

PROAST Manual

Menu version

Table of Contents

- 1. INTRODUCTION..... 2
- 2. GENERAL SET-UP.....5.
- 3. EXAMPLE ANALYSES WITH PROAST (menu version).....7
 - 3.1 Illustrative example: Continuous data.....7
 - 3.2 Illustrative example: Quantal data.....11
 - 3.3 Illustrative example: Quantal data, with model averaging.....17
 - 3.4 Illustrative example: Continuous data with sex as covariate, and two endpoints considered.....18
 - 3.5 Illustrative example: Quantal data with sex as covariate, and two endpoints considered.....21
- 4. CREATING YOUR OWN DATAFILE IN THE PROAST FORMAT.....24
 - 4.1 Data format.....24
 - 4.2 Importing your data into R.....35
- 5. COMPREHENSIVE DESCRIPTION OF PROAST.....37
 - 5.1. Continuous individual data (data type = 1).....40
 - 5.1.0 Single model or set of models (continuous individual data).....41
 - 5.1.1 Data specification (Change settings) for continuous individual data.....41
 - 5.1.2 Main menu (continuous individual data, data type = 1).....47
 - 5.2. Continuous summary data (data type = 10).....61
 - 5.2.0 Single model or set of models (continuous summary data).....61
 - 5.2.1 Data specification (change settings) for data type = 10.....61
 - 5.2.2 Main menu (continuous summary data).....62
 - 5.3. Quantal data.....63
 - 5.3.0 Single model or set of models (quantal data).....63
 - 5.3.1 Data specification (Change settings) for quantal data.....63
 - 5.3.2 Main menu for quantal data.....64
 - 5.4 Binary data (data type = 2).....69
 - 5.4.1 Data specification (binary data).....72
 - 5.4.2 Main menu (binary data).....72
 - 5.5 Ordinal data (data type = 3).....73
 - 5.5.0 Single model or set of models (ordinal data).....75
 - 5.5.1 Data specification (change settings) for ordinal data.....76
 - 5.5.2 Main menu (ordinal data).....76
 - 5.6 Clustered continuous data, individual or clustered (dtype = 5, 15).....80
 - 5.7 Clustered quantal data (data type = 6).....81
- 6. SPECIFIC APPLICATIONS.....83
 - 6.1 Dose addition.....83
 - 6.2 Two-step analysis of data with two independent variables.....86
 - 6.3 CxT models.....86
- 7. THE R OBJECT WITH THE SAVED RESULTS.....87
- 8. THE PLOT LEGENDS.....90
- REFERENCES.....93

1. INTRODUCTION

Main purpose

PROAST has been developed for the analysis of toxicological dose-response data by the BMD (Benchmark dose) approach. However, it has a much wider range of applicability, and can be used for (nonlinear) regression in general (by replacing the term dose by any other independent variable, etc).

Environment

PROAST works under R (downloadable freeware, see section 2).

PROAST platforms

PROAST can be run as a stand-alone package under R, or by means of a web application (see the PROAST website).

The MENU and GUI version of PROAST (under R)

The MENU version of PROAST is completely based on multiple choice questions which the user needs to answer. The GUI (graphical user interface) version is more user-friendly, but the MENU option has, so far, more options.

This manual describes the MENU version and statistical background; the latter obviously holds for the GUI version (and web applications) as well. Users who prefer the GUI version are referred to the GUI manual for PROAST.

PROAST vs. BMDS

There are some differences in statistical assumptions and methods between BMDS (developed by USEPA) and PROAST. The aim is to make the PROAST consistent with the BMDS software to the extent possible.

An important difference with the BMDS software is that PROAST can be used to analyse different subgroups (e.g. sexes, compounds, studies) of dose-response data in a single analysis. PROAST then compares the dose-responses of the subgroups, and evaluates in which sense they differ and in which sense they do not differ from each other. For instance, sexes may differ in background response but otherwise their dose-responses may be similar. When the dose-responses related to different subgroups are at least partially similar, a combined analysis will result in higher statistical precision (more effective use of the available data).

Some important differences with BMDS that may lead to different results are:

- BMDS uses the normal distribution as the default for continuous data, PROAST uses the lognormal as the default
- BMDS puts constraints on the shape parameter in most quantal models, as the default setting, while PROAST does not.

Some further specifics of PROAST

- The data types that can be handled by PROAST are: continuous, quantal and ordinal data.
- Nested data (such as litter effects in developmental studies) are covered as well.
- For continuous data, PROAST focuses on four families of (nested) models: the exponential, the Hill, the inverse exponential and the LN model.
- For quantal data, PROAST uses the seven models recommended by EFSA (and also available in the BMDS software). In addition, two so-called latent variable models are available.
- The confidence intervals for the BMD or CED can be assessed by the Maximum Likelihood Profile method (this is also the method used by BMDS), or by the Bootstrap method.
- For continuous data, the default is analysis after log-transformation. Although generally not needed, it is possible to omit the transformation, or to apply the square root transformation. Plots of the regression residuals can be produced to check the assumptions of normality and homogeneity (for any of these transformations).
- Model averaging is implemented based on bootstrap sampling.
- Mixture studies can be analyzed by PROAST, by fitting DR models based on dose addition. In this way, it can be examined to what extent the mixture responses are predicted by dose addition.

Some major changes in EFSA guidance

In 2017, EFSA issued an update of its BMD guidance, which differs in some respects from the earlier guidance of 2009. As PROAST follows the EFSA guidance, PROAST has been updated concoringly, so that results from earlier versions (before 60.x) that were based on the same data may no longer coincide with those in newer versions, in particular in cases where various model families are fitted in a single run.

The two most important changes as compared to the earlier EFSA guidance and earlier PROAST versions are:

- Selecting or accepting models is no longer based on the log-likelihoods (using the likelihood ratio test) but rather on the AIC. At this point, a default value of 2 units difference between AICs is considered as the critical value by EFSA. However, PROAST allows the user to changing this critical value.
- The approach of finding the “minimal” model from the nested exponential or Hill families of models (in the case of continuous data) was abandoned. Instead, only two of the models (3 and 5, see below) are considered, and from these the one with lowest AIC is selected as the model for calculating the BMD confidence interval.
- In later versions (starting 66.x) model averaging (based on bootstrap sampling) is implemented, making model selection obsolete. Hence, the previous two bullets are no longer relevant when model averaging is applied (which is the preferred method according to EFSA).

ATTENTION

The development of PROAST (and the BMD approach as such) is going fast, and we are not always able to fully update the manual at the same time. Therefore, the manual may not always completely match with the latest PROAST version, and we kindly ask the user to be a bit flexible in matching the manual with the actual software.

2. GENERAL SET-UP

The figure below illustrates the main schedule for how one may navigate through the various options. The quickest analysis is given by selecting “set of models”. This is also the route that is followed by the GUI version of PROAST.

The choice to select a single model or a set of models can be made at the very beginning of the PROAST session (second question):

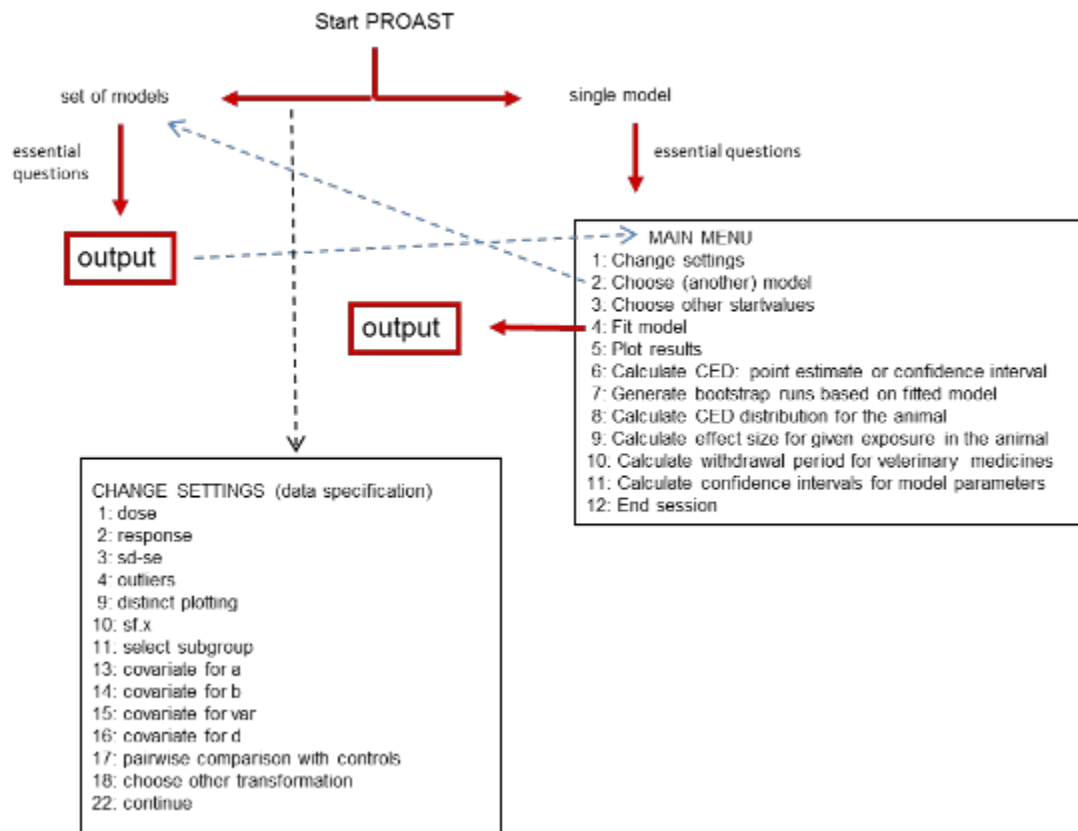
Do you want to fit a single model or fit various nested families of models?

- 1: single model*
- 2: change settings first*
- 3: select model 3 or 5 from various families of models*
- 4: select model 3 from various nested families of models*
- 5: select model 5 from various nested families of models*
- 6: select model 15 in terms of RPF*

The most usual answer is option 3 (see below). After that some questions will appear, and after answering those, the analysis will start, resulting in the output consisting of plots and numerical results in the Console window of R.

However, the user may also want to change a setting before applying the set of models, such as selecting a subgroup from the complete dataset. In that case, option 2 (*change settings first*) should be selected, which offers the opportunity to adjust settings 4, 8 or 18 in the *CHANGE SETTINGS* list.

When the option *single model* was selected, the users will arrive at the MAIN MENU (after some essential questions). Here, the user should choose a specific model (see option 2), and then choose option 4 to fit the model. Option 5 (*Plot results*) is useful to change or improve the plotted results. By choosing option 1 the user can always return to *CHANGE SETTINGS*, e.g. to choose another response variable, or to select another subgroup. Also note that option 2 in the main menu offers the possibility to go to *set of models*.



3. EXAMPLE ANALYSES WITH PROAST (menu version)

In this section various example analyses are shown, as a quick introduction to PROAST (menu version). These can be mimicked by the user. (These examples could also be run with the GUI).

The first two examples relate to a single dose-response dataset (in one sex). The third and fourth examples show a somewhat more advanced analysis, where sex is included as a covariate in a combined analysis of the dose-response data, and where more than one endpoint is analysed (consecutively).

Before mimicking these analyses, make sure that you have done the following:

R users:

Load the PROAST package (as usual)

Make the example data available by typing

```
> data(das1)
```

and

```
> data(das4)
```

and

```
> data(das11)
```

You may check if the data are available in the R workspace by typing

```
> ls()
```

NOTE:

During any PROAST session it is convenient to make the commands (console) window and the graphical window both visible at all times. You can do that as soon as the first plot appears. Adjust the size of the Console window to around half of the R window, and use the remaining space for the graphical window.

3.1 Illustrative example: Continuous data

In this example, the nested family of exponential models will be fitted to dose-response data on male kidney weights from a 28-day study, which used 7 dose groups.

Note: the input that should be given by the user is printed in bold.

First delete all existing plots (if applicable, in the case that plots were previously generated in R) by typing

```
> graphics.off()
```

This is to make sure that your plot numbers are the same as used below.

Type

> **f.proast(das11)**

Note 1: *das11* is the name of the data file. If the file *das11* cannot be found, first type:

> data(das11)

Note 2: You can always interrupt the analysis by using the Escape button, e.g. when you accidentally type a wrong answer to any of the PROAST questions.

Note 3: The first thing you are recommended to do is to move/resize the console (command) window to the left half of the R window. In this way the plots that are going to be produced will have room on the right half within the R window, so that you can fully see it when it appears. It is important that you can see both the console and the graphical window completely, since the analysis with PROAST is interactive, and the intermediate results will be given in both these two windows. If you use PROAST for the first time in a particular working directory, you may want to change the size of the graphical window. If so, type

> f.proast(das11, resize=T)

You can now interactively change the size of the graphical window, which will be remembered for later sessions.

Next, answer to the consecutive questions as follows:

What type of response data do you want to consider?

- 1: *continuous, individual data*
 - 2: *binary*
 - 3: *ordinal*
 - 4: *quantal*
 - 5: *clustered continuous, individual data*
 - 6: *clustered quantal*
 - 7: *continuous, summary data*
 - 8: *clustered continuous, summary data*
 - 9: *quantal, CxT*
 - 10: *other*
- Selection: **1**

Do you want to fit a single model or fit various nested families of models?

- 1: *single model*
- 2: *change settings first*
- 3: *select model 3 or 5 from various families of models*
- 4: *select model 3 from various nested families of models*
- 5: *select model 5 from various nested families of models*

6: *select model 15 in terms of RPF*

Selection: **3**

Q1: *Which variable do you want to consider as independent variable?*
(e.g. dose, age)

1 : *dose*
2 : *sex*
3 : *cage*
4 : *water*
5 : *food*
6 : *Group*
7 : *OECD*
8 : *ALT*
9 : *number*
10 : *BWDay0*
11 : *BWDay28*
12 : *BWabs*
13 : *relKidney*
14 : *relThym*
15 : *relBrain*
16 : *relSpleen*
17 : *relLiver*
18 : *Kidneys*
19 : *Thymus*
20 : *Brain*
21 : *Spleen*
22 : *Liver*
Selection: **1**

1 *dose*
2 *sex*
3 *cage*
4 *water*
5 *food*
6 *Group*
7 *OECD*
8 *ALT*
9 *number*
10 *BWDay0*
11 *BWDay28*
12 *BWabs*
13 *relKidney*
14 *relThym*
15 *relBrain*
16 *relSpleen*

17 relLiver

18 Kidneys

19 Thymus

20 Brain

21 Spleen

22 Liver

Give number(s) of the response(s) you want to analyse

Selection: **13** (= relative kidney weights)

Give number of factor serving as potential covariate (e.g.sex)

-- type 0 if none ---

1: dose	2: sex	3: cage	4: water	5: food
6: Group	7: OECD	8: ALT	9: number	10: BWDay0
11: BWDay28	12: BWabs	13: relKidney	14: relThym	15: relBrain
16: relSpleen	17: relLiver	18: Kidneys	19: Thymus	20: Brain
21: Spleen	22: Liver			

Selection: **0**

Give value for CES (always positive)

type 0 to avoid calculation of CIs > **0.05**

Note: CES stands for Critical Effect Size, i.e., the Benchmark Response (BMR) defined as a percent change in average response compared to the response in the controls

1

Do you want to calculate the BMD confidence interval by model averaging?

1: no

2: yes

Selection: **2**

Give number of bootstrap runs for calculating BMD confidence interval based on MA (e.g. 200) > **200**

ATTENTION: the constraints on parameter d in the exponential and Hill model are set at 0.25 and 4

type 0 if you want to change these constraints, otherwise enter any other number >**1**

Now the analysis starts.

When finished, seven plots are produced.

The last plot (number 7) shows the bootstrap curves based on model averaging. Plot 5 and 6 show the data with the selected models, one model for each sub-model family (inverse exponential, exponential, Hill, and NL) with parameters a , CED, c and d . The subtitle of the plot indicates the selected model.

The parameter b (see legend) was back-calculated from the CED and the other model parameters.

The graphical window (plot 1-4) before/behind the one just discussed shows the various (nested) models that have been fitted. In the first subplot (plot 1), related to the full model the responses are connecting by (dashed) straight lines. The full model consists of the observed means (with the assumed lognormal distribution for the within-group variation) without assuming any particular dose-response relationship.

The commands window produces some additional output, including a table of log-likelihood values with the number of parameters estimated for each model, followed by the BMD (=CED) confidence interval.

After answering the question
give name for file to store results (or type 0 if none)
 you will see the PROAST main menu. Type **13** to end the session.

Copying graphical output into Word

Graphical output in 1R can be copied and pasted into a word document. Right click on the graphical window to copy the plot (as metafile, usually).

Saving graphical output in file

In R the graphical window can be saved in various formats under *File* in the main bar of the R window. Note that the quality differs among formats, and some of them are not accepted by Journals.

