

Maximum Likelihood Estimation

for GAUSS™

Version 5.0

Aptech Systems, Inc.

Information in this document is subject to change without notice and does not represent a commitment on the part of Aptech Systems, Inc. The software described in this document is furnished under a license agreement or nondisclosure agreement. The software may be used or copied only in accordance with the terms of the agreement. The purchaser may make one copy of the software for backup purposes. No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose other than the purchaser's personal use without the written permission of Aptech Systems, Inc. ©Copyright 1988-2001 by Aptech Systems, Inc., Maple Valley, WA. All Rights Reserved.

GAUSS, GAUSS Engine, GAUSS Light are trademarks of Aptech Systems, Inc. All other trademarks are the properties of their respective owners.

Documentation Version: June 6, 2001

Part Number: 001414

Contents

1	Installation	1
1.1	UNIX	1
1.1.1	Download	1
1.1.2	Floppy	1
1.1.3	Solaris 2.x Volume Management	2
1.2	Windows/NT/2000	3
1.2.1	Download	3
1.2.2	Floppy	3
1.3	Differences Between the UNIX and Windows/NT/2000 Versions	3
2	Maximum Likelihood Estimation	5
2.1	Getting Started	5
2.1.1	README Files	5
2.1.2	Setup	5
2.2	The Log-likelihood Function	6
2.3	Algorithm	7
2.3.1	Derivatives	7

2.3.2	The Secant Algorithms	8
2.3.3	Berndt, Hall, Hall, and Hausman’s (BHHH) Method	9
2.3.4	Polak-Ribiere-type Conjugate Gradient (PRCG)	9
2.3.5	Line Search Methods	9
2.3.6	Random Search	11
2.3.7	Weighted Maximum Likelihood	11
2.3.8	Active and Inactive Parameters	11
2.3.9	Example	12
2.4	Managing Optimization	13
2.4.1	Scaling	13
2.4.2	Condition	13
2.4.3	Starting Point	17
2.4.4	Diagnosis	17
2.5	Gradients	18
2.5.1	Analytical Gradient	18
2.5.2	User-Supplied Numerical Gradient	19
2.5.3	Analytical Hessian	19
2.5.4	User-Supplied Numerical Hessian	21
2.5.5	Switching Algorithms Automatically	21
2.6	FASTMAX – Fast Execution MAXLIK	22
2.6.1	Undefined Function Evaluation	22
2.7	Inference	23
2.7.1	Wald Inference	23
2.7.2	Profile Likelihood Inference	25

2.7.3	Profile Trace Plots	28
2.7.4	Bootstrap	29
2.7.5	Pseudo-Random Number Generators	30
2.7.6	Bayesian Inference	31
2.8	Run-Time Switches	33
2.9	Calling MAXLIK Recursively	34
2.10	Using <code>_MAXLIK</code> Directly	35
2.11	Error Handling	35
2.11.1	Return Codes	35
2.11.2	Error Trapping	36
2.12	References	36
3	Maximum Likelihood Reference	39
	FASTMAX	40
	FASTBayes	48
	FASTBoot	51
	FASTPfClimits	54
	FASTProfile	56
	MAXLIK	59
	MAXBayes	70
	MAXBoot	73
	MAXBlimits	75
	MAXCLPrt	76
	MAXDensity	78
	MAXHist	80
	MAXProfile	82
	MAXPfClimits	85
	MAXPrt	87
	MAXSet	88
	MAXTlimits	89

4	Event Count and Duration Regression	91
4.1	Getting Started	92
4.1.1	README Files	92
4.1.2	Setup	92
4.2	About the COUNT Procedures	93
4.2.1	Inputs	94
4.2.2	Outputs	94
4.2.3	Global Control Variables	95
4.2.4	Statistical Inference	97
4.2.5	Problems with Convergence	98
4.3	Annotated Bibliography	100
5	Count Reference	103
	CountCLPrt	104
	CountPrt	105
	CountSet	106
	Expgam	107
	Expon	112
	Hurdlep	116
	Negbin	120
	Pareto	126
	Poisson	131
	Supreme	136
	Supreme2	140
	Index	145

Chapter 1

Installation

1.1 UNIX

If you are unfamiliar with UNIX, see your system administrator or system documentation for information on the system commands referred to below. The device names given are probably correct for your system.

1.1.1 Download

1. Copy the `.tar.gz` file to `/tmp`.
2. Unzip the file.

```
gunzip appxxx.tar.gz
```

3. `cd` to the **GAUSS** or **GAUSS Engine** installation directory. We are assuming `/usr/local/gauss` in this case.

```
cd /usr/local/gauss
```

4. Untar the file.

```
tar xvf /tmp/appxxx.tar
```

1.1.2 Floppy

1. Make a temporary directory.

```
mkdir /tmp/workdir
```

1. INSTALLATION

2. `cd` to the temporary directory.

```
cd /tmp/workdir
```

3. Use `tar` to extract the files.

```
tar xvf device_name
```

If this software came on diskettes, repeat the `tar` command for each diskette.

4. Read the README file.

```
more README
```

5. Run the `install.sh` script in the work directory.

```
./install.sh
```

The directory the files are install to should be the same as the install directory of **GAUSS** or the **GAUSS Engine**.

6. Remove the temporary directory (optional).

The following device names are suggestions. See your system administrator. If you are using Solaris 2.x, see Section 1.1.3.

Operating System	3.5-inch diskette	1/4-inch tape	DAT tape
Solaris 1.x SPARC	<code>/dev/rfd0</code>	<code>/dev/rst8</code>	
Solaris 2.x SPARC	<code>/dev/rfd0a</code> (vol. mgt. off)	<code>/dev/rst12</code>	<code>/dev/rmt/11</code>
Solaris 2.x SPARC	<code>/vol/dev/aliases/floppy0</code>	<code>/dev/rst12</code>	<code>/dev/rmt/11</code>
Solaris 2.x x86	<code>/dev/rfd0c</code> (vol. mgt. off)		<code>/dev/rmt/11</code>
Solaris 2.x x86	<code>/vol/dev/aliases/floppy0</code>		<code>/dev/rmt/11</code>
HP-UX	<code>/dev/rfloppy/c20Ad1s0</code>		<code>/dev/rmt/0m</code>
IBM AIX	<code>/dev/rfd0</code>	<code>/dev/rmt.0</code>	
SGI IRIX	<code>/dev/rdisk/fds0d2.3.5hi</code>		

1.1.3 Solaris 2.x Volume Management

If Solaris 2.x volume management is running, insert the floppy disk and type

```
volcheck
```

to signal the system to mount the floppy.

The floppy device names for Solaris 2.x change when the volume manager is turned off and on. To turn off volume management, become the superuser and type

```
/etc/init.d/volmgt off
```

To turn on volume management, become the superuser and type

```
/etc/init.d/volmgt on
```


1. INSTALLATION

1.2 Windows/NT/2000

1.2.1 Download

Unzip the .zip file into the **GAUSS** or **GAUSS Engine** installation directory.

1.2.2 Floppy

1. Place the diskette in a floppy drive.
2. Call up a DOS window
3. In the DOS window log onto the root directory of the diskette drive. For example:

```
A:<enter>
cd\

```

4. Type: **ginstall** *source_drive* *target_path*

source_drive Drive containing files to install
with colon included

For example: **A:**

target_path Main drive and subdirectory to install
to without a final \

For example: **C:\GAUSS**

A directory structure will be created if it does not already exist and the files will be copied over.

<i>target_path</i> \ src	source code files
<i>target_path</i> \ lib	library files
<i>target_path</i> \ examples	example files

1.3 Differences Between the UNIX and Windows/NT/2000 Versions

- If the functions can be controlled during execution by entering keystrokes from the keyboard, it may be necessary to press *Enter* after the keystroke in the UNIX version.

1. *INSTALLATION*

- On the Intel math coprocessors used by the Windows/NT/2000 machines, intermediate calculations have 80-bit precision, while on the current UNIX machines, all calculations are in 64-bit precision. For this reason, **GAUSS** programs executed under UNIX may produce slightly different results, due to differences in roundoff, from those executed under Windows/NT/2000.

Chapter 2

Maximum Likelihood Estimation

MaxLik

written by

Ronald Schoenberg

This module contains a set of procedures for the solution of the constrained maximum likelihood problem

2.1 Getting Started

GAUSS 3.6.18+ is required to use these routines.

2.1.1 README Files

The file **README.ml** contains any last minute information on this module. Please read it before using the procedures in this module.

2.1.2 Setup

In order to use the procedures in the *MAXIMUM LIKELIHOOD* Module, the **MAXLIK** library must be active. This is done by including `maxlik` in the **library** statement at the top of your program or command file:

2. MAXIMUM LIKELIHOOD ESTIMATION

```
library maxlik,pgraph;
```

This enables **GAUSS** to find the *MAXIMUM LIKELIHOOD* procedures. If you plan to make any right hand references to the global variables (described in the *REFERENCE* section), you also need the statement:

```
#include maxlik.ext;
```

Finally, to reset global variables in succeeding executions of the command file the following instruction can be used:

```
maxset;
```

This could be included with the above statements without harm and would insure the proper definition of the global variables for all executions of the command file.

The version number of each module is stored in a global variable:

__ml__ver 3×1 matrix, the first element contains the major version number of the *MAXIMUM LIKELIHOOD* Module, the second element the minor version number, and the third element the revision number.

If you call for technical support, you may be asked for the version number of your copy of this module.

2.2 The Log-likelihood Function

MAXIMUM LIKELIHOOD is a set of procedures for the estimation of the parameters of models via the maximum likelihood method with general constraints on the parameters, along with an additional set of procedures for statistical inference.

MAXIMUM LIKELIHOOD solves the general maximum likelihood problem

$$L = \sum_{i=1}^N \log P(Y_i; \theta)^{w_i}$$

where N is the number of observations, $P(Y_i, \theta)$ is the probability of Y_i given θ , a vector of parameters, and w_i is the weight of the i -th observation.

The *MAXIMUM LIKELIHOOD* procedure **MAXLIK** finds values for the parameters in θ such that L is maximized. In fact **MAXLIK** minimizes $-L$. It is important to note, however, that the user must specify the log-probability to be *maximized*. **MAXLIK** transforms the function into the form to be minimized.

MAXLIK has been designed to make the specification of the function and the handling of the data convenient. The user supplies a procedure that computes $\log P(Y_i; \theta)$, i.e., the log-likelihood, given the parameters in θ , for either an individual observation or set of observations (i.e., it must return either the log-likelihood for an individual observation or a vector of log-likelihoods for a matrix of observations; see discussion of the global variable **__row** below). **MAXLIK** uses this procedure to construct the function to be minimized.

2. MAXIMUM LIKELIHOOD ESTIMATION

2.3 Algorithm

MAXIMUM LIKELIHOOD finds values for the parameters using an iterative method. In this method the parameters are updated in a series of iterations beginning with a starting values that you provide. Let θ_t be the current parameter values. Then the succeeding values are

$$\theta_{t+1} = \theta_t + \rho\delta$$

where δ is a $k \times 1$ direction vector, and ρ a scalar step length.

Direction

Define

$$\begin{aligned}\Sigma(\theta) &= \frac{\partial^2 L}{\partial\theta\partial\theta'} \\ \Psi(\theta) &= \frac{\partial L}{\partial\theta}\end{aligned}$$

The direction, δ is the solution to

$$\Sigma(\theta_t)\delta = \Psi(\theta_t)$$

This solution requires that Σ be positive definite.

Line Search

The line search finds a value of ρ that minimizes or decreases $L(\theta_t + \rho\delta)$.

2.3.1 Derivatives

The minimization requires the calculation of a Hessian, Σ , and the gradient, Ψ . **MAXLIK** computes these numerically if procedures to compute them are not supplied.

If you provide a proc for computing Ψ , the first derivative of L , **MAXLIK** uses it in computing Σ , the second derivative of L , i.e., Σ is computed as the Jacobian of the gradient. This improves the computational precision of the Hessian by about four places. The accuracy of the gradient is improved and thus the iterations converge in fewer iterations. Moreover, the convergence takes less time because of a decrease in function calls - the numerical gradient requires k function calls while an analytical gradient reduces that to one.

2.3.2 The Secant Algorithms

The Hessian may be very expensive to compute at every iteration, and poor start values may produce an ill-conditioned Hessian. For these reasons alternative algorithms are provided in **MAXLIK** for updating the Hessian rather than computing it directly at each iteration. These algorithms, as well as step length methods, may be modified during the execution of **MAXLIK**.

Beginning with an initial estimate of the Hessian, or a conformable identity matrix, an update is calculated. The update at each iteration adds more “information” to the estimate of the Hessian, improving its ability to project the direction of the descent. Thus after several iterations the secant algorithm should do nearly as well as Newton iteration with much less computation.

There are two basic types of secant methods, the BFGS (Broyden, Fletcher, Goldfarb, and Shanno), and the DFP (Davidon, Fletcher, and Powell). They are both rank two updates, that is, they are analogous to adding two rows of new data to a previously computed moment matrix. The Cholesky factorization of the estimate of the Hessian is updated using the functions **cholup** and **choldn**.

In addition, **MAXLIK** includes a scoring method, BHHH (Berndt, Hall, Hall, and Hausman). This method computes the gradient of the likelihood by observation, i.e., the Jacobian, and estimates Σ as the cross-product of this Jacobian.

Secant Methods (BFGS and DFP)

BFGS is the method of Broyden, Fletcher, Goldfarb, and Shanno, and DFP is the method of Davidon, Fletcher, and Powell. These methods are complementary (Luenberger 1984, page 268). BFGS and DFP are like the NEWTON method in that they use both first and second derivative information. However, in DFP and BFGS the Hessian is approximated, reducing considerably the computational requirements. Because they do not explicitly calculate the second derivatives they are sometimes called *quasi-Newton* methods. While it takes more iterations than the NEWTON method, the use of an approximation produces a gain because it can be expected to converge in less overall time (unless analytical second derivatives are available in which case it might be a toss-up).

The secant methods are commonly implemented as updates of the *inverse* of the Hessian. This is not the best method numerically for the BFGS algorithm (Gill and Murray, 1972). This version of **MAXLIK**, following Gill and Murray (1972), updates the Cholesky factorization of the Hessian instead, using the functions **cholup** and **choldn** for BFGS. The new direction is then computed using **cholsol**, a Cholesky solve, as applied to the updated Cholesky factorization of the Hessian and the gradient.

2. MAXIMUM LIKELIHOOD ESTIMATION

2.3.3 Berndt, Hall, Hall, and Hausman's (BHHH) Method

BHHH is a method proposed by Berndt, Hall, Hall and Hausman (1974) for the maximization of log-likelihood functions. It is a *scoring* method that uses the cross-product of the matrix of first derivatives to estimate the Hessian matrix.

This calculation can be time-consuming, especially for large data sets, since a gradient matrix exactly the same size as the data set must be computed. For that reason BHHH cannot be considered a preferred choice for an optimization algorithm.

