

A short introduction to Biogeme

Michel Bierlaire

April 20, 2023

Report TRANSP-OR 230620
Transport and Mobility Laboratory
School of Architecture, Civil and Environmental Engineering
Ecole Polytechnique Fédérale de Lausanne
`transp-or.epfl.ch`

SERIES ON BIOGEME

This is an updated version of Bierlaire (2020), adapted for Biogeme 3.2.11.

The package Biogeme (`biogeme.epfl.ch`) is designed to estimate the parameters of various models using maximum likelihood estimation. It is particularly designed for discrete choice models. In this document, we present step by step how to specify a simple model, estimate its parameters and interpret the output of the package. We assume that the reader is already familiar with discrete choice models (Ben-Akiva and Lerman, 1985), and has successfully installed Biogeme.

Biogeme is a Python package written in Python and C++, that relies on the Pandas library for the management of the data. This is the standard mode of operations of more and more data scientists. This document has been written using Biogeme 3.2.11.

1 Data

Biogeme assumes that a Pandas database is available, containing only numerical entries. Each column corresponds to a variable, each row to an observation.

If you are not familiar with Pandas, prepare a file that contains in its first line a list of labels corresponding to the available data, and that each subsequent line contains the exact same number of numerical data, each row corresponding to an observation. Delimiters can be tabs or spaces.

The data file used for this example is `swissmetro.dat`. It can be downloaded from the “Data” section of `biogeme.epfl.ch`.

2 Python

Biogeme is a package of the Python programming language. Therefore, estimating a model amounts to writing a script in Python. Online tutorials and documentation about Python can easily be found. Although it is not necessary to master the Python language to specify models for Biogeme, it would definitely help to learn at least the basics. In this Section, we report some useful information when using the package Biogeme.

- Two versions of Python are commonly used: 2 and 3. Biogeme works only with Python version 3.
- Python is available on Linux, MacOSX and Windows. Biogeme is platform independent.
- The syntax of Python is case sensitive. It means that `varname` and `Varname`, for instance, would represent two different entities.

- The indentation of the code is important in Python. It is advised to use a text editor that has a “Python mode” to help managing these indentations.
- A Python statement must be on a single line, except if it is surrounded by parentheses. Sometimes, for the sake of readability, it is convenient to split the statement on several lines. In that case, make sure to include the parentheses. There are several examples below, for instance in the specification of the utility functions.

3 The model

The model is a logit model with 3 alternatives: *train*, *Swissmetro* and *car*. The utility functions are defined as:

```
V1 = (
    ASC_TRAIN +
    B_TIME * TRAIN_TT_SCALED +
    B_COST * TRAIN_COST_SCALED
)
V2 = (
    ASC_SM +
    B_TIME * SM_TT_SCALED +
    B_COST * SM_COST_SCALED
)
V3 = (
    ASC_CAR +
    B_TIME * CAR_TT_SCALED +
    B_COST * CAR_CO_SCALED
)
```

where

- TRAIN_TT_SCALED,
- TRAIN_COST_SCALED,
- SM_TT_SCALED,
- SM_COST_SCALED,
- CAR_TT_SCALED,
- CAR_CO_SCALED

are variables, and

- ASC_TRAIN,
- ASC_SM,
- ASC_CAR,
- B_TIME,
- B_COST

are parameters to be estimated. Note that it is not possible to identify all alternative specific constants ASC_TRAIN, ASC_SM, ASC_CAR from data. Consequently, ASC_SM is normalized to 0.

The availability of an alternative i is determined by the variable y_i , $i=1, 2, 3$, which is equal to 1 if the alternative is available, and 0 otherwise. The probability of choosing an available alternative i is given by the logit model:

$$P(i|\{1, 2, 3\}; x, \beta) = \frac{y_i e^{V_i(x, \beta)}}{y_1 e^{V_1(x, \beta)} + y_2 e^{V_2(x, \beta)} + y_3 e^{V_3(x, \beta)}}. \quad (1)$$

Given a data set of N observations, the log likelihood of the sample is

$$\mathcal{L} = \sum_n \log P(i_n|\{1, 2, 3\}; x_n, y_N, \beta) \quad (2)$$

where i_n is the alternative actually chosen by individual n , and x_n are the explanatory variables associated with individual n .

4 Data preparation

It is advised to perform the data preparation in a separate file. For instance, the file `swissmetro_data.py` is reported in Section A.1.

The file can contain comments, designed to document the specification. Single-line comments are included using the characters `#`, consistently with the Python syntax. All characters after this command, up to the end of the current line, are ignored by Python. Multiple lines comments are created by adding a delimiter (`"""`) at the beginning and the end of the comment. In our example, the file starts with comments describing the name of the file, its author and the date when it was created. A short description of its content is also provided.

```
"""File swissmetro_data.py
:author: Michel Bierlaire, EPFL
```

```
:date: Mon Mar 6 15:17:03 2023
```

```
Data preparation for Swissmetro, and definition of the variables  
"""
```

These comments are completely ignored by Python. However, it is recommended to use many comments to describe the content of the script, for future reference, or to help other persons to understand it.