3.2 Illustrative example: Quantal data

In this example, a set of models will be fitted to a quantal dataset. For those models that are considered to provide similar fits (based on the AIC criterion) the BMD confidence interval is calculated.

The dataset used contains the observed responses (forestomach and liver tumors) for both males and females. In this example we will only consider the males, and this subset of the data can be selected within PROAST as shown below.

Note: the input that should be given by the user is printed in bold.

Type

> f.proast(das4)

Note: if the data file *das4* is not found, you should first type:
 > data(das4)

What type of response data do you want to consider?

- 1: continuous, individual data
 - 2: binary
 - 3: ordinal
 - 4: quantal
 - 5: clustered continuous, individual data
 - 6: clustered quantal
 - 7: continuous, summary data
 - 8: clustered continuous, summary data
 - 9: quantal, CxT
 - 10: other
- Selection: **4**

Do you want to fit a set of models, or choose a single model?

- 1: single model
 - 2: set of models
 - 3: change settings first
- Selection: **3**

Which setting(s) do you want to change?

- 1:
 - 2:
 - 3:
 - 4: outliers
 - 5:
 - 6:
 - 7:
 - 8:
 - 9: distinct plotting
 - 10: scaling for x
 - 11: select subgroup
 - 12:
 - 13:
 - 14:
 - 15:
 - 16:
 - 17:
 - 18:
 - 19: constraints on shape parameter
 - 20: set critical difference in AIC
 - 21:
 - 22:
 - 23: relax fit conditions
 - 24: time variable
 - 25: continue
- Selection: **11**

*Q11a: Give number of (another) factor for which you want to select data
 --- type 0 if you want to analyze all (remaining) data ---*

- 1: dose.kg.bw
- 2: forestomach
- 3: liver
- 4: sample.size

5: sex

Selection: 5

Q11b: Give number(s) associated with the level(s) you want to select

1 : 1

2 : 2

Selection > 1

*Q11a: Give number of (another) factor for which you want to select data
--- type 0 if you want to analyze all (remaining) data ---*

1: dose.kg.bw

2: forestomach

3: liver

4: sample.size

5: sex

Selection: 0

Note: this question is repeated, since in some cases you may want to select a subset for another factor as well

Which setting(s) do you want to change?

1:

2:

3:

4: outliers

5:

6:

7:

8:

9: distinct plotting

10: scaling for x

11: select subgroup

12:

13:

14:

15:

16:

17:

18:

19: constraints on shape parameter 20: set critical difference in AIC

21:

22:

23: relax fit conditions

24: time variable

25: continue

Selection: 25 (= continue)

Do you want to fit a set of models, or choose a single model?

1: single model

2: set of models

Selection: 2

Q1: Which variable do you want to consider as independent variable?

(e.g. dose, age)

- 1 dose.kg.bw
 - 2 forestomach
 - 3 liver
 - 4 sample.size
 - 5 sex
- Selection: **1**

- 1 dose.kg.bw
- 2 forestomach
- 3 liver
- 4 sample.size
- 5 sex

Which response(s) you want to analyse
by set of models > **2**

- 1 dose.kg.bw
- 2 forestomach
- 3 liver
- 4 sample.size
- 5 sex

Enter column number(s) with the associated sample sizes > **4**
Give number of factor serving as potential covariate (e.g.sex)
-- type 0 if none ---

- 1: dose.kg.bw
 - 2: forestomach
 - 3: liver
 - 4: sample.size
 - 5: sex
- Selection: **0**

What type of Benchmark response do you want to consider?
Type 0 if you do not need CIs

- 1: ED50
 - 2: Additional risk, i.e. $P[\text{BMD}] - P[0]$
 - 3: Extra risk, i.e. $(P[\text{BMD}] - P[0]) / (1 - P[0])$
 - 4: CED for latent variable
- Selection: **3**

Give value for the BMR,

in terms of extra risk > **0.10**

Do you want to calculate the BMD confidence interval by model averaging?

1: no

2: yes

Selection: **1**

Now the calculations start. The last graphical window is the important one, and shows all the models fitted.

(The previous graphical window shows (some of the) members of the exponential latent variable model fitted to the data, but this window can normally be ignored.)

You will see the following table in the commands window:

```

      model No.par loglik  AIC accepted BMDL BMDU BMD conv
1      null    1 -138.59 279.18      NA NA NA
2      full    4 -47.00 102.00      NA NA NA
3 two.stage    3 -47.00 100.00   yes 1.25 2.86 2.24 yes
4 log.logist   3 -47.65 101.30   yes 1.86 3.22 2.53 yes
5   Weibull    3 -47.00 100.00   yes 1.43 3.18 2.28 yes
6 log.prob     3 -47.25 100.50   yes 1.91 3.16 2.54 yes
7   gamma     3 -47.02 100.04   yes 1.54 3.16 2.38 yes
8 logistic     2 -48.90 101.80   yes 2.44 4.00 3.15 yes
9  probit      2 -48.37 100.74   yes 2.29 3.71 2.94 yes
10 LVM: Expon. m3- 3 -47.02 100.04   yes 1.43 3.20 2.33 yes
11 LVM: Hill m3-  3 -47.07 100.14   yes 1.58 3.20 2.40 yes

```

BMR: 0.1 extra risk

constraint on steepness parameter: 0.01

no litter effects

critical AIC value: 2

PROAST version: 66.30

This table can be transported into an editable table in Word as follows.

Browse to your working directory. Click on “last.BMDtable”, and open with notepad.

Copy the content of table.tmp into your Word document. Then, in the Word document, select the table (including the header, starting with ‘model’), and click on Insert > Table > Insert table. You should now have a regular Word table, as shown below, which you can adjust in the usual way.

das4 , forestomach :

model	No.par	log-likelihood	accepted	AIC	BMDL	BMDU	BMD

null	1	-138.59		279.18	NA	NA	NA
full	4	-47		102	NA	NA	NA
two.stage	3	-47	yes	100	1.25	2.86	2.24
log.logist	3	-47.65	yes	101.3	1.86	3.22	2.53
Weibull	3	-47	yes	100	1.43	3.18	2.28
log.prob	3	-47.25	yes	100.5	1.91	3.16	2.54
gamma	3	-47.02	yes	100.04	1.54	3.16	2.38
logistic	2	-48.9	yes	101.8	2.44	4	3.15
probit	2	-48.37	yes	100.74	2.29	3.71	2.94
LVM: Expon. m3-	3	-47.02	yes	100.04	1.43	3.2	2.33
LVM: Hill m3-	3	-47.07	yes	100.14	1.58	3.2	2.4

PROAST version: 66.30

no covariate

BMR: 0.1 extra risk

lower constraint on steepness parameter: 0.01

Then, after the question

give name for file to store results (or type 0 if none) > **forestom.fit**

PROAST will store the results in a file in the R workspace (as an R object) with the name provided. If you stop the PROAST analysis here, you may proceed this analysis later on, by typing

```
> f.proast(, forestom.fit)
```

and you are back where you just left PROAST.

The next questions can be used to continue with one specific model, e.g. for plotting or further analysis.

*Which model do you want to continue with?
type 0 for all models*

1: two.stage

2: log.logist

3: Weibull

4: log.prob

5: gamma

6: logistic

7: probit

8: EXP

9: HILL

Type 4 if you want for instance the log.prob model.

Type 13 to end the session.

3.3 Illustrative example: Quantal data, with model averaging

To see how model averaging can be applied in an analysis of quantal data, just repeat the previous example, but now after the question

Do you want to calculate the BMD confidence interval by model averaging?

1: no

2: yes

select **2 (yes)**.

This will prompt the question

Do you want to include the latent variable models in model averaging?

1: no

2: yes

and, if you want to follow EFSA guidance, answer 2 (yes). Then, you will be asked to enter the number of model averaging bootstrap runs you want to be applied. Usually, 200 runs will give a reasonably precise answer. The more runs, the more precise the answer will be, but the calculation will take more time, of course.

After having fitted the suite of models, as usual, model averaging starts. You will see a new plot with the data, to which a new dashed curve is added, representing the model average curve for that bootstrap runs. The small circles close to the x-axis are the CEDs associated with each curve. The distribution of these points defines the final model average confidence interval for the BMD.

3.4 Illustrative example: Continuous data with sex as covariate, and two endpoints considered

In PROAST, the user may appoint a factor (e.g. sex) to be handled as a covariate in the model. PROAST will then explore if particular parameters in the model should receive different values, depending on the level of the covariate (e.g. males vs. females of the covariate sex). For instance, the parameter reflecting the background response (e.g. BW in the controls) may differ between males and females. Further, one of the sexes may be more sensitive to the compound than the other (reflected by different values for the slope parameter b).

It is optional to additionally apply the family of Hill models, which consists of a set of five models that are analogous to the five exponential models (i.e., they contain the same parameters which have similar meanings). Comparing the selected exponential model with the selected Hill model provides some information on model uncertainty.

A whole list of continuous endpoints can be given as an input, and PROAST then automatically produces the results for the selected model of each family of models for each endpoint consecutively.

The following example can be mimicked as an illustration.

Type

```
> f.proast(das1)
```

```
Note : das1 is the name of the data file. If the file das1 cannot be found, first type:
> data(das1)
```

What type of response data do you want to consider?

```
1: continuous, individual data
2: binary
3: ordinal
4: quantal
5: clustered continuous, individual data
6: clustered quantal
7: continuous, summary data
8: clustered continuous, summary data
9: quantal, CxT
10: other
Selection: 1
```

Do you want to fit a single model or fit various nested families of models?

```
1: single model
```

2: *change settings first*
 3: *select model 3 or 5 from various families of models*
 4: *select model 3 from various nested families of models*
 5: *select model 5 from various nested families of models*
 6: *select model 15 in terms of RPF*
 Selection: **3**

Q1: *Which variable do you want to consider as independent variable?*
(e.g. dose, age)

1 : *Dose*
 2 : *LDH*
 3 : *UBH*
 4 : *BW*
 5 : *sex*
 6 : *food*
 Selection: **1**

1 *Dose*
 2 *LDH*
 3 *UBH*
 4 *BW*
 5 *sex*
 6 *food*
Give number(s) of the response(s) you want to analyse
 Selection: **c(2,4)**

The last answer means that you want to analyse both response number 2 and response number 4. Here, c(..) is an R function (c stands for concatenate) which puts the two numbers together in a single vector, to be treated as a single answer.

Give number of factor serving as potential covariate (e.g.sex)
 -- type 0 if none ---

1: *Dose*
 2: *LDH*
 3: *UBH*
 4: *BW*
 5: *sex*
 6: *food*
 Selection: **5**

Do you want to adjust CES to within group SD?

1: *no*
 2: *yes*

Selection: **1**

Note: this option relates to a specific application, not discussed in the manual.

*Give value for CES (always positive)
type 0 to avoid calculation of CIs > 0.05*

Do you want to interrupt calculations after each endpoint?

1: no
2: yes (opportunity to save results per endpoint)

Selection: **2**

Do you want to calculate the BMD confidence interval by model averaging?

1: no
2: yes

Selection: **1**

Which models to you want to be fitted?

1 : Exponential model only
2 : Exponential and Hill model
3 : previous option with inverse exponential model added
4 : previous option with lognormal DR model added

Selection:**4**

ATTENTION: the constraints on parameter d in the exponential and Hill model are set at 0.25 and 4

type 0 if you want to change these constraints,
otherwise enter any other number >1

Now the analysis starts. It will produce the four plots for LDH (one for each model).

*give name for file to store results (or type 0 if none) > **das1.ldh.fit***

PROAST will then produce the four plots for BW (one for each model)

*give name for file to store results (or type 0 if none) > **BW***

After that, PROAST will produce a plot with the CED confidence intervals, both for the female subgroup, and (after pressing a key to continue) for the male subgroup. For each endpoint the upper CI relates to the exponential, the second to the Hill model, the third to

the Inverse Exponential model, and the fourth to the LN model. This option in PROAST is particularly useful in analyzing a larger number of endpoints from a given study in a single run (you may use option not to interrupt the calculations after each endpoint). Of course, the endpoints must have the same data type, such as continuous or quantal.

3.5 Illustrative example: Quantal data with sex as covariate, and two endpoints considered

Type

```
> f.proast(das4)
```

Note: *das4* is the name of the data file. If the file *das4* cannot be found, first type:

```
> data(das4)
```

What type of response data do you want to consider?

- 1: continuous, individual data
 - 2: binary
 - 3: ordinal
 - 4: quantal
 - 5: clustered continuous, individual data
 - 6: clustered quantal
 - 7: continuous, summary data
 - 8: clustered continuous, summary data
 - 9: quantal, CxT
 - 10: other
- Selection: 4

Do you want to fit a set of models, or choose a single model?

- 1: single model
 - 2: set of models
 - 3: change settings first
- Selection: 2

*Q1: Which variable do you want to consider as independent variable?
(e.g. dose, age)*

- 1 : dose.kg.bw
- 2 : forestomach
- 3 : liver
- 4 : sample.size

5 : sex

Selection: 1

1 dose.kg.bw

2 forestomach

3 liver

4 sample.size

5 sex

Which response(s) you want to analyse by set of models > 2:3

Note: the notation 2:3 means all response from 2 up to 3. This may be helpful for larger numbers of responses, if they form a non-interrupted list. Note that 2:3 and c(2,3) are equivalent.

1 dose.kg.bw

2 forestomach

3 liver

4 sample.size

5 sex

Enter column number(s) with the associated sample sizes >

Selection: 4

Note: when sample sizes differ between the endpoints indicated in the previous question, you need to provide the associated column numbers using c(...)

Give number of factor serving as potential covariate (e.g.sex)

-- type 0 if none ---

1: dose.kg.bw

2: forestomach

3: liver

4: sample.size

5: sex

Selection: 5

What type of Benchmark response do you want to consider?

type 0 if you do not need CIs

1: ED50

2: Additional risk, i.e. $P[\text{BMD}] - P[0]$

3: Extra risk, i.e. $(P[\text{BMD}] - P[0]) / (1 - P[0])$

4: CED for latent variable

Selection: 3

Give value for the BMR,

in terms of extra risk > **0.10**

Do you want to calculate the BMD confidence interval by model averaging?

1: no

2: yes

Selection: **1**

Do you want to interrupt calculations after each endpoint?

1: yes

2: no

Selection: **2**

Here the analysis starts. When it is finished you will see two plots, one for each endpoint, with all models fitted.

4. CREATING YOUR OWN DATAFILE IN THE PROAST FORMAT

The first part of this section describes how to format your dataset to make it suitable for analysis, with various examples. The second part shows how to import the dataset into R.

4.1 Data format

Attention:

Make sure that the decimal point in your excel sheet is always a point (dot), not a comma, as is the habit in some countries; in Windows this can be changed by clicking on: start – control panel – region and language – change number format - additional settings.

You can create the data file in any suitable application, e.g. in Excel. The format of the data however needs to obey some strict rules, as will now be discussed.

The information on dose, response, and any other relevant factor is presented by consecutive columns, where the rows relate either to an individual animal (or other experimental unit), or to a whole dose group. The latter may occur with quantal data, or with continuous data where the responses are given as group means. When the rows relate to dose groups, the group sizes are required as an additional column. For continuous group means, an additional column is needed, representing either the SDs or the SEMs for that group.

Each column requires a header (see below for format restrictions).

No empty cells are allowed. Missing values need to be represented by “NA” (without the quotes).

Do not enter numerical values with only one significant figure (like 0.1). Use at least 3 significant figures (e.g. 0.124).

Do not use spaces in any of the cells. The entries should be single strings, which may include dots, or underscores, but not special characters like * or /.

Keep the entries in the cells as short as possible, in particular the ones that may be used in the analysis. The reason is that the entries may be used in labeling model parameters, which are presented in the legends of the plots. Long entries will make the legends with the plots harder to read (or the text will not fit in the window).

If needed, the meaning of the brief entries in the data sheet may be explained in an associated sheet called “key”.

Make sure that the dimensions related to the entries in any quantitative column (like dose) are exactly the same (e.g. mg / kg body weight).

Make sure that all entries in a column that will or may be used as quantitative information in the analysis are indeed quantitative values. For example, indicating group size as 5-7 (varying between 5 and 7) is not allowed.

Tip: In Excel, numerical values are aligned to the right, and non-numerical entries to the left. However, alignment might have been changed by the user, which might mislead you. To avoid that, select all cells of the sheet, click on the button ‘alignment’ (between ‘font’ and ‘number’ under the home tab), and select ‘general’ in the window “horizontal:” Now, you can see which entries are numerical or not, from the way they are aligned.