2.3.4 Polak-Ribiere-type Conjugate Gradient (PRCG)

The conjugate gradient method is an improvement on the steepest descent method without the increase in memory and computational requirements of the secant methods. Only the gradient is stored, and the calculation of the new direction is different:

$$d_{t+1} = -g_{t+1} + \beta_t d_t$$

where t indicates t -th iteration, d is the direction, g is the gradient. The conjugate gradient method used in **MAXLIK** is a variation called the Polak-Ribiere method where

$$\beta_t = \frac{(g_{t+1} - g_t)' g_{t+1}}{g_t' g_t}$$

The Newton and secant methods require the storage on the order of the Hessian in memory, i.e., $8k^2$ bytes of memory, where k is the number of parameters. For a very large problem this can be prohibitive. For example, 200 parameters will require 3.2 megabytes of memory, and this doesn't count the copies of the Hessian that may be generated by the program. For large problems, then, the PRCG and STEEP methods may be the only alternative. As described above, STEEP can be very inefficient in the region of the minimum, and therefore the PRCG is the method of choice in these cases.

2.3.5 Line Search Methods

Given a direction vector d , the updated estimate of the parameters is computed

$$\theta_{t+1} = \theta_t + \rho \delta$$

where ρ is a constant, usually called the *step length*, that increases the descent of the function given the direction. **MAXLIK** includes a variety of methods for computing ρ . The value of the function to be minimized as a function of ρ is

$$L(\theta_t + \rho \delta)$$

2. MAXIMUM LIKELIHOOD ESTIMATION

Given θ and d , this is a function of a single variable ρ . Line search methods attempt to find a value for ρ that decreases m . STEPBT is a polynomial fitting method, BRENT and HALF are iterative search methods. A fourth method called ONE forces a step length of 1. The default line search method is STEPBT. If this, or any selected method, fails, then BRENT is tried. If BRENT fails, then HALF is tried. If all of the line search methods fail, then a random search is tried (provided `__max__RandRadius` is greater than zero).

STEPBT

STEPBT is an implementation of a similarly named algorithm described in Dennis and Schnabel (1983). It first attempts to fit a quadratic function to $m(\theta_t + \rho\delta)$ and computes an ρ that minimizes the quadratic. If that fails it attempts to fit a cubic function. The cubic function more accurately portrays the F which is not likely to be very quadratic, but is, however, more costly to compute. STEPBT is the default line search method because it generally produces the best results for the least cost in computational resources.

BRENT

This method is a variation on the *golden section* method due to Brent (1972). In this method, the function is evaluated at a sequence of test values for ρ . These test values are determined by extrapolation and interpolation using the constant, $(\sqrt{5} - 1)/2 = .6180\dots$. This constant is the inverse of the so-called “golden ratio” $((\sqrt{5} + 1)/2 = 1.6180\dots$ and is why the method is called a golden section method. This method is generally more efficient than STEPBT but requires significantly more function evaluations.

HALF

This method first computes $m(x + d)$, i.e., sets $\rho = 1$. If $m(x + d) < m(x)$ then the step length is set to 1. If not, then it tries $m(x + .5d)$. The attempted step length is divided by one half each time the function fails to decrease, and exits with the current value when it does decrease. This method usually requires the fewest function evaluations (it often only requires one), but it is the least efficient in that it is not very likely to find the step length that decreases m the most.

BHHHStep

This is a variation on the golden search method. A sequence of step lengths are computed, interpolating or extrapolating using a golden ratio, and the method exits when the function decreases by an amount determined by `__max__Interp`.

2. MAXIMUM LIKELIHOOD ESTIMATION

2.3.6 Random Search

If the line search fails, i.e., no ρ is found such that $m(\theta_t + \rho\delta) < m(\theta_t)$, then a search is attempted for a random direction that decreases the function. The radius of the random search is fixed by the global variable, `__max__RandRadius` (default = .01), times a measure of the magnitude of the gradient. **MAXLIK** makes `__max__MaxTry` attempts to find a direction that decreases the function, and if all of them fail, the direction with the smallest value for m is selected.

The function should never increase, but this assumes a well-defined problem. In practice, many functions are not so well-defined, and it often is the case that convergence is more likely achieved by a direction that puts the function somewhere else on the hyper-surface even if it is at a higher point on the surface. Another reason for permitting an increase in the function here is that halting the minimization altogether is only alternative if it is not at the minimum, and so one might as well retreat to another starting point. If the function repeatedly increases, then you would do well to consider improving either the specification of the problem or the starting point.

2.3.7 Weighted Maximum Likelihood

Weights are specified by setting the **GAUSS** global, `__weight` to a weighting vector, or by assigning it the name of a column in the **GAUSS** data set being used in the estimation. Thus if a data matrix is being analyzed, `__weight` must be assigned to a vector.

MAXLIK assumes that the weights sum to the number of observations, i.e, that the weights are frequencies. This will be an issue only with statistical inference. Otherwise, any multiple of the weights will produce the same results.

2.3.8 Active and Inactive Parameters

The **MAXLIK** global `__max__Active` may be used to fix parameters to their start values. This allows estimation of different models without having to modify the function procedure. `__max__Active` must be set to a vector of the same length as the vector of start values. Elements of `__max__Active` set to zero will be fixed to their starting values, while nonzero elements will be estimated.

This feature may also be used for model testing. `__max__NumObs` times the difference between the function values (the second return argument in the call to **MAXLIK**) is chi-squared distributed with degrees of freedom equal to the number of fixed parameters in `__max__Active`.

2.3.9 Example

This example estimates coefficients for a tobit model:

```

library maxlik;
#include maxlik.ext;
maxset;

proc lpr(x,z);
  local t,s,m,u;
  s = x[4];
  if s <= 1e-4;
    retp(error(0));
  endif;
  m = z[.,2:4]*x[1:3,.];
  u = z[.,1] ./= 0;
  t = z[.,1]-m;
  retp(u.*(-(t.*t)/(2*s)-.5*ln(2*s*pi)) + (1-u).*(ln(cdfnc(m/sqrt(s)))));
endp;

x0 = { 1, 1, 1, 1 };
__title = "tobit example";

{x,f,g,cov,ret} = maxlik("tobit",0,&lpr,x0);
call maxprt(x,f,g,cov,ret);

```

The output is:

```

=====
                                tobit example
=====
MAXLIK Version 5.0.0                                5/30/2001  1:11 pm
=====
                                Data Set:  tobit
-----

return code =      0
normal convergence

Mean log-likelihood      -1.13291
Number of cases          100

Covariance matrix of the parameters computed by the following method:
Inverse of computed Hessian

Parameters   Estimates   Std. err.  Est./s.e.  Prob.      Gradient
-----
P01          0.0104      0.0845     0.123      0.4510     -0.0000

```

2. MAXIMUM LIKELIHOOD ESTIMATION

P02	-0.2081	0.0946	-2.200	0.0139	-0.0000
P03	-0.0998	0.0801	-1.245	0.1065	-0.0000
P04	0.6522	0.0999	6.531	0.0000	-0.0000

Correlation matrix of the parameters

1.000	0.035	0.155	-0.090
0.035	1.000	-0.204	0.000
0.155	-0.204	1.000	-0.030
-0.090	0.000	-0.030	1.000

Number of iterations 17
Minutes to convergence 0.03200

2.4 Managing Optimization

The critical elements in optimization are scaling, starting point, and the condition of the model. When the data are scaled, the starting point is reasonably close to the solution, and the data and model go together well, the iterations converge quickly and without difficulty.

For best results therefore, you want to prepare the problem so that model is well-specified, the data scaled, and that a good starting point is available.

The tradeoff among algorithms and step length methods is between speed and demands on the starting point and condition of the model. The less demanding methods are generally time consuming and computationally intensive, whereas the quicker methods (either in terms of time or number of iterations to convergence) are more sensitive to conditioning and quality of starting point.

2.4.1 Scaling

For best performance, the diagonal elements of the Hessian matrix should be roughly equal. If some diagonal elements contain numbers that are very large and/or very small with respect to the others, **MAXLIK** has difficulty converging. How to scale the diagonal elements of the Hessian may not be obvious, but it may suffice to ensure that the constants (or “data”) used in the model are about the same magnitude.

2.4.2 Condition

The specification of the model can be measured by the condition of the Hessian. The solution of the problem is found by searching for parameter values for which the

2. MAXIMUM LIKELIHOOD ESTIMATION

gradient is zero. If, however, the Jacobian of the gradient (i.e., the Hessian) is very small for a particular parameter, then **MAXLIK** has difficulty determining the optimal values since a large region of the function appears virtually flat to **MAXLIK**. When the Hessian has very small elements, the inverse of the Hessian has very large elements and the search direction gets buried in the large numbers.

Poor condition can be caused by bad scaling. It can also be caused by a poor specification of the model or by bad data. Bad models and bad data are two sides of the same coin. If the problem is highly nonlinear, it is important that data be available to describe the features of the curve described by each of the parameters. For example, one of the parameters of the Weibull function describes the shape of the curve as it approaches the upper asymptote. If data are not available on that portion of the curve, then that parameter is poorly estimated. The gradient of the function with respect to that parameter is very flat, elements of the Hessian associated with that parameter is very small, and the inverse of the Hessian contains very large numbers. In this case it is necessary to respecify the model in a way that excludes that parameter.

computer arithmetic

Computer arithmetic is fundamentally flawed by the fact that the computer number is finite (see Higham, 1996, for a general discussion). The standard double precision number in PC's carries about 16 decimal significant places. A simple operation can destroy nearly all of those places. The most destructive operation on a computer is addition and subtraction. Numbers are stored in a computer in the form of an **abscissa** and an **exponent**, e.g., 1.234567890123456e+02. There are about 16 decimal places of precision on most computers. The problem occurs when adding numbers that are of very different size. Before adding the number must be transformed so that the exponents are the same. For example consider adding 1.2345678901232456e-07 to 1.0000000000000000e+00:

```
1.0000000000000000e+00
0.0000001234567890e+00
-----
1.0000001234567890e+00
```

As you can see eight places were lost in the smaller number. If the exponent in the smaller number was 16 all of the places in that number would be lost.

This problem is due to the finiteness of the computer number, not to the implementation of the operators. It is an inherent problem in all computers and the only solution, adding more bits to the computer number, is only temporary because sooner or later a problem will arise where that quantity of bits won't be enough. The first lesson to be learned from this is to avoid operations combining very small numbers

2. MAXIMUM LIKELIHOOD ESTIMATION

with relatively large numbers. And for very small numbers, 1 can be a large number, as the example shows.

The standard method for evaluating the precision lost in computing a matrix inverse is the ratio of the largest to the smallest eigenvalue of the matrix. This quantity is sometimes called the condition number. The log of the condition number to the base 10 is approximately the number of decimal places lost in computing the inverse. A condition number greater than $1e16$ therefore indicates that all of the 16 decimal places are lost that are available in the standard double precision floating point number.

The BFGS optimization method in **MAXLIK** has been successful primarily because its method of generating an approximation to the Hessian encourages better conditioning. The implementation of the NEWTON method involves a numerical calculation of the Hessian. A numerical Hessian, like all numerical derivatives, are computed by first computing a difference, the most destructive operation as we've seen, and then compounding that by dividing the difference by a very small quantity. In general, when using double precision with 16 places of accuracy, about four places are lost in calculating a first derivative and another four with the second derivative. The numerical Hessian therefore begins with a loss of eight places of precision. If there are any problems computing the function itself, or if the model itself contains any problems of condition, there may be nothing left at all.

The BFGS method avoids much of the problems in computing a numerical Hessian. It produces an approximation by building information slowly with each iteration. Initially the Hessian is set to the identity matrix, the matrix with the best condition but the least information. Information is increased at each iteration with a method that guarantees a positive definite result. This provides for stabler, though slower, progress towards convergence.

The implementation of has been designed to minimize the damage to the precision of the optimization problem. The BFGS method avoids a direct calculation of the numerical Hessian, and uses sophisticated techniques for calculating the direction that preserve as much precision as possible. However, all of this can be defeated by a poorly scaled problem or a poorly specified model. When the objective function being optimized is a log-likelihood, the inverse of the Hessian is an estimate of the covariance matrix of the sampling distribution of the parameters. The condition of the Hessian is related to (i) the scaling of the parameters, and (ii) the degree with which there are linear dependencies in the sampling distribution of the parameters.

Scaling

Scaling is under the direct control of the investigator and should never be an issue in the optimization. It might not always be obvious how to do it, though. In estimation problems scaling of the parameters is usually implemented by scaling the data. in regression models this is simple to accomplish, but in more complicated models it might

2. MAXIMUM LIKELIHOOD ESTIMATION

be more difficult to do. It might be necessary to experiment with different scaling to get it right. The goal is to optimize the condition of the Hessian. The definition of the condition number implies that we endeavor to minimize the difference of the largest to the smallest eigenvalue of the Hessian. A rule of thumb for this is to scale the Hessian so that the diagonal elements are all about the same magnitude.

If the scaling of the Hessian proves too difficult, an alternative method is to scale the parameters directly in the procedure computing the log-likelihood. Multiply or divide the parameter values being passed to the procedure by setting quantities before their use in the calculation of the log-likelihood. Experiment with different values until the diagonal elements of the Hessian are all about the same magnitude.

Linear dependencies or nearly linear dependencies in the sampling distribution.

This is the most common difficulty in estimation and arises because of a discrepancy between the data and the model. If the data do not contain sufficient information to “identify” a parameter or set of parameters, a linear dependency is generated. A simple example occurs in regressors that cannot be distinguished from the constant because its variation is too small. When this happens, the sampling distribution of these two parameters becomes highly collinear. This collinearity will produce an eigenvalue approaching zero in the Hessian, increasing the number of places lost in the calculation of the inverse of the Hessian, degrading the optimization.

In the real world the data we have available will frequently fail to contain the information we need to estimate all of the parameters of our models. This means that it is a constant struggle to a well-conditioned estimation. When the condition sufficiently deteriorates to the point that the optimization fails, or the statistical inference fails through a failure to invert the Hessian, either more data must be found, or the model must be re-specified. Re-specification means either the direct reduction of the parameter space, that is, a parameter is deleted from the model, or some sort of restriction is applied to the parameters.

Diagnosing the linear dependency.

At times it may be very difficult to determine the cause of the ill-conditioning. If the Hessian being computed at convergence for the covariance matrix of the parameters fails to invert, try the following: first generate the pivoted QR factorization of the Hessian,

$$\{ R, E \} = \text{qre}(H);$$

The linearly dependent columns of H are pivoted to the end of the R matrix. E contains the new order of the columns of H after pivoting. The number of linearly

2. MAXIMUM LIKELIHOOD ESTIMATION

dependent columns is found by looking at the number of nearly zero elements at the end of the diagonal of R .

We can compute a coefficient matrix of the linear relationship of the dependent columns on the remaining columns by computing $R_{11}^{-1}R_{12}$ where R_{11} is that portion of the R matrix associated with the independent columns and R_{12} the independent with dependent. Rather than use the inverse function in **GAUSS**, we use a special solve function that takes advantage of the triangular shape of R_{11} . Suppose that the last two elements of R are nearly zero, then

```
r0 = rows(R);
r1 = rows(R) - 1;
r2 = rows(R) - 2;
B = utrisol(R[1:r2,r1:r0],R[1:r2,1:r2]);
```

B describes the linear dependencies among the columns of H and can be used to diagnose the ill-conditioning in the Hessian.

2.4.3 Starting Point

When the model is not particularly well-defined, the starting point can be critical. When the optimization doesn't seem to be working, try different starting points. A closed form solution may exist for a simpler problem with the same parameters. For example, ordinary least squares estimates may be used for nonlinear least squares problems or nonlinear regressions like probit or logit. There are no general methods for computing start values and it may be necessary to attempt the estimation from a variety of starting points.