The file must start by loading the Python libraries needed by Biogeme. For data preparation, the following libraries must be loaded:

- `pandas`, the generic package for data management,
- `biogeme.database`, the Biogeme module for data management.

It is custom in Python to use shortcuts to simplify the syntax. Here, we use `pd` for `pandas`, and `db` for `biogeme.database`. Finally, we need to import some mathematical expressions useful to build the model specification. In the data preparation file, we use only the expression `Variable` that defines the variables available in the database.

```
import pandas as pd  
import biogeme.database as db  
from biogeme.expressions import Variable
```

The next step consists in preparing the Pandas database. For instance, if the data file is separated by tabs, you can use the following statements:

```
df = pd.read_csv('swissmetro.dat', sep='\t')  
database = db.Database('swissmetro', df)
```

The first statement reads the data from the file, using tabs as delimiters. It stores it in a Pandas data structure. The second statement prepares the database for Biogeme. Clearly, if you prefer to create your Pandas database in another way, it is possible. In that case, you still have to use the second statement to transfer the Pandas database to Biogeme.

The name of the columns in the database characterize the variables for your model. In order to make them available as a Python variable, the following statements must be included:

```
PURPOSE = Variable('PURPOSE')  
CHOICE = Variable('CHOICE')  
GA = Variable('GA')  
LUGGAGE = Variable('LUGGAGE')  
TRAIN_CO = Variable('TRAIN_CO')  
CAR_AV = Variable('CAR_AV')  
SP = Variable('SP')  
TRAIN_AV = Variable('TRAIN_AV')
```

```

TRAIN_TT = Variable('TRAIN_TT')
SM_TT = Variable('SM_TT')
CAR_TT = Variable('CAR_TT')
CAR_CO = Variable('CAR_CO')
SM_CO = Variable('SM_CO')
SM_AV = Variable('SM_AV')
MALE = Variable('MALE')
GROUP = Variable('GROUP')
TRAIN_HE = Variable('TRAIN_HE')
SM_HE = Variable('SM_HE')
INCOME = Variable('INCOME')

```

Although not formally necessary, it is highly recommended to use the exact same name for the Python variable (on the left hand side) and the Biogeme variable (on the right hand side).

It is possible to tell Biogeme to ignore some observations in the data file. A boolean expression must be defined, that is evaluated for each observation in the data file. Each observation such that this expression is “true” is discarded from the sample. In our example, the modeler has developed the model only for work trips, so that every observation such that the trip purpose is not 1 or 3 is removed.

Observations such that the dependent variable `CHOICE` is 0 are also removed. The convention is that “false” is represented by 0, and “true” by 1, so that the multiplication sign ‘*’ can be interpreted as a “and”, and the addition sign ‘+’ as a “or”. Note also that the result of the ‘+’ can be 2, so that we test if the result is equal to 0 or not. The exclude condition in our example is therefore interpreted as: either (`PURPOSE` different from 1 and `PURPOSE` different from 3), or `CHOICE` equal to 0.

```

exclude = ((PURPOSE != 1) * (PURPOSE != 3) + (CHOICE == 0)) > 0
database.remove(exclude)

```

- We have conveniently used an intermediary Python variable `exclude` in this example. It is not necessary. The above statement is completely equivalent, but may be less readable:

```

database.remove(
    (
        (PURPOSE != 1) *
        (PURPOSE != 3) +
        (CHOICE == 0)
    ) > 0
)

```

- The same result can be obtained using Pandas directly, using the following syntax:

```

remove = (
    (
        (database.data.PURPOSE != 1) &
        (database.data.PURPOSE != 3)
    ) |
    (database.data.CHOICE == 0)
)
database.data.drop(
    database.data[remove].index,
    inplace=True
)

```

Pandas provides more powerful tools to manage the database. If you need to perform sophisticated data manipulations, it is advised to use Pandas instead of Biogeme for these purposes. Refer to the online Pandas documentation and the many tutorials available online.

It is possible to define new variables in addition to the variables defined in the data files.

```

SM_COST = SM_CO * (GA == 0)
TRAIN_COST = TRAIN_CO * (GA == 0)
CAR_AV_SP = CAR_AV * (SP != 0)
TRAIN_AV_SP = TRAIN_AV * (SP != 0)

```

When boolean expressions are involved, the value True is represented by 1, and the value False is represented by 0. Therefore, a multiplication involving a boolean expression is equivalent to a “and” operator. The above code is interpreted in the following way:

- CAR_AV_SP is equal to CAR_AV if SP is different from 0, and is equal to 0 otherwise. TRAIN_AV_SP is defined similarly.
- SM_COST is equal to SM_CO if GA is equal to 0, that is, if the traveler does not have a yearly pass (called “Generalabonnement” in German). If the traveler possesses a yearly pass, then GA is different from 0, and the variable SM_COST is zero. The variable TRAIN_COST is defined in the same way.