Data types

The analysis of dose-response data depends on the type of data you are considering. PROAST uses a code (number) for distinguishing the various data types (used in the legend of the plots):

Table 1. Data types and code used in PROAST

Data type	Code in PROAST
continuous data (e.g. weights, concentrations, counts), assuming lognormal distribution	1
binary data, i.e., yes or no response in individual animals, given as 0 and 1	2
ordinal data (e.g. histopathological scores of severity)	3
quantal data, i.e. number of responding animals per number of animals in the group	4
continuous data that are nested, e.g. fetal weights within litters	5
quantal data that are nested, e.g. number of fetuses affected within litters	6
Summary statistics for continuous data, assuming lognormal distribution	10
Summary statistics for clustered continuous data, assuming lognormal distribution	15
Continuous data, assuming normality	25
Continuous data, assuming normality after square root transformaton	26
mean (continuous) response in a dose group, assuming normal distribution	250
Continuous data, assuming normality after square root transformaton	26
Summary statistics for continuous data, assuming normal distribution	250
Summary statistics for continuous data,	260

assuming normal distribution after square root transformation	
---	--

Example formats of various data types

Example of data set (continuous, data type = 1) in PROAST format:

In this example, the dose is PCB.dose, the response is EROD, while there is one additional factor (sex).

PCB.dose	EROD	sex
2	0.33	m
4	0	m
8	0.12	m
16	0.16	m
32	0.2	m
64	0.65	m
125	0.92	m
256	2.05	m
512	2.61	m
1024	4.62	m
2048	6.59	m
4096	5.27	m
2	0.15	f
4	0.06	f
8	0.06	f
16	0.08	f
32	0.1	f
64	NA	f
125	0.47	f
256	1.03	f
512	1.35	f
1024	2.3	f
2048	3.3	f
4096	3	f

Example of data set (continuous, data type = 1) with detection limits in PROAST format:

When (continuous) data contain zero's these usually indicate an observation below the detection limit. When all observations can be assumed to have the same detection limit, this value can be entered interactively during the PROAST session. If not, the different detection limits need to be given in the data sheet, with a separate column for the various values, as illustrated below.

dose	IL_18	compound	detlim_IL18
0	120.34	1	100
0	0	1	100
0	0	1	100
1.4	1830	1	100
14	2776	1	100
14	2177	1	100
140	6467	1	100
140	6096	1	100
140	5040	1	100
1400	4415	1	100
1400	5319	1	100
1400	5271	1	100
0	120.34	2	10
0	0	2	10
0	0	2	10
1.4	296.38	2	10
1.4	0	2	10
1.4	0	2	10
14	32.97	2	10
14	0	2	10
14	0	2	10
140	136.77	2	10
140	123.38	2	10
1400	574.93	2	10
1400	753.43	2	10
1400	1128.3	2	10

Example of data set (continuous summary data, data type = 10) in PROAST format:

Note that continuous summary data must include a measure for the within-group variation (SD or SEM) and the group size.

dose	sex	mean.bw	sd.bw	n.bw
0	1	704	124.7	33
0.5	1	739	140.5	35
3.5	1	742	97.7	40
25	1	646	119.4	41
50	1	572	97	49
0	2	496	105.7	37
0.5	2	477	132.6	33
3.5	2	480	106.8	32
25	2	402	106.8	27
50	2	361	81.1	20

Example of data set (quantal, data type = 4) in PROAST format:

dose	sex	response	group.size
0	1	0	20
0.5	1	2	18
3.5	1	5	19
25	1	6	20
50	1	12	18
0	2	1	20
0.5	2	0	19
3.5	2	3	20
25	2	7	19
50	2	10	19

Example of data set (binary, data type = 2) in PROAST format:

Here, for each animal (experimental unit) it is indicated if it responded or not, where 0 is a non-response.

dose	response	sex
2	0	1
4	0	1
8	0	1
16	1	1
32	0	1
64	1	1
125	1	1
256	0	1
512	0	1
1024	1	1
2048	1	1
4096	1	1
2	1	2
4	0	2
8	0	2
16	0	2
32	1	2
64	0	2
125	1	2
256	1	2
512	1	2
1024	0	2
2048	1	2
4096	1	2

Example of data set (ordinal, data type = 3) in PROAST format:

Ordinal data are scores of severity categories in each animal (e.g. normal, minimal, mild, moderate, severe). They should be provided as score values (e.g. 0, 1, 2, ...), where 0 is 'normal'. If you use other scores, PROAST will try to translate them into scores 0, 1, 2, etc, where the 0 score reflects "normal". This translation will be shown in the Console window. You should keep in mind when interpreting the PROAST output, which relates to the translated scores(' temp.scores'). However, PROAST will always assume that higher scores relate to more serious effects.

These scores need to be provided for each individual animal separately, just like in binary data, see previous example. The difference with binary data is that in ordinal data a lesion can be scored by more than two values (0 and 1).

The appropriate format for ordinal data is like that for binary data, where each animal is associated with a particular score. Usually, however, the data are available as incidences related to each score. These can be easily transferred into individual data by using the PROAST function **f.datatransfer.ord** (see section 4.2). To make this function work, make sure that the incidences are available as in the following example, where S0, S1, S2, ... are the severity scores, and the numbers in these columns reflect the number of animals with that score. The dose *must* be in the first column, and the columns for any additional factors *must* come after the columns with the scores.

Dose	S0	S1	S2	S3	sex	other factors...
0	5	0	0	0	f	
20	5	0	0	0	f	
100	3	1	1	0	f	
500	1	2	2	0	f	
0	5	0	0	0	m	
20	2	2	1	0	m	
100	3	1	1	0	m	
500	1	0	3	1	m	

Example of data set (clustered individual continuous, data type = 5) in PROAST format:

Here, a column is needed indicating the factor defining the clustering (litter effect), in this example a dam number.

dose	dam	foetalweight
0	2	4.11
0	2	4.22
0	2	4.27
0	2	4.26
0	2	4.03
0	2	4.13
0	2	4.07
75	6	3.96
75	6	3.97
75	6	4.05
75	6	3.39
75	6	3.27
350	7	4.86
350	7	4.31
350	7	4.78
350	7	4.62
350	7	4.63
350	7	4.53
350	7	4.3
350	7	4.26
350	7	4.44
580	9	3.62
580	9	4.32
580	9	4.18
580	9	3.84
580	9	3.86
580	9	3.64
580	9	4.03
580	9	3.71

Example of data set (clustered quantal, data type = 6) in PROAST format:

In clustered quantal data, each response incidence relates to a particular dam (or level of the factor defining the clustering). Therefore, there is no need for a column defining the clustering factor.

dose	sex	response	group.size
0	1	0	20
0	1	2	18
0	1	5	19
0	1	6	20
5	1	12	18
5	1	1	20
5	1	0	19
10	1	3	20
10	1	7	19
10	1	10	19

Example of data set (clustered continuous summary statistics, data type = 15) in PROAST format:

dose	meanresp	sd	N
750	3.96	NA	1
350	4.52	0.215	9
580	4.04	0.370	3
970	3.06	0.466	2
0	4.12	0.176	5
450	4.67	0.176	8
0	4.28	0.115	3
350	4.51	0.173	8
0	4.89	0.135	3
750	3.48	0.09192	2
750	4.28	0.108	4
270	4.12	0.235	5
450	4.49	0.335	5
580	4.75	0.231	7
970	2.34	NA	1

Note that in this case there is no need to indicate the specific litter, just as in the case of clustered quantal data, as each row relates to a litter.

Example of mixture dataset

In this case, the doses for each single chemical are entered in a separate column. The column “Dosing type” denotes if the row relates to a single chemical or a mixture. The next columns provide the responses, and the covariates, if applicable. This example relates to individual continuous data, but similar formats apply for other data types, using the instructions for that data type above.

IE.dose	CA.dose	Dosing type	cells.ALN	IL10.ml	IFNg.ml
0	0	IE	8029000	337	58
0	0	IE	6675200	382	86
0	0	IE	7413000	422	119
0	0	IE	6443850	447	104
0	0	CA	7171500	306	14 ¹⁾
0	0	CA	6647900	546	161
0	0	CA	6144133	365	93
0	0	CA	7710500	532	230
0.04	0	IE	8036000	521	255
0.04	0	IE	20458667	727	482
0.04	0	IE	8596000	538	260
0.04	0	IE	9068500	408	198
0	0.32	CA	12453000	1028	2167
0	0.32	CA	13804000	1013	1334
0	0.32	CA	6221250	951	1282
0	0.32	CA	12971000	938	1167
0.02	0.02	mix	8092000	714	574
0.02	0.02	mix	5680500	562	280
0.02	0.02	mix	7045500	800	577
0.02	0.02	mix	9464000	616	399
0.04	0.12	mix	16597000	1393	1426
0.12	0.04	mix	15053500	1562	1570
0.24	0.08	mix	16961000	1932	3180
0.24	0.08	mix	14686000	1590	1816
0.24	0.08	mix	17199000	1431	1158
0.24	0.08	mix	13953333	1382	1811
0.08	0.24	mix	11791500	1228	1800
0.08	0.24	mix	17031000	969	1382

0.08	0.24	mix	13349000	982	965
0.08	0.24	mix	14703500	1089	1425

4.2 Importing your data into R

Saving the excel sheet

After having completed your data set in the PROAST format you need to save it in ‘text’ format. Within Excel this is done by saving the sheet as ‘text, tab delimited’, in the window following “save as”. Thus, the file will get the extension ‘.txt’. Choose a name for the data in text format, say *name.txt* (where name is replaced by a string that makes the dataset identifiable).

Now, open R, and type

```
> getwd()
```

which will show you the directory that R considers as the current working directory. If this is not the same directory as the one where you saved your data, you need to change that. Click on *File* in the left upper corner of the R window, and select *Change dir ...*

Then browse to the directory where you have saved your dataset. You can check this by typing

```
> getwd()
```

once more, or by typing

```
> dir()
```

which will show you the files present in the current working directory, including your dataset.

You are now ready to import the dataset into R by

```
> f.scan()
```

which shows you the list of files in the current working directory. Enter the number with the data file you want to import, and next provide a name to be used for the data object in R.

Alternatively, you can use:

```
> name <- f.scan('name.txt')
```

Either way, the function `f.scan` reads the text file ‘name.txt’, and transforms the information into the R object ‘name’. Of course, you can use any name instead of ‘name’.

Error messages after f.scan()

If you get the message that R cannot read the file, this is probably because it is not located in the appropriate directory. Check if the dataset is available by typing

```
> dir()
```

If you get the message “a reading error”, this might be due to, a.o.:

- empty cells
- some name in the headers is not a “single word”, e.g., it contains a space or not-allowed special character.
- some entry is not a single word or not a single number, e.g. it contains a space or not-allowed special character.

R indicates the location of the error by a line number.

One option to find an error in the dataformat is by selecting the first few rows and save them in a test file. Then try to read the test file by f.scan. By selecting varying numbers of rows you may be able to locate the error. (E.g., select half of the file, check, and if not yet OK, take half of the remaining rows, etc.)

Finally, you can check if the data are available in your R Workspace in the right format by typing the name of the imported dataset.

Special case: ordinal data

When you have ordinal data represented by incidences for each severity score, import that dataset with f.scan (see the required format discussed in section 4.1). Assuming the name of the imported dataset is ord.incid, you can transfer the data into the required format with a score for each individual animal by

```
> ord.indiv <- f.datatransfer.ord(ord.incid)
```

The individual scores will now be available in ord.indiv. Of course, if you have another number of severity categories, you should adjust the number 6.

To check if the data are in the appropriate format, type ord.indiv (or other name that you used), and you will see the data in the new format. This dataset is now suitable for analysis by PROAST.

5. COMPREHENSIVE DESCRIPTION OF PROAST

The use of PROAST for automated model selection has been illustrated in section 3. In this section various (other) options in PROAST will be discussed in more detail.

Arguments of the function `f.proast()`

A PROAST session (in the MENU version) is invoked by typing

```
> f.proast(datafilename)
```

where `datafilename` is the first argument of the PROAST function.

The function `f.proast` has a few other arguments, that might be used:

Second argument

The second argument is the name of an R object with the stored results of a previous analysis. So, typing

```
> f.proast( , resultsfilename)
```

will show the the results of that previous analysis. You can then continue the analysis starting from there.

Third argument

The third argument might be useful when at some point in answering the PROAST questions something went wrong, and the PROAST session was interrupted (because the PROAST session crashed, or because you pressed *Esc*). If you type

```
> f.proast(datafilename, er=T)
```

the analysis will restart, skipping the questions that were already answered. However, it may occur that PROAST gives the message that the option cannot be used and then you need to start all over.

Fourth argument

When you type

```
> f.proast(datafilename, resize=T)
```

PROAST allows you to resize your graphical window. Once you have done this, the size of the graphical window that you chose will be maintained (as long as you will be working in the current working directory).

Fifth argument

This is a more advanced option, and will be rarely used. It might be useful when you have fitting problems. When you type

```
> f.proast(datafilename, scale.ans = T)
```

you get the opportunity to change the scaling of the parameters that is in the optimization algorithm. See R manual on the function `nlminb`.

Entering data type

When starting a PROAST analysis by typing `f.proast(...)` the first question is:

What type of response data do you want to consider?

- 1: *continuous, individual data*
- 2: *binary*
- 3: *ordinal*
- 4: *quantal*
- 5: *clustered continuous, individual data*
- 6: *clustered quantal*
- 7: *continuous, summary data*
- 8: *clustered continuous, summary data*
- 9: *quantal, CxT*
- 10: *other*

Here, you need to indicate the appropriate data type or your dataset. See section 4 for the required data format.

When you select option 10 (*other*) you can directly enter the PROAST code for that data type (see Table 1), including data types that are not listed here (e.g. 25 = continuous data assuming normal distribution).

Make sure that you enter the appropriate answer here, otherwise your analysis will not be valid: the dose-response analysis that is required depends on the data type that you are analyzing. Therefore, the various options in PROAST will be discussed for the different data types separately.

Single or set of models

With the second question you make clear if you want to fit a specific single model, or if you want to do an automated runs of a sequence of models. This question depends on the data type, and will be discussed in the appropriate section.

NOTE: The current GUI version only provides for the automated analysis.

During a PROAST session (menu version) it is possible to change from automated run fitting a set of models to an analysis where a single model is fitted, and *vice versa*.

Although the dose-response analysis is quite different for the different data types, the general framework of the menu version of PROAST is similar. Basically, PROAST consists of:

- a main menu, which gives various options to be performed (e.g. choosing a model, fitting a model, plotting the results).

- a set of questions that are used to specify the additional settings (e.g., selection of a sub-part of the data, selection of a covariate). These questions start with the letter “Q” followed by a number.

It is recommended to first go through the description of continuous data (data type=1), since many options are similar for the other data types (quantal, or ordinal).

For each data type, an example dataset is available under the names: das1, das2, das3, das4, das5, das6, das10, and das15, for continuous, binary, ordinal, quantal, clustered continuous, clustered quantal, continuous summary, and clustered continuous summary data, respectively.

5.1. Continuous individual data (data type = 1)

This section describes the use of PROAST for continuous data, with an observation for each individual subject (i.e. data type = 1, see table 1). An analysis for this type of data will be started after choosing option 1 after the first question when typing `f.proast(...)`.

It should be noted that in PROAST the default assumption for continuous data (i.e. data with a natural lower bound of zero, and with a dimension, such as g, L, s, m, or any combination) is that they are lognormally distributed. Therefore, the analysis is performed on log-scale, i.e., the response data, as well as the dose-response model, is log-transformed, and back-transformed after the statistical analysis. In this way, the dose-response model, and the meaning of its parameters, remains as it is. Note, however, that the parameter *var* reflects the scatter (residual variance) around the fitted curve on the (natural) log-scale. PROAST also provides a value for the residual variation in terms of a Coefficient of Variation (CV), which is shown in the console window after a model has been fitted.

Further, it is important to note that the fitted curves in the final plots relate to the median (= geometric mean) at each dose. Similarly, if group means are plotted, these reflect geometric means, not arithmetic means. See e.g. Slob (1994) for some arguments of choosing the geometric mean (median) as the appropriate measure of central tendency.

It is possible to deviate from the lognormal default, and omit any transformation, or apply a square root transformation (see Q19).

For continuous data a zero in the observations normally means that the real value is somewhere below the detection limit, not that it is truly zero. PROAST offers the possibility to enter a detection limit via question Q6 (see below), which will be used in the analysis, assuming that the observed “zero” actually means any (positive) value below the detection limit that was entered in Q6. It is also possible to have multiple detection limits, for instance when the data include different studies (see section 3 for an example dataset in the PROAST format).

Although counts are discrete rather than continuous data, they may also be analysed as data type 1 in many cases, since they can often be approximated by a lognormal distribution. Zero counts might also be considered as an observation below a limit of detection in specific cases. For example, when the number of leucocytes is counted out of 100 white blood cells, zero counts will occur more often than when 1000 white blood cells are counted. So, in a sense, the second case has a lower “limit of detection” (lower probability of finding zero counts) than the first case.

As an alternative, the user can add a small number to the counts, but this will probably lead to a distorted lognormal distribution (see Q6-7 below). The assumption of lognormal distribution can be checked after the model has been fitted (see option 5 of main menu).