2.4.4 Diagnosis

When the optimization is not proceeding well, it is sometimes useful to examine the function, the gradient Ψ , the direction δ , the Hessian Σ , the parameters θ_t , or the step length ρ , during the iterations. The current values of these matrices can be printed out or stored in the global **_max_Diagnostic** by setting **_max_Diagnostic** to a nonzero value. Setting it to 1 causes **MAXLIK** to print them to the screen or output file, 2 causes **MAXLIK** to store them in **_max_Diagnostic**, and 3 does both.

When you have selected **_max_Diagnostic** = 2 or 3, **MAXLIK** inserts the matrices into **_max_Diagnostic** using the **vput** command. The matrices are extracted using the **vread** command. For example,

```
_max_Diagnostic = 2;
call MAXPrt(maxlik("tobit",0,&lpr,x0));
h = vread(_max_Diagnostic,"hessian");
d = vread(_max_Diagnostic,"direct");
```

2. MAXIMUM LIKELIHOOD ESTIMATION

The following table contains the strings to be used to retrieve the various matrices in the **vread** command:

θ	“params”
δ	“direct”
Σ	“hessian”
Ψ	“gradient”
ρ	“step”

When nested calls to **MAXLIK** are made, i.e., when the procedure for computing the log-likelihood itself calls its own version of **MAXLIK**, **_max_Diagnostic** returns the matrices of the outer call to **MAXLIK** only.

2.5 Gradients

2.5.1 Analytical Gradient

To increase accuracy and reduce time, you may supply a procedure for computing the gradient, $\Psi(\theta) = \partial L / \partial \theta$, analytically.

This procedure has two input arguments, a $K \times 1$ vector of parameters and an $N_i \times L$ submatrix of the input data set. The number of rows of the data set passed in the argument to the call of this procedure may be less than the total number of observations when the data are stored in a **GAUSS** data set and there was not enough space to store the data set in RAM in its entirety. In that case subsets of the data set are passed to the procedure in sequence. The gradient procedure must be written to return a gradient (or more accurately, a “Jacobian”) with as many rows as the input submatrix of the data set. Thus the gradient procedure returns an $N_i \times K$ matrix of gradients of the N_i observations with respect to the K parameters. The **MAXLIK** global, **_max_GradProc** is then set to the pointer to that procedure. For example,

```
library maxlik;
#include maxlik.ext;
maxset;

proc lpsn(b,z); /* Function - Poisson Regression */
  local m;
  m = z[.,2:4]*b;
  retp(z[.,1].*m-exp(m));
endp;

proc lgd(b,z); /* Gradient */
  retp((z[.,1]-exp(z[.,2:4]*b)).*z[.,2:4]);
```


2. MAXIMUM LIKELIHOOD ESTIMATION

```
endp;

x0 = { .5, .5, .5 };
_max_GradProc = &lgd;
_max_GradCheckTol = 1e-3;

{ x,f0,g,h,retcode } = MAXLIK("psn",0,&lpsn,x0);
call MAXPrt(x,f0,g,h,retcode);
```

In practice, unfortunately, much of the time spent on writing the gradient procedure is devoted to debugging. To help in this debugging process, **MAXLIK** can be instructed to compute the numerical gradient along with your prospective analytical gradient for comparison purposes. In the example above this is accomplished by setting **_max_GradCheckTol** to 1e-3.

2.5.2 User-Supplied Numerical Gradient

You may substitute your own numerical gradient procedure for the one used by **MAXLIK** by default. This is done by setting the **MAXLIK** global, **_max_UserGrad** to a pointer to the procedure.

MAXLIK includes some numerical gradient functions in `gradient.src` which can be invoked using this global. One of these procedures, **gradre**, computes numerical gradients using the Richardson Extrapolation method. To use this method set

```
_max_UserNumGrad = &gradre;
```

2.5.3 Analytical Hessian

You may provide a procedure for computing the Hessian, $\Sigma(\theta) = \partial^2 L / \partial \theta \partial \theta'$. This procedure has two arguments, the $K \times 1$ vector of parameters, an $N_i \times L$ submatrix of the input data set (where N_i may be less than N), and returns a $K \times K$ symmetric matrix of second derivatives of the objection function with respect to the parameters.

The pointer to this procedure is stored in the global variable **_max_HessProc**.

In practice, unfortunately, much of the time spent on writing the Hessian procedure is devoted to debugging. To help in this debugging process, **MAXLIK** can be instructed to compute the numerical Hessian along with your prospective analytical Hessian for comparison purposes. To accomplish this **_max_GradCheckTol** is set to a small nonzero value.

2. MAXIMUM LIKELIHOOD ESTIMATION

```

library maxlik;
#include maxlik.ext;

proc lnk(b,z);
  local dev,s2;
  dev = z[.,1] - b[1] * exp(-b[2]*z[.,2]);
  s2 = dev' dev/rows(dev);
  retp(-0.5*(dev.*dev/s2 + ln(2*pi*s2)));
endp;

proc grdlk(b,z);
  local d,s2,dev,r;
  d = exp(-b[2]*z[.,2]);
  dev = z[.,1] - b[1]*d;
  s2 = dev' dev/rows(dev);
  r = dev.*d/s2;
/*  retp(r~(-b[1]*z[.,2].*r));          correct gradient */
  retp(r~(z[.,2].*r));          /* incorrect gradient */
endp;

proc hslk(b,z);
  local d,s2,dev,r, hss;
  d = exp(-b[2]*z[.,2]);
  dev = z[.,1] - b[1]*d;
  s2 = dev' dev/rows(dev);
  if s2 <= 0;
    retp(error(0));
  endif;
  r = z[.,2].*d.*(b[1].*d - dev)/s2;
  hss = -d.*d/s2~r~(-b[1].*z[.,2].*r);
  retp(xpnd(sumc(hss)));
endp;

maxset;
_max_HessProc = &hslk;
_max_GradProc = &grdlk;
_max_GradCheckTol = 1e-3;

startv = { 2, 1 };

{ x,f0,g,cov,retcode } = MAXLIK("nlls",0,&lnk,startv);
call MAXPrt(x,f0,g,cov,retcode);

```

The gradient is incorrectly computed, and **MAXLIK** responds with an error message. It is clear that the error is in the calculation of the gradient for the second parameter.

```

analytical and numerical gradients differ
      numerical      analytical
-0.015387035      -0.015387035
 0.031765317      -0.015882659

```

2. MAXIMUM LIKELIHOOD ESTIMATION

```
=====
analytical Hessian and analytical gradient
=====
MAXLIK Version 5.0.0                    5/30/2001 10:10 am
=====
Data Set:  nlls
-----

return code =      7
function cannot be evaluated at initial parameter values

Mean log-likelihood          1.12119
Number of cases              150

The covariance of the parameters failed to invert

Parameters    Estimates      Gradient
-----
P01           2.000000     -0.015387
P02           1.000000     -0.015883

Number of iterations      .
Minutes to convergence    .
```

2.5.4 User-Supplied Numerical Hessian

You may substitute your own numerical Hessian procedure for the one used by **MAXLIK** by default. This done by setting the **MAXLIK** global, **_max_UserHess** to a pointer to the procedure. This procedure has three input arguments, a pointer to the log-likelihood function, a $K \times 1$ vector of parameters, and an $N_i \times K$ matrix containing the data. It must return a $K \times K$ matrix which is the estimated Hessian evaluated at the parameter vector.

2.5.5 Switching Algorithms Automatically

The global variable **_max_Switch** can be used to switch algorithms automatically during the iterations. If **_max_Switch** has one column, the algorithm is switched once during the iterations, and if it has two columns it is switched back and forth. The conditions for the switching is determined by the elements of **_max_Switch** in the second through fourth rows. If these are rows are not supplied default values are entered. The first row contains the algorithm numbers to switch to, or if two columns to switch to and from. The algorithm switches if the log-likelihood function improves by less than the quantity in the second row, or if the number of iterations exceeds the

quantity in the third row, or if the line search changes by less than the quantity in the fourth row.

If only the first row is specified in the command file, that is, if only the algorithm numbers are entered, the second, third and fourth rows are set by default to .001, 10, .001 respectively.

2.6 FASTMAX – Fast Execution MAXLIK

Depending on the type of problem **FASTMAX**, the fast version of **MAXLIK**, can be called with speed-ups from 10 percent to 500 percent over the regular version of **MAXLIK**. This is achieved at the expense of losing some features, in particular, it won't print any iteration information to the screen, the globals cannot be modified on the fly, it can't print or store diagnostic information. Moreover, the dataset must be entirely storable in RAM.

The gain in time depends on the type of problem. The greatest speedup occurs with problems that are function call intensive. The speedup will be less if gradients and/or Hessians are provided. The least speedup occurs for problems where convergence is quick, and the most where convergence is slow. Thus **FASTMAX** will least affect a bootstrap or profile likelihood estimation for models that converge quickly, and most affect those that don't.

FASTMAX is most useful for problems that will be repeated in some way such as in a Monte Carlo study or a bootstrap. The initial runs would use **MAXLIK** where monitoring the progress is most important, and subsequent runs would use **FASTMAX**.

FASTMAX has the same arguments and returns as **MAXLIK** and thus to call it you may change the name **MAXLIK** in your command file to **FASTMAX**. **FASTMAX** does require that the dataset be storable in memory in its entirety, however, and if that isn't possible **FASTMAX** will fail.

In a similar way, for the fast versions of **MAXBOOT**, **MAXPROFILE**, and **MAXBAYES**, change the calls to **FASTBOOT**, **FASTPROFILE**, and **FASTBAYES**, respectively. No changes in input or output arguments are necessary.

2.6.1 Undefined Function Evaluation

On occasion the log-likelihood function will evaluate to an undefined value, for example, the log-likelihood procedure may attempt to take the log of a negative quantity for one or more observations. If you have written your procedure to return a scalar missing value when this happens, **MAXLIK** will succeed in recovering in most

2. MAXIMUM LIKELIHOOD ESTIMATION

cases. That is, depending on circumstances it will find another set of parameter values or use a different line search method.

If you are using **FASTMAX**, you can try a different strategy. Write your procedure to enter a missing value in the log-likelihood vector for that observation for which the calculation is undefined. **FASTMAX** will compute gradients and function values by list-wise deletion. In other words it will compute the function and gradient from the available observations.

2.7 Inference

MAXLIK includes four classes of methods for analyzing the distributions of the estimated parameters:

- Wald
- Profile likelihood
- Bootstrap
- Bayesian

The Wald type statistical inference is the most commonly used method which relies on a quadratic approximation to the log-likelihood surface, and uses an estimate of the covariance matrix of the parameters for computing standard errors and confidence limits. **MAXLIK** provides three methods for estimating the covariance matrix, the inverse of the Hessian, the inverse of the cross-products of the first derivatives, and the quasi-maximum likelihood (or QML) estimate which is computed from both the Hessian and the cross-product of the first derivatives.

The bootstrap and Bayesian methods both produce simulated “data” sets of the parameters from which kernel density plots, histograms, surface plots, and confidence limits may be computed.

The profile likelihood method computes confidence limits directly from the log-likelihood surface. Profile likelihood confidence limits are to be preferred to Wald confidence limits when the quadratic approximation is poor which is likely to be the case in particular for nonlinear models. The profile likelihood inference package includes a procedure for computing confidence limits as well as likelihood profile traces and profile t traces used for evaluating the shape of the log-likelihood surface.

2.7.1 Wald Inference

An argument based on a Taylor-series approximation to the likelihood function (e.g., Amemiya, 1985, page 111) shows that

$$\hat{\theta} \rightarrow N(\theta, A^{-1}BA^{-1})$$

2. MAXIMUM LIKELIHOOD ESTIMATION

where

$$A = E \left[\frac{\partial^2 L}{\partial \theta \partial \theta'} \right]$$

$$B = E \left[\left(\frac{\partial L}{\partial \theta} \right)' \left(\frac{\partial L}{\partial \theta} \right) \right]$$

Estimates of A and B are

$$\hat{A} = \frac{1}{N} \sum_i^N \frac{\partial^2 L_i}{\partial \theta \partial \theta'}$$

$$\hat{B} = \frac{1}{N} \sum_i^N \left(\frac{\partial L_i}{\partial \theta} \right)' \left(\frac{\partial L_i}{\partial \theta} \right)$$

Assuming the correct specification of the model $\text{plim}(A) = \text{plim}(B)$ and thus

$$\hat{\theta} \rightarrow N(\theta, \hat{A}^{-1})$$

When `__max__CovPar` = 1, \hat{A}^{-1} , the inverse of the Hessian, is returned as the covariance matrix of the parameters.

When `__max__CovPar` = 2, **MAXLIK** returns \hat{B}^{-1} , the cross-product of the first derivatives computed by observation (i.e., the “Jacobian” of the log-likelihood) as the covariance matrix of the parameters.

When `__max__CovPar` is set to 3, **MAXLIK** returns $\hat{A}^{-1} \hat{B} \hat{A}^{-1}$, the QML covariance matrices of the parameters.

When the QML method has been selected, the covariance matrices computed from the Hessian and the cross-product of first derivatives will both be returned in the global variables, `__max__HessCov` and `__max__XprodCov`, respectively. A rough measure of the misspecification in the model may be gauged from the extent to which the covariance matrices computed from the Hessian and the cross-product of first derivatives diverge. A method for computing a statistic to measure this divergence (thereby providing a test for misspecification) has been developed by White (1981,1982).

The QML covariance matrix is expensive to compute since it requires the calculation of both the matrix of second derivatives and the first derivatives by case. The expense will usually be worth it, however, because this matrix will always generate the correct standard errors (unless there is a misspecification in the model that renders the parameter estimates inconsistent in which case no method will produce correct standard errors). To determine whether either the Hessian or the cross-product covariance matrix of parameters are sufficiently correct by themselves it would be necessary to compute them both anyway.

2. MAXIMUM LIKELIHOOD ESTIMATION

When Computing the Covariance Matrix of the Parameters Fails

The computation of the covariance matrix of the parameters may fail if there is not enough information in the data to identify the model parameters, or if the model specification includes parameters that cannot be identified for any set of data. In these cases there may be some utility in a collinearity analysis of the matrix used in the computation of the covariance matrix of the parameters. This matrix is stored in the global variable `_max_FinalHess` before the inversion attempt. If the inversion fails (of the Hessian if `_max_CovPar = 1`, or of the cross-product of the first derivatives if `_max_CovPar = 2`), **MAXLIK** will return a missing code for the covariance matrix and the user can then retrieve the matrix stored in `_max_FinalHess` for a collinearity analysis. Linear dependencies in this matrix will indicate which parameters are not identified and an analysis of these linear dependencies may suggest tactics for respecifying the model.

MaxLik

2.7.2 Profile Likelihood Inference

Wald confidence limits for parameters assume the appropriateness of the quadratic approximation to the log-likelihood surface. For some models, in particular nonlinear models, this approximation may not be satisfactory. In this case, the profile likelihood confidence limit would be preferred.

The profile likelihood confidence region is defined as the set of points (Cook and Wiesberg, 1990, Meeker and Escobar, 1995):

$$\{\theta \mid \sqrt{2(L(\hat{\theta}) - L(\theta))} \geq \chi^2_{(1-\alpha; k)}\}$$

where

$$L(\theta) = \sum_{i=1}^N \log P(Y_i; \theta)$$

and K is the length of θ .