Variables can be also be rescaled. For numerical reasons, it is good practice to scale the data so that the values of the estimated parameters are around 1. A previous estimation with the unscaled data has generated parameters around -0.01 for both cost and time. Therefore, time and cost are divided by 100.

```

TRAIN_TT_SCALED = TRAIN_TT / 100
TRAIN_COST_SCALED = TRAIN_COST / 100

```

```

SM_TT_SCALED = SM_TT / 100
SM_COST_SCALED = SM_COST / 100
CAR_TT_SCALED = CAR_TT / 100
CAR_CO_SCALED = CAR_CO / 100

```

The Python syntax presented above is well suited for readability of the code. However, the calculations that are involved will be redone again and again each time the variable is needed, that is for each observation, and for each iteration of the estimation algorithm.

Therefore, it is advised to create new columns in the database that store the new variables, so that they are calculated once for all. This can be done using the following syntax:

```

SM_COST = database.DefineVariable(
    'SM_COST', SM_CO * (GA == 0)
)
TRAIN_COST = database.DefineVariable(
    'TRAIN_COST', TRAIN_CO * (GA == 0)
)
CAR_AV_SP = database.DefineVariable(
    'CAR_AV_SP', CAR_AV * (SP != 0)
)
TRAIN_AV_SP = database.DefineVariable(
    'TRAIN_AV_SP', TRAIN_AV * (SP != 0)
)
TRAIN_TT_SCALED = database.DefineVariable(
    'TRAIN_TT_SCALED', TRAIN_TT / 100
)
TRAIN_COST_SCALED = database.DefineVariable(
    'TRAIN_COST_SCALED', TRAIN_COST / 100
)
SM_TT_SCALED = database.DefineVariable(
    'SM_TT_SCALED', SM_TT / 100
)
SM_COST_SCALED = database.DefineVariable(
    'SM_COST_SCALED', SM_COST / 100
)
CAR_TT_SCALED = database.DefineVariable(
    'CAR_TT_SCALED', CAR_TT / 100
)
CAR_CO_SCALED = database.DefineVariable(
    'CAR_CO_SCALED', CAR_CO / 100
)

```


5 Model specification: Biogeme

The file `b011logit.py` is reported in Section A.2. We describe here its content. The objective is to provide to Biogeme the formula of the log likelihood function to maximize, using a syntax based on the Python programming language, and extended for the specific needs of Biogeme.

Like any Python script, the file must start by loading the Python libraries needed by Biogeme (after the comments describing the content of the file). The following libraries must be loaded:

- `biogeme.biogeme`, the core of Biogeme, that we rename `bio`,
- `biogeme.models`, containing the specification of useful models, such as the logit model.

We need also to import some mathematical expressions useful to build the model specification. In this specification file, we use only the expression `Beta` that defines the unknown parameters to be estimated. Other expressions such as `log` or `exp` can also be used.

```
import biogeme.biogeme as bio
from biogeme import models
from biogeme.expressions import Beta
```

Then, we need to import from the data preparation file (described in Section 4) the variables that are used for the model specification, as well as the database object.

```
from swissmetro_data import (
    database ,
    CHOICE ,
    SM_AV ,
    CAR_AV_SP ,
    TRAIN_AV_SP ,
    TRAIN_TT_SCALED ,
    TRAIN_COST_SCALED ,
    SM_TT_SCALED ,
    SM_COST_SCALED ,
    CAR_TT_SCALED ,
    CAR_CO_SCALED ,
)
```

The next statements use the function `Beta` to define the parameters to be estimated. For each parameter, the following information must be mentioned:

1. the name of the parameter,
2. the default value,

3. a lower bound (or `None`, if no bound is specified),
4. an upper bound, (or `None`, if no bound is specified),
5. a flag that indicates if the parameter must be estimated (0) or if it keeps its default value (1).

```
ASC_CAR = Beta('ASC_CAR', 0, None, None, 0)
ASC_TRAIN = Beta('ASC_TRAIN', 0, None, None, 0)
ASC_SM = Beta('ASC_SM', 0, None, None, 1)
B_TIME = Beta('B_TIME', 0, None, None, 0)
B_COST = Beta('B_COST', 0, None, None, 0)
```

- In Python, case sensitivity is enforced, so that `varname` and `Variance` would represent two different variables. In our example, the default value of each parameter is 0. If a previous estimation had been performed before, we could have used the previous estimates as default value.
- For the parameters that are estimated by Biogeme, the default value is used as the starting value for the optimization algorithm. For the parameters that are not estimated, the default value is used throughout the estimation process. In our example, the parameter `ASC_SM` is not estimated (as specified by the 1 in the fifth argument on the corresponding line), and its value is fixed to 0.
- A lower bound and an upper bound may be specified. If no bound is meaningful, use `None`.
- As for the definition of the variables, nothing prevents to write

```
car_cte = Beta('ASC_CAR', 0, None, None, 0)
```

and to use `car_cte` later in the specification. We **strongly** advise against this practice, and suggest to use the exact same name for the Python variable on the left hand side, and for the Biogeme variable, appearing as the first argument of the function, as illustrated in this example.