5.1.0 Single model or set of models (continuous individual data)

The menu version of PROAST allows for fitting a single model, selected by the user from a list of models, or for fitting a set of models in an automated single run. This choice can be made after the second question in the PROAST session:

Do you want to fit a single model or select a model from a nested family of models?

1: single model

2: select model 3 or 5 from nested (exponential/Hill) family of models

3: change settings first

4: select model 3 from nested (exponential/Hill) family of models

5: select model 5 from nested (exponential/Hill) family of models

By choosing option 1 (single model) the user has more freedom and more options for doing the analysis as compared to options 2, 4 and 5, which select an appropriate model after having fitted a set of models. However, the automated runs are much faster, and in many cases are fit for purpose.

Risk assessors who follow EFSA guidance (2016) should select option 2 here. The ensuing analysis will compare the fits of models 3 and 5 (possibly with covariates) and select the model with the lowest AIC. The BMD confidence intervals will be calculated based on that selected model (for both exponential and Hill).

Options 4 and 5 may be useful for specific applications.

If you first want to change the settings, such as selecting a subset from the data, choose option 3. The possibilities to change the settings are discussed in the next session.

5.1.1 Data specification (*Change settings*) for continuous individual data

When choosing option 1 of the main menu, a list of options for changing the settings appears. The questions discussed here are numbered Q1 – Q24, and these can be used to change the dose or the response or to change additional settings.

Q1: Which variable do you want to consider as independent variable?

Here, any column can be chosen that the user wants to consider as the independent variable, i.e. the variable to be plotted on the x-axis. For instance, in inhalation studies it will be concentration rather than dose. Or, the interest may be in a change in response with time instead of dose.

When the interest is in both dose and time, this may be achieved in a two-step procedure or by fitting a model with two independent variables. See section 6.

Q2: Which response variable do you want to analyse?

Choose the column with the response of interest.

Q4: If you want to remove outliers, type number of factor to indicate them

--- type 0 if you do not want to remove any data ---

Normally, this question will be relevant only after a first analysis (with no outliers removed) has been done. When the result of that first analysis indicates that some data point appears to be an outlier, you may wish to exclude it. However, this option may also be helpful for selecting a subset of the complete dataset, in particular if you want to exclude some levels of a particular factor (e.g, exclude observation day 0 from a range of different observation days).

Type the number of the column that you want to use for removing data. When the column represents a factor (i.e. with nonnumerical values) all values will appear as a column on the screen preceded by a number. In that case, select the number(s) preceding the values you want to exclude. When the column represents a numerical variable (i.e. with numerical values only) these values will appear as such, and you need to enter the value(s) to be removed themselves.

NOTE:

When you enter multiple numbers, you need to concatenate them, to make clear that they should be consider as a single answer. This can be done by using the R function `c()`.

When you type

```
> 3,6,10
```

In the Console window you will get an error message, but when you type

```
> c(3,6,10)
```

you get the three numbers back as a single string (vector). Consecutive numbers can be indicated by separating the first and last by a colon, e.g.

```
> 3:7
```

is read as the string (vector) 3 4 5 6 7

After having indicated which data you wanted to delete, PROAST will show the values for the response variable that are deleted.

Note: When you want to get the removed outliers back in the dataset, first refresh the complete dataset by typing 0 after invoking this question, and then invoke this question again.

Q5: Which variable do you want as denominator ? (0 = NONE)

This question may be helpful, for instance, if your data contain organ weights that need to be divided by the body weights in the same dataset. Simply enter the column number that is associated with body weights.

*Q6: Give global value of detection limit,
or type 0 for individual values >*

Enter the value of the detection limit, if it is the same for all observations. If you do not know it, give an estimate, which can be based on the lowest observation (which was printed just preceding this question).

It could also occur that the detection limit is not the same for each observation (e.g., when analyzing various studies in a combined analysis). This can be handled by PROAST as well. In that case you need to have a separate column in the data file, giving the detection limit for that particular row. When you answer Q6 by typing 0, the following question will arise:

Give column with individual detection limits

and here you can enter the number of the column containing the detection limits.

Note: observations below the detection limits will be plotted as being equal to half the detection limit. The detection limit is indicated by a horizontal dashed line in the plots. For the situation of various detection limits, two horizontal dashed lines are plotted, indicating the maximum and the minimum detection limit from the list.

Q7: Give small value to be added to the data >

Here you can enter a small value added to all response data.

Note: If you have zero values in the response data, the following question will automatically appear:

Q6-7: What do you want to do?

1: Give detection limit

2: Add nonzero value to observations

and by answering 1 or 2, you will be directed to Q6 or Q7, respectively.

Q9: Which variable do you want to consider for distinct plotting?

(type 0 if none)

As an example, when you have males and females in the data set, it may be useful to plot these separately, so that you can visually inspect if these data differ systematically. That will help in further deciding how to proceed with the analysis. Of course, you may also be interested in comparing different exposure durations, different compounds, etc.

Note: if you have indicated that a particular factor (e.g. sex) should be considered as a covariate in the analysis (see Q13-16), then distinct plotting will be done without answering this question.

This option is particularly useful in dose addition modeling, see section 6.

Q10: Give scaling factor for dose >

The dose (or any other independent variable selected in question 1) will be divided by this scaling factor. E.g. when you enter 1, the original values will be used. To prevent numerical problems in the analysis, it may help to scale the values for the dose such that they are not too far away from unity. For instance, when the largest dose were 5000 (mg/kg), you may enter the value 1000 at this question, so that the largest value in the analysis will be 5 (g/kg). Note: by using a factor of 1000 (and not say 100 or 10) you just change the units of the dose, which is convenient for reporting the results, etc. But otherwise, you can use any value here.

Q11: Give number of factor for which you want to select data

--- type 0 if you want to analyze all (remaining) data ---

Here you can indicate for which factor you want to restrict the data to be analyzed. Type the number of the column that you want to use for selecting data. When the column represents a factor (i.e. with nonnumerical values) all values will appear as a column on the screen preceded by a number. In that case, select the number(s) preceding the values you want to exclude. When the column represents a numerical variable (i.e. with numerical values only) these values will appear as such, and you need to enter the value(s) to be removed themselves.

For factors with more than two levels, you may want to select various levels. This can be done by the concatenate function `c()`, e.g. you may enter

```
> c(2,4)
```

and the levels associated with 2 and 4 will be selected both. See note in the discussion of Q4 in section 5.1.1

The question will be repeated until you enter a zero; in this way you can select factor levels for different factors (e.g. sex = males with exposure duration = 2).

Note: When you want to make a completely different selection of the data (including data that were no longer present), first refresh the original complete dataset by typing 0 after invoking this question, and then invoke this question again.

Q13: Give number of factor serving as covariate with respect to parameter a

--- type 0 if none ---

Here you may enter the number in the list related to the factor that you want to include as a covariate for parameter *a* in model fitting. In most of the models, parameter *a* represents the response level in the controls. So, when you enter the factor sex as a covariate here, two different values for *a* (one for males, one for females) will be estimated in fitting the model to the data of both sexes combined.

Q14: Give number of factor serving as covariate with respect to parameter b

--- type 0 if none ---

This question is analogous to question 13, but now different values for parameter *b* will be estimated for each level of the covariate entered here.

Q15: Give number of factor serving as covariate with respect to parameter var

--- type 0 if none ---

This question is analogous to question 13, but now different values for the variance parameter will be estimated for each level of the covariate entered here. The variance parameter (*var*) relates to the (natural) log of the observations.

Note: with respect to dose (more generally, the independent variable) *var* is assumed to be homogeneous (which is equivalent to a homogeneous Coefficient of Variation among doses). In most cases this is a reasonable assumption. It can be further checked by looking at residual plots (see section 5.1.2, option 5).

*Q16: Give number of factor serving as covariate with respect to parameter d
--- type 0 if none ---*

This question will normally not be used by risk assessors. It may be used to test if this parameter may be assumed equal in different subgroups. If so, the dose-responses are parallel on log-scale.

*Q17: Give number of factor serving as covariate with respect to parameter c
--- type 0 if none ---*

In specific applications it may be useful to attach the covariate to parameter c, e.g. to check if the assumption of constant c is reasonable.

Q18: Do you want pairwise comparison of dose groups against controls?

By answering yes, you will get the results from pairwise hypothesis testing of the dose groups against the controls. Standard t-tests are used (no compensation for familywise error rate), based on the residual MS of the associated one-way ANOVA (on log-scale, obviously).

The output includes the following columns:

N: the number of subjects in that group

GM: the geometric mean response in that group

ES-%: the effect size, as a percent change in the response compared to the controls

L05-ES: the lower confidence limit of ES (for 90% two-sided confidence)

L95-ES: the upper confidence limit of ES (for 90% two-sided confidence)

Note: When you want to compare within a subpopulation (e.g. only females), you need to select that subpopulation by question 11, before invoking this question.

Q19: Which transformation do you want?

Here you can choose between the options:

1: no transformation

2: sqrt-transformation

3: log-transformation

After choosing the transformation, PROAST uses the following codes in plotting results:

dtype = 1 : log-transformation

dtype = 25: no transformation

dtype = 26: square root transformation

Q21: Give value for right censoring >

This question is preceded by a list of the values for the response variable. You may enter a value here for upper censoring the responses. This is useful e.g. in time-to-tumor observations, where the study was ended at a fixed time, so that for the animals that were tumor-free at the end of the study it is only known that time-to-tumor was larger than that point in time.

Note: right censoring can be seen as a detection limit in the other tail of the distribution.

Q22: What type of regression do you want?

1: error in response only (default)

2: error in both response and dose (minimize sum of products)

Minimizing sum of product rather than sum of squares is more appropriate for data where errors occur in both x and y (e.g. in correlating NOAELs to BMDLs, or lengths against body weights).

Q23: How strict do you want the fitting conditions to be?

1: mild conditions

2: moderate conditions

3: strict conditions (R default)

In cases where a large number of parameters need to be estimated, or when confidence intervals need to be calculated for multiple subgroups, computing time may become lengthy. By choosing milder conditions in the optimization algorithm (the R function `nlminb`) computing time can be decreased. This might however result in less precise fits, depending on the situation.

Q24: Give column number of time variable

Here, enter the number of the column representing time (or, more generally, the second independent variable), if you want to fit with two independent variables (CxT). For continuous data this is model 48 in the list of models.

5.1.2 Main menu (continuous individual data, data type = 1)

The main menu contains a number of possible actions that can be done with PROAST. Users who are only interested in doing a BMD analysis (or, more generally, a regression analysis), will normally only use the first 6 options.

Options 7 and 8 relate to the possibility of calculating BMD confidence intervals using the bootstrap method.

Options 9 can be used to calculate the predicted response (according to the fitted model) at a given dose.

Option 10 relates to the specific problem of calculating withdrawal periods (waiting times) for veterinary medicines.

Option 11 calculates confidence intervals for any estimated parameter(s) in the model.

Option 12 may be used to end the session. (Note that the Escape button can also be used to end a session at any time).

Option 1: Change settings

By selecting this option you will get a list of the questions Q1-Q24, discussed in section 5.1.1. In this way, you can always change any of the settings during your analysis.

Option 2: Choose (another) model

This option may be used to specify a particular model that you want to fit to the data. Note, however, that from here, you can directly go to automatic model fitting, by choosing 38, 39, or 40 from this list (see section 3 for examples and description of that option).

Once you choose one of the (individual) models from this list, a dashed line should appear in the graphical window. Note that this line does not yet represent the best fit! It is only a first guess, based on some rough calculations. For obtaining the best fit, go to option 4 (*fit model*). This will normally be the next thing you want to do.

However, when the first guess (as indicated by the dashed line being far from the data, or not visible at all) is poor, it may be necessary to first go to option 3 (*choose other start values*). In general, you would only change the start values when directly fitting the model resulted in a problematic or suspicious outcome. This will however occur only incidentally.

The first five models of the list of models are five members of a *nested* family of exponential models. In general, two models are nested if one of the models can be deduced from the other by giving one of the parameters a specific value (see figure 1 below).

The five models have the following default constraints for the parameters:

$$\text{model 1: } y = a$$

$$\text{with } a > 0$$

model 2: $y = a \exp(b x)$	with $a > 0$
model 3: $y = a \exp(b x^d)$	with $a > 0, d > 0$
model 4: $y = a [c - (c - 1) \exp(-b x)]$	with $a > 0, b > 0, c > 0$
model 5: $y = a [c - (c - 1) \exp(-b x^d)]$	with $a > 0, b > 0, c > 0, d > 0$

Parameter a represents the background response (response at $x=0$).

Parameter b reflects the potency of the compound.

Parameter c is used for curves that level off. It is the maximum fold-change relative to the background response value ($= a$). When $c > 1$ the curve increases, when $c < 1$ the curve decreases.

Parameter d is a shape parameter. It is sometimes argued that the lower bound of parameter d should be 1, because values of d smaller than 1 imply that the slope of the curve at dose zero is infinite. This is often considered to be biologically implausible. However, when talking about biological plausibility, the log-dose scale is more appropriate. The curve of the model with $d < 1$ has slope zero when plotted against log-dose. Therefore, all positive values of parameter d are allowed. See also the updated EFSA guidance on the BMD (2016). Of course, the user can always change this constraint (see below), but setting $d > 1$ will result in a too optimistic BMDL, and is not justified.




Fig. 1. Graphical representation of the relationship between the nested models of the exponential family of models.

Models 6 up to 10 represent the same five models in another form. Here, the dose is not multiplied with a parameter b , but divided by it. Note however, that in models 8 and 10 the power d relates to (x/b) , while in models 3 and 5 it relates to x only.

Models 12 up to 15 constitute yet another representation of the same five models. Here, the models are reparameterized such that the CED (= BMD, see discussion of option 6 below) is one of the model's parameters (it replaces parameter b). This representation should be used if one wishes to assess a confidence interval for the CED by the likelihood profile method (see below). If you choose one of the models 12-15 you need to enter a value for the CES (= BMR) associated with the CED (see option 6, discussed below).

It may happen that numerical problems (e.g., no convergence, local optimum) disappear when using another representation of the same model. For instance, in cases with high dose values (left unscaled by Q10 in change settings), the second representation (divide dose by b) might behave numerically better than the first (multiply dose by b).

Model 6 is another representation of model 5, used for specific applications, not discussed here.

In model 11 no particular function for the dose-response relationship is assumed: the parameters of the model consist of the (geometric) means of the individual observations (plus the pooled within group variance), without assuming any specific dose-response relationship. Model 11 is also referred to as the “full” (or “saturated”) model, as it perfectly “fits” the data. This model may be used as a reference in the goodness-of-fit test, by statistically testing if the full model gives a significantly better fit than any other model. Or, as recommended by the EFSA updated guidance, the AIC associated with the full model may be used as a reference to the AICs of the fitted models, and lead to an alert if the AIC of the full model is substantially smaller than all others. However, one should always realize that experiments are never perfectly randomized or may contain outliers, while these statistical criteria assume that the data only contain random errors according to the assumed distribution.

Models 17-20 form an alternative nested family of models: the Hill family of models. These models, while having different mathematical expressions, are analogous to the exponential models 2-5, in the sense that they contain the same parameters, with the same interpretation.

Models 22-25 are the same Hill models as models 17-20 but now parameter b has been replaced by CED. Just as the exponential models 12-15, models 22-25 should be chosen if the user wishes to calculate the confidence interval for the CED by the profile likelihood method (see option 6).

The Hill family of models may be used in addition to the exponential family of models to evaluate “model uncertainty”. Model uncertainty refers to the phenomenon that different models providing an acceptable description of the data may result in dissimilar results (such as dissimilar BMD confidence intervals). This will happen if the data contain relatively little dose-response information. For data that contain relatively good dose-response information, the results from different models will be more similar. This type of uncertainty in the data is called model uncertainty and may be evaluated by comparing the results from different models with similarly good fits.

The other models will only be partly discussed here, since for describing dose-response data the exponential and Hill models work satisfactorily in general. However, in specific situations one might prefer one of these other models. For instance, in describing internal dose data as a function of external dose, one would need a model where $y=0$ when $x=0$ (which does not hold for the exponential and Hill models just discussed). In that case, you may choose model 26, 30 or 43.

Models 26 and 30 represent two parameterizations of the power model (this was done because covariates in PROAST can only relate to parameters a and b).

Models 21, 32, 33 and 45 are (nested) extensions of the exponential models. These models have the flexibility to be nonmonotone, or to have more than one point of inflexion.

Model 28 represents a model with a dose threshold as one of its parameters (in particular, a). This model is not recommended, as it is very unstable and results may be unreliable (see e.g., Slob 1999). The notation “ $(x > a)$ ” in the formule of model 28 should be read as a logical expression, being equal to 1 when true, and equal to 0 when not true.

Model 34 resrepresents a third degree polynomial. By fixing the relevant parameters at zero, it can also be used for fitting a second order polynomial, or a simple straight line. See discussion of option 4 below, for how to fix parameters.