For individual parameters this method is implemented in **MAXLIK** in the following way: define

$$G(\phi) = \min(\text{Logl}(\theta) \mid \eta'_i \theta = \phi) \quad (2.1)$$

where η_i is a conformable vector of zeros with a one in position i .

Then the lower profile likelihood confidence limit at the $1 - \alpha$ interval are the values of ϕ such that

$$G(\phi) = \chi^2_{(1-\alpha; k)}.$$

and the upper limit is found by redefining Equation 2.1 as a maximum.

2. MAXIMUM LIKELIHOOD ESTIMATION

Example

This examples illustrates and compares Wald confidence limits and profile likelihood confidence limits:

```
library maxlik;
#include maxlik.ext;
maxset;

proc lpr(x,z);
  local t,s,m,u;
  s = x[4];
  if s <= 1e-4;
    retp(error(0));
  endif;
  m = z[.,2:4]*x[1:3,.];
  u = z[.,1] ./= 0;
  t = z[.,1]-m;
  retp(u.*(-(t.*t)/(2*s)-.5*ln(2*s*pi)) +
      (1-u).*(ln(cdfnc(m/sqrt(s))))
    );
endp;

x0 = { 1, 1, 1, 1 };

{f,f,g,cov,ret} = maxlik("tobit",0,&lpr,x0);

__title = "Wald Confidence Limits";
cl1 = maxtlimits(x,cov);
call maxclprt(x,f,g,cl1,ret);

__title = "Profile Likelihood Confidence Limits";
cl2 = maxpflclimits(x,f,"tobit",0,&lpr);
call maxclprt(x,f,g,cl2,ret);
```

The output is:

```
=====
                                Wald Confidence Limits
=====
MAXLIK Version 5.0.0                                5/30/2001  1:16 pm
=====
                                Data Set:  tobit
-----

return code =      0
normal convergence
```


2. MAXIMUM LIKELIHOOD ESTIMATION

Mean log-likelihood -1.13291
 Number of cases 100

Parameters	Estimates	0.95 confidence limits		Gradient
		Lower Limit	Upper Limit	
P01	0.0104	-0.1573	0.1781	-0.0000
P02	-0.2081	-0.3958	-0.0203	-0.0000
P03	-0.0998	-0.2588	0.0593	-0.0000
P04	0.6522	0.4540	0.8505	-0.0000

Number of iterations 17
 Minutes to convergence 0.03200



```

=====
                          Profile Likelihood Confidence Limits
=====
MAXLIK Version 5.0.0                               5/30/2001  1:16 pm
=====
                          Data Set:  tobit
=====
    
```

return code = 0
 normal convergence

Mean log-likelihood -1.13291
 Number of cases 100

Parameters	Estimates	0.95 confidence limits		Gradient
		Lower Limit	Upper Limit	
P01	0.0104	-0.1560	0.1720	-0.0000
P02	-0.2081	-0.3918	-0.0245	-0.0000
P03	-0.0998	-0.2562	0.0549	-0.0000
P04	0.6522	0.4928	0.8885	-0.0000

Number of iterations 17
 Minutes to convergence 0.03200

In this example, the model is conditionally linear and we see that the Wald and profile likelihood limits are quite similar.

2.7.3 Profile Trace Plots

MAXProfile generates profile t plots as well as plots of the likelihood profile traces for all of the parameters in the model in pairs. The profile t plots are used to assess the nonlinearity of the distributions of the individual parameters, and the likelihood profile traces are used to assess the bivariate distributions. The input and output arguments to **MAXProfile** are identical to those of **MAXLIK**. But in addition to providing the maximum likelihood estimates and covariance matrix of the parameters, a series of plots are printed to the screen using **GAUSS**' Publication Quality Graphics. A screen is printed for each possible pair of parameters. There are three plots, a profile t plot for each parameter, and a third plot containing the likelihood profile traces for the two parameters.

The discussion in this section is based on Bates and Watts (1988), pages 205-216, which is recommended reading for the interpretation and use of profile t plots and likelihood profile traces.

The Profile t Plot

Define

$$\tilde{\theta}_k = (\tilde{\theta}_1, \tilde{\theta}_2, \dots, \tilde{\theta}_{k-1}, \theta_k, \tilde{\theta}_{k+1}, \dots, \tilde{\theta}_K)$$

This is the vector of maximum likelihood estimates *conditional* on θ_k , i.e., where θ_k is fixed to some value. Further define the profile t function

$$\tau(\theta_k) = \text{sign}(\theta_k - \hat{\theta}_k)(N - K) \sqrt{2 [L(\tilde{\theta}_k) - L(\hat{\theta}_k)]}$$

For each parameter in the model, τ is computed over a range of values for θ_k . These plots provide exact likelihood intervals for the parameters, and reveal how nonlinear the estimation is. For a linear model, τ is a straight line through the origin with unit slope. For nonlinear models, the amount of curvature is diagnostic of the nonlinearity of the estimation. High curvature suggests that the usual statistical inference using the t-statistic is hazardous.

The Likelihood Profile Trace

The likelihood profile traces provide information about the bivariate likelihood surfaces. For nonlinear models the profile traces are curved, showing how the parameter

2. MAXIMUM LIKELIHOOD ESTIMATION

estimates affect each other and how the projection of the likelihood contours onto the (θ_k, θ_ℓ) plane might look. For the (θ_k, θ_ℓ) plot, two lines are plotted, $L(\tilde{\theta}_k)$ against θ_k and $L(\tilde{\theta}_\ell)$ against θ_ℓ .

If the likelihood surface contours are long and thin, indicating the parameters to be collinear, the profile traces are close together. If the contours are fat, indicating the parameters to be more uncorrelated, the profile traces tend to be perpendicular. And if the contours are nearly elliptical, the profile traces are straight. The surface contours for a linear model would be elliptical and thus the profile traces would be straight and perpendicular to each other. Significant departures of the profile traces from straight, perpendicular lines, therefore, indicate difficulties with the usual statistical inference.

To generate profile t plots and likelihood profile traces from the example in Section 2.3.9, it is necessary only to change the call to **MAXLIK** to a call to **MAXProfile**:

```
call MAXPrt(MAXProfile("tobit",0,&lpr,x0));
```

MAXProfile produces the same output as **MAXLIK** which can be printed out using a call to **MAXPRT**.

For each pair of parameters a plot is generated containing an xy plot of the likelihood profile traces of the two parameters, and two profile t plots, one for each parameter.

2.7.4 Bootstrap

The bootstrap method is used to generate empirical distributions of the parameters, thus avoiding the difficulties with the usual methods of statistical inference described above.

MAXBoot

Rather than randomly sample with replacement from the data set, **MAXBoot** performs **__max_NumSample** weighted maximum likelihood estimations where the weights are Poisson pseudo-random numbers with expected value equal to the the number of observations. **__max_NumSample** is set by the **MAXBoot** global variable. The default is 100 re-samplings. Efron and Tibshirani (1993:52) suggest that 100 is satisfactory, and rarely are more than 200 needed.

The mean and covariance matrix of the bootstrapped parameters is returned by **MAXBoot**. In addition **MAXBoot** writes the bootstrapped parameter estimates to a **GAUSS** data set for use with **MAXHist**, which produces histograms and surface plots, **MAXDensity**, which produces kernel density plots, and **MAXLimits**, which produces confidence limits based on the bootstrapped coefficients. The data set name can be specified by the user in the global **__max_BootFname**. However, if not specified, **MAXBoot** selects a temporary filename.

MAXDensity

MAXDensity is a procedure for computing kernel type density plots. The global, **__max__Kernel** permits you to select from a variety of kernels, normal, Epanechnikov, biweight, triangular, rectangular, and truncated normal. For each selected parameter, a plot is generated of a smoothed density. The smoothing coefficients may be specified using the global, **__max__Smoothing**, or **MAXDensity** will compute them.

MAXHist

MAXHist is a procedure for visually displaying the results of the bootstrapping in univariate histograms and bivariate surface plots for selected parameters. The univariate discrete distributions of the parameters used for the histograms are returned by **MAXHist** in a matrix.

Example

To bootstrap the example in Section 2.3.9, the only necessary alteration is the change the call to **MAXLIK** to a call to **MAXBoot**:

```

_max_BootFname = "bootdata";

call MAXPrt(maxlikboot("tobit",0,&lpr,x0));

call MAXDensity("bootdata",0);
call MAXHist("bootdata",0);

```

2.7.5 Pseudo-Random Number Generators

Pseudo-Random numbers are generated by **MAXLIK** and **FASTMAX** in the random line search, by **MAXBoot** and **FASTBoot** for re-sampling, and by **MAXBayes** and **FASTBayes** also for re-sampling. There two types of pseudo-random generators, the linear congruential (LC) and another based on Marsaglia's **Kiss-Monster** algorithm (KM). The LC generators are faster but have shorter period (2^{32}), whereas the KM generators are slower but have much longer periods (2^{3859}).

The global variable **__max__RandType** chooses between these. By default the LC generators are used.

The seed for these generators is kept in the global **__max__State**. The default value is 345678. You may set this to any integer value in your command file.

2. MAXIMUM LIKELIHOOD ESTIMATION

2.7.6 Bayesian Inference

The **MAXLIK** proc **MAXBayes** generates a simulated posterior of the parameters of a maximum likelihood estimation using the weighted likelihood bootstrap method described in Newton and Raftery (1994). In this method, a weighted bootstrap is conducted using weighted Dirichlet random variates for weights. After generating the weighted bootstrapped parameters, “Importance” weights are computed:

$$r(\hat{\theta}) = \pi(\hat{\theta})e^{L(\hat{\theta})}/\hat{g}(\hat{\theta})$$

where $\pi(\hat{\theta})$ is the prior distribution of the parameters, and $\hat{g}(\hat{\theta})$ is a normal kernel density estimate of the of the parameters using Terrell’s (1990) method of maximum smoothing. The SIR algorithm, described in Rubin (1988), is applied to the bootstrapped parameters using these importance weights.

The Dirichlet variates are weighted to generate over-dispersion in order to make sure they have coverage with respect to the posterior distribution. This weight is stored in the **MAXLIK** global, **_max_BayesAlpha**, and is set to 1.4 by default. See Newton and Raftery (1994) for a discussion of this weight.

Example

This example computes ordinary maximum likelihood estimates, and then calls **MAXBayes** which generates a simulated posterior. The call to **MAXDensity** produces kernel density plots and returns the data used in the plots. This information is used to determine the modes of the simulated posterior distributions and **MAXPrt** prints that information to output.

```
library maxlik,pgraph;
#include maxlik.ext;
#include pgraph.ext;
graphset;
maxset;

proc lpr(x,z);
  local t,s,m,u;
  s = x[4];
  if s <= 1e-4;
    retp(error(0));
  endif;
  m = z[.,2:4]*x[1:3,.];
  u = z[.,1] ./= 0;
  t = z[.,1]-m;
  retp(u.*(-(t.*t)/(2*s)-.5*ln(2*s*pi)) + (1-u).*(ln(cdfnc(m/sqrt(s))))));
endp;
```

2. MAXIMUM LIKELIHOOD ESTIMATION

```
start = { 1, 1, 1, 1 };

__title = "Maximum Likelihood Estimates";
{x0,f,g,cov,ret} = maxlik("tobit",0,&lpr,start);
call maxprt(x0,f,g,cov,ret);

_max_BootFname = "bayes";
_max_NumSample = 500;

{x1,f,g,cov,ret} = maxBayes("tobit",0,&lpr,x0);

{ px,py,smth } = maxDensity("bayes",0);
x_mode = diag(px[maxindc(py),.]);

__title = "modal Bayesian estimates";
call maxprt(x_mode,f,g,cov,ret);
```

```
=====
                          Maximum Likelihood Estimates
=====
MAXLIK Version 5.0.0                                5/30/2001 11:18 am
=====
                          Data Set:  tobit
-----
```

```
return code = 0
normal convergence
```

```
Mean log-likelihood      -1.13291
Number of cases         100
```

```
Covariance matrix of the parameters computed by the following method:
Inverse of computed Hessian
```

Parameters	Estimates	Std. err.	Est./s.e.	Prob.	Gradient
P01	0.0104	0.0873	0.119	0.4525	0.0000
P02	-0.2081	0.0946	-2.200	0.0139	0.0000
P03	-0.0998	0.0800	-1.247	0.1062	0.0000
P04	0.6522	0.0999	6.531	0.0000	0.0000

```
Correlation matrix of the parameters
```

1.000	0.030	0.151	-0.092
0.030	1.000	-0.205	0.000
0.151	-0.205	1.000	-0.029
-0.092	0.000	-0.029	1.000

2. MAXIMUM LIKELIHOOD ESTIMATION

Number of iterations 17
Minutes to convergence 0.01462

```
=====
                        modal Bayesian estimates
=====
MAXLIK Version 5.0.0                                5/30/2001 11:20 am
=====
                        Data Set:  tobit
=====
```

return code = 0
normal convergence

Mean log-likelihood -0.0117326
Number of cases 100

Covariance matrix of the parameters computed by the following method:
Bayesian covariance matrix

Parameters	Estimates	Std. err.	Est./s.e.	Prob.	Gradient
P01	0.1729	0.1661	1.041	0.1488	0.0000
P02	-0.2054	0.1930	-1.065	0.1435	0.0000
P03	-0.1425	0.1735	-0.821	0.2057	0.0000
P04	0.6598	0.2329	2.833	0.0023	0.0000

Correlation matrix of the parameters

1.000	-0.165	0.362	0.465
-0.165	1.000	-0.473	0.026
0.362	-0.473	1.000	0.322
0.465	0.026	0.322	1.000

Number of iterations 7
Minutes to convergence 0.00354

2.8 Run-Time Switches

If the user presses **Alt-H** during the iterations, a help table is printed to the screen which describes the run-time switches. By this method, important global variables may be modified during the iterations.

Alt-G	Toggle <code>__max__GradMethod</code>
Alt-V	Revise <code>__max__GradTol</code>
Alt-O	Toggle <code>__output</code>
Alt-M	Maximum Tries
Alt-I	Compute Hessian
Alt-E	Edit Parameter Vector
Alt-C	Force Exit
Alt-A	Change Algorithm
Alt-J	Change Line Search Method
Alt-H	Help Table

The algorithm may be switched during the iterations either by pressing **Alt-A**, or by pressing one of the following:

Alt-1	Steepest Descent (STEEP)
Alt-2	Broyden-Fletcher-Goldfarb-Shanno (BFGS)
Alt-3	Davidon-Fletcher-Powell (DFP)
Alt-4	Newton-Raphson (NEWTON) or (NR)
Alt-5	Berndt, Hall, Hall & Hausman (BHHH)
Alt-6	Polak-Ribiere Conjugate Gradient (PRCG)

The line search method may be switched during the iterations either by pressing **Alt-S**, or by pressing one of the following:

Shift-1	no search (1.0 or 1 or ONE)
Shift-2	cubic or quadratic method (STEPBT)
Shift-3	step halving method (HALF)
Shift-4	Brent's method (BRENT)
Shift-5	BHHH step method (BHHHSTEP)

2.9 Calling MAXLIK Recursively

The procedure that computes the log-likelihood may itself call **MAXLIK**. This version of **MAXLIK** nested inside the procedure is actually a separate copy of **MAXLIK** with its own set of globals and must have its own log-likelihood function (or otherwise you would have infinite recursion).