We now write the specification of the utility functions. Note the use of parentheses to split the specification on several lines, improving readability.

```
V1 = (
    ASC_TRAIN +
    B_TIME * TRAIN_TT_SCALED +
```

```

    B_COST * TRAIN_COST_SCALED
)
V2 = (
    ASC_SM +
    B_TIME * SM_TT_SCALED +
    B_COST * SM_COST_SCALED
)
V3 = (
    ASC_CAR +
    B_TIME * CAR_TT_SCALED +
    B_COST * CAR_CO_SCALED
)

```

We need to associate each utility function with the identifier, of the alternative, using the same numbering convention as in the data file. In this example, the convention is described in Table 1.

Train	1
Swissmetro	2
Car	3

Table 1: Numbering of the alternatives

To do this, we use a Python dictionary:

```
V = {1: V1, 2: V2, 3: V3}
```

We use also a dictionary to describe the availability conditions of each alternative:

```
av = {1: TRAIN_AV_SP, 2: SM_AV, 3: CAR_AV_SP}
```

We now define the choice model. The function `models.loglogit` provides the logarithm of the choice probability of the logit model. It takes three arguments:

1. the dictionary describing the utility functions,
2. the dictionary describing the availability conditions,
3. the alternative for which the probability must be calculated.

In this example, we obtain

```
logprob = models.loglogit(V, av, CHOICE)
```

We are now ready to create the BIOGEME object, using the following syntax:

```
the_biogeme = bio.BIOGEME(database, logprob)
```

Avoid using the name `biogeme` for the variable storing the Biogeme object, as this name refers to the package. Prefer `the_biogeme`, for instance. The constructor accepts two mandatory arguments:

- the database object containing the data,
- the formula for the contribution to the log likelihood of each row in the database.

It is advised to give a name to the model using the following statement:

```
the_biogeme.modelName = 'b01logit'
```

This name will be used to name the output files, generated once the model is estimated. In particular, a file named `b01logit.html` is generated, that can be opened in any browser.

It is custom to calculate the likelihood of the model where all coefficients are zero. This requires the knowledge of the choice set of each individual. Therefore, it relies on the availability conditions. This can be done using the following syntax:

```
the_biogeme.calculateNullLoglikelihood(av)
```

The estimation of the model parameters is then performed using the following statement.

```
results = the_biogeme.estimate()
```

6 Running Biogeme

The script is executed like any python script. Typically, by typing

```
python b01logit.py
```

in a terminal, or by typing “shift-return” in a Jupyter notebook.

Two files are generated:

- `b01logit.html` reports the results of the estimation in HTML format, and can be opened in your favorite browser.
- `b01logit.pickle` is a snapshot of the results of the estimation, and can be used in another Python script.

In order to avoid erasing previously generated results, the name of the files may vary from one run to the next. Therefore, it is important to verify the latest files created in the directory.

You can also print the name of the files that were actually created using the following Python statement:

```
print(f'HTML file: {results.data.htmlFileName}')  
print(f'Pickle file: {results.data.pickleFileName}')
```

7 Biogeme: the report file

The report file generated by Biogeme gathers information about the result of the estimation. First, some information about the version of Biogeme, and some links to relevant URLs is provided. Next, the name of the report file and the name of the database are reported.

The estimation report follows, including

- The number of parameters that have been estimated.
- The sample size, that is, the number of rows in the data file that have not been excluded.
- The number of excluded observations.
- `Null log likelihood` is the log likelihood \mathcal{L}^0 of the sample for a model where all parameters are zero (if its calculation has been requested),
- `Init log likelihood` is the log likelihood \mathcal{L}^i of the sample for the model defined with the default values of the parameters. Note that, if the default values of the parameters are all zero, it coincides with the null loglikelihood.
- `Final log likelihood` is the log likelihood \mathcal{L}^* of the sample for the estimated model.
- `Likelihood ratio test for the null model` is

$$-2(\mathcal{L}^0 - \mathcal{L}^*) \tag{3}$$

where \mathcal{L}^0 is the log likelihood of the null model as defined above, and \mathcal{L}^* is the log likelihood of the sample for the estimated model.

- `Rho-square for the null model` is

$$\rho^2 = 1 - \frac{\mathcal{L}^*}{\mathcal{L}^0}. \tag{4}$$

- `Rho-square-bar for the null model` is

$$\bar{\rho}^2 = 1 - \frac{\mathcal{L}^* - K}{\mathcal{L}^0}. \tag{5}$$

where K is the number of estimated parameters.

- Likelihood ratio test for the init. model is

$$-2(\mathcal{L}^i - \mathcal{L}^*) \quad (6)$$

where \mathcal{L}^i is the log likelihood of the init model as defined above, and \mathcal{L}^* is the log likelihood of the sample for the estimated model.

- Rho-square for the init. model is

$$\rho^2 = 1 - \frac{\mathcal{L}^*}{\mathcal{L}^i}. \quad (7)$$

- Rho-square-bar for the init. model is

$$\bar{\rho}^2 = 1 - \frac{\mathcal{L}^* - K}{\mathcal{L}^i}. \quad (8)$$

where K is the number of estimated parameters.

