SSE user guide

PsN 4.6.0

Revised 2016-05-10

1 Introduction

SSE – Stochastic Simulation and Estimation – is a tool for model comparison and hypothesis testing. First, using the input model, a number of simulated datasets are generated. Then the input model and a set of alternative models are fitted to the simulated data. Finally, a set of statistical measures are computed for the parameter estimates and objective function values of the various models.

Example calls:

```
sse run1.mod -samples=1000 -alternative_models=alt1.mod,alt2.mod
```

2 Input and options

A model file, the simulation model, is required on the command-line.

-samples = *N*
: The number of simulated datasets to generate is a required option. N must be greater than or equal to 1.

-alternative_models = *alt1,alt2,alt3*
: A comma-separated list of model files with alternative models is optional. If not given, only the input model will be fitted to the simulated data.

-estimate_simulation

This option is set by default, and can be disabled with -no-estimate_simulation. By default, the simulated datasets are estimated with the input model. The OFV-values from these estimations are used as reference when computing statistics on results from estimations with the alternative models. If -no-estimate_simulation (no option value) is set, the input model will not be used to estimate (only to simulate). Instead, either a given reference OFV-value will be used (see option -ref_ofv) or the OFV-values from the first alternative model will be treated as the reference ones. Regardless of how option -no-estimate_simulation is set, the initial estimates in the simulation model will be used as reference when computing bias, rmse and rsebias for theta, omega and sigma. If simulating with uncertainty then the reference values will be different for each simulated dataset.

-ref_ofv = *X*

A reference OFV-value to use when computing statistics for the results from the alternative models. This option is only allowed together with -no-estimate_simulation.

-parallel_simulation = *X*

The number of parallel processes to start for the simulation step (not the estimation step) on a parallel computer. Default is the same setting as option -threads.

-keep_tables

By default, all pre-existing $TABLE will be deleted from the simulation and alternative models, to save disk space. If option -keep_tables is set, PsN will instead keep $TABLE and number the file names according to the sample number.

-append_columns = *col1,col2,...*

Optional, only allowed when -add_models is not set. A comma-separated list of columns to append to the simulated data sets. These variables must be defined in the simulation model, otherwise all the simulation models will crash with an NMtran-error. PsN will not check that the variables are defined, but

simply append them to the $TABLE generating the simulated datasets. This option cannot be used in combination with -add_models, since in that case no simulations will be run. The user must ensure that the estimation models have these columns properly included in $INPUT, PsN will not add them to $INPUT.

-add_models

This option tells PsN to add the alternative models listed with option -alternative_models to an old sse run. All models given via option -alternative_models will be estimated from scratch, so the alternatives from the old run should not be listed in the input again. The -add_models option requires that the old sse run has been completed without errors. It is necessary to also set the general PsN option -directory to the sse directory of the old run, e.g. -directory=sse_dir50. The user must ensure that the -samples option in the new sse run is equal to or smaller than in the old run. If some simulated datasets are missing it will result in an error. The simulation model must still be given as input on the command-line. If the option -estimate_simulation is set (it is set by default) old estimation results will be reloaded if they exist, otherwise the simulation model will be estimated. The numbering of the extra models will start at the number of the last old alternative plus 1. Results for the added models are in file sse_results_add1.csv

-rawres_input = *filename*

A simple way to simulate with uncertainty. Instead of using identical parameter estimates for simulation of each new dataset, take parameter estimates from a raw_results.csv file, e.g. from a bootstrap run or the intial_estimates.csv file from a previous sse run with $PRIOR in the simulation model. The raw results file must be comma-separated and contain at least as many samples as the input -samples to sse, the labels for THETA/OMEGA/SIGMA in the file must match the labels in the simulation model given as input to sse, the theta columns must be directly followed by the omega columns which must be directly followed by the sigma columns, and the column

header must be model either in the first column just as in a bootstrap raw_results file or in the second or third column as in a sse raw_results file. If a column header contains a comma, e.g. OMEGA(2,2), then that header must be enclosed in double quotes. This is done automatically in PsN raw results files. Note that is is possible to generate a file with initial parameter estimates outside of PsN, as long as the file follows the format rules.