When calling **MAXLIK** recursively, the following considerations apply:

- Variable selection (as opposed to case selection) can be done on any level by means of the second argument in the call to each copy of **MAXLIK**.
- Data sets can be opened by nested copies of **MAXLIK**. If a nested copy of **MAXLIK** is going to use the data set opened by the outer copy of **MAXLIK**, then pass a null string in the first argument in the call. If it is

2. MAXIMUM LIKELIHOOD ESTIMATION

going to analyze a different data set from the outer copy, then pass it the data set name in a string. You may also load and store a data set in memory in the command file and pass it in the first argument in the nested call to **MAXLIK**.

- Before the call to the nested copy of **MAXLIK**, the global variables should be reset by calling **MAXCLR**. You must not use **MAXSET** because that will clear information about the data sets opened and processed in the outer copy. The only differences between **MAXSET** and **MAXCLR** are references to these globals.
- You may also want to disable the keyboard control of the nested copies. This is done by setting the global `__max_Key = 0` after the call to **MAXCLR** and before the call to the nested **MAXLIK**.

2.10 Using `__MAXLIK` Directly

When **MAXLIK** is called, it directly references all the necessary globals and passes its 4 arguments and the values of the globals to a function called `__maxlik`. When `__maxlik` returns, **MAXLIK** then sets the output globals to the values returned by `__maxlik` and returns 5 arguments directly to the user. `__maxlik` makes no global references to matrices or strings (except to `__max_eps2` which is set to the cube of machine precision), and all procedures it references have names that begin with an underscore “_”.

`__maxlik` can be used directly in situations where you do not want any of the global matrices and strings in your program. If **MAXLIK**, **MAXPRT**, **MAXSET**, and **MAXCLR** are not referenced, the global matrices and strings in `maxlik.dec` will not be included in your program.

The documentation for **MAXLIK**, the globals it references, and the code itself should be sufficient documentation for using `__maxlik`.

2.11 Error Handling

2.11.1 Return Codes

The fourth argument in the return from **MAXLIK** contains a scalar number that contains information about the status of the iterations upon exiting **MAXLIK**. The following table describes their meanings:

2. MAXIMUM LIKELIHOOD ESTIMATION

0	normal convergence
1	forced exit
2	maximum iterations exceeded
3	function calculation failed
4	gradient calculation failed
5	Hessian calculation failed
6	line search failed
7	function cannot be evaluated at initial parameter values
8	error with gradient
9	gradient vector transposed
10	secant update failed
11	maximum time exceeded
12	error with weights
20	Hessian failed to invert
34	data set could not be opened
99	termination condition unknown

2.11.2 Error Trapping

Setting the global `___output = 0` turns off all printing to the screen. Error codes, however, still are printed to the screen unless error trapping is also turned on. Setting the trap flag to 4 causes **MAXLIK** to *not* send the messages to the screen:

```
trap 4;
```

Whatever the setting of the trap flag, **MAXLIK** discontinues computations and returns with an error code. The trap flag in this case only affects whether messages are printed to the screen or not. This is an issue when the **MAXLIK** function is embedded in a larger program, and you want the larger program to handle the errors.

2.12 References

- Amemiya, Takeshi, 1985. *Advanced Econometrics*. Cambridge, MA: Harvard University Press.
- Bates, Douglas M. and Watts, Donald G., 1988. *Nonlinear Regression Analysis and Its Applications*. New York: John Wiley & Sons.
- Berndt, E., Hall, B., Hall, R., and Hausman, J. 1974. "Estimation and inference in nonlinear structural models". *Annals of Economic and Social Measurement* 3:653-665.

2. MAXIMUM LIKELIHOOD ESTIMATION

- Brent, R.P., 1972. *Algorithms for Minimization Without Derivatives*. Englewood Cliffs, NJ: Prentice-Hall.
- Cook, R.D. and Weisberg, S., 1990. "Confidence Curves in Nonlinear Regression", *Journal of the American Statistical Association*, 85: 544-551.
- Dennis, Jr., J.E., and Schnabel, R.B., 1983. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Englewood Cliffs, NJ: Prentice-Hall.
- Efron, Bradley, Robert J. Tibshirani, 1993. *An Introduction to the Bootstrap*. New York: Chapman & Hall.
- Gill, P. E. and Murray, W. 1972. "Quasi-Newton methods for unconstrained optimization." *J. Inst. Math. Appl.*, 9, 91-108.
- Judge, G.G., R.C. Hill, W.E. Griffiths, H. Lütkepohl and T.C. Lee. 1988. *Introduction to the Theory and Practice of Econometrics*. 2nd Edition. New York:Wiley.
- Judge, G.G., W.E. Griffiths, R.C. Hill, H. Lütkepohl and T.C. Lee. 1985. *The Theory and Practice of Econometrics*. 2nd Edition. New York:Wiley.
- Meeker, W.Q and Escobar, L.A., 1995. "Teaching about approximate confidence regions based on maximum likelihood estimate", *The American Statistician*, 49: 48-53.
- Newton, M.A. and A.E. Raftery, 1994. "Approximate Bayesian inference with the weighted likelihood bootstrap", *J.R. Statist. Soc. B*, 56: 3-48.
- Rubin, D.B., 1988. "Using the SIR algorithm to simulate posterior distributions", in *Bayesian Statistics 3*, J.M. Bernardo, M.H. DeGroot, D.V. Lindley, and A.F.M. Smith (eds.), pp. 395-402.
- Terrell, G.R., 1990. "The maximal smoothing principle in density estimation", *Journal of the American Statistical Association*, 85: 470-477.
- White, H. 1981. "Consequences and detection of misspecified nonlinear regression models." *Journal of the American Statistical Association* 76:419-433.
- White, H. 1982. "Maximum likelihood estimation of misspecified models." *Econometrica* 50:1-25.

2. *MAXIMUM LIKELIHOOD ESTIMATION*

Chapter 3

Maximum Likelihood Reference

■ Purpose

Computes estimates of parameters of a maximum likelihood function

■ Library

maxlik

■ Format

{ *x,f,g,cov,retcode* } = **FASTMAX**(*data,vars,&fct,start*)

■ Input

data $N \times NV$ matrix, data

vars $NV \times 1$ character vector, labels of variables selected for analysis

– or –

$NV \times 1$ numeric vector, indices of variables selected for analysis.

vars may be a character vector containing either the standard labels created by **FASTMAX** (i.e., either V1, V2, ..., or V01, V02, See discussion of the global variable **__vpad** below, or the user-provided labels in **__altnam**).

&fct a pointer to a procedure that returns either the log-likelihood for one observation or a vector of log-likelihoods for a matrix of observations (see discussion of the global variable **__row** in global variable section below).

start $K \times 1$ vector, start values.

■ Output

x $K \times 1$ vector, estimated parameters

f scalar, function at minimum (the mean log-likelihood)

g $K \times 1$ vector, gradient evaluated at *x*

h $K \times K$ matrix, covariance matrix of the parameters (see discussion of the global variable **__max__CovPar** below).

retcode scalar, return code. If normal convergence is achieved, then *retcode* = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:

0 normal convergence

- 1 forced exit.
- 2 maximum iterations exceeded.
- 3 function calculation failed.
- 4 gradient calculation failed.
- 5 Hessian calculation failed.
- 6 line search failed.
- 7 function cannot be evaluated at initial parameter values.
- 8 error with gradient
- 9 gradient vector transposed
- 10 secant update failed
- 11 maximum time exceeded
- 12 error with weights
- 34 data set could not be opened.
- 99 termination condition unknown.

■ Globals

The global variables used by **FASTMAX** can be organized in the following categories according to which aspect of the optimization they affect:

Options `_max_Options`

Descent and Line Search `_max_Algorithm`, `_max_Delta`, `_max_LineSearch`,
`_max_Maxtry`, `_max_Extrap`, `_max_Interp`, `_max_RandRadius`,
`_max_Switch`, `_max_RandType`, `_max_State`,

Covariance Matrix of Parameters `_max_CovPar`, `_max_XprodCov`,
`_max_HessCov`, `_max_FinalHess`

Gradient `_max_GradMethod`, `_max_GradProc`, `_max_UserNumGrad`,
`_max_HessProc`, `_max_GradStep`,

Terminations Conditions `_max_GradTol`, `_max_MaxIters`, `_max_MaxTime`

Data `_max_NumObs`, `__weight`,

Parameters `_max_Active`, `_max_ParNames`

Miscellaneous `__title`, `_max_IterData`,

The list below contains an alphabetical listing of each global with a complete description.

__max__Active vector, defines fixed/active coefficients. This global allows you to fix a parameter to its starting value. This is useful, for example, when you wish to try different models with different sets of parameters without having to re-edit the function. When it is to be used, it must be a vector of the same length as the starting vector. Set elements of `'max'Active` to 1 for an active parameter, and to zero for a fixed one.

__max__Algorithm scalar, selects optimization method:

- 1 STEEP - Steepest Descent
- 2 BFGS - Broyden, Fletcher, Goldfarb, Shanno method
- 3 DFP - Davidon, Fletcher, Powell method
- 4 NEWTON - Newton-Raphson method
- 5 BHHH - Berndt, Hall, Hall, Hausman method
- 6 PRCG - Polak-Ribiere Conjugate Gradient

Default = 3

__max__CovPar scalar, type of covariance matrix of parameters

- 0 not computed
- 1 computed from Hessian calculated after the iterations
- 2 computed from cross-product of Jacobian
- 3 Quasi-maximum likelihood (QML) covariance matrix of the parameters

Default = 1;

__max__Delta scalar, floor for eigenvalues of Hessian in the NEWTON algorithm. When nonzero, the eigenvalues of the Hessian are augmented to this value.

__max__GradTol scalar, convergence tolerance for gradient of estimated coefficients. When this criterion has been satisfied **FASTMAX** exits the iterations. Default = 1e-5.

__max__Extrap scalar, extrapolation constant in BRENT. Default = 2.

__max__FinalHess $K \times K$ matrix, the Hessian used to compute the covariance matrix of the parameters is stored in **__max__FinalHess**. This is most useful if the inversion of the hessian fails, which is indicated when **FASTMAX** returns a missing value for the covariance matrix of the parameters. An analysis of the Hessian stored in **__max__FinalHess** can then reveal the source of the linear dependency responsible for the singularity.

__max__GradMethod scalar, method for computing numerical gradient.

- 0 central difference
- 1 forward difference (default)

__max_GradProc scalar, pointer to a procedure that computes the gradient of the function with respect to the parameters. For example, the statement:

```
__max_GradProc=&gradproc;
```

tells **FASTMAX** that a gradient procedure exists as well where to find it. The user-provided procedure has two input arguments, an $K \times 1$ vector of parameter values and an $N \times K$ matrix of data. The procedure returns a single output argument, an $N \times K$ matrix of gradients of the log-likelihood function with respect to the parameters evaluated at the vector of parameter values.

For example, suppose the log-likelihood function is for a Poisson regression, then the following would be added to the command file:

```
proc lgd(b,z);
  retp((z[.,1]-exp(z[.,2:4]*b)).*z[.,2:4]);
endp;
```

```
__max_GradProc = &lgd;
```

Default = 0, i.e., no gradient procedure has been provided.

__max_GradStep scalar, increment size for computing gradient. When the numerical gradient is performing well, set to a larger value (1e-3, say). Default is the cube root of machine precision.

__max_HessCov $K \times K$ matrix. When **__max_CovPar** is set to 3 the information matrix covariance matrix of the parameters, i.e., the inverse of the matrix of second order partial derivatives of the log-likelihood by observations, is returned in **__max_HessCov**.

__max_HessProc scalar, pointer to a procedure that computes the hessian, i.e., the matrix of second order partial derivatives of the function with respect to the parameters. For example, the instruction:

```
__max_HessProc = &hessproc;
```

tells **FASTMAX** that a procedure has been provided for the computation of the hessian and where to find it. The procedure that is provided by the user must have two input arguments, a $K \times 1$ vector of parameter values and an $N \times P$ data matrix. The procedure returns a single output argument, the $K \times K$ symmetric matrix of second order derivatives of the function evaluated at the parameter values.

__max_Interp scalar, interpolation constant in BRENT. Default = .25.

__max_IterData 3x1 vector, contains information about the iterations.

The first element contains the # of iterations, the second element contains the elapsed time in minutes of the iterations, and the third element contains a character variable indicating the type of covariance matrix of the parameters.

__max_LineSearch scalar, selects method for conducting line search. The result of the line search is a *step length*, i.e., a number which reduces the function value when multiplied times the direction..

- 1 step length = 1.
- 2 cubic or quadratic step length method (STEPBT)
- 3 step halving (HALF)
- 4 Brent's step length method (BRENT)
- 5 BHHH step length method (BHHHSTEP)

Default = 2.

Usually **__max_LineSearch** = 2 is best. If the optimization bogs down, try setting **__max_LineSearch** = 1, 4 or 5. **__max_LineSearch** = 3 generates slower iterations but faster convergence and **__max_LineSearch** = 1 generates faster iterations but slower convergence.

When any of these line search methods fails, **FASTMAX** attempts a random search of radius **__max_RandRadius** times the truncated log to the base 10 of the gradient when **__max_RandRadius** is set to a nonzero value.

__max_MaxItrs scalar, maximum number of iterations.

__max_MaxTime scalar, maximum time in iterations in minutes. This global is most useful in bootstrapping. You might want 100 re-samples, but would be happy with anything more than 50 depending on the time it took. Set **__max_NumSample** = 100, and **__max_MaxTime** to maximum time you would be willing to wait for results. Default = 1e+5, about 10 weeks.

__max_MaxTry scalar, maximum number of tries to find step length that produces a descent.

__max_NumObs scalar, number of cases in the data set that was analyzed.

__max_Options character vector, specification of options. This global permits setting various **FASTMAX** options in a single global using identifiers. The following

```
__max_Options = { bfgs stepbt forward };
```

sets to the default values, i.e. the descent method to BFGS, the line search method to STEPBT, the numerical gradient method to central differences.

The following is a list of the identifiers:

Algorithms STEEP, BFGS, DFP, NEWTON, BHHH, PRCG

Line Search ONE, STEPBT, HALF, BRENT, BHHHSTEP

Covariance Matrix NOCOV, INFO, XPROD, HETCON

Gradient method CENTRAL, FORWARD

__max__ParNames $K \times 1$ character vector, parameter labels.

__max__RandRadius scalar, if set to a nonzero value (1e-2, say) and all other line search methods fail then **FASTMAX** attempts **__max__MaxTry** tries to find a random direction within radius determined by **__max__RandRadius** that is a descent. Default = 1e-2.

__max__RandType scalar, if nonzero, pseudo-random numbers of the linear congruential type are generated, otherwise, they are generated by Marsaglia's **Kiss-Monster** method. The latter method is slower but has a much larger period. Random numbers are generated for the random line search.

__max__State scalar or vector, state vector for pseudorandom number generators containing seed. By default it is set to 345678. If you wish to select a seed, set to a different value.

__max__Switch 4×1 or 4×2 vector, controls algorithm switching. If **__max__Switch** is 4×1 , set its elements in the following way,

- 1 , algorithm number to switch to
- 2 , **FASTMAX** will switch to algorithm in the first element when the function value is less than the value entered here
- 3 **FASTMAX** switches if the number of iterations exceeds the number entered here
- 4 **FASTMAX** switches if line search step changes less than the amount entered here

If **__max__Switch** is 4×2 , **FASTMAX** switches between the algorithms in column 1 and column 2 subject to the conditions specified for the 4×1 vector.