- Akaike Information Criterion is:

$$2K - 2\mathcal{L}^*, \quad (9)$$

where K is the number of estimated parameters.

- Bayesian Information Criterion is:

$$-2\mathcal{L}^* + K \ln(N), \quad (10)$$

where K is the number of estimated parameters, and N is the sample size.

- Final gradient norm is the gradient of the log likelihood function computed for the estimated parameters.
- Nbr of threads is the number of processors used by Biogeme to calculate the log likelihood at each iteration.
- Algorithm is the optimization algorithm used to solve the maximum likelihood estimation problem.
- Proportion analytical hessian is the proportion of iterations where the analytical second derivatives matrix (called “hessian”) has been calculated.

- **Relative projected gradient** is the norm of the projected gradient, scaled to account for the level of magnitude of the log likelihood. This quantity is used as stopping criterion for the algorithm.
- **Relative change** is the norm of the relative change between two consecutive iterates of the algorithm. This quantity is also used as a stopping criterion for the algorithm.
- **Number of iterations** is the number of iterations performed by the optimization algorithms.
- **Number of function evaluations** reports the number of times that the log likelihood function has been calculated.
- **Number of gradient evaluations** reports the number of times that the gradient of the log likelihood function has been calculated.
- **Number of hessian evaluations** reports the number of times that the second derivatives matrix (or hessian) of the log likelihood function has been calculated.
- **Cause of termination** provides the reason why the optimization algorithm has stopped.
- **Optimization time** is the actual time used by the algorithm.

The following section reports the estimates of the parameters of the utility function, together with some statistics. For each parameter β_k , the following is reported:

- The name of the parameter.
- The estimated value β_k .
- The robust standard error σ_k^R of the estimate, calculated as the square root of the k th diagonal entry of the robust estimate of the variance covariance matrix. (see Appendix B).
- The robust t statistics, calculated as $t_k^R = \beta_k / \sigma_k^R$.
- The robust p value, calculated as $2(1 - \Phi(t_k^R))$, where $\Phi(\cdot)$ is the cumulative density function of the univariate normal distribution.

The last section reports, for each pair of parameters k and ℓ ,

- the name of β_k ,

- the name of β_ℓ ,
- the entry $\Sigma_{k,\ell}$ of the Rao-Cramer bound (see Appendix B),
- the correlation between β_k and β_ℓ , calculated as

$$\frac{\Sigma_{k,\ell}}{\sqrt{\Sigma_{k,k}\Sigma_{\ell,\ell}}}, \quad (11)$$

- the t statistics, calculated as

$$t_{k,\ell} = \frac{\beta_k - \beta_\ell}{\sqrt{\Sigma_{k,k} + \Sigma_{\ell,\ell} - 2\Sigma_{k,\ell}}}, \quad (12)$$

- the p value, calculated as $2(1 - \Phi(t_{k,\ell}))$, where $\Phi(\cdot)$ is the cumulative density function of the univariate standard normal distribution,
- the entry $\Sigma_{k,\ell}^R$ of Σ^R , the robust estimate of the variance covariance matrix (see Appendix B),
- the robust correlation between β_k and β_ℓ , calculated as

$$\frac{\Sigma_{k,\ell}^R}{\sqrt{\Sigma_{k,k}^R \Sigma_{\ell,\ell}^R}}, \quad (13)$$

- the robust t statistics, calculated as

$$t_{k,\ell}^R = \frac{\beta_k - \beta_\ell}{\sqrt{\Sigma_{k,k}^R + \Sigma_{\ell,\ell}^R - 2\Sigma_{k,\ell}^R}}, \quad (14)$$

- the robust p value, calculated as $2(1 - \Phi(t_{k,\ell}^R))$, where $\Phi(\cdot)$ is the cumulative density function of the univariate standard normal distribution,

The final lines report the value of the smallest and the largest eigenvalues, as well as the ratio between the two, called the “condition number”. If the smallest eigenvalue is close to zero, it is a sign of singularity, that may be due to a lack of variation in the data or an unidentified model.

8 The results as Python variables

The estimation function returns an object that contains the results of the estimation as well as the associated statistics. It provides some functions for reporting. For instance, the statement

```
print(results.shortSummary())
```

generates the following output:

```
Results for model b01logit
Nbr of parameters:          4
Sample size:                6768
Excluded data:             3960
Null log likelihood:       -6964.663
Final log likelihood:      -5331.252
Likelihood ratio test (null): 3266.822
Rho square (null):         0.235
Rho bar square (null):     0.234
Akaike Information Criterion: 10670.5
Bayesian Information Criterion: 10697.78
```

The values of the estimated parameters, as well as the corresponding statistics, can be organized in a Pandas dataframe. For instance, the following statements

```
pandas_results = results.getEstimatedParameters()
print(pandas_results)
```

generate the following output:

	Value	Rob. Std err	Rob. t-test	Rob. p-value
ASC_CAR	-0.154633	0.058163	-2.658590	0.007847
ASC_TRAIN	-0.701187	0.082562	-8.492857	0.000000
B_COST	-1.083790	0.068225	-15.885521	0.000000
B_TIME	-1.277859	0.104254	-12.257120	0.000000

If `results` is the object returned by the estimation function, the results of the estimation can be accessed in `results.data`:

- `results.data.modelName`: the model name.
- `results.data.userNotes`: optional notes provided by the user to include in the report.
- `results.data.nparam`: the number K of estimated parameters.
- `results.data.betaValues`: a Numpy array containing the estimated values of the parameters, in an arbitrary order.
- `results.data.betaNames`: a list containing the name of the estimated parameters, in the same order as the values above.
- `results.data.nullLogLike`: the value \mathcal{L}^0 is the null log likelihood.
- `results.data.initLogLike`: the value \mathcal{L}^i is the initial log likelihood.

- `results.data.betas`: a list of objects corresponding to the parameters. Each of these objects contains the following entries, which should be self explanatory.
 - `beta.name`,
 - `beta.value`,
 - `beta.stdErr`,
 - `beta.lb`,
 - `beta.ub`,
 - `beta.tTest`,
 - `beta.pValue`,
 - `beta.robust_stdErr`,
 - `beta.robust_tTest`,
 - `beta.robust_pValue`,
 - `beta.bootstrap_stdErr`,
 - `beta.bootstrap_tTest`,
 - `beta.bootstrap_pValue`.
- `results.data.logLike`: the value \mathcal{L}^* of the log likelihood at the final value of the parameters.
- `results.data.g`: the gradient of the log likelihood at the final value of the parameters.
- `results.data.H`: the second derivatives matrix of the log likelihood at the final value of the parameters.
- `results.data.bhhh`: the BHHH matrix (19) at the final value of the parameters.
- `results.data.dataname`: the name of the database.
- `results.data.sampleSize`: the sample size N .
- `results.data.numberOfObservations`: the number of rows in the data file. If the data is not panel, it is the same as the sample size.
- `results.data.monteCarlo`: a boolean that is True if the model involves Monte-Carlo simulation for the calculation of integrals.

- `results.data.numberOfDraws`: number of draws used for Monte-Carlo simulation.
- `results.data.typesOfDraws`: type of draws used for Monte-Carlo simulation.
- `results.data.excludedData`: number of excluded data.
- `results.data.dataProcessingTime`: time needed to process the data before estimation.
- `results.data.drawsProcessingTime`: time needed to generate the draws for Monte-Carlo simulation.
- `results.data.optimizationTime`: time used by the optimization algorithm.
- `results.data.gradientNorm`: norm of the gradient of the log likelihood at the final value of the parameters.
- `results.data.optimizationMessages`: informations returned by the optimization routine.
- `results.data.numberOfWorkers`: number of processors used.
- `results.data.htmlFileName`: name of the HTML file.
- `results.data.pickleFileName`: name of the Pickle file.
- `results.data.latexFileName`: name of the \LaTeX file.
- `results.data.F12FileName`: name of the ALogit F12 file.
- `results.data.bootstrap`: a boolean that is True if the calculation of statistics using bootstrapping has been requested.
- `results.data.bootstrapTime`: the time needed for calculating the statistics with bootstrapping, if applicable.

In addition the robust variance-covariance matrix can be obtained using

```
results.data.secondOrderTable
```

References

- Ben-Akiva, M. E. and Lerman, S. R. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand*, MIT Press, Cambridge, Ma.
- Berndt, E. K., Hall, B. H., Hall, R. E. and Hausman, J. A. (1974). Estimation and inference in nonlinear structural models, *Annals of Economic and Social Measurement* **3/4**: 653–665.
- Bierlaire, M. (2020). A short introduction to pandasbiogeme, *Technical Report TRANSP-OR 200605*, Transport and Mobility Laboratory, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.
- Kauermann, G. and Carroll, R. (2001). A note on the efficiency of sandwich covariance matrix estimation, *Journal of the American Statistical Association* **96**(456).
- White, H. (1982). Maximum likelihood estimation of misspecified models, *Econometrica* **50**: 1–25.