Model 44 is a representation of model 2 such that the parameter estimates for a and b are uncorrelated, used for specific applications (in the package “milkings”).

Model 46 is equivalent to model 15, except that the different CEDs (related to different subgroups) are expressed as RPFs (relative potency factors) relative to the CED of one of the subgroups (e.g. the reference chemical). This model allows calculation of confidence intervals for the RPFs. NOTE: the RPF is defined such that values lower than unity realte to lower potency (higher BMD) than the reference.

Model 47 is equivalent to model 6, also used for specific applications (not dicussed here).

Model 48 describes the response as a function of two independent variables, such as dose and time.

Option 3: Choose other startvalues

Normally, you will choose option 4 after having appointed a model in option 2. However, it may happen that a problem arises in fitting the model, e.g. graphical inspection shows that the fitted model does not agree with the data very well. In that case, it may be helpful to change the start values as automatically produced by the software, and try to get better values. Or, it may be needed to explore whether or not different start values result in the same fit (if not, this indicates that there are local optima in the likelihood function). Changing start values may be done as follows.

First, answer the question

Which parameter(s) do you want to change ?

by entering the relevant number, then, after

Please, give new value for $a-f >$

enter another value for that parameter (in this example $a-f$ represents parameter a for females). By visual inspection of the plot you will learn if the change was an improvement or not. In addition, you may compare the old with the new log-likelihood value. In some situations the log-likelihood value may be reported to be “NA” (i.e. not available), or $-\text{Inf}$, or another nonnumerical value. This may happen when the curve is too far away from the data, or when one of the parameters is NA.

Note: as long as the log-likelihood is reported on your screen as some nonnumerical value, you need to adjust the start values; otherwise fitting the model (option 4) will (most likely) not succeed.

When you think the start values for the parameters result in a curve that is not too far away from that data, enter the number associated with *Please, CONTINUE*

Option 4: Fit model

After choosing option 4, the constraints (lower and upper bound) you get the question:

Do you want to change the parameter constraints ?

1: no, fit model

2: yes

If you answer by ‘yes’ you will see a list of the parameters to be estimated (ignoring possible covariates), with the current lower and upper constraint, i.e., the fit algorithm is not allowed to consider parameter values outside that range. You can now change the lower and/or upper constraints for each parameter.

Note that Inf ($-\text{Inf}$) stands for infinite (minus infinite).

You may also fix a parameter to a specific value, by setting the lower and upper bounds at the same value. For example, one may fit a straight line to the data by choosing model 34 (third order polynomial) and then fix parameters c and d to zero, by setting both their constraints at zero.

The output given in the command window consists of the estimated parameter values, the log-likelihood value, the AIC, and the CV (coefficient of variation) of the residual variation of the individual observations around the fitted model. Note that the estimate of the parameter var relates to the variance of the natural logarithms of the data. It is related to the CV by:

$$CV = \sqrt{\exp(var) - 1}$$

ATTENTION: In incidental cases the result from the fit algorithm depends on the start values used, i.e. when the model is refitted but with different start values, the best fit may differ from the previous one. In that situation, there is apparently some ambiguity on the best fit, due to relatively poor data (e.g. large noise, small effects, few dose groups)

and/or too many parameters in the model. This problem is indicated by the term “local optima”, i.e., the fit algorithm may end up in an optimum, while there are other (possibly better) optima. In many cases, local optima are discovered (and cured) when you calculate confidence intervals for the parameters based on the profile likelihood method (see option 6 or 12 of the main menu).

After the model has been fitted, you get the opportunity to store the fit results (or, actually, the current state of the analysis) in an R file:

give name for file to store results (or type 0 if none)

Storing the results may be useful for further analysis (including producing plots) in the future, especially when the calculations took some time. You can give it any name that facilitates its identification. For instance, you may use names like ‘tcdd.m5’ to indicate the name of the compound, and the model fitted.

When, at some point, you have stopped the current PROAST session you may invoke the stored results by typing

```
> f.proast( , tcdd.m5)
```

and continue the analysis at any later point in time.

Note: do not forget the comma before the name of the stored results, to indicate that it is the second argument of `f.proast`.

No convergence

It may happen that the fit algorithm reports “no convergence”. This can have various reasons. The most common ones are the following:

- The start values were too bad, i.e. the curve was too far away from the data at the start of the fit algorithm.
- The start values were too good. When you have, for instance, already fitted the model, and try to fit it again, PROAST will use the best estimates of the parameters as starting values. In such a situation it might happen that the fit algorithm creates only small changes in the parameters, without noting any changes in the fit criteria (the log-likelihood). Just change one of the start values a bit, and convergence will result.
- The data may not be informative enough to estimate (all) the parameters. For instance, there may be only two or three different response levels, while three or more regression parameters are estimated. Or, the noise within dose groups may be large compared to the (mean) differences in response. Non-convergence might occur when the number of regression parameters to be estimated is larger than the number of dose groups (or actually when it is larger than the number of clearly different response levels), or when the data leave ambiguity about the dose-response relationship.

Option 5: Plot results

This option may be used to change the plot of the data with the fitted model. The following question will be asked.

Which plot type do you want ?

- 1: y vs. x
- 2: y vs. $\log(x)$
- 3: $\log(y)$ vs. x
- 4: $\log(y)$ - $\log(x)$
- 5: y vs. \sqrt{x}
- 6: $\log(y)$ vs. \sqrt{x}
- 7: \sqrt{y} vs. x
- 8: \sqrt{y} vs. $\log(x)$
- 9: qq-plot of residuals
- 10: residuals against dose
- 11:
- 12: change plot limits
- 13: shift datapoints at same dose
- 14: use current plot for identifying points
- 15: back to main menu

Selection:

In general, option 4 is most appropriate, as biological measures (including dose) have a ratio scale, and with log-scales visual distances relate to relative differences (ratios, or percentage), which is the appropriate way of comparing values in biology. If you prefer to use original scales, e.g. for direct reading of the values, make sure you have at least had a look at the plot on log-scales. Further, keep in mind that the visual shape of a dose-response curve has no meaning, and any interpretation of the shape of a curve is spurious (Slob and Setzer, 2014).

After choosing one of the scale options (1-8) you will get, if you had different subgroups, the question:

Do you want color plot?

Here you may choose to make a black-and-white or a color plot.

Do you want to plot the group means?

You may choose to plot the group (geometric) means together with the individual data, or leave them out. Or, you may want to plot the means only (without the individual data). In the latter case, you will be asked if you want to plot the confidence intervals for the responses (recommended in most cases).

If you are not happy with the limits of the scales in the plot, choose option 12 (*change plot limits*). You can then choose other values for the limits of dose, or response. Note

that, when changing the lower bound of the x-axis, you also need to adjust the dummy value for dose zero in case dose is plotted on log-scale.

Note that, in all these plotting options, the same data, and same fitted model are plotted, just on other scales. So, the data are not transformed (which would imply another assumed statistical distribution), they are just plotted on other scales. Similarly, the plotted curves are completely equivalent, and just visualized in another way..

Option 9 produces a QQ-plot, i.e. a plot of the observed quantiles to the theoretical quantiles (according to the assumed distribution). When the data are a random sample from the assumed distribution, the QQ-plot should show up as a straight line. Note however, that the tails will always deviate from the theoretical straight line, as they are more sensitive to random sampling error.

After creating a QQ-plot you may get a warning that the data appear to contain outliers. These outliers are identified based on Grubb's test. You get the opportunity to remove the outliers. If so, you may refit the model, and compare the results.

When you choose one of plot options 9, 10, 11 the resulting plot has a second label on the x-axis, to inform you on the data-transformation used, by the following codes:

dtype = 1 : log-transformation

dtype = 25: no transformation

dtype = 26: square root transformation

Option 14 (active after replotting the data) gives the opportunity to identify individual data points: by clicking on them a label will appear next to it, giving the value of a particular factor associated with that point. The factor can be chosen by the user to be any of the columns available in the dataset. For instance, if the data contain identification numbers for the individual animals in the study, the user may tell PROAST to use that column in the dataset for labeling the data points that are appointed in the plot. This is a convenient way for identifying outliers. You can click on more than one point, if you want to identify various points. Click on the right mouse button and click on "Stop" if you are done.

Further plot options

If you have selected the option “*Back to main menu*” you will still get another two plotting options:

Further plot options:

1: *plot separate curves for each level of the covariate*

2: *produce normalized plots of all curves*

3: *continue*

The first of these is convenient for situations where you have analysed a large number of subgroups, since then it may be difficult to see the fits in each subgroup separately.

The second option may be used if you want to normalize the response to the background response in each subgroup. If only parameter *a* depends on the covariate, you will get a single curve here. If parameter *b* also differs among subgroups the dose-responses will be parallel on log-scale (i.e. the horizontal distance between the curves is the same for every value of *y*).

Option 6: Calculate CED point estimate or confidence interval

The Benchmark response in continuous data is defined as a percent change in mean response (relative to the controls). To distinguish this metric from the BMR in quantal data (usually defined as extra risk) it is often called CES (critical effect size¹) rather than BMR. More precisely, the CES and CED are defined by

$$CES = \frac{f(CED) - f(0)}{f(0)} = \frac{f(CED)}{f(0)} - 1$$

After the question

give value for CES = f(CED)/f(0) - 1 >

enter the desired value, where a 5% change should be entered as 0.05, meaning that the (predicted) response variable differs by 5% from that in the controls.

Note 1: in some applications much larger effect levels are relevant, e.g. the stimulation index of three in immunotoxicology. Here, a relevant change is a factor of 3 compared to the controls, and thus one should enter the value 2 here.

Note 2: if you enter the value 0 after this question, you will get the question:

give the value of BMR in terms of SD

*f(d)-f(0) = BMR * SD >*

¹ The word ‘critical’ in critical effect size does not need to be taken literally as the borderline between adverse and nonadverse. One could also define the CES as a value that may be regarded as a nonadverse effect size, without saying that larger values are necessarily adverse.

and when you enter the value 1 here, you can apply the BMR 1 SD approach. The resulting BMD is associated with a 1 SD change in means (on log-scale). If you also want to calculate a confidence interval for that BMD, you need to use the bootstrap method in PROAST (see below).

The calculated value for the CED is printed on your screen in original dose units, while in the legend to the plot it is printed in the same dose units as used in the plot (which may differ by the dose scaling factor, entered in Question 10 in section 5.1.1). The dose scaling factor is indicated in the legend of the plot as *scaling factor on x*.

The confidence interval for the CED may be calculated by the profile likelihood method (also called the likelihood ratio method), if one of models 12-15, or 22-25 (“in terms of CED”) had been chosen. This method is usually much faster than the bootstrap method (see option 7 below), while the results are similar (e.g. Moerbeek et al. 2003). In PROAST this method will be offered as soon as you have selected one of the models 12-15, or 22-25. So, if you have decided that model 3 is the most appropriate model for a particular dataset, you have to refit model 13 to the data. Then, by choosing option 6 you will get the question

do you want C.I. for CED?

After the calculations, you will see two plots. The lower panel shows the log-likelihood profile, where each dot represents an evaluated value, which relates to one of the curves in the upper panel. The evaluated points in the lower panel are described by a spline function fitted through them. This fitted function is cut off by a horizontal line, which depends on the confidence level (default in PROAST is two-sided 90%). The cut-off points on the CED-axis constitute the confidence interval for the CED.

One should always visually inspect the plot of the log-likelihood profiles for having irregular shapes. Under normal circumstances (i.e. the data contain sufficient information to determine the dose-response relationship), the log-likelihood profile should be a smoothly increasing and then decreasing curve. In problematic data sets (i.e. the data do not contain sufficient information to estimate the parameters of the model) the picture may be dramatically different, e.g. over a long range it remains constant and then suddenly drops. Such cases indicate problems with the data (the data may not be good enough to adequately estimate all parameters in the model).

Note: option 12 of the main menu allows you to calculate CIs for any of the parameters in the model, including the CEDs.

Option 7: Generate bootstrap runs based on fitted model

Option 7 from the main menu allows you to quantify the (statistical) uncertainty in the estimated CED using the bootstrap method (Slob 2001; Moerbeek et al. 2003), by producing an uncertainty distribution around the CED. This distribution would be needed

in a probabilistic hazard characterization (see e.g. Slob and Pieters, 1998; Chiu and Slob 2015), but it can also be used as an alternative method to option 6 (the likelihood profile method) for assessing confidence intervals around the CED.

The basic idea behind the bootstrap method is simple. To obtain the uncertainty distribution of an estimated CED a large number of artificial data sets is generated, each of them obtained by random sampling from a lognormal distribution with geometric mean (=median) defined by the fitted regression model, and (geometric) standard-deviation as estimated by the fitted model. Each artificial data set is based on the original experimental design, i.e. the same dose groups and number of subjects in each group. By sampling data sets we assume the error variance known, which ignores the estimation error involved. To account for the uncertainty in the estimated residual variance (on log-scale), it is, in each simulation run, multiplied by a correction factor $df / rchisq(df)$, where df equals the total number of data minus the number of regression parameters estimated, and $rchisq(df)$ denotes a random drawing from the chisquare distribution with df degrees of freedom (with a new drawing for each run).

Please, give the number of bootstrap runs you want to be done >

Enter the number of runs you wish to be done. It may be useful to first enter a small number (e.g. five), and then after the next question:

do you want plots of each sampling run?

1: no

2: yes (only for small number of runs!)

select *yes*. Each run will now be plotted, and after that you need to press a key to get the next plot (so that you can examine them, and check if the bootstrapping works properly, i.e. similar but different data and curves).

When you actually start the bootstrap analysis with a larger² number of runs (e.g. 500), you should not use the plot option, for obvious reasons.

At the end of the bootstrap calculations, you will be asked:

give name for file to store bootstrap results (or type 0 if none) >

Give a convenient name, e.g. TCDD.m5.bt, to indicate that model 5 was fitted to TCDD data, and bootstrap calculations have been done. These results may be used at any time after having ended the current PROAST session by typing

```
> f.proast( , TCDD.m5.bt)
```

and continue the analysis.

² Unfortunately, no general rule can be given for the number of bootstrap runs needed: it depends on the data. One might check the adequacy of the number of runs by repeating the bootstrap analysis with a larger number of runs. Or, if feasible, with the likelihood profile method in main menu option 6. See also Moerbeek et al. (2003).

Option 8: Calculate CED distribution for the animal

When you have performed a bootstrap analysis, option 8 allows you to compose a histogram of the bootstrap CEDs, as well as to estimate a confidence interval for the CED based on that histogram.

First, you will get a message how many of the bootstrap runs resulted in convergence (a large number of non-converged runs indicates that the data were poor relative to the number of parameters).

Next, the question

Do you want to change the Critical Effect Size?

current value: 0.05

1: no

2: yes

Selection:

gives you the opportunity to change the CES. If you answer by *yes*, simply enter the CES after the next question (see main menu option 6 for definition of CES).

Note: the value of CES cannot be changed here when you had fitted a model in terms of the CED (models 12-15, or 22-25).

PROAST will now generate a plot with the histogram of the CED, as well of that of the 10-log of the CED.

Note: this histogram should not be interpreted as variation among animals!
It reflects the scientific (statistical) uncertainty in the point estimate of the CED, due to the statistical limitations in the data. The distribution will be wider when based on fewer animals and/or larger noise in the data.

After

Press enter key to continue ... > You will see the plot of the data and fitted model, with the confidence bounds added in the legend.

Option 10: Calculate withdrawal period for veterinary medicines

Option 11 is a specific application of PROAST used for calculating withdrawal periods based on declining concentrations of veterinary medicines in e.g. meat, as a function of time. The aim is to estimate a point in time where the concentration in the matrix considered in 95% of the animals is below a given reference value, with 95% confidence. Not further discussed here.

Option 11: Calculate confidence interval for model parameter(s)

Option 12 may be used to calculate the confidence interval for any of the model parameters that were estimated in fitting the model, based on the likelihood profile method. After the question

Give number(s) of the (consecutive) parameter(s)

you may select a single parameter, or a set of parameters, using the concatenate function in R (see NOTE in discussion of Q4 in section 5.1.1).

5.2. Continuous summary data (data type = 10)

Dose-response modeling is possible for summarized dose-response data in terms of the group means if the associated standard deviations (or standard errors) and the group sizes are available (see Slob 2002 for statistical details). For this type of data the use of PROAST is largely the same as for data type = 1, except a few additional questions in the *Change settings* option of PROAST. These will be discussed in the next section.

5.2.0 Single model or set of models (continuous summary data)

See section 5.1.0, not different.

5.2.1 Data specification (*change settings*) for data type = 10

First, start PROAST, for instance with the example data set *das10*:

> **f.proast(das10)**

NOTE: if the data file *das10* is not found, you should first type:
> data(das10)

Then, the question

What type of response data do you want to consider?

should be answered by 7 (continuous, summary data). Most of the questions that will be given are the same as described in the previous paragraph for continuous individual data, but there are a few additional questions that need to be answered:

Q3a: Give number of the variation statistic associated to these endpoints.....