Thus if **__max__Switch** is a 4×1 vector, **FASTMAX** will switch algorithms no more than once during the iterations, whereas if it is 4×2 it may switch back and forth between the two algorithms throughout the iterations.

- __title** string title of run
- __weight** vector, frequency of observations. By default all observations have a frequency of 1. zero frequencies are allowed. It is assumed that the elements of **__weight** sum to the number of observations.
- __max_XprodCov** $K \times K$ matrix. When **__max_CovPar** is set to 3 the cross-product matrix covariance matrix of the parameters, i.e., the inverse of the cross-product of the first derivatives of the log-likelihood computed by observations, is returned in **__max_XprodCov**.

■ Remarks

Writing the Log-likelihood Function

The user must provide a procedure for computing the log-likelihood for a matrix of observations. The procedure must have two input arguments: first, a vector of parameter values, and second, the data matrix. The output argument is the log-likelihood for the observations in the second argument evaluated at the parameter values in the first argument. Suppose that the function procedure has been named *pfct*, the following considerations apply:

The format of the procedure is:

```
logprob = pfct(x,y);
```

where

x column vector of parameters of model

y data

The output from the procedure *pfct* is the vector of log-likelihoods for a set of observations.

Supplying an Analytical GRADIENT Procedure

To decrease the time of computation, the user may provide a procedure for the calculation of the gradient of the log-likelihood. The global variable **__max_GradProc** must contain the pointer to this procedure. Suppose the name of this procedure is *gradproc*. Then,

```
g = gradproc(x,y);
```

where the input arguments are

3. MAXIMUM LIKELIHOOD REFERENCE

x vector of coefficients

y matrix, dataset.

and the output argument is

g row vector of gradients of log-likelihood with respect to coefficients, or a matrix of gradients (i.e., a Jacobian)

It is important to note that the gradient is row oriented. **_max_GradProc** must return a matrix of first derivatives in which rows are associated with observations and columns with coefficients.

Providing a procedure for the calculation of the first derivatives also has a significant effect on the calculation time of the Hessian. The calculation time for the numerical computation of the Hessian is a quadratic function of the size of the matrix. For large matrices, the calculation time can be very significant. This time can be reduced to a linear function of size if a procedure for the calculation of analytical first derivatives is available. When such a procedure is available, **FASTMAX** automatically uses it to compute the numerical Hessian.

The major problem one encounters when writing procedures to compute gradients and Hessians is in making sure that the gradient is being properly computed. For best results use **MAXLIK** with **_max_GradCheckTol** set to a nonzero value to ensure that they are being calculated correctly.

Supplying an Analytical HESSIAN Procedure.

Selection of the NEWTON algorithm becomes feasible if the user supplies a procedure to compute the Hessian. If such a procedure is provided, the global variable **_max_HessProc** must contain a pointer to this procedure. Suppose this procedure is called *hessproc*, the format is

```
h = hessproc(x,y);
```

The input arguments are

x $K \times 1$ vector of coefficients

y matrix containing data set

and the output argument is

h $K \times K$ matrix of second order partial derivatives evaluated at the coefficients in x .

In practice much of the time spent on writing the Hessian procedure is devoted to debugging. To help in this debugging process, use the **MAXLIK** procedure with **_max_GradCheckTol** is set to a small nonzero value.

■ Source

`fastmax.src`

■ Purpose

Computes a simulated posterior of the parameters of a maximum likelihood function using **FASTMAX**

■ Library

maxlik

■ Format

$\{ x, f, g, cov, retcode \} = \text{FASTBayes}(data, vars, \&fct, start)$

■ Input

data $N \times NV$ matrix, dataset

vars $NV \times 1$ character vector, labels of variables selected for analysis
 – or –
 $NV \times 1$ numeric vector, indices of variables selected for analysis.
vars may be a character vector containing either the standard labels created by **FASTBayes** (i.e., either V1, V2,..., or V01, V02,..... See discussion of the global variable **___vpad** below, or the user-provided labels in **___altnam**).

&fct a pointer to a procedure that returns the log-likelihood for a vector of log-likelihoods for a matrix of observations.

start $K \times 1$ vector, start values.

■ Output

x $K \times 1$ vector, means of simulated posterior

f scalar, mean weighted bootstrap log-likelihood

g $K \times 1$ vector, means gradient of weighted bootstrap

h $K \times K$ matrix, covariance matrix of simulated posterior

retcode scalar, return code. If normal convergence is achieved, then *retcode* = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:

0 normal convergence

1 forced exit.

2	maximum iterations exceeded.
3	function calculation failed.
4	gradient calculation failed.
5	Hessian calculation failed.
6	line search failed.
7	function cannot be evaluated at initial parameter values.
8	error with gradient
9	gradient vector transposed
10	secant update failed
11	maximum time exceeded
12	error with weights
34	data set could not be opened.
99	termination condition unknown.

■ Globals

The **FASTMAX** procedure global variables are also applicable.

__max__BayesAlpha scalar, exponent of the Dirichlet random variates used for weights for the weighted bootstrap. See Newton and Raftery, “Approximate Bayesian Inference with the Weighted Likelihood Bootstrap”, J.R. Statist. Soc. B (1994), 56:3-48. Default = 1.4.

__max__BootFname string, file name of **GAUSS** data set (do not include .DAT extension) containing bootstrapped parameter estimates. If not specified, **FASTBayes** selects a temporary name.

__max__MaxTime scalar, maximum amount of time spent in re-sampling. Default = 1e5 (about 10 weeks).

__max__NumSample scalar, number of samples to be drawn. Default = 100.

__max__PriorProc scalar, pointer to proc for computing prior. This proc takes the parameter vector as its only argument, and returns a scalar probability. If a proc is not provided, a uniform prior is assumed.

__max__RandType scalar, if nonzero, pseudo-random numbers of the linear congruential type are generated, otherwise, they are generated by Marsaglia’s **Kiss-Monster** method. The latter method is slower but has a much larger period. Random numbers are generated for the random line search.

__max_State scalar or vector, state vector for pseudorandom number generators containing seed. By default it is set to 345678. If you wish to select a seed, set to a different value.

■ Remarks

FASTBayes generates **__max_NumSample** simulations from the posterior distribution of the parameters using a weighted likelihood bootstrap method. The simulation is put into a **GAUSS** data set. The file name of the data set is either the name found in the global **__max_BootFname**, or a temporary name. If **FASTBayes** selects a file name, it returns that file name in **__max_BootFname**.

The simulated parameters in this data set can be used as input to the procedures **MAXHist** and **MAXDensity** for further analysis.

The output from **MAXDensity** can also be used to compute modal estimates of the parameters.

■ Source

`fastbayes.src`

■ Purpose

Computes bootstrapped estimates of parameters of a maximum likelihood function using **FASTMAX**.

■ Library

maxlik

■ Format

$\{ x, f, g, cov, retcode \} = \mathbf{FASTBoot}(data, vars, \&fct, start)$

■ Input

<i>data</i>	$N \times NV$ matrix, dataset
<i>vars</i>	$NV \times 1$ character vector, labels of variables selected for analysis – or – $NV \times 1$ numeric vector, indices of variables selected for analysis. <i>vars</i> may be a character vector containing either the standard labels created by FASTBoot (i.e., either V1, V2, ..., or V01, V02, See discussion of the global variable ___vpad below, or the user-provided labels in ___altnam).
<i>&fct</i>	a pointer to a procedure that returns the log-likelihood for a matrix of observations
<i>start</i>	$K \times 1$ vector, start values.

■ Output

<i>x</i>	$K \times 1$ vector, means of re-sampled parameters
<i>f</i>	scalar, mean re-sampled function at minimum (the mean log-likelihood)
<i>g</i>	$K \times 1$ vector, means of re-sampled gradients evaluated at the estimates
<i>h</i>	$K \times K$ matrix, covariance matrix of the re-sampled parameters
<i>retcode</i>	scalar, return code. If normal convergence is achieved, then <i>retcode</i> = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:
0	normal convergence
1	forced exit.

- 2 maximum iterations exceeded.
- 3 function calculation failed.
- 4 gradient calculation failed.
- 5 Hessian calculation failed.
- 6 line search failed.
- 7 function cannot be evaluated at initial parameter values.
- 8 error with gradient
- 9 gradient vector transposed
- 10 secant update failed
- 11 maximum time exceeded
- 12 error with weights
- 34 data set could not be opened.
- 99 termination condition unknown.

■ Globals

The **FASTMAX** procedure global variables are also applicable.

__max__BootFname string, file name of **GAUSS** data set (do not include .DAT extension) containing bootstrapped parameter estimates. If not specified, **FASTBoot** selects a temporary name.

__max__MaxTime scalar, maximum amount of time spent in re-sampling. Default = 1e5 (about 10 weeks).

__max__NumSample scalar, number of samples to be drawn. Default = 100.

__max__RandType scalar, if nonzero, pseudo-random numbers of the linear congruential type are generated, otherwise, they are generated by Marsaglia's **Kiss-Monster** method. The latter method is slower but has a much larger period. Random numbers are generated for the random line search.

__max__State scalar or vector, state vector for pseudorandom number generators containing seed. By default it is set to 345678. If you wish to select a seed, set to a different value.

■ Remarks

FASTBoot generates **__max__NumSample** random samples of size **__max__NumObs** from the data set with replacement and calls **FASTMAX**. **FASTBoot** returns the mean

vector of the estimates in the first argument and the covariance matrix of the estimates in the third argument.

A **GAUSS** data set is also generated containing the bootstrapped parameter estimates. The file name of the data set is either the name found in the global **__max_BootFname**, or a temporary name. If **FASTBoot** selects a file name, it returns that file name in **__max_BootFname**. The coefficients in this data set may be used as input to the procedures **MAXHist** and **MAXDensity** for further analysis.

■ Source

`fastboot.src`

■ Purpose

Computes profile likelihood confidence limits using **FASTMAX**

■ Library

maxlik

■ Format

`cl = FASTPflimits(b,f,data,vars,&fct)`

■ Input

<i>b</i>	$K \times 1$ vector, maximum likelihood estimates
<i>f</i>	scalar, function at minimum (mean log-likelihood)
<i>data</i>	$N \times NV$ matrix, data
<i>vars</i>	$NV \times 1$ character vector, labels of variables selected for analysis – or – $NV \times 1$ numeric vector, indices of variables selected for analysis. <i>vars</i> may be a character vector containing either the standard labels created by MAXPflimits (i.e., either V1, V2, ..., or V01, V02, ...). See discussion of the global variable ___vpad below, or the user-provided labels in ___altnam .
<i>&fct</i>	a pointer to a procedure that returns a vector of log-likelihoods for a matrix of observations.

■ Output

<i>cl</i>	$K \times 2$ vector, upper (first column) and lower (second column) confidence limits for the parameters in <i>b</i>
-----------	----------------------------------------------------------------------------------------------------------------------

■ Globals

The **FASTMAX** procedure global variables are also applicable.

___max_Alpha (1-**___max_Alpha**)% confidence limits are computed. The default is .05

___max_NumObs scalar, number of observations. Must be set. If the call to **MaxPflimits** comes after a call to **MAXLIK**, it will be set by **MAXLIK**.

___max_Select selection vector for selecting parameters for analysis. For example,

```
_max_Select = { 1, 3, 4 };
```

selects the 1st, 3rd, and 4th parameters for limits.

■ Remarks

FASTPflimits computes profile likelihood confidence limits given a maximum likelihood estimation. b and f should be returns from a call to **FASTMAX**. This will also properly set up **_max_NumObs** for **FASTPflimits**.

FASTPflimits solves for the confidence limits as a parametric likelihood problem. Thus it itself calls **FASTMAX** several times for each confidence limit.

■ Source

fastpflcl.src

■ Purpose

Computes profile t plots and likelihood profile traces for maximum likelihood models using **FASTMAX**

■ Library

maxlik

■ Format

$\{ x, f, g, cov, retcode \} = \mathbf{FASTProfile}(data, vars, \&fct, start)$

■ Input

data $N \times NV$ matrix, dataset

vars $NV \times 1$ character vector, labels of variables selected for analysis
 – or –
 $NV \times 1$ numeric vector, indices of variables selected for analysis.
vars may be a character vector containing either the standard labels created by **FASTProfile** (i.e., either V1, V2, ..., or V01, V02, ...). See discussion of the global variable **__vpad** below, or the user-provided labels in **__altnam**).

&fct a pointer to a procedure that returns a vector of log-likelihoods for a matrix of observations

start $K \times 1$ vector, start values.

■ Output

x $K \times 1$ vector, means of re-sampled parameters

f scalar, mean re-sampled function at minimum (the mean log-likelihood)

g $K \times 1$ vector, means of re-sampled gradients evaluated at the estimates

h $K \times K$ matrix, covariance matrix of the re-sampled parameters

retcode scalar, return code. If normal convergence is achieved, then *retcode* = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:

0 normal convergence

1 forced exit.

2	maximum iterations exceeded.
3	function calculation failed.
4	gradient calculation failed.
5	Hessian calculation failed.
6	line search failed.
7	function cannot be evaluated at initial parameter values.
8	error with gradient
9	gradient vector transposed
10	secant update failed
11	maximum time exceeded
12	error with weights
34	data set could not be opened.
99	termination condition unknown.

■ Globals

The **FASTMAX** procedure global variables are also relevant.

__max__NumCat scalar, number of categories in profile table. Default = 16.

__max__Increment $K \times 1$ vector, increments for cutting points, default is $2 * \text{__max__Width} * \text{std dev} / \text{__max__NumCat}$. If scalar zero, increments are computed by **FastProfile**.

__max__Center $K \times 1$ vector, value of center category in profile table. Default values are coefficient estimates.

__max__Select selection vector for selecting coefficients to be included in profiling, for example

```
__max__Select = { 1, 3, 4 };
```

selects the 1st, 3rd, and 4th parameters for profiling.

__max__Width scalar, width of profile table in units of the standard deviations of the parameters. Default = 2.

■ Remarks

For each pair of the selected parameters, three plots are printed to the screen. Two of the are the profile t trace plots that describe the univariate profiles of the parameters, and one of them is the profile likelihood trace describing the joint distribution of the two parameters. Ideally distributed parameters would have univariate profile t traces that are straight lines, and bivariate likelihood profile traces that are two straight lines

intersecting at right angles. This ideal is generally not met by nonlinear models, however, large deviations from the ideal indicate serious problems with the usual statistical inference.

- **Source**

`fastprof.src`

■ Purpose

Computes estimates of parameters of a maximum likelihood function.

■ Library

maxlik

■ Format

$\{ x, f, g, cov, retcode \} = \text{MAXLIK}(\text{dataset}, \text{vars}, \&fct, \text{start})$

■ Input

<i>dataset</i>	string containing name of GAUSS data set – or – $N \times NV$ matrix, data
<i>vars</i>	$NV \times 1$ character vector, labels of variables selected for analysis – or – $NV \times 1$ numeric vector, indices of variables selected for analysis. If <i>dataset</i> is a matrix, <i>vars</i> may be a character vector containing either the standard labels created by MAXLIK (i.e., either V1, V2, ..., or V01, V02, ...). See discussion of the global variable __vpad below, or the user-provided labels in __altnam).
<i>&fct</i>	a pointer to a procedure that returns either the log-likelihood for one observation or a vector of log-likelihoods for a matrix of observations (see discussion of the global variable __row in global variable section below).
<i>start</i>	$K \times 1$ vector, start values.