A Complete specification files

A.1 swissmetro_data.py

```
1 """ File swissmetro_data.py
2
3 :author: Michel Bierlaire, EPFL
4 :date: Mon Mar 6 15:17:03 2023
5
6 Data preparation for Swissmetro, and definition of the variables
7 """
8
9 import pandas as pd
10 import biogeme.database as db
11 from biogeme.expressions import Variable
12
13 # Read the data
14 df = pd.read_csv('swissmetro.dat', sep='\t')
15 database = db.Database('swissmetro', df)
16
17 PURPOSE = Variable('PURPOSE')
18 CHOICE = Variable('CHOICE')
19 GA = Variable('GA')
20 LUGGAGE = Variable('LUGGAGE')
21 TRAIN_CO = Variable('TRAIN_CO')
22 CAR_AV = Variable('CAR_AV')
23 SP = Variable('SP')
24 TRAIN_AV = Variable('TRAIN_AV')
25 TRAIN_TT = Variable('TRAIN_TT')
26 SM_TT = Variable('SM_TT')
27 CAR_TT = Variable('CAR_TT')
28 CAR_CO = Variable('CAR_CO')
29 SM_CO = Variable('SM_CO')
30 SM_AV = Variable('SM_AV')
31 MALE = Variable('MALE')
32 GROUP = Variable('GROUP')
33 TRAIN_HE = Variable('TRAIN_HE')
34 SM_HE = Variable('SM_HE')
35 INCOME = Variable('INCOME')
36 # Removing some observations can be done directly using pandas.
37 # remove = (((database.data.PURPOSE != 1) &
38 #           (database.data.PURPOSE != 3)) |
39 #           (database.data.CHOICE == 0))
40 # database.data.drop(database.data[remove].index, inplace=True)
41 # Here we use the "biogeme" way:
42 exclude = ((PURPOSE != 1) * (PURPOSE != 3) + (CHOICE == 0)) > 0
43 database.remove(exclude)
44
```

```

45
46 # Definition of new variables
47 SMLCOST = database.DefineVariable(
48     'SM_COST', SMLCO * (GA == 0)
49 )
50 TRAIN_COST = database.DefineVariable(
51     'TRAIN_COST', TRAIN_CO * (GA == 0)
52 )
53 CAR_AV_SP = database.DefineVariable(
54     'CAR_AV_SP', CAR_AV * (SP != 0)
55 )
56 TRAIN_AV_SP = database.DefineVariable(
57     'TRAIN_AV_SP', TRAIN_AV * (SP != 0)
58 )
59 TRAIN_TT_SCALED = database.DefineVariable(
60     'TRAIN_TT_SCALED', TRAIN_TT / 100
61 )
62 TRAIN_COST_SCALED = database.DefineVariable(
63     'TRAIN_COST_SCALED', TRAIN_COST / 100
64 )
65 SM_TT_SCALED = database.DefineVariable(
66     'SM_TT_SCALED', SM_TT / 100
67 )
68 SMLCOST_SCALED = database.DefineVariable(
69     'SM_COST_SCALED', SMLCOST / 100
70 )
71 CAR_TT_SCALED = database.DefineVariable(
72     'CAR_TT_SCALED', CAR_TT / 100
73 )
74 CAR_CO_SCALED = database.DefineVariable(
75     'CAR_CO_SCALED', CAR_CO / 100
76 )

```

A.2 b01logit.py

```

1 """ File b01logit.py
2
3 :author: Michel Bierlaire, EPFL
4 :date: Sun Apr 9 17:02:18 2023
5
6 Example of a logit model.
7 Three alternatives: Train, Car and Swissmetro
8 SP data
9 """
10 import biogeme.biogeme as bio
11 from biogeme import models
12 from biogeme.expressions import Beta
13 from swissmetro_data import (
14     database,

```

```

15     CHOICE,
16     SMAV,
17     CAR_AV_SP,
18     TRAIN_AV_SP,
19     TRAIN_TT_SCALED,
20     TRAIN_COST_SCALED,
21     SM_TT_SCALED,
22     SM_COST_SCALED,
23     CAR_TT_SCALED,
24     CAR_CO_SCALED,
25 )
26
27 # Parameters to be estimated
28 ASC_CAR = Beta('ASC_CAR', 0, None, None, 0)
29 ASC_TRAIN = Beta('ASC_TRAIN', 0, None, None, 0)
30 ASC_SM = Beta('ASC_SM', 0, None, None, 1)
31 B_TIME = Beta('B_TIME', 0, None, None, 0)
32 B_COST = Beta('B_COST', 0, None, None, 0)
33
34
35 # Definition of the utility functions
36 V1 = (
37     ASC_TRAIN +
38     B_TIME * TRAIN_TT_SCALED +
39     B_COST * TRAIN_COST_SCALED
40 )
41 V2 = (
42     ASC_SM +
43     B_TIME * SM_TT_SCALED +
44     B_COST * SM_COST_SCALED
45 )
46 V3 = (
47     ASC_CAR +
48     B_TIME * CAR_TT_SCALED +
49     B_COST * CAR_CO_SCALED
50 )
51
52 # Associate utility functions with the numbering of alternatives
53 V = {1: V1, 2: V2, 3: V3}
54
55 # Associate the availability conditions with the alternatives
56 av = {1: TRAIN_AV_SP, 2: SMAV, 3: CAR_AV_SP}
57
58 # Definition of the model. This is the contribution of each
59 # observation to the log likelihood function.
60 logprob = models.loglogit(V, av, CHOICE)
61
62 # Create the Biogeme object
63 the_biogeme = bio.BIOGEME(database, logprob)

```



```
64 the_biogeme.modelName = 'b01logit'
65
66 # Calculate the null log likelihood for reporting.
67 the_biogeme.calculateNullLoglikelihood(av)
68
69 # Estimate the parameters
70 results = the_biogeme.estimate()
71 print(results.shortSummary())
72
73 # Get the results in a pandas table
74 pandas_results = results.getEstimatedParameters()
75 print(pandas_results)
```

