters, or values per uncertainty parameter. Another solution to this memory allocation limitation is to use the LFT in Windows to set up the limiting factor experiments and then execute them in a 64-bit Linux installation of *inSTREAM*, which is much less subject to memory limitations. The LFT can be started in Windows and used to set all the LFT parameter values. The "Start Limiting Factor Experiments" button is then clicked, after which the "Terminate All LFT Experiment Runs" button is clicked to stop the runs. This will create all the separate directories and input files for each experiment; these can be copied to Linux for execution there. An especially large swap file in the Linux installation can offset the graduate consumption of memory by *inSTREAM*.)

## LFT Parameter Index

### **B**

*baseFlowRangeHigh* .....3, 11  
*baseFlowRangeLow* .....3, 11

### **F**

*foodAvailabilityRatioHigh* .....4, 12  
*foodAvailabilityRatioLow* .....4, 12

### **G**

*gravelAvailabilityRangeHigh* .....4, 5, 12  
*gravelAvailabilityRangeLow* .....4, 5, 12

### **H**

*hidingCoverRangeHigh* .....5, 12  
*hidingCoverRangeLow* .....5, 12  
*hidingCoverThreshold* .....5, 12

### **L**

LFT\_Output.rpt .....15

### **N**

*numScenarios* .....3, 11, 18  
*numUncertaintyParameters* .....13  
*numUncertaintyParams* .....8, 18  
*numUncertaintyScenarios* .....8, 13, 18

### **P**

*piscivoryRiskRatioHigh* .....6, 12  
*piscivoryRiskRatioLow* .....6, 12

### **R**

*reddScourRatioHigh* .....8, 13  
*reddScourRatioLow* .....8, 13  
*resultsAgeThreshold* .....3, 11  
*resultsCensusDay* .....3, 11

### **S**

*shelterAvailabilityRangeHigh* .....5, 12  
*shelterAvailabilityRangeLow* .....5, 12  
*summerTemperatureEndDay* .....4, 11  
*summerTemperatureRangeHigh* .....4, 11  
*summerTemperatureRangeLow* .....4, 11  
*summerTemperatureStartDay* .....4, 11

### **U**

*uncertaintyParam1* .....8, 13  
*uncertaintyParam1HighFactor* .....8, 9, 13  
*uncertaintyParam1LowFactor* .....8, 9, 13  
*uncertaintyParam2* .....8, 13  
*uncertaintyParam2HighFactor* .....9, 13  
*uncertaintyParam2LowFactor* .....9, 13  
*uncertaintyParam3* .....8, 13  
*uncertaintyParam3HighFactor* .....9, 13  
*uncertaintyParam3LowFactor* .....9, 13

### **W**

*winterTemperatureEndDay* .....4, 11  
*winterTemperatureRangeHigh* .....4, 11  
*winterTemperatureRangeLow* .....4, 11  
*winterTemperatureStartDay* .....4, 11