Here you enter the column that contains the standard deviations or standard errors.

Q3b: Do you have standard deviations or standard errors associated with the means?

1: standard deviatons

2: standard errors

It is important to give the right answer here, otherwise you will get spurious results.

Q3c: Give number(s) of the group size(s) associated to these endpoints

Enter the column with the sample sizes of each group

After the question to choose a model, you will get a plot of the data (the geometric means, with the associated confidence intervals) against the x-variable (e.g., dose). It should be noted that the original data (arithmetic means) are translated by PROAST into geometric

means. The geometric means will be lower than the original arithmetic means, while also the relative distances between the data points may change. The assumption of lognormality rather than normality of biological data implies that the geometric means give a more realistic picture than the arithmetic means.

5.2.2 Main menu (continuous summary data)

The main menu of PROAST similarly applies to data type 10 just as it does to data type 1, and the user is referred to section 5.1 for a description of the options in the main menu.

5.3. Quantal data

This section describes the use of PROAST for quantal data.

5.3.0 Single model or set of models (quantal data)

Just as in the case of continuous data, you can choose to fit a single or a set of models at the second question in the PROAST session:

Do you want to fit a set of models, or choose a single model?

1: single model

2: set of models

3: change settings first

The choice of options 1 and 2 are discussed at the end of the discussion of option 2 in section 5.3.2.

If you choose option 3, there are a number of settings that you may change, as discussed in the next section.

5.3.1 Data specification (*Change settings*) for quantal data

Most questions in the *Data Specification* part are the same or similar to those for continuous data, and the user is referred to section 5.1. 1. There are a few differences, however:

Q15: covariate for theta

--- type 0 if none ---

This question will only appear when you selected an LVM model. It relates to a specific application and will not be discussed here.

Q19: constraints on shape parameter >

Here you can change the (lower) constraint on the steepness parameter. Normally, no constraint is recommended by EFSA, but to compare results with BMDS calculations, which still uses this constraint as a default.

Q20: Give value for critical difference in AICs among models (default = 2) >

EFSA (2016) recommends a critical difference in AICs of 2 units. However, this value is not strict, and may be changed (increased), for instance, when none or only few of the models fitted would be accepted by the default AIC difference.

5.3.2 Main menu for quantal data

The options in the main menu that are available for quantal data are roughly similar to the ones for continuous data (see section 5.1.2), but there are a few differences, which will now be discussed.

Option 2: Choose (another) model

For quantal data two types of models are available in PROAST. The first type of models comprise the more usual models applied to quantal data, and therefore indicates as classical models. The latent variable models (explained below) have a better biological basis, but are still hardly used. So, after the question

Choose type of models

1: *classical models*

2: *latent variable models*

3:

4: *set of models*

you may select a single classical, or a single latent variable model, or choose the automated procedure of fitting a set of models.

Classical models

Note that some of the available classical models are nested to each other, e.g., model 2 is nested within model 3, and model 3 is nested within model 4. Thus, one may choose model 3, fix parameter c at 0 (see option 4 below), and effectively fit model 2.

Model 5 (the log-logistic model) and model 7 (the Hill model) are equivalent when parameter c is fixed at unity. When parameter c is left free (i.e. between a small positive number and 1), one may use model 7 for data that level off at a level below 100% response.

Most classical models include a parameter for the background response (usually denoted by a), reflecting the probability of response at dose zero. Theoretically, this probability can be very small, but not zero (since a is the fraction of responses in the infinite population). Therefore, its value should be estimated from the (dose-response) data. For data sets that have zero or very few observed positive responses at the lower doses, this should lead to an estimate of a that is very close to zero. If not, this indicates that the model does not adequately describe the observed dose-response relationship.

In principle, the user could fix parameter a to be zero in fitting the model (by setting the lower and upper constraint to zero), but that will lead to an error: the assumed binomial distribution is not defined for probability zero. Briefly, parameter a should always be estimated as a nonzero parameter in fitting the model to the complete dose-response dataset.

Most of the models in this list occur in two parameterizations: one with a parameter b and one where parameter b is replaced by the BMD. The latter models should be chosen

if you wish to calculate a confidence interval for the BMD using the profile-likelihood method. In the former models, a confidence interval for the BMD can only be calculated with the bootstrap method (which takes more time).

Latent variable models (LVMs)

When you choose *latent variable models*, you will get a list of models that is the same as the first part of the list of models used for continuous data. LVMs assume that the observed incidences originate from an underlying continuous response, which is not directly observed. Instead, each animal (experimental unit) is observed to have a response below or above a certain cut-off value, resulting in yes/no responses for all animals. This cut-off value is (normally) unknown, as it relates to the invisible latent variable. In fitting the model to the data, it can however be estimated (see e.g., Appel et al., 2001; Slob and Pieters, 1998; Slob, 1999). Thus, the models, defined for continuous response data, can be similarly used for quantal data. The LVMs similarly apply to ordinal data (see section 5.4). Fig. 4 in section 5.5 illustrates the idea of LVMs graphically for the case of ordinal data with three severity categories.

The parameters of a quantal LVM consist of the parameters of the underlying continuous model, plus on additional parameter, representing the cut-off between response and non-response, in PROAST denoted as *th*. Further, the parameter *var* in continuous models is now called *sigma* (interpreted as the square root of *var*). Not all these parameters can be estimated independently, due to perfect or very high correlations. In PROAST the parameter *th* is fixed at zero, and the value of *sigma* is fixed at 0.25. The user can change these fixed values, but that would not change the model fit.

A more extensive discussion about latent variable models is given in section 5.5, where they are applied to ordinal data. Note that quantal data are equivalent to ordinal data for the case of just one severity category. For instance, ordinal data with just two possible scores (no-response and mild response) reduce to quantal data.

When you choose any model in terms of CED, you will be able to calculate confidence intervals for the CED (BMD) by the likelihood profile method, just as in the case of continuous data (see option 6). In this case, you need to define the BMR associated with the CED/BMD in the model, by answering the question:

What type of Benchmark response do you want to consider?

- 1: ED50
- 2: Additional risk, i.e. $P[BMD] - P[0]$
- 3: Extra risk, i.e. $(P[BMD]-P[0])/(1-P[0])$
- 4: CED for latent variable
- 5: Relative risk

The default choice for risk assessors (EFSA guidance) would be 3 here: extra risk (with default value 0.1).

However, WHO-IPCS (2014) recommends to use the ED50 as the BMR, for quantal endpoints that can be considered to be “deterministic”. Therefore, option 1 may be

appropriate, in particular in the case of probabilistic risk assessment. The ED50 can be interpreted as the dose at which the average animal in the study changes from non-responding into responding. It can be regarded as the CED in the underlying continuous response with CES equal to the borderline between response and non-response (= *th*).

Option 4 and 5 relate to specific applications, not discussed here.

After selecting one of the models, a dashed line should appear in the graphical window. Note that this line does not yet represent the best fit! It is only a first guess, based on some rough calculations. To obtain the best fit, go to option 4 (*fit model*). However, when fitting the model leads to problems, it may be necessary to change the start values by hand (see option 3).

Latent variable model 32 is the one that you need for estimating an RPF (or various of them), as this model is LVM model 15 reparameterized with one CED (for the reference compound), while the CEDs for the other chemicals are expressed as RPF relative to the reference CED. In this way, the confidence intervals for the RPFs (and for the reference CED) will be calculated when selecting option 6 from the main menu.

Set of models

When you select “set of models” PROAST consecutively fits the exponential and Hill LVMs, and six classical models, as illustrated in sections 3.2 and 3.4. In the case of quantal data PROAST always applies model 3 as the LVM for the exponential and Hill models, as this model assumes the response to increase to 100% at sufficiently high doses (unlike model 5), just like the classical models that are applied here.

Option 3: Choose other startvalues

Similar to continuous data, see section 5.1.2

Option 4: Fit model

Similar to continuous data, see section 5.1.2

Option 5: Plot results

This option differs from that for continuous data in a number of respects.

The first question now is:

Do you want confidence intervals in plot?

Confidence intervals are useful in interpreting observed incidences, but you might wish to delete them in some cases to improve the visibility of the dose-response information in the plot.

After the question about color plots, you will get the list of plot options, which for quantal data is:

What plot do you want ?

- 1: *y vs. x*
- 2: *y vs. log(x)*
- 3: *log(y) vs. x*
- 4: *log(y) vs. log(x)*
- 5: *arcsin.sqrt(y) vs. x*
- 6: *arcsin.sqrt(y) vs. log(x)*
- 7: *latent variable vs. x*
- 8: *latent variable vs. log-x*
- 9: *log(latent variable) vs. x*
- 10: *log(latent variable) vs. log-x*
- 11: *qqplot for betadistr*
- 12: *change plot limits*
- 13: *back to main menu*

In this case, y relates to the response in terms of incidences.

Options 5 and 6 plot the incidence on a scale that makes the scatter homogenous (to correct for the fact that binomial variation increases when the incidence approaches 50%). This scale may be particularly useful when you have replicated observations at each dose (x).

Options 7 up to 10 only apply to LVMs and will not be visible when a classical model was fitted.

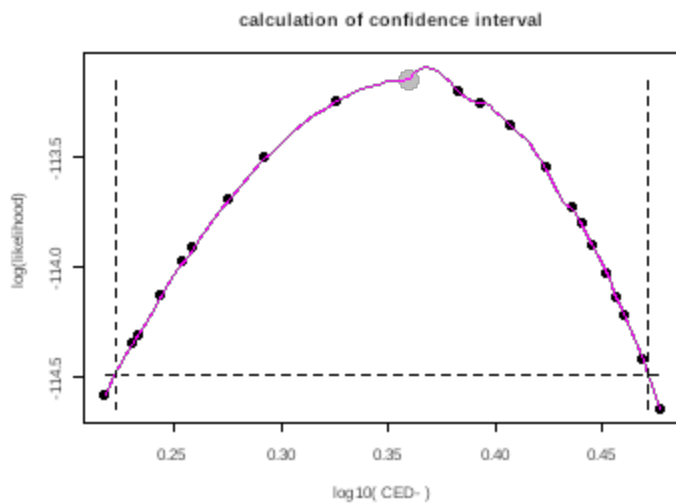
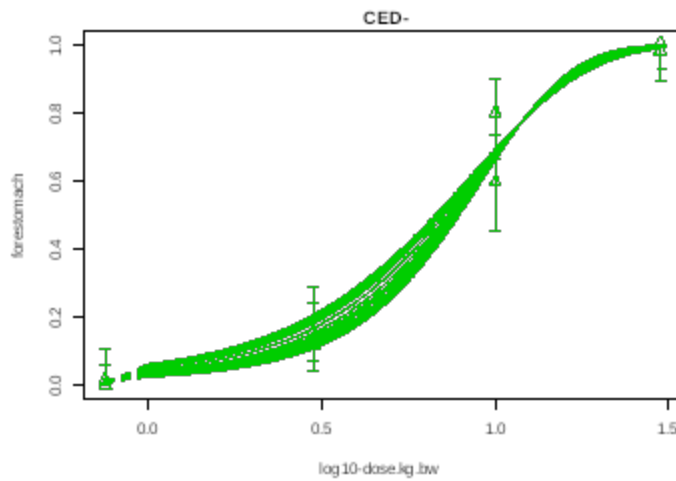
Option 11 will only be visible if you had quantal clustered data (see section 5.7). It will produce the QQ-plot for the observed frequencies (per litter) assuming a Beta distribution.

Option 6: Calculate BMD/CED: point estimate or confidence interval

This option allows you to calculate a point estimate of the BMD and/or a BMD confidence interval. Note that BMD and CED are both used in the PROAST software, in the case of quantal data.

The user is reminded that a point estimate of the BMD is not very meaningful, and a confidence interval is always required to evaluate in what range the true BMD might be, given the data available.

When you have selected a model that includes the BMD/CED as a model parameter, you will be asked if you want to calculate a confidence interval. In that case, you had already defined the BMR (see option 2 above). If you indicate to calculate the confidence interval(s) for the BMD(s) you will get a plot of the log-likelihood profile as shown in the figure below (lower panel), see discussion of option 6 for continuous data.



After having checked to log-likelihood profile, you can indicate if you want to include confidence intervals for the responses in the final dose-response plot. The legend of this plot now includes the confidence interval(s) just obtained.

Finally, you may store the results, by using the same name as after fitting the model (if applicable). All that has changed is the additional information on the BMD confidence interval(s).

If, on the other hand, you had selected a model where the potency parameter b is not replaced by the BMD, you will get the question:

What type of Benchmark response do you want to consider?

- 1: ED50
- 2: Additional risk, i.e. $P[\text{BMD}] - P[0]$
- 3: Extra risk, i.e. $(P[\text{BMD}] - P[0]) / (1 - P[0])$
- 4:

5:

The default choice for risk assessors (EFSA guidance) would be 3 here: extra risk (with default value 0.1).

However, WHO-IPCS (2014) recommends to use the ED50 as the BMR, for quantal endpoints that can be considered to be “deterministic”. Therefore, option 1 may be appropriate, in particular in the case of probabilistic risk assessment.

The ED50 can be interpreted as the dose at which the average animal in the study changes from non-responding into responding. It can be regarded as the CED in the underlying continuous response with CES equal to the borderline between response and non-response (= *th*).

Note: in the PROAST output you may encounter “CES = 0”, which indicates that the value for the CED relates to the ED50.

When you had used a dose scaling factor (Q10 in change settings) the calculated value for the BMD / CED is printed on your screen (R Console) in original dose units, while in the legend to the plot it is printed in the same dose units as used in the plot (which may differ by the dose scaling factor, if it is unequal to one). The factor you have used is shown in the legend of the plot (“scaling factor on x”).

The profile likelihood method is usually much faster than the bootstrap method (see option 7 below), while the results are similar (Moerbeek et al. 2003).

Option 7: Generate bootstrap distribution for CED.animal

Similar to continuous data, see section 5.1.2

Option 8: Calculate CED distribution for the animal

Similar to continuous data, see section 5.1.2

5.4 Binary data (data type = 2)

Binary data are yes/no responses for each individual subject. In principle, binary and quantal data are equivalent. For instance, the binary data

dose	response
0	0
0	1
0	0
0	0
10	1
10	1
10	0

10	1
----	---

are equivalent with the following quantal data:

f.proast(dose	response	group size
0	1	4
10	3	4

Therefore, when there are various animals (experimental units) in each dose group, binary data can be rendered into quantal data. However, it may occur that group sizes are small, or even equal to one. For instance, in epidemiological data each individual may be associated with a personal exposure estimate. Such data can be subjected to dose-response modeling in exactly the same way as quantal data (using the same models). The main difference is that the usual way of plotting the fitted model together with the observed incidences is not informative for visual inspection of the fit. As an illustration, consider Fig. 2. In this plot, the response is either zero or one, indicated by the small circles (which are artificially distorted to improve readability of individual points). According to the fitted model, the probability of response increases with dose, but this plot does not give visual information on the adequacy of the model for describing the data. Therefore, PROAST provides a special way of plotting for these data: both the data and the model are plotted cumulatively against dose. This is illustrated in Fig. 3.




Fig. 2. Curve fitted to binary data, with yes/no (0/1) observations per individual on y-axis. Both the zero and one observations are distorted by a small random error, to make overlapping points at the same dose visible.




Fig. 3. Same data and fitted model (left panel) as shown in Fig. 2, but now plotted cumulatively, illustrating that in this way the fit of the model can be made visual. The right panel shows the fitted null model (no dose-response), to illustrate that this model shows a poorer fit.

5.4.1 Data specification (binary data)

Similar to quantal data. There is however one additional option:

- Q12 is added here (*lump responses per dose*). It gives you the opportunity to render data, given as binary in the data file, into quantal. Of course, the analysis will produce exactly the same results for both ways of representing the data. However, the cumulative plots can no longer be produced, which is a disadvantage when group sizes are small. However, for larger group sizes this option may be useful, as it results in the usual plots for quantal data.

5.4.2 Main menu (binary data)

Similar to quantal data, see section 5.3.2.

5.5 Ordinal data (data type = 3)

Note that PROAST assumes that the severity categories are always scored as: 0, 1, 2, ... for increasing severities. Empty categories are not allowed. The numbering of the scores as used in PROAST is given after you have specified the data in the beginning of the PROAST session. If needed PROAST will translate the scores into a more appropriate format, which will be shown (as *temp.scores*) in the console window, e.g.:

```

orig.scores temp.scores
[1,] "0"      "0"
[2,] "1"      "1"
[3,] "2"      "2"
[4,] "4"      "3"
[5,] "5"      "4"
[6,] "6"      "5"

```

Theoretical background: latent variable models

Ordinal data indicate the severity of the response as ordered categories (e.g., none, minimal, moderate, severe). In PROAST, dose-response models for ordinal data always are “latent variable” models. These models assume that with increasing doses the severity of the response increases, but gradually, just like a continuous response. However, this continuous response cannot be quantified, and instead, it is measured by discrete severity classes. One may imagine that the underlying continuous response is cut off into severity classes, where the bounds between the classes are determined by the experimental observer (e.g. the histopathologist). In the case of histopathological data these bounds often relate to a “multidimensional” observation, i.e. various characteristics of the cell are taken into account. In other cases, the observed variable might be one-dimensional, but hard to measure quantitatively in each experimental unit (e.g. hypertrophy).