-offset_rawres = *N*

Only relevant in combination with -rawres_input. Default 1. The number of result lines to skip in the input raw results file before starting to read final parameter estimates. In a regular bootstrap raw_results file, and also in an initial_estimates.csv file from an sse run, the first line of estimates refers to the input model with the full dataset, so therefore the default offset is 1.

-in_filter = *comma-separated list of conditions*

Only relevant in combination with -rawres_input. Default not used. The parameter estimates lines in the file can be filtered on values in the different columns. When specifying which column(s) the filtering should be based on, the exact column name must be used, e.g. minimization_successful. Filtering can only be based on columns with numeric values. The allowed relations are .gt. (greater than), .lt. (less than) and .eq. (equal to). If the value in the filter column is 'NA' then that parameter set will be skipped, regardless of the defined filter relation. Conditions are separated with commas. If the remaining number of lines after filtering is smaller than -samples, sse will stop with an error message. Then the user must either change the filtering rules or change -samples. If the user has created a file with parameter estimates outside of PsN, filtering can be done on any numeric column in that file. Do not set column headers containing .eq. or .lt. or .gt.in the user-generated file as this would interfere with the in_filter option syntax.

Example (there must be no linebreaks in the actual command):

```
-in_filter=minimization_successful.eq.1,
```

```
significant_digits.gt.3.5
```

-random_estimation_inits

This option can only be used in combination with -rawres_input. It turns off simulation with uncertainty and instead uses the parameter values from the rawres_input file as initial values for the estimation step. When this option is not used, the estimation initial values will be the same as the values used for simulating data.

-out_filter = *comma-separated list of conditions*

Default not used. The user may choose to only compute results based on estimations which fulfill certain conditions. The default is to only skip runs where the ofv cannot be read from the lst-file or is equal to 0. Filtering of output can be done on any numeric column in a standard sse raw_results file, for example minimization_successful, significant_digits and covariance_step_successful. The allowed relations are .gt. (greater than), .lt. (less than) and .eq. (equal to). If the value in the filter column is 'NA' then that parameter set will be skipped, regardless of the defined filter relation. Conditions are separated with commas. If the remaining number of estimation results after filtering is less than 2, sse will stop with an error message.

Example:

```
-out_filter=minimization_successful.eq.1,
significant_digits.gt.3.5
```

-recompute = *raw results filename, including directory name*

Default not set. Setting this option makes PsN recompute output statistics based on the specified raw_results file. Note that the filename must be given including the directory name. The user may change the -out_filter settings for the recomputation. Apart from -out_filter, the input model must be set, and -samples. Alternative models are not needed, information about them will be read from the raw results file. Option

-directory will be ignored, instead the directory specified as part of the file path will be used.

Example:

```
-recompute=sse_dir12/raw_results_run1.csv
```

2.1 Some important common PsN options

For a complete list see common_options.pdf, or psn_options -h on the commandline.

-h or -?

Print the list of available options and exit.

-help

With -help all programs will print a longer help message. If an option name is given as argument, help will be printed for this option. If no option is specified, help text for all options will be printed.

-directory = *'string'*

Default sse_dirN, where N will start at 1 and be increased by one each time you run the script. The directory option sets the directory in which PsN will run NONMEM and where PsN-generated output files will be stored. You do not have to create the directory, it will be done for you. If you set -directory to a the name of a directory that already exists, PsN will run in the existing directory.

-seed = *'string'*

You can set your own random seed to make PsN runs reproducible. The random seed is a string, so both -seed=12345 and -seed=JustinBieber are valid. It is important to know that because of the way the Perl pseudo-random number generator works, for two similar string seeds the random sequences may be identical. This is the case e.g. with the two different seeds 123 and 122. Setting the same seed guarantees the same sequence, but setting two slightly different seeds does not guarantee two different random sequences, that must be verified.

-clean = *'integer'*

Default 1. The clean option can take four different values:

0 Nothing is removed

1 NONMEM binary and intermediate files except INTER are removed, and files specified with option -extra_files.

2 model and output files generated by PsN restarts are removed, and data files in the NM_run directory, and (if option -nmqual is used) the xml-formatted NONMEM output.

3 All NM_run directories are completely removed. If the PsN tool has created modelfit_dir:s inside the main run directory, these will also be removed.