■ Output

<i>x</i>	$K \times 1$ vector, estimated parameters
<i>f</i>	scalar, function at minimum (the mean log-likelihood)
<i>g</i>	$K \times 1$ vector, gradient evaluated at <i>x</i>
<i>h</i>	$K \times K$ matrix, covariance matrix of the parameters (see discussion of the global variable __max__CovPar below).
<i>retcode</i>	scalar, return code. If normal convergence is achieved, then <i>retcode</i> = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:

0	normal convergence
1	forced exit.
2	maximum iterations exceeded.
3	function calculation failed.
4	gradient calculation failed.
5	Hessian calculation failed.
6	line search failed.
7	function cannot be evaluated at initial parameter values.
8	error with gradient
9	gradient vector transposed
10	secant update failed
11	maximum time exceeded
12	error with weights
34	data set could not be opened.
99	termination condition unknown.

■ Globals

The global variables used by **MAXLIK** can be organized in the following categories according to which aspect of the optimization they affect:

Options `_max_Options`

Descent and Line Search `_max_Algorithm, _max_Delta, _max_LineSearch, _max_Maxtry, _max_Extrap, _max_Interp, _max_RandRadius, _max_UserSearch, _max_Switch, _max_RandType, _max_State,`

Covariance Matrix of Parameters `_max_CovPar, _max_XprodCov, _max_HessCov, _max_FinalHess`

Gradient `_max_GradMethod, _max_GradProc, _max_UserNumGrad, _max_HessProc, _max_UserNumHess, _max_GradStep, _max_GradCheckTol`

Terminations Conditions `_max_GradTol, _max_MaxIters, _max_MaxTime`

Data `_max_Lag, _max_NumObs, __weight, __row, __rowfac`

Parameters `_max_Active, _max_ParNames`

Miscellaneous `__title, _max_IterData, _max_Diagnostic`

The list below contains an alphabetical listing of each global with a complete description.

__max__Active vector, defines fixed/active coefficients. This global allows you to fix a parameter to its starting value. This is useful, for example, when you wish to try different models with different sets of parameters without having to re-edit the function. When it is to be used, it must be a vector of the same length as the starting vector. Set elements of `'max'Active` to 1 for an active parameter, and to zero for a fixed one.

__max__Algorithm scalar, selects optimization method:

- 1 STEEP - Steepest Descent
- 2 BFGS - Broyden, Fletcher, Goldfarb, Shanno method
- 3 DFP - Davidon, Fletcher, Powell method
- 4 NEWTON - Newton-Raphson method
- 5 BHHH - Berndt, Hall, Hall, Hausman method
- 6 PRCG - Polak-Ribiere Conjugate Gradient

Default = 3

__max__CovPar scalar, type of covariance matrix of parameters

- 0 not computed
- 1 computed from Hessian calculated after the iterations
- 2 computed from cross-product of Jacobian calculated after iterations
- 3 Quasi-maximum likelihood (QML) covariance matrix of the parameters

Default = 1;

__max__Delta scalar, floor for eigenvalues of Hessian in the NEWTON algorithm. When nonzero, the eigenvalues of the Hessian are augmented to this value.

__max__Diagnostic scalar.

- 0 nothing is stored or printed
- 1 current estimates, gradient, direction, function value, Hessian, and step length are printed to the screen
- 2 the current quantities are stored in **__max__Diagnostic** using the **vput** command. Use the following strings to extract from **__max__Diagnostic** using **vread**:

function	“function”
estimates	“params”
direction	“direct”
Hessian	“hessian”
gradient	“gradient”
step	“step”

When **_max_Diagnostic** is nonzero, **__output** is forced to 1.

- _max_GradTol** scalar, convergence tolerance for gradient of estimated coefficients. When this criterion has been satisfied MAXLIK exits the iterations. Default = 1e-5.
- _max_Extrap** scalar, extrapolation constant in BRENT. Default = 2.
- _max_FinalHess** $K \times K$ matrix, the Hessian used to compute the covariance matrix of the parameters is stored in **_max_FinalHess**. This is most useful if the inversion of the hessian fails, which is indicated when **MAXLIK** returns a missing value for the covariance matrix of the parameters. An analysis of the Hessian stored in **_max_FinalHess** can then reveal the source of the linear dependency responsible for the singularity.
- _max_GradCheckTol** scalar. Tolerance for the deviation of numerical and analytical gradients when proc's exist for the computation of analytical gradients or Hessians. If set to zero, the analytical gradients will not be compared to their numerical versions. When adding procedures for computing analytical gradients it is highly recommended that you perform the check. Set **_max_GradCheckTol** to some small value, 1e-3, say when checking. It may have to be set larger if the numerical gradients are poorly computed to make sure that **MAXLIK** doesn't fail when the analytical gradients are being properly computed.
- _max_GradMethod** scalar, method for computing numerical gradient.
- 0 central difference
 - 1 forward difference (default)
- _max_GradProc** scalar, pointer to a procedure that computes the gradient of the function with respect to the parameters. For example, the statement:
- ```
_max_GradProc=&gradproc;
```
- tells **MAXLIK** that a gradient procedure exists as well where to find it. The user-provided procedure has two input arguments, an  $K \times 1$  vector of parameter values and an  $N \times K$  matrix of data. The procedure returns a single output argument, an  $N \times K$  matrix of gradients of the log-likelihood function with respect to the parameters evaluated at the vector of parameter values.

For example, suppose the log-likelihood function is for a Poisson regression, then the following would be added to the command file:

```
proc lgd(b,z);
 retp((z[.,1]-exp(z[.,2:4]*b)).*z[.,2:4]);
endp;

_max_GradProc = &lgd;
```

Default = 0, i.e., no gradient procedure has been provided.

**\_max\_GradStep** scalar, increment size for computing gradient. When the numerical gradient is performing well, set to a larger value (1e-3, say). Default is the cube root of machine precision.

**\_max\_HessCov**  $K \times K$  matrix. When **\_max\_CovPar** is set to 3 the information matrix covariance matrix of the parameters, i.e., the inverse of the matrix of second order partial derivatives of the log-likelihood by observations, is returned in **\_max\_HessCov**.

**\_max\_HessProc** scalar, pointer to a procedure that computes the hessian, i.e., the matrix of second order partial derivatives of the function with respect to the parameters. For example, the instruction:

```
_max_HessProc = &hessproc;
```

tells MAXLIK that a procedure has been provided for the computation of the hessian and where to find it. The procedure that is provided by the user must have two input arguments, a  $K \times 1$  vector of parameter values and an  $N \times P$  data matrix. The procedure returns a single output argument, the  $K \times K$  symmetric matrix of second order derivatives of the function evaluated at the parameter values.

**\_max\_Interp** scalar, interpolation constant in BRENT. Default = .25.

**\_max\_IterData** 3x1 vector, contains information about the iterations.

The first element contains the # of iterations, the second element contains the elapsed time in minutes of the iterations, and the third element contains a character variable indicating the type of covariance matrix of the parameters.

**\_max\_Lag** scalar, if the function includes lagged values of the variables **\_max\_Lag** may be set to the number of lags. When **\_max\_Lag** is set to a nonzero value then **\_\_row** is set to 1 (that is, the function must be evaluated one observation at a time), and **MAXLIK** passes a matrix to the user-provided function and gradient procedures. The first row in this matrix is the  $(i - \text{\_max\_Lag})$ -th observation and the last row is the  $i$ -th observation. The read loop begins with the  $(\text{\_max\_Lag}+1)$ -th observation. Default = 0.

**\_\_max\_LineSearch** scalar, selects method for conducting line search. The result of the line search is a *step length*, i.e., a number which reduces the function value when multiplied times the direction..

- 1 step length = 1.
- 2 cubic or quadratic step length method (STEPBT)
- 3 step halving (HALF)
- 4 Brent's step length method (BRENT)
- 5 BHHH step length method (BHHHSTEP)

Default = 2.

Usually **\_\_max\_LineSearch** = 2 is best. If the optimization bogs down, try setting **\_\_max\_LineSearch** = 1, 4 or 5. **\_\_max\_LineSearch** = 3 generates slower iterations but faster convergence and **\_\_max\_LineSearch** = 1 generates faster iterations but slower convergence.

When any of these line search methods fails, **MAXLIK** attempts a random search of radius **\_\_max\_RandRadius** times the truncated log to the base 10 of the gradient when **\_\_max\_RandRadius** is set to a nonzero value. If **\_\_max\_UserSearch** is set to 1, **MAXLIK** enters an interactive line search mode.

**\_\_max\_MaxIters** scalar, maximum number of iterations.

**\_\_max\_MaxTime** scalar, maximum time in iterations in minutes. This global is most useful in bootstrapping. You might want 100 re-samples, but would be happy with anything more than 50 depending on the time it took. Set **\_\_max\_NumSample** = 100, and **\_\_max\_MaxTime** to maximum time you would be willing to wait for results. Default = 1e+5, about 10 weeks.

**\_\_max\_MaxTry** scalar, maximum number of tries to find step length that produces a descent.

**\_\_max\_NumObs** scalar, number of cases in the data set that was analyzed.

**\_\_max\_Options** character vector, specification of options. This global permits setting various **MAXLIK** options in a single global using identifiers. The following

```
__max_Options = { bfgs stepbt forward screen };
```

sets to the default values, i.e. the descent method to BFGS, the line search method to STEPBT, the numerical gradient method to central differences, and `OUTPUT = 2`.

The following is a list of the identifiers:

**Algorithms** STEEP, BFGS, DFP, NEWTON, BHHH, PRCG

**Line Search** ONE, STEPBT, HALF, BRENT, BHHHSTEP

**Covariance Matrix** NOCOV, INFO, XPROD, HETCON

**Gradient method** CENTRAL, FORWARD

**Output method** NONE, FILE, SCREEN

**\_\_\_output** scalar, determines printing of intermediate results. Generally when **\_\_\_output** is nonzero, i.e., where there some kind of printing during the iterations, the time of the iterations is degraded.

- 0 nothing is written
- 1 serial ASCII output format suitable for disk files or printers
- 2 output is suitable for screen only. ANSI.SYS must be active.
- $\geq 5$  same as **\_\_\_output** = 1 except that information is printed only every **\_\_\_output**-th iteration.

When **\_\_max\_\_Diagnostic** is nonzero, **\_\_\_output** is forced to 1.

**\_\_max\_\_ParNames**  $K \times 1$  character vector, parameter labels.

**\_\_max\_\_RandRadius** scalar, if set to a nonzero value (1e-2, say) and all other line search methods fail then **MAXLIK** attempts **\_\_max\_\_MaxTry** tries to find a random direction within radius determined by **\_\_max\_\_RandRadius** that is a descent. Default = 1e-2.

**\_\_max\_\_RandType** scalar, if nonzero, pseudo-random numbers of the linear congruential type are generated, otherwise, they are generated by Marsaglia's **Kiss-Monster** method. The latter method is slower but has a much larger period. Random numbers are generated for the random line search.

**\_\_max\_\_State** scalar or vector, state vector for pseudorandom number generators containing seed. By default it is set to 345678. If you wish to select a seed, set to a different value.

**\_\_max\_\_Switch**  $4 \times 1$  or  $4 \times 2$  vector, controls algorithm switching. If **\_\_max\_\_Switch** is  $4 \times 1$ , set its elements in the following way,

- 1 , algorithm number to switch to
- 2 , **MAXLIK** will switch to algorithm in the first element when the function value is less than the value entered here
- 3 **MAXLIK** switches if the number of iterations exceeds the number entered here
- 4 **MAXLIK** switches if line search step changes less than the amount entered here

If **\_\_max\_Switch** is  $4 \times 2$ , **MAXLIK** switches between the algorithms in column 1 and column 2 subject to the conditions specified for the  $4 \times 1$  vector.

Thus if **\_\_max\_Switch** is a  $4 \times 1$  vector, **MAXLIK** will switch algorithms no more than once during the iterations, whereas if it is  $4 \times 2$  it may switch back and forth between the two algorithms throughout the iterations.

**\_\_max\_UserNumGrad** scalar, pointer to user provided numerical gradient procedure.  
The instruction

```
__max_UserNumGrad = &userproc;
```

tells **MAXLIK** that a procedure for computing the numerical gradients exists. The user-provided procedure has three input arguments, a pointer to a function that computes the log-likelihood function, a  $K \times 1$  vector of parameter values, and an  $K \times P$  matrix of data. The procedure returns a single output argument, an  $N \times K$  matrix of gradients of each row of the input data matrix with respect to each parameter.

**MAXLIK** includes a procedure, **GRADRE**, for computing numerical derivatives using the Richardson Extrapolation method. It is invoked by setting the global to a pointer to this function:

```
__max_UserNumGrad = &gradre;
```

**\_\_row** scalar, specifies how many rows of the data set are read per iteration of the read loop. See the *REMARKS* Section for a more detailed discussion of how to set up your log-likelihood to handle more than one row of your data set. By default, the number of rows to be read is calculated by **MAXLIK**.