B Estimation of the variance-covariance matrix

Under relatively general conditions, the asymptotic variance-covariance matrix of the maximum likelihood estimates of the vector of parameters $\theta \in \mathbb{R}^K$ is given by the Cramer-Rao bound

$$-\mathbb{E} [\nabla^2 \mathcal{L}(\theta)]^{-1} = \left\{ -\mathbb{E} \left[\frac{\partial^2 \mathcal{L}(\theta)}{\partial \theta \partial \theta^\top} \right] \right\}^{-1}. \quad (15)$$

The term in square brackets is the matrix of the second derivatives of the log likelihood function with respect to the parameters evaluated at the true parameters. Thus the entry in the k th row and the ℓ th column is

$$\frac{\partial^2 \mathcal{L}(\theta)}{\partial \theta_k \partial \theta_\ell}. \quad (16)$$

Since we do not know the actual values of the parameters at which to evaluate the second derivatives, or the distribution of \mathbf{x}_{in} and \mathbf{x}_{jn} over which to take their expected value, we estimate the variance-covariance matrix by evaluating the second derivatives at the estimated parameters $\hat{\theta}$ and the sample distribution of \mathbf{x}_{in} and \mathbf{x}_{jn} instead of their true distribution. Thus we use

$$\mathbb{E} \left[\frac{\partial^2 \mathcal{L}(\theta)}{\partial \theta_k \partial \theta_\ell} \right] \approx \sum_{n=1}^N \left[\frac{\partial^2 (\mathbf{y}_{\text{in}} \ln P_n(\mathbf{i}) + \mathbf{y}_{\text{jn}} \ln P_n(\mathbf{j}))}{\partial \theta_k \partial \theta_\ell} \right]_{\theta=\hat{\theta}}, \quad (17)$$

as a consistent estimator of the matrix of second derivatives.

Denote this matrix as $\hat{\mathbf{A}}$. Note that, from the second order optimality conditions of the optimization problem, this matrix is negative semi-definite, which is the algebraic equivalent of the local concavity of the log likelihood function. If the maximum is unique, the matrix is negative definite, and the function is locally strictly concave.

An estimate of the Cramer-Rao bound (15) is given by

$$\hat{\Sigma}_\theta^{\text{CR}} = -\hat{\mathbf{A}}^{-1}. \quad (18)$$

If the matrix $\hat{\mathbf{A}}$ is negative definite then $-\hat{\mathbf{A}}$ is invertible and the Cramer-Rao bound is positive definite.

Another consistent estimator of the (negative of the) second derivatives matrix can be obtained by the matrix of the cross-products of first derivatives as follows:

$$-\mathbb{E} \left[\frac{\partial^2 \mathcal{L}(\theta)}{\partial \theta \partial \theta^\top} \right] \approx \sum_{n=1}^n \left(\frac{\partial \ell_n(\hat{\theta})}{\partial \theta} \right) \left(\frac{\partial \ell_n(\hat{\theta})}{\partial \theta} \right)^\top = \hat{\mathbf{B}}, \quad (19)$$

where

$$\left(\frac{\partial \ell_n(\hat{\theta})}{\partial \theta}\right) = \frac{\partial}{\partial \theta}(\log P(i_n | \mathcal{C}_n; \hat{\theta})) \quad (20)$$

is the gradient vector of the likelihood of observation n . This approximation is employed by the BHHH algorithm, from the work by Berndt et al. (1974). Therefore, an estimate of the variance-covariance matrix is given by

$$\hat{\Sigma}_{\theta}^{\text{BHHH}} = \hat{\mathbf{B}}^{-1}, \quad (21)$$

although it is rarely used. Instead, $\hat{\mathbf{B}}$ is used to derive a third consistent estimator of the variance-covariance matrix of the parameters, defined as

$$\hat{\Sigma}_{\theta}^{\text{R}} = (-\hat{\mathbf{A}})^{-1} \hat{\mathbf{B}} (-\hat{\mathbf{A}})^{-1} = \hat{\Sigma}_{\theta}^{\text{CR}} (\hat{\Sigma}_{\theta}^{\text{BHHH}})^{-1} \hat{\Sigma}_{\theta}^{\text{CR}}. \quad (22)$$

It is called the *robust* estimator, or sometimes the *sandwich* estimator, due to the form of equation (22). Biogeme reports statistics based on both the Cramer-Rao estimate (18) and the robust estimate (22).

When the true likelihood function is maximized, these estimators are asymptotically equivalent, and the Cramer-Rao bound should be preferred (Kauermann and Carroll, 2001). When other consistent estimators are used, the robust estimator must be used (White, 1982). Consistent non-maximum likelihood estimators, known as pseudo maximum likelihood estimators, are often used when the true likelihood function is unknown or difficult to compute. In such cases, it is often possible to obtain consistent estimators by maximizing an objective function based on a simplified probability distribution.