-nm_version = *'string'*

Default is 'default'. If you have more than one NONMEM version installed you can use option -nm_version to choose which one to use, as long as it is defined in the [nm_versions] section in psn.conf, see psn_configuration.pdf for details. You can check which versions are defined, without opening psn.conf, using the command

```
psn -nm_versions
```

-threads = *'integer'*

Default 5 (if default PsN config file is used). Use the threads option to enable parallel execution of multiple models. This option decides how many models PsN will run at the same time, and it is completely independent of whether the individual models are run with serial NONMEM or parallel NONMEM. If you want to run a single model in parallel you must use options -parafile and -nodes. On a desktop computer it is recommended to not set -threads higher the number of CPUs in the system plus one. You can specify more threads, but it will probably not increase the performance. If you are running on a computer cluster, you should consult your system administrator to find out how many threads you can specify.

-version

Prints the PsN version number of the tool, and then exit.

2.2 Auto-generated R-plots from PsN

PsN can automatically generate R plots to visualize results for sse, using a default template found in the R-scripts subdirectory of the installation directory. The user can also create a custom template, see more details in section "Auto-generated R-plots from PsN" in common_options.pdf.

Parametric power estimation (PPE) allows to generate data for full power versus sample size curve(s) based on a SSE with full and reduced models performed with a single sample size. The procedure is described in [1].

The default sse template will perform all required PPE calculations in R and output the resulting power curve(s). For the template to work, the R libraries ggplot2 and plyr need to be installed. Furthermore, it is recommended to indicate the reduced model associated with a specific full model by adding the suffix 'red' (i.e. 'run1red.mod' is the reduced model associated with the full model 'run1.mod'). If no full and reduced model pair with the expected naming scheme is found, matching will occur by position (i.e., estimation model 1 will be taken to be the full model and estimation model 2 the associated reduced model, same for 3 and 4, etc.).

-rplots = *level*

-rplots<0 means R script is not generated
-rplots=0 (default) means R script is generated but not run
-rplots=1 means basic plots are generated
-rplots=2 means basic and extended plots are generated

Troubleshooting

If no pdf was generated even if a template file is available and the appropriate options were set, check the .Rout-file in the main run directory for error messages. If no .Rout-file exists then check that R is properly installed, and that either command 'R' is available or that R is configured in psn.conf.

Basic plots

The basic PPE plot will be generated if option -rplots is set >0, and the general rplots conditions fulfilled, see above. The plot shows, for each of the provided full/reduced model pairs, a color-coded power versus sample size curve together with a shaded region indicating the uncertainty in the power estimate due to Monte-Carlo noise (-samples option in PsN). By default power curves are drawn from a sample size of 1 to the samples size of 99% power.

Extended plots

Extended PPE plots, will be generated if -rplots>1, and are on the second page of the pdf-output. The extended plots are intended as a diagnostic to evaluate the assumptions of the PPE algorithm. They show separately for each full/reduced model pair the empirical cumulative distribution function (ECDF) of the Δ-OFVs together with a shaded area indicating the 95% confidence region for the ECDF provided the underlying distributional assumptions of the PPE algorithm are true. The results of the algorithm should be re-evaluated through an alternative method for power estimation when the ECDF curve lies outside the confidence interval.

2.3 Special cases

To simulate odd-type data, relevant code (with a simulation block) is needed in the input model file, as well as a $SIMULATION row. It should contain one ordinary (from a normal distribution) seed number and one extra seed number (from a uniform distribution) "$SIMULATION (11111) (11111 UNIFORM)". The random number generator handled by PsN will replace the seed numbers in a controlled manner (derived from "–seed=" if provided).

Please note that NONPARAMETRIC is not yet completely supported, since it requires additional results handling features.

2.4 Simulation with uncertainty

Simulation with uncertainty is supported using three methods. One method is to use options -rawres_input and -offset_rawres, see above. The second method is to define $PRIOR NWPRI in the input model, including the PLEV option required when simulating from a prior. Please note that the

update_inits script has some functionality for automatically adding $PRIOR NWPRI to a model, based on output from an estimation. However the update_inits feature is still experimental and the generated model needs to be manually checked. The third method is to define $PRIOR TNPRI in the input model, including the PLEV option required when simulating from a prior. If no $SIM record is present in the input model containing $PRIOR, PsN will create one and add option TRUE=PRIOR. If PLEV option is not set in $PRIOR, NONMEM will halt saying that the value of PLEV is inappropriate. When $PRIOR is present in the input model, PsN will store the initial estimates sampled from the prior in the file initial_estimates.csv. This file is formatted so that it can be used with option -rawres_input in a later sse run. The first row, model 0, in initial_estimates.csv is the initial estimates of the input model, and these will be skipped if -offset_rawres is equal to the default value.