**\_\_rowfac** scalar, “row factor”. If **MAXLIK** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
__rowfac = 0.8;
```

causes **GAUSS** to read in 80% of the rows of the **GAUSS** data set that were read when **MAXLIK** failed due to insufficient memory.

This global has an affect only when **\_\_row** = 0. Default = 1.

**\_\_title** string title of run

**\_\_max\_UserNumHess** scalar, pointer to user provided numerical Hessian procedure.  
The instruction

```
__max_UserHess = &hessproc;
```



tells **MAXLIK** that a procedure for computing the numerical Hessian exists. The user-provided procedure three input arguments, a pointer to a function that computes the log-likelihood function, a  $K \times 1$  vector of parameter values, and an  $N \times P$  matrix of data. The procedure returns a single output argument, a  $K \times K$  Hessian matrix of the function with respect to the parameters.

- \_\_max\_UserSearch** scalar, if nonzero and if all other line search methods fail **MAXLIK** enters an interactive mode in which the user can select a line search parameter
- \_\_weight** vector, frequency of observations. By default all observations have a frequency of 1. zero frequencies are allowed. It is assumed that the elements of **\_\_weight** sum to the number of observations.
- \_\_max\_XprodCov**  $K \times K$  matrix. When **\_\_max\_CovPar** is set to 3 the cross-product matrix covariance matrix of the parameters, i.e., the inverse of the cross-product of the first derivatives of the log-likelihood computed by observations, is returned in **\_\_max\_XprodCov**.

## ■ Remarks

### Writing the Log-likelihood Function

The user must provide a procedure for computing the log-likelihood for either one observation, or for a matrix of observations. The procedure must have two input arguments: first, a vector of parameter values, and second, one or more rows of the data matrix. The output argument is the log-likelihood for the observation or observations in the second argument evaluated at the parameter values in the first argument. Suppose that the function procedure has been named *pfct*, the following considerations apply:

The format of the procedure is:

```
logprob = pfct(x,y);
```

where

- x* column vector of parameters of model
- y* one or more rows of the data set (if the data set has been transformed, or if *vars*  $\neq$  0, i.e., there is selection, then *y* is a transformed, selected observation)
- if **\_\_row** = *n*, then *n* rows of the data set are read at a time
- if **\_\_row** = 0, the maximum number of rows that fit in memory is computed by **MAXLIK**.

The output from the procedure *pfct* is the log-likelihood for a single observation or a vector of log-likelihoods for a set of observations. If it is not possible to compute the log-likelihood for a set of observations, then either `___row` may be set to 1 to force **MAXLIK** to send one observation at a time to *pfct* or the procedure computing the function may contain a loop. If possible, *pfct* should be written to compute a vector of log-likelihoods for a set of observations because this speeds up the computations significantly. If `_max_Lag`  $\geq 1$ , then `___row` is forced to 1.

Setting `___row` = 0 causes **MAXLIK** to send the entire matrix to *pfct* if it is stored entirely in memory, or to compute the maximum number of rows if it is a **GAUSS** data set stored on disk (Note that even if the data starts out in a **GAUSS** data set, **MAXLIK** determines whether the data set will fit in memory, and if it does, then it reads the data set into an array in memory). If you are getting *insufficient memory* messages, then set `___rowfac` to a positive value less than 1.

#### Supplying an Analytical GRADIENT Procedure

To decrease the time of computation, the user may provide a procedure for the calculation of the gradient of the log-likelihood. The global variable `_max_GradProc` must contain the pointer to this procedure. Suppose the name of this procedure is *gradproc*. Then,

```
g = gradproc(x,y);
```

where the input arguments are

*x*                vector of coefficients  
*y*                one or more rows of data set.

and the output argument is

*g*                row vector of gradients of log-likelihood with respect to coefficients, or a matrix of gradients (i.e., a Jacobian) if the data passed in *y* is a matrix (unless `_max_Lag`  $\geq 1$  in which case the data passed in *y* is a matrix of lagged values but a row vector of gradients is passed back in *g*).

It is important to note that the gradient is row oriented. Thus if the function that computes the log-likelihood returns a scalar value (`___row` = 1), then a row vector of the first derivatives of the log-likelihood with respect to the coefficients must be returned, but if the procedure that computes the log-likelihood returns a column vector, then `_max_GradProc` must return a matrix of first derivatives in which rows are associated with observations and columns with coefficients.

Providing a procedure for the calculation of the first derivatives also has a significant effect on the calculation time of the Hessian. The calculation time for the numerical

computation of the Hessian is a quadratic function of the size of the matrix. For large matrices, the calculation time can be very significant. This time can be reduced to a linear function of size if a procedure for the calculation of analytical first derivatives is available. When such a procedure is available, **MAXLIK** automatically uses it to compute the numerical Hessian.

The major problem one encounters when writing procedures to compute gradients and Hessians is in making sure that the gradient is being properly computed. **MAXLIK** checks the gradients and Hessian when **\_max\_GradCheckTol** is nonzero. **MAXLIK** generates both numerical and analytical gradients, and viewing the discrepancies between them can help in debugging the analytical gradient procedure.

#### Supplying an Analytical HESSIAN Procedure.

Selection of the NEWTON algorithm becomes feasible if the user supplies a procedure to compute the Hessian. If such a procedure is provided, the global variable **\_max\_HessProc** must contain a pointer to this procedure. Suppose this procedure is called *hessproc*, the format is

```
h = hessproc(x,y);
```

The input arguments are

|          |                                     |
|----------|-------------------------------------|
| <i>x</i> | $K \times 1$ vector of coefficients |
| <i>y</i> | one or more rows of data set        |

and the output argument is

|          |                                                                                                     |
|----------|-----------------------------------------------------------------------------------------------------|
| <i>h</i> | $K \times K$ matrix of second order partial derivatives evaluated at the coefficients in <i>x</i> . |
|----------|-----------------------------------------------------------------------------------------------------|

In practice much of the time spent on writing the Hessian procedure is devoted to debugging. To help in this debugging process, **MAXLIK** can be instructed to compute the numerical Hessian along with your prospective analytical Hessian for comparison purposes. To accomplish this **\_max\_GradCheckTol** is set to a small nonzero value.

#### ■ Source

```
maxlik.src
```

## ■ Purpose

Computes a simulated posterior of the parameters of a maximum likelihood function

## ■ Library

maxlik

## ■ Format

$\{ x, f, g, cov, retcode \} = \text{MAXBayes}(\text{dataset}, \text{vars}, \&fct, \text{start})$

## ■ Input

*dataset* string containing name of **GAUSS** data set  
 – or –  
 $N \times NV$  matrix, data

*vars*  $NV \times 1$  character vector, labels of variables selected for analysis  
 – or –  
 $NV \times 1$  numeric vector, indices of variables selected for analysis.  
 If *dataset* is a matrix, *vars* may be a character vector containing either the standard labels created by **MAXBayes** (i.e., either V1, V2, ..., or V01, V02, .... See discussion of the global variable **\_\_\_vpad** below, or the user-provided labels in **\_\_\_altnam**).

*&fct* a pointer to a procedure that returns either the log-likelihood for one observation or a vector of log-likelihoods for a matrix of observations (see discussion of the global variable **\_\_\_row** in global variable section below).

*start*  $K \times 1$  vector, start values.

## ■ Output

*x*  $K \times 1$  vector, means of simulated posterior

*f* scalar, mean weighted bootstrap log-likelihood

*g*  $K \times 1$  vector, means gradient of weighted bootstrap

*h*  $K \times K$  matrix, covariance matrix of simulated posterior

*retcode* scalar, return code. If normal convergence is achieved, then *retcode* = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:

0 normal convergence

- 1 forced exit.
- 2 maximum iterations exceeded.
- 3 function calculation failed.
- 4 gradient calculation failed.
- 5 Hessian calculation failed.
- 6 line search failed.
- 7 function cannot be evaluated at initial parameter values.
- 8 error with gradient
- 9 gradient vector transposed
- 10 secant update failed
- 11 maximum time exceeded
- 12 error with weights
- 34 data set could not be opened.
- 99 termination condition unknown.

## ■ Globals

The **MAXLIK** procedure global variables are also applicable.

**\_\_max\_\_BayesAlpha** scalar, exponent of the Dirichlet random variates used for weights for the weighted bootstrap. See Newton and Raftery, “Approximate Bayesian Inference with the Weighted Likelihood Bootstrap”, J.R. Statist. Soc. B (1994), 56:3-48. Default = 1.4.

**\_\_max\_\_BootFname** string, file name of **GAUSS** data set (do not include .DAT extension) containing bootstrapped parameter estimates. If not specified, **MAXBayes** selects a temporary name.

**\_\_max\_\_MaxTime** scalar, maximum amount of time spent in re-sampling. Default = 1e5 (about 10 weeks).

**\_\_max\_\_NumSample** scalar, number of samples to be drawn. Default = 100.

**\_\_max\_\_PriorProc** scalar, pointer to proc for computing prior. This proc takes the parameter vector as its only argument, and returns a scalar probability. If a proc is not provided, a uniform prior is assumed.

## ■ Remarks

**MAXBayes** generates **\_\_max\_\_NumSample** simulations from the posterior distribution of the parameters using a weighted likelihood bootstrap method. The simulation is put into a **GAUSS** data set. The file name of the data set is either the name found in the

## MAXBayes

### 3. *MAXIMUM LIKELIHOOD REFERENCE*

global `__max_BootFname`, or a temporary name. If **MAXBayes** selects a file name, it returns that file name in `__max_BootFname`.

The simulated parameters in this data set can be used as input to the **MAXLIK** procedures **MAXHist** and **MAXDensity** for further analysis.

The output from **MAXDensity** can also be used to compute modal estimates of the parameters.

#### ■ **Source**

`maxbayes.src`

## ■ Purpose

Computes bootstrapped estimates of parameters of a maximum likelihood function.

## ■ Library

maxlik

## ■ Format

$\{ x, f, g, cov, retcode \} = \text{MAXBoot}(dataset, vars, \&fct, start)$

## ■ Input

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>dataset</i>  | string containing name of <b>GAUSS</b> data set<br>– or –<br>$N \times NV$ matrix, data                                                                                                                                                                                                                                                                                                                                                                               |
| <i>vars</i>     | $NV \times 1$ character vector, labels of variables selected for analysis<br>– or –<br>$NV \times 1$ numeric vector, indices of variables selected for analysis.<br>If <i>dataset</i> is a matrix, <i>vars</i> may be a character vector containing either the standard labels created by <b>MAXBoot</b> (i.e., either V1, V2, ..., or V01, V02, ...). See discussion of the global variable <b>___vpad</b> below, or the user-provided labels in <b>___altnam</b> ). |
| <i>&amp;fct</i> | a pointer to a procedure that returns either the log-likelihood for one observation or a vector of log-likelihoods for a matrix of observations (see discussion of the global variable <b>___row</b> in global variable section below).                                                                                                                                                                                                                               |
| <i>start</i>    | $K \times 1$ vector, start values.                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## ■ Output

|                |                                                                                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>x</i>       | $K \times 1$ vector, means of re-sampled parameters                                                                                                                           |
| <i>f</i>       | scalar, mean re-sampled function at minimum (the mean log-likelihood)                                                                                                         |
| <i>g</i>       | $K \times 1$ vector, means of re-sampled gradients evaluated at the estimates                                                                                                 |
| <i>h</i>       | $K \times K$ matrix, covariance matrix of the re-sampled parameters                                                                                                           |
| <i>retcode</i> | scalar, return code. If normal convergence is achieved, then <i>retcode</i> = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination: |
|                | 0      normal convergence                                                                                                                                                     |

- 1 forced exit.
- 2 maximum iterations exceeded.
- 3 function calculation failed.
- 4 gradient calculation failed.
- 5 Hessian calculation failed.
- 6 line search failed.
- 7 function cannot be evaluated at initial parameter values.
- 8 error with gradient
- 9 gradient vector transposed
- 10 secant update failed
- 11 maximum time exceeded
- 12 error with weights
- 34 data set could not be opened.
- 99 termination condition unknown.

## ■ Globals

The **MAXLIK** procedure global variables are also applicable.

**\_\_max\_\_BootFname** string, file name of **GAUSS** data set (do not include .DAT extension) containing bootstrapped parameter estimates. If not specified, **MAXBoot** selects a temporary name.

**\_\_max\_\_MaxTime** scalar, maximum amount of time spent in re-sampling. Default = 1e5 (about 10 weeks).

**\_\_max\_\_NumSample** scalar, number of samples to be drawn. Default = 100.

## ■ Remarks

**MAXBoot** generates **\_\_max\_\_NumSample** random samples of size **\_\_max\_\_NumObs** from the data set with replacement and calls **MAXLIK**. **MAXBoot** returns the mean vector of the estimates in the first argument and the covariance matrix of the estimates in the third argument.

A **GAUSS** data set is also generated containing the bootstrapped parameter estimates. The file name of the data set is either the name found in the global **\_\_max\_\_BootFname**, or a temporary name. If **MAXBoot** selects a file name, it returns that file name in **\_\_max\_\_BootFname**. The coefficients in this data set may be used as input to the **MAXLIK** procedures **MAXHist** and **MAXDensity** for further analysis.

## ■ Source

maxboot.src



## ■ Purpose

Generates histograms and surface plots from **GAUSS** data sets

## ■ Library

maxlik

## ■ Format

*cl* = MAXBlimits(*dataset*)

## ■ Input

*dataset*      string containing name of **GAUSS** data set  
                   – or –  
                   N×K matrix, data

## ■ Output

*cl*              $K \times 2$  matrix, lower (first column) and upper (second column) confidence limits of the selected parameters

## ■ Globals

**\_max\_Alpha** (1-**\_max\_Alpha**)% confidence limits are computed. The default is .05

**\_max\_Select** selection vector for selecting coefficients to be included in profiling, for example

```
_max_Select = { 1, 3, 4 };
```

selects the 1st, 3rd, and 4th parameters for profiling.

## ■ Remarks

**MAXBlimits** sorts each column of the parameter data set and computes (1-**\_max\_Alpha**)% confidence limits by measuring back **\_max\_Alpha**/2 times the number of rows from each end of the columns. The confidence limits are the values in those elements. If amount to be measured back from each end of the columns doesn't fall exactly on an element of the column, the confidence limit is interpolated from the bordering elements.

## ■ Source

maxblim.src

## ■ Purpose

Formats and prints the output from a call to **MAXLIK** along with confidence limits

## ■ Library

maxlik

## ■ Format

$\{ x, f, g, cl, retcode \} = \text{MAXCLPrt}(x, f, g, cl, retcode);$

## ■ Input

|           |                                                                                       |
|-----------|---------------------------------------------------------------------------------------|
| $x$       | $K \times 1$ vector, parameter estimates                                              |
| $f$       | scalar, value of function at minimum                                                  |
| $g$       | $K \times 1$ vector, gradient evaluated at $x$                                        |
| $cl$      | $K \times 2$ matrix, lower (first column) and upper (second column) confidence limits |
| $retcode$ | scalar, return code.                                                                  |

## ■ Output

The input arguments are returned unchanged.

## ■ Globals

**\_\_header** string. This is used by the printing procedure to display information about the date, time, version of module, etc. The string can contain one or more of the following characters:

|     |                                   |          |
|-----|-----------------------------------|----------|
| “t” | print title (see <b>__title</b> ) |          |
| “l” | bracket title with lines          |          |
| “d” | print date and time               | Example: |
| “v” | print version number of program   |          |
| “f” | print file name being analyzed    |          |

```
__header = "tld";
```

Default = “tldvf”.

**\_\_title** string, message printed at the top of the screen and printed out by **MAXCLPrt**. Default = “”.

**■ Remarks**

Confidence limits computed by **MAXBlimits** or **MAXTlimits** may be passed in the fourth argument in the call to **MAXCLPrt**:

```
{ b,f,g,cov,ret } = MAXBoot("tobit",0,&lpr,x0);
cl = MAXBLimit(_max_BootFname,0);
call MAXCLPrt(b,f,g,cl,ret);
```

**■ Source**

maxlik.src

### ■ Purpose

Generates histograms and surface plots from **GAUSS** data sets

### ■ Library

maxlik

### ■ Format

$\{ px, py, smth \} = \text{MAXDensity}(dataset, vars)$

### ■ Input

*dataset*      string containing name of **GAUSS** data set  
                   – or –  
                   N×K matrix, data

*vars*             $K \times 1$  character vector, labels of variables selected for analysis  
                   – or –  
                    $K \times 1$  numeric vector, indices of variables selected for analysis.

If *dataset* is a matrix, *vars* may be a character vector containing either the standard labels created by **MAXDensity** (i.e., either V1, V2, ..., or V01, V02, .... See discussion of the global variable **\_\_\_vpad** below, or the user-provided labels in **\_\_\_altnam**).

### ■ Output

*px*             **\_\_max\_\_NumPoints** × K matrix, abscissae of plotted points

*py*             **\_\_max\_\_NumPoints** × K matrix, ordinates of plotted points

*smth*             $K \times 1$  vector, smoothing coefficients

### ■ Globals

The **MAXLIK** procedure global variables are also applicable.

**\_\_max\_\_Kernel**  $K \times 1$  character vector, type of kernel:

**NORMAL** normal kernel  
**EPAN** Epanechnikov kernel  
**BIWGT** biweight kernel  
**TRIANG** triangular kernel  
**RECTANG** rectangular kernel

**TNORMAL** truncated normal kernel

If **\_\_max\_\_Kernel** is scalar, the kernel is the same for all parameter densities. Default = NORMAL.