These bounds between severity categories thus cannot be determined from direct measurements and will be estimated in fitting the model, as unknown parameters. Note that their values usually cannot be interpreted in a biological sense, and their values are in that sense arbitrary. The only restriction is that the bounds should be strictly increasing or decreasing, since the severity scores are ordered as well.

The following plot illustrates the idea of a latent variable model for ordinal data, where the underlying dose-response is drawn as a straight line (for illustration purposes only). The boundaries between the severity scores cut off fractions of the assumed distribution, reflecting the probability of response of each category. These fractions are “fitted” to the observed fractions, when fitting the model.




Fig. 4. Schematic representation of a latent variable model (LVM) for ordinal data with three severity categories (including ‘no response’). The decreasing line represents the dose-response of the latent variable (where lower values are worse). The model is fitted to the data by matching the model-predicted frequencies (such as the 50% and 45% in the figure) to the observed frequencies in the dataset.

Latent variable models are also applied in the EPA software CATREG (which stands for categorical regression). However, in CATREG, the focus is on the assumed distribution, while the underlying continuous dose-response is assumed to be linear. Thus the various models used in CATREG vary in the assumed distribution (e.g., lognormal for the log-probit model, log-logistic for the logit model).

In PROAST the focus is on the underlying continuous dose-response relationship, while the distribution is always assumed to be lognormal. Further, PROAST defines the BMDs as the ED50s for each severity category, so they indicate the doses where the average animal changes from one severity category to another.

The latent variable models contain the regression parameters as given in the dose-response model, a number of parameters ($th1$, $th2$, $th3$, ...) representing the boundaries between the severity categories, and a parameter (σ) representing the scatter in the data:

		constraints
Regression parameters	a b, \dots	lower = 1; upper depends on number of categories See continuous data, assuming decreasing function
First boundary parameter, defined for ln of the latent variable	th1	Fixed at 0 (= 1 on back-transformed latent variable scale)
Further boundary parameters, defined as a difference with th1 (on ln-scale)	th2, th3,	< 0
SD parameter, relates to the normal distribution of the ln of the latent variable	sigma	Fixed at 0.25

Parameters $th1$ and $sigma$ are fixed (by default) to prevent overparameterization of the model (some of the parameters show very high or even perfect correlation).

To ensure that the boundaries between the categories form a strictly decreasing series of values they are defined as the difference with parameter $th1$, using the constraint that all $th2, th3, \dots$ are negative. As a result, the cut-off values between categories are forced to have decreasing order (in fitting the model to real data, this might not be the case without this constraint).

Note that in binary or quantal data, there is only one single severity category, and hence the latent variable models for these data only contain boundary parameter $th1$.

5.5.0 Single model or set of models (ordinal data)

Just as for the other data types you have the choice between fitting a single model or a set of models. In both these options, the general approach is similar to that for continuous data. However, if you choose set of models in the case of ordinal data, model 3 for both the exponential and Hill model families will be used for calculating the BMD confidence intervals. After the analysis, the Console window will show the confidence intervals for the BMD (ED50) each severity category (for exponential and Hill model, consecutively).

When you have chosen to fit a single model, and have specified the dose and response, a plot is produced showing the severity scores on the y-axis. The numbers in the plot represent the number of data points with that severity category (at that dose).

5.5.1 Data specification (*change settings*) for ordinal data

The questions for changing the settings are similar to those for continuous data, see section 5.3.1. However, there is an additional option:

Q3: *give severity category for dichotimization*

1: 1

2: 2

3: 3

4: 4

5: 5

This question allows you to dichotomize the ordinal data into quantal data, with respect to the severity category indicated here. Of course, this is not a good idea in general, as it reduces the information in the data, but it may in some cases be helpful in comparing results.

5.5.2 Main menu (ordinal data)

Option 2 Choose (another) model

You will see a list of 32 models, which is a subset from the list of models for continuous data (see section 5.1.2, option 2 for a discussion of these models).

After having selected a (single) model, you will see a plot, representing the model for the start values of the parameters. The plot shows the observed and predicted data both plotted cumulatively. See binary data (section 5.4) for an explanation. In ordinal data, these cumulative plots are shown for each severity category (indicated by the numbers in the plot). This plot gives an idea of the adequacy of the start values, and may offer some help in finding better ones, if needed.

Just as for the other data types, if you want to calculate confidence intervals for the CEDs (defined as the ED50s for each category) using the likelihood profile method, you need to choose a latent variable model in terms of CED. If you have selected a model where the CED is one of the parameters (models 12-15 or 22-25), you will see the question

give adverse category for CED >

The BMD confidence interval will only be calculated for the BMD associated to a single severity category, and here you can enter the one you find most relevant.

Option 4: Fit model

After fitting the model, you will see a plot as in Fig. 5, where both the observations and the fitted model are plotted cumulatively, for each separate severity category (just as in Fig. 3, where this is done for one single severity category, in binary data).

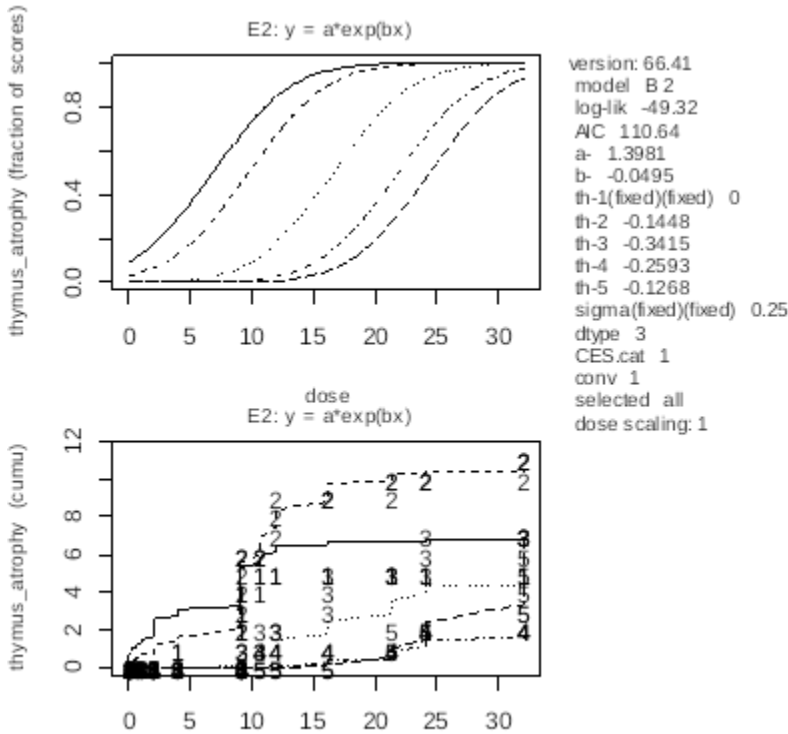


Fig. 5. Graphical output after fitting a model to ordinal data. Upper panel: Cumulative plot of observations and fitted model, for each severity category. The cumulative observations are indicated by the relevant severity category. Data used: das3 in the PROAST package.

Option 6 Calculate BMD/CED: point estimate or confidence interval

In PROAST, the BMD for ordinal data is defined as the ED50 associated with a particular severity score, i.e. the severity score is considered as the BMR. In this way, each BMD may be regarded as the dose where the *average* animal would respond exactly on the borderline between the two consecutive categories. For instance, the ED50 associated with category 2 is defined as the dose where the average animal responds on the borderline between categories 1 and 2. This dose is analogous to the CED in continuous data, where it represents the dose at which the average animal’s response is equal to the CES. In ordinal data the CES is defined as the transient from one severity category to the next.

Thus, when selecting option 6 of the main menu, the resulting BMDs are the ED50s for each of the severity categories.

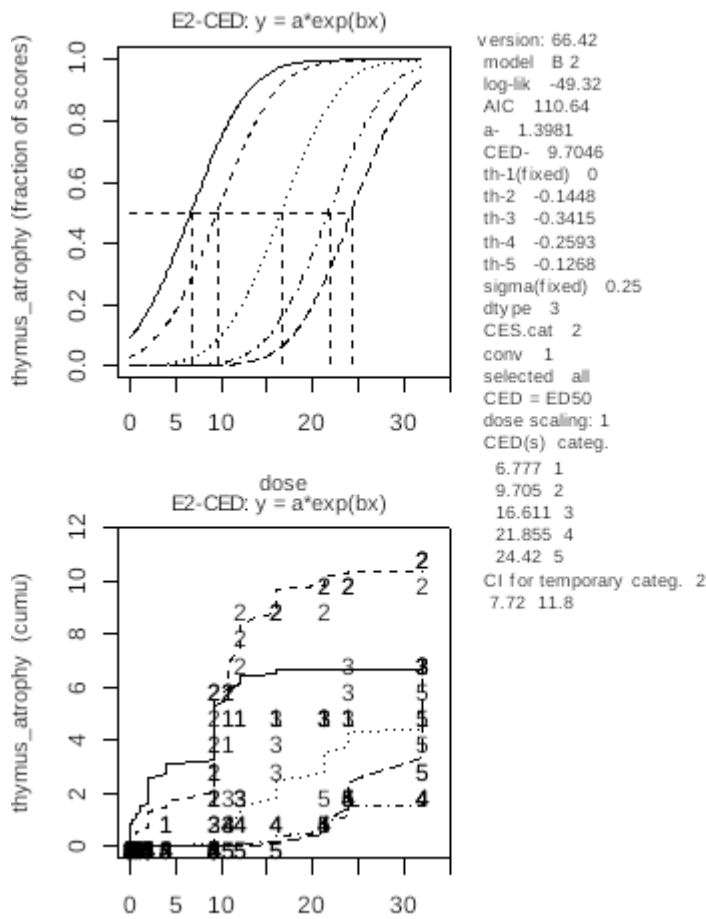
In the PROAST output, CES=0 means that the critical effect size is defined as transients between severity categories, and that the associated CEDs are ED50s.

Which dose scale do you want for plotting?

- 1: original scale
- 2: log-scale

This question enables you to change the dose-scale in the plots that will come next. These plots contain two different plots:

Upper panel: the probability of response for each category. To be more precise, each curve represents the probability of that category, or any lower category. In this way, the curves behave as quantal dose-response curves, i.e. between 0% and 100%.
 Lower panel: the cumulative data and cumulative model, for each of the severity categories.



5.6 Clustered continuous data, individual or clustered (dtype = 5, 15)

PROAST allows for a dose-response analysis where the individual observations per dose group are not independent, such as in developmental studies where the observations related to the embryos or the pups will show correlations within litters (litter effects).

PROAST will fit the model, as if there are no litter effects, and after that estimate the within and between litter (or other subgroup) variances (from the regression residuals in the case of individual data, or from the group SDs/SEMs in the case of summary data). Both variance components are reported in the Console window, followed by a P-value, which indicates if the litter effect is statistically significant. If the variance between litters is much smaller than within litters, the litter effect will have little impact on the BMD confidence interval and might be ignored. You could also compare the log-likelihoods associated with the model fits with or without litter effects (i.e. as dtype = 1 vs. dtype = 5) : if they differ by more than 1.92 units, the litter effect is significant at $P = 0.05$ (likelihood ratio test).

For clustered continuous data, the confidence interval for the CED cannot be calculated by the profile likelihood method in the current PROAST version, and instead is assessed by the bootstrap method (option 7 from main menu). This procedure takes the litter effects (i.e., the between and within variance components) into account.

If you choose the option “set of models” you are asked to enter the number of bootstrap runs that you want to do, in advance of the analysis.

If you have indicated that you have clustered continuous data, then PROAST will ask:

*Give column number of the nested factor
(type 0 to skip nested analysis)*

where you may indicate the column in the data matrix relating to the factor defining the clustering in the data (e.g. the numbers of the dams).

Note that with clustered continuous summary data there is no need to have a column defining the clusters, as they are already defined by the summary statistics (just like in clustered quantal data, see next section).

5.7 Clustered quantal data (data type = 6)

In PROAST clustering in quantal data is taken into account as follows. It is assumed that the response at each dose (= value of x) has a betabinomial rather than binomial distribution. The betabinomial distribution results from assuming a Beta distribution for the expected response for a particular cluster (e.g. litter) and a binomial distribution for the experimental units (e.g. embryos) within the cluster (litter). In this way all models available for quantal models, as discussed in section 5.4, can be similarly applied to quantal data with clustering.

The difference between fitting a model as if the data were not clustered is just one additional parameter, called *alfa* in PROAST. This parameter *alfa* is one of the parameters of the Beta distribution. In PROAST *alfa* is defined such that the lower its value the wider the Beta distribution (higher variation).

The addition of the parameter *alfa* can be regarded as a nested extension of the dose-response model. Therefore, the log-likelihood value associated with the best fit of the betabinomial dose-response model (with *alfa*) can be compared to the equivalent binomial model (without *alfa*) in the usual way (repeat the same analysis, but select data type *quantal data*, option 4 from the first PROAST question).

The plot produced after fitting a model shows the observed incidences in the single litters (clusters) as well as a larger marks, indicating the group means (see Fig. 6 for an illustration). These group means are calculated (automatically) by fitting the full model to the data, with *alfa* as an additional parameter.

In option 5 of the main menu there is an additional option (11) to plot the QQ-plot of the observed frequencies in each litter, assuming a Beta distribution.




Fig. 6. Example of a fitted model to clustered quantal data. Here, the larger marks denote the group mean incidences, the smaller mark the incidences in the individual litters.

6. SPECIFIC APPLICATIONS

6.1 Dose addition

Data from mixture studies can be analysed in PROAST, based on dose addition. The complete dataset is analysed (including both the single doses and the mixtures). PROAST estimates the RPF (or, in case on more than two chemicals in the studies, multiple RPFs), using the complete dataset, and then predicts the response for all applied doses or combinations of doses. The responses are plotted against the dose that results from dose addition (using the estimated RPF), resulting in a single curve. The plots show different marks indicating which responses relate to single doses (of a particular compound) and which to the mixtures. Note that the fitted curve is based on the assumption that dose addition holds, in which case the plotted responses of each type should randomly scatter around the fitted curve.

With dose addition, it is possible to fit a single selected model or perform an automated sequence of various model fits, just as in the usual situation of one compound with administered doses.

Data format

Datasets from mixture studies should include a column representing the dose for each individual chemical, be it applied as a single dose or as a mixture. Note that the first three letters in the headings of the dose columns will be used in the PROAST output, so make sure that they are discriminative. An additional column indicates which single chemical was applied, or that a mixture was applied. In this example, aaa and bbb represent two chemicals.

Example of dataset for a mixture experiment in the PROAST format.

aaa-dose	bbb-dose	Dosing.type	Resp1	Resp2
0	0	aaa	8029000	337
0	0	aaa	6675200	382
0	0	aaa	6144133	365
0	0	aaa	7710500	532
0.01	0	aaa	6905033	451
0.01	0	aaa	11172000	532
0.01	0	aaa	6534500	496
0.01	0	aaa	9646000	408
0	0.02	bbb	9789500	409
0	0.02	bbb	8564500	529
0	0.02	bbb	7042000	534

0	0.02	bbb	5351500	474
0.02	0.02	mix	8092000	714
0.02	0.02	mix	5680500	562
0.02	0.02	mix	7045500	800
0.02	0.02	mix	9464000	616
0.04	0.04	mix	7154000	950
0.04	0.04	mix	8876000	756

Applying dose addition analysis in PROAST

The analysis of a mixture study proceeds just as a single chemical study, except that you need to answer Q1

Q1: Which variable do you want to consider as independent variable?

(e.g. dose, age)

by a multiple answer, by indicating the various columns of the chemicals involved. Since the answer needs to be understood by PROAST as a single answer, you need to use the concatenate function here (see NOTE in discussion of Q4 in section 5.1.1). For example, when the first two columns of the datasheet represent the doses of two compounds, you need to provide the answer c(1,2).

Furthermore, after Q1, PROAST will ask which column represents the labels for single doses and mixtures:

Q9: Which column defines the single doses and mixtures? (type 0 if none)

and here you need to enter the appropriate column number (called dosing type in the above data example).