If $PRIOR TNPRI is set, then NONMEM requires $MSFI in the first $PROB. PsN will automatically copy the msf-file to the run directory, meaning that it is not necessary to set PsN option -extra_files for the msf-file.

3 Output

The output from each SSE run is collected in a new directory with a name of the form sse_dirX, where X is an integer. The most recent SSE run is found in the directory with the highest number. The only exception is when the user restarts a run from a specified existing directory using the -directory option, then the results will be saved in that directory.

When $PRIOR is present in the input model, PsN will store the initial estimates sampled from the prior in the file initial_estimates.csv, see Simulation with uncertainty above.

SSE creates two files summarizing the results that can be opened in e.g. Excel. The raw_results.csv file is a standard PsN file containing raw result data for each estimation run (termination status, parameter estimates, uncertainty estimates, etc. . .). The sse_results.csv file is a file specific to the SSE routine containing summary statistics and comparisons common to all estimation runs.

The header line for the reference/true values uses complete indexes for OMEGA and SIGMA, e.g OMEGA(2,1), SIGMA(3,3). The header lines for the input model and the alternative models use naïve numbering, for example

if OMEGA(1,1), OMEGA(2,2), OMEGA(3,2) and OMEGA(3,3) are defined in the modelfile, they will in the sse output be numbered OM_1, OM2, OM_3 and OM_4. See also section Principles of statistical analyses.

3.1 Explanation for data items in sse_results.csv:

The 'true' values in the formulas below are always the initial parameter estimates of the simulation model. When simulating without uncertainty the true value is constant over i, but when simulating with uncertainty the true value varies with index i. If for any i the true value is equal to 0 in a computation of a relative value, then that sample is skipped and N is adjusted. *sd = standard deviation of estimated parameter* x_i.

$$Skewness = \frac{N}{(N-1)\cdot(N-2)} sd^{-3} \sum_{i=1}^{N} (x_i - \bar{x})^3$$

$$Kurtosis = \frac{N\cdot(N+1)}{(N-1)\cdot(N-2)\cdot(N-3)} sd^{-4} \sum_{i=1}^{N} (x_i - \bar{x})^4 - \frac{3\cdot(N-1)^2}{(N-2)\cdot(N-3)}$$

$$rmse = \sqrt{\frac{1}{N}\sum_i (est_i - true_i)^2}$$

$$relative_rmse = 100\% \sqrt{\frac{1}{N}\sum_i \frac{(est_i - true_i)^2}{true_i^2}}$$

$$bias = \frac{1}{N}\sum_i (est_i - true_i)$$

$$relative\_bias = 100\% \frac{1}{N}\sum_i \frac{est_i - true_i}{true_i}$$

$$relative\_absolute\_bias = 100\% \frac{1}{N}\sum_i \frac{est_i - true_i}{|true_i|}$$

$$Rse_{bias} = relativestandarderror = 100\% \frac{sd}{True\sqrt{N}}$$

Standard error CI_{bias} = parametric confidence intervals = mean bias +/- Z*rse where Z = [2.58, 1.96, 1.645] for CI = [99%, 95%, 90%] Items under the OFV statistics heading by default refer to likelihood ratio test comparisons between the input model and the alternative models. The exceptions are described under section Principles of statistical analyses below.

4 Known problems

The program will use the wrong simulated data for reestimation if the simulation model file has more than one $INPUT statement, the table file with simulated data used for reestimation will only include the items from the last $INPUT. Note: It works fine to split $INPUT over several lines, as long as the text $INPUT only appears once.

The simulated datasets are $TABLE output from NONMEM, and NONMEM formats and rounds off values when printing tables. This leads to two known problems.

1. In NONMEM6 1013201 is rounded to 1013200 (five significant digits), and if this makes a significant change to the model estimation, for example if the value is a covariate, then the sse results will be wrong. In NONMEM7 it is possible to set the FORMAT or RFORMAT option in $TABLE to make sure no important information is lost. With NONMEM6 the user must make sure the rounding to five significant digits does not harm the results. PsN cannot detect this problem.

2. A column of integers in the input dataset, e.g. OCC for occasion, will be formatted as a floating point number in table output. When the table output is used as input in the estimation step of sse, statements like IF (OCC .EQ. 1) will not work, because to NONMEM 1 is not equal to 1.000000e00. As a workaround the user must write code that will work both for integers and floating point numbers, for example (IF OCC .GT. 0 AND OCC .LT. 2). Later versions of NONMEM support ID.EQN.1, which will be true both for 1 and 1.000e00.

5 Recovering a crashed/stopped sse

If the sse run is halted before the simulations are finished and option -clean is <=2, it is possible to reuse the already finished runs. Run the same command as before, set option -directory if it was not set in the original call. Please note that PsN will print a message "Starting NNN NONMEM executions" that indicates that all simulations are rerun, but the check for reusable results is done after this message is printed. Some copying of results is done, and some existing result files will be overwritten with identical copies.

If the sse run is halted after the simulations are finished and option -clean is <=2, PsN can reuse the results from both simulations and finished estimations. Run the same command again with -directory set. Do not set the rerun option. PsN will print that all NONMEM executions are started even if results are reused, see above.

If the sse run is stopped during the estimation phase when -clean=3, partial results can still be recoved but it requires a number of steps.

1. Change directory to the m1 subdirectory of the sse run directory
2. Run execute on all .mod-files in the m1-subdirectory that do not already have a .lst-file. Note that execute accepts multiple .mod-files in a single run.
3. Run the rawresults script (no options) in the m1 directory after all execute runs have finished. This creates a raw_results file based on the output in the m1 subdirectory.
4. Change directory to outside the main sse directory and run sse with option -recompute (see help text for this option) to compute the final results.

6 Technical overview of algorithm

The program creates N (N= number of samples) simulation model files that are modified copies of the input model file given on the command line. Then the program creates N*M (M=number of alternative model files + 1 (for the input model file)) estimation model files. Each model file has one estimation copy per simulated dataset. N simulated datasets are generated by running NONMEM once for each simulation model file, and then each simulated dataset is used for parameter estimation M times, once with each model (simulation + alternatives). The estimated parameters and OFV values from the M*N runs are collected in raw_results.csv together with other NONMEM outputs, and statistical measures are written to sse_results.csv.

6.1 Important features of the N simulation model files generated from the input model

- In $DATA record IGNORE and ACCEPT statements are kept intact.

- If the input model contains a $SIMULATION record, the record is kept intact, except that the random seeds are set to different values in each copy, option NSUBS/SUBPROBLEMS is removed, and it is made sure that TRUE=PRIOR is set if the input model has $PRIOR. If there is no existing $SIMULATION record, the program creates a new one according to $SIMULATION (<random seed>) ONLYSIMULATION and adds TRUE=PRIOR if the input model has $PRIOR.
- If any $ESTIMATION record contains LIKELIHOOD, -2LOGLIKELIHOOD, -2LLIKELIHOOD or LAPLACIAN, then ONLYSIMULATION NOPREDICTION is set in $SIMULATION.
- All $ESTIMATION records are removed.
- Any $COVARIANCE record is removed.
- If option -keep_tables is not set, old $TABLE records are removed. Otherwise old $TABLE records are kept, and the filename in option FILE=<filename> will have trailing numberes removed and then be numbered -sim-i, where i is the order number of the simulation model file.
- In the newly generated $TABLE record for the simulated dataset the option FILE=<simdata> is set, where <simdata> is a unique name for each simulation model file.
- If option -rawres_input is set, then replace the THETA/OMEGA/SIGMA parameter values in the ith simulation model with the final estimates from the ith results line in the input raw_results file.

6.2 Important features of the N estimation model files generated from the input model and of the estimation model files generated from alternative model files (N files per alternative model)

- The N model files generated from the input model are numbered with the number of the simulated dataset used as input in the modelfile. The files generated from the alternative models are numbered with X-Y, where X is the order number of the alternative model and Y is the order number of the simulated dataset.