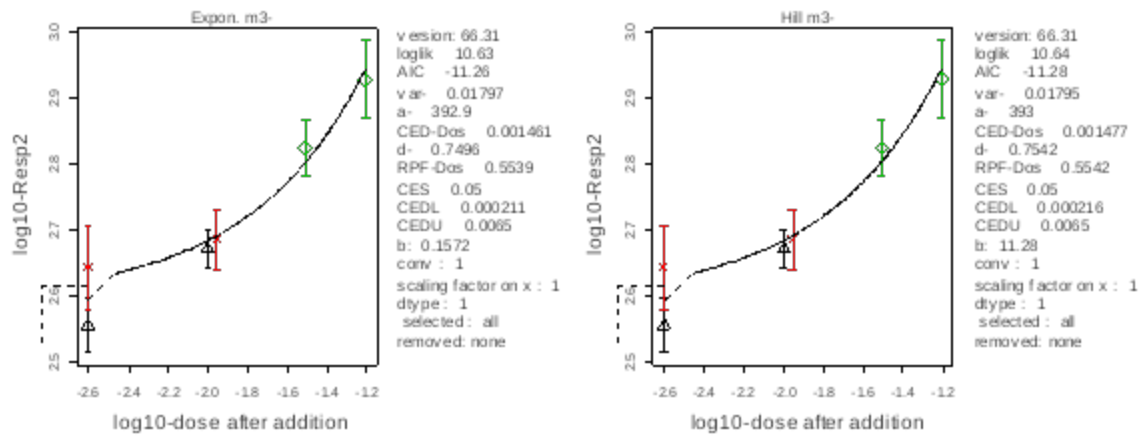
The next question is:

Do you want to fit dose addition model, or parallel curves model

1: dose addition model

2: parallel curves model

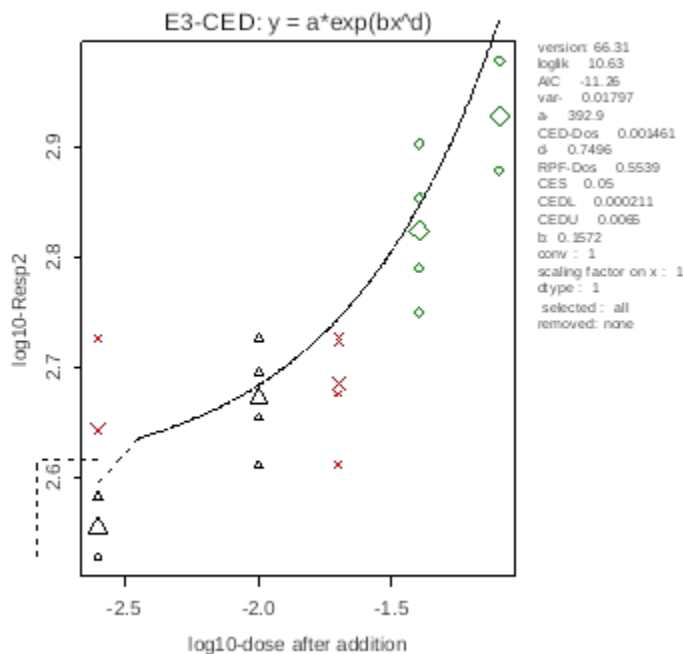
If your data also contains mixtures, you need to select option 1. If your data do not contain mixtures (e.g. you have removed them in *Change settings*), then you could also select option 2. In that case, the model will be fitted with compound as a covariate. You may also want to choose model 46 (which is exponential model 5, with one CED for the reference compound, and RPFs for the other compounds). If you fit this model, you get an estimate for the RPFs of the other compounds. By selecting option 6 from the main menu, you get the confidence intervals of the CED (of the reference) and the RPFs (of the other compounds).



After fitting the model(s) you get a plot that uses different marks (and colors) for the responses associated with the single chemicals, and those associated with the mixtures (i.e. the column indicated after question Q9). The Console window provides information what the symbols represent, e.g.

The colors in the plot relate to the following subgroups:

color	mark	subgroup
1 black	upward triangle	Chem1
2 red	cross	Chem2
3 green	diamond	mix



If you are not sure which colors/marks belong to which levels of dosing.type, use option 14 of the plot menu (*use current plot for identifying points*), to make that clear (see section 5.1.2, discussion of option 5).

Furthermore, if you did an automated run of various models, you will see in the Console window the confidence intervals for the CED of the reference compound and the RPF for the other compound(s), per model that was fitted (and accepted according to a maximum difference of 5 units with the model with minimal AIC, in the case of quantal data).

6.2 Two-step analysis of data with two independent variables

When you have data with two independent variables, such as time and dose, the analysis can be done in two steps. As an example, consider body weight data that were measured over time (in the same animals), while the animals were distributed over different dose groups.

The dataset should now have an additional column with time of observation given in each row, as well as an ID for the relevant animal.

As a first step in the analysis choose time as the independent variable (instead of dose). Then, you may choose the automatic sequence of modeling (option 2 or 5 after second PROAST question), indicating ID of the animal as the covariate. You may omit the Hill model here. Save the results of the calculations in a file in the R workspace with a suitable name, say “name.fit”.

When the resulting model shows different curves for each animal, with the CED depending on animal, you can now examine the dose-response of the CED with dose by typing:

```
> f.subseq.anal(name.fit).
```

Here, indicate that you want to consider dose as the independent variable, and that you used “ID” as the covariate in the first step, and finally indicate if there is a covariate that you want to consider in this second step of the analysis (e.g. sex). After that, the PROAST analysis proceeds as usual, starting with Q1:

*Which variable do you want to consider as independent variable?
(e.g. dose, age)*

Proceed as usual, and you will get an analysis of the CED as a function of dose.

6.3 CxT models

These models are still under development, both for quantal and continuous data. For quantal data the (probit) model is similar to the Ten Berge models. It also allows for inclusion of a covariate.

Note that the analysis discussed in section 6.2 is an alternative analysis for CxT data, and in a sense more transparent, in particular if you want to include a covariate in the analysis.

7. THE R OBJECT WITH THE SAVED RESULTS

In R pieces of information are called “objects”. Objects include datasets and functions. A function is an object that is able to perform some action. For example, the function ‘ls’ prints a list of the objects that are present in a particular directory. In general, to execute an action associated with a function, one has to type the function name together with brackets. E.g., to actually see the objects in a directory type

```
> ls()
```

Typing ls without brackets will show you the content of the function itself, which is not of interest for the user.

Objects

In giving names to an object (e.g. for saving results) you may use all alphanumeric symbols and ‘.’ (the dot). The underscore ‘_’ is also allowed. Avoid other symbols, for instance ‘-’ is read as minus. Spaces are not allowed as well.

Further analysis of PROAST results

At various places you are asked if you want to save the results at that point. If you answered by giving a name, PROAST creates an object with the results and other relevant information. For instance, suppose you have saved results from a bootstrap analysis (option 7 of the main menu) of the data object “das1” in the object “das1.bt”. You can invoke these results in another PROAST session by typing

```
> f.proast( , das1.bt)
```

and you will arrive in the PROAST main menu, with all results stored in das1.bt now known to PROAST. From here on you can proceed the analysis (e.g. choose another model) or reproduce/change the plot, etc.

The saved object contains a list of other objects, the names of which can be made visible by typing

```
> names(das1.bt)
```

or, to get them in alphabetical order:

```
> sort(names(das1.bt))
```

You will now see the names of the elements in the list. You can make the content of a particular element visible by pasting its name after the name of the object, separated by the \$ sign. For example, typing

```
> das1.bt$par.boot
```

gives you the values of the model parameters estimated in each bootstrap run.

Some examples of other elements in the saved object are:

\$BMD	the estimate (MLE) of the BMD (in the case of quantal data)
\$BMR	the BMR that was used (in the case of quantal data)
\$CED	the estimate (MLE) of the CED
\$CES	the BMR that was used
\$CED.boot	the estimate (MLE) of the CED in each bootstrap run
\$cond.ans:	indicates the optimization criterion used in fitting the model. Normally, this value is 3, but if you relaxed the criterion (question Q23 of change settings) the value may be 2 or 1.
\$data.0	The dataset that was analysed in the current analysis. This may be a subset of the original dataset (= \$odt)
\$detlim	The value of the detection limit that was used
\$date	The data and time that the analysis was performed
\$gui	If TRUE the analysis was done by the GUI version, otherwise by the MENU version
\$loglik	The log-likelihood value of the fitted model
\$MA\$confint.ma	Confidence interval based on model averaging (quantal data)
\$MLE	The estimates (MLEs) of the model parameters
\$model.ans	The number of the fitted model, according to PROAST numbering
\$modelname	The fitted model
\$odt	The original (complete) dataset, used when calling f.proast()
\$odt.name	The name of the original (complete) dataset, used when calling f.proast()
\$PRversion	Version number of PROAST that was used
\$regr.resid	The regression residuals of the fitted model
\$scores.mtr	Ordinal data only: The original scores as used in the dataset, and the scores as used in the analysis
\$text.par	The names of the model parameters that were estimated

When you saved the results after an automatic run of a set of models, the results for each model are stored in a sublist, with names:

EXP, HILL

in the case of continuous or ordinal data

EXP, HILL, log.logist, log.prob, two.stage, Weibull, gamma, logit, probit.

in the case of quantal data

For example, when `das1` was analysed by set of models, and stored in `das1.som` (`som = set of models`) then

```
> names(das1.som$EXP)
```

shows the elements with fit results related to the exponential model. For example,

```
> das1.som$EXP$modelname
```

gives the selected model from the exponential models, and

```
> das1.som$EXP$loglik
```

gives the associated log-likelihood. Note that

```
> das1.som$logliks
```

(note the `s` at the end) gives the list of log-likelihoods for all models that were considered in that family of models.

Of course, you can apply all R functions to the elements in the saved object. For example, you can make a histogram of the bootstrap CEDs by typing

```
> hist(das1.bt$CED.boot)
```

or calculate a specific quantile, say the 10th, by typing

```
> quantile(das1.bt$CED.boot, 0.10)
```

Exporting particular results for analysis outside R can be done as follows. First type

```
> sink('name.txt')
```

which means that all output that is normally written on your screen, will be written in a text file called `'name.txt'` (note the quotation marks in using `sink`). Then type the name of the object you want to be written in this text file, e.g.,

```
> das1.bt$CED.boot
```

and finally type

```
> sink()
```

to restore the situation that all output is written on your screen again. You will find the file (`name.txt`) in your current working directory, containing the bootstrap results for the CED.

8. THE PLOT LEGENDS

The legends on the right side of the PROAST plots give a summary of the essential information related to the plotted results.

An example for continuous data is given in Fig. 7. In this example LDH is analysed as a function of dose, using sex as a covariate (data from das1).




Fig. 7. Example of PROAST plot with legend, for continuous data.

The header of the plot indicates which model was fitted. The meaning of the terms used in the legend on the right is as follows.

version	The PROAST version that was used for this analysis
loglik	The log-likelihood for the fitted model shown
a-f and a-m	The estimated values (MLEs) of parameter a for females and parameter a for males. Here, f and m are the two levels of the factor (i.e. sex in this case) as they occurred in the original dataset
CED	Estiamted value of the CED (MLE)
d	Estiamted value of parameter d
CES	CES used in the analysis
CEDL and CEDU	Lower and upper bound of the CED, obtained by the profile likelihood method. If they were derived by bootstrapping, they are indicated as CEDL.bt and CEDU.bt
b	Value of parameter b. Note that b is in this case not a parameter in the model, but was re-calculated based on the estimated parameters just mentioned.
conv	if 1: converged was reached, if 0: no convergence
scaling factor on x	The value entered after Q10 (change settings), i.e. the factor by which the doses were divided before analysis.
dtype	Code for data type, see table 1
selected:	Column name (here: food) and the levels (here: levels 1 and 2 of factor food) that were selected
removed:	Column name (here: LDH) and the levels that were removed (here: value 6080 of response LDH)
fact1	The factor that was used as a covariate relative to parameter a
fact2	The factor that was used as a covariate relative to parameter b
fact3	The factor that was used as a covariate relative to parameter var
fact4	The factor that was used as a covariate relative to parameter c
fact5	The factor that was used as a covariate relative to parameter d

An example for quantal data is given in Fig. 8.

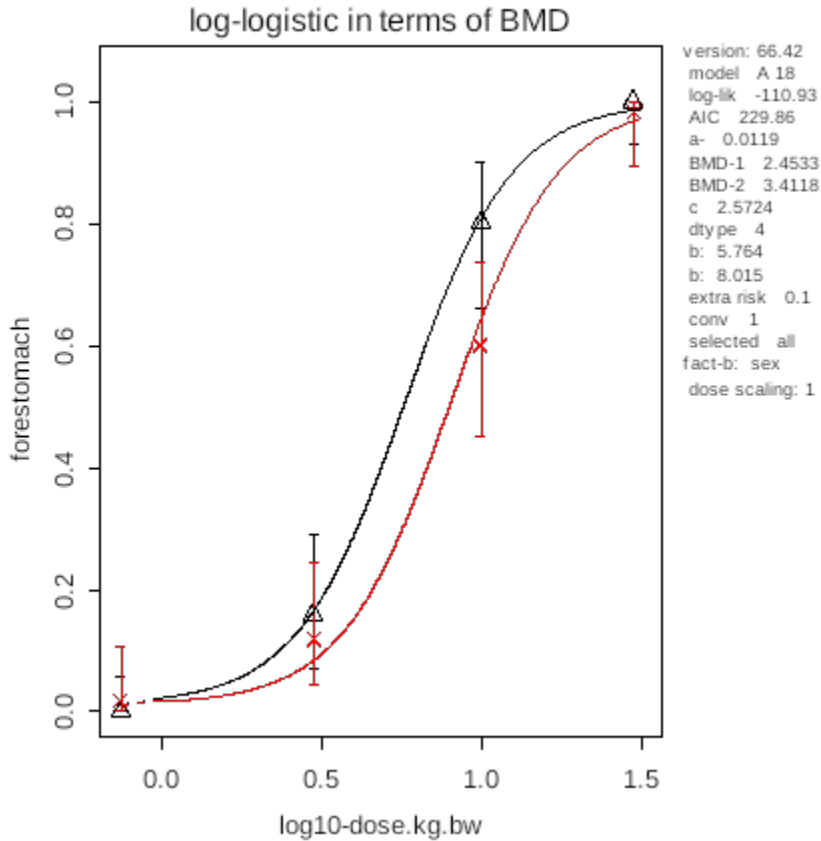


Fig. 8. Example of PROAST plot with legend, for quantal data.

The meaning of the terms is similar to those in the legends for continuous data, but there are a few differences:

model	Coded by a letter A (=classical model) or B (=LVM), followed by the model number in the list of models.
ces.ans	Indicates the BMR type: = 1: ED50 = 2: additional risk = 3: extra risk
CES	Value of the BMR. If the BMR relates to the ED50, CES in the legend is 0.
th	Value of the parameter theta (only if LVM was fitted)
sigma	Value of the parameter sigma (only if LVM was fitted)

REFERENCES

- Appel, M.J., Bouman, H.G.M., Pieters, M.N., Slob, W. (2001). Evaluation of the applicability of the Benchmark approach to existing toxicological data. RIVM/TNO Report No. 601930 001.
- Chiu, W.A. and Slob, W. (2015). A Unified Probabilistic Framework for Dose-Response Assessment of Human Health Effects. *Environmental health perspectives*, 123(12), 1241-1251.
- EFSA (2009). Guidance of the Scientific Committee on a request from EFSA on the use of the benchmark dose approach in risk assessment. *The EFSA Journal* (2009) 1150, 1-72
- ISO TC 147. Water Quality – Guidance document on the statistical analysis of ecotoxicity data. April 2003.
- Moerbeek, M., Piersma, A.H., Slob, W. (2003). A comparison of three methods for calculating confidence intervals for the benchmark dose. *Risk Anal* 24:31-40.
- Slob, W. (1994). Uncertainty analysis in multiplicative models. *Risk Analysis* 14: 571-576
- Slob, W. and Pieters, M.N. (1998). A probabilistic approach for deriving acceptable human intake limits and human health risks from toxicological studies: general framework. *Risk Analysis* 18: 787-798.
- Slob, W. (1999). Thresholds in Toxicology and Risk Assessment. *Internat. J. Toxicol.* 18:259-268
- Slob, W. (2002). Dose-response modeling of continuous endpoints. *Toxicol. Sci.* 66, 298-312.
- Slob, W., Moerbeek, M., Rauniomaa E., Piersma, A.H. (2005). A statistical evaluation of toxicity study designs for the estimation of the Benchmark dose in continuous endpoints. *Toxicol. Sci.*, 84: 167-185
- Slob, W., Setzer, R.W. (2014). Shape and steepness of toxicological dose-response relationships of continuous endpoints. *Crit Rev Toxicol* 44, 270-297.
- WHO-IPCS (2014). Guidance document on evaluating and expressing uncertainty in hazard characterization. WHO-IPCS harmonization project document no. 11. ISBN 978 92 4 150761 5.
http://www.who.int/ipcs/methods/harmonization/areas/hazard_assessment/en/