- In $DATA record, IGNORE is set to @ (i.e. ignore text lines), replacing any old IGNORE=<single character>. All old IGNORE=(list) and ACCEPT statements are kept. The user must be cautious regarding statistics if IGNORE or ACCEPT statements differ between models.
- Any $SIMULATION record is removed from the first $PROBLEM.
- All $ESTIMATION records are kept. If MSFO=<filename> is specified, <filename> will first have any trailing numbers removed, and then <filename> will have the same numbering appended as the name of the modelfile itself (see above).
- If option -keep_tables is not set, old $TABLE records are removed. Otherwise, in $TABLE records of the first $PROBLEM, the filename in FILE=<filename> will have trailing numbers removed and then be numbered as the modelfile, see above.
- Input is set to <simdata> from one of the simulation model files.
- If an alternative model has a second $PROBLEM (this is not allowed for the input model unless the first $PROB has $PRIOR TNPRI), then it is assumed that the first $PROB has MSFO=<filename> in $ESTIMATION and that the second $PROBLEM has an $MSFI record. PsN will set the same numbered filename in MSFO and MSFI. It is also assumed that the second $PROBLEM has a $TABLE record (it is not checked that -keep_tables is set, the user must do that). In the first $TABLE of the second $PROBLEM PsN will set FILE=simtabX-Y.dat, where X-Y is numbers as above. The filename in $DATA of the second $PROB will be set to the same <simdata> as in the first $PROBLEM, but IGNORE statements will not be changed. This means that the user must ensure that there is an IGNORE=@ in $DATA of the second $PROBLEM. PsN will not change the random seeds in $SIMULATION, but since the simulated datasets used as input are different, the table output will also be different.
- If the input model is to be estimated and it has $PRIOR set, then option PLEV is removed from $PRIOR.

6.3 Principles of statistical analyses

All estimated parameters values; thetas, omegas and sigmas, are compared with the values used for generating the simulated datasets. The matching is based on numbering, so the value of theta 3 used for simulation will be compared with all estimated theta 3. It is up to the user to ensure that the matching is correct, i.e. that the simulation value of CL is always compared with the estimated values of CL. If for example OMEGA(1,1), OMEGA(2,2), OMEGA(3,2) and OMEGA(3,3) are defined in the modelfile, they will in the sse output be numbered OM_1, OM2, OM_3 and OM_4, i.e. values are numbered naively without considering if they are on the diagonal or not. Same for SIGMA. If OMEGA/SIGMA is diagonal in the original model but of block form in an alternative model, the matching will be incorrect. To avoid that error, use block form also in the original model, setting off-diagonal elements to 0 (NONMEM will keep these fixed). If there are more parameters in an alternative model than in the simulation or vice versa, no matching can be done and the spaces for the comparison statistics will be empty. If parameters are different from a model to another one, numbering should be done considering that, or if not comparisons should not be interpreted.

SSE does not check whether minimization was successful or not. Statistical computations include also parameter estimates from NONMEM runs terminated with e.g. rounding errors, unless the option -out_filter is used, see above. SSE will skip all runs where the ofv cannot be read from the lst-file or is equal to 0. This means that e.g. the mean may be based on different number of runs (samples) for different alternative models. If the ofv is available but a parameter value is missing, then SSE will print a warning and compute the statistics without that value.

If a run neither minimizes nor terminates, but gets into an infinite loop, it won't be killed by PsN; so, in order for the sse results to be computed, a manual kill can be contemplated.

In sse_results.csv no individual parameter estimates are reported. Mean, median, sd and other measures are the means etc. over the simulated datasets.

The OFV values are treated differently. By default, the OFV value obtained when estimating parameters using the input model is used as the reference for each simulated dataset, and the OFV values obtained when estimating the alternative models are compared with this reference. There are two alternative methods. The first is to set the option -no-estimate_simulation

without setting -ref_ofv. In this case the OFV-values from the first alternative model are used as references. The second alternative is to set -no-estimate_simulation together with -ref_ofv=X, and then X is used as the reference OFV value for all alternative models.

References

[1] S. Ueckert, M. O. Karlsson, and A. C. Hooker. "Accelerating Monte Carlo power studies through parametric power estimation". In: *J Pharmacokinet Pharmacodyn 43(2):223-234* (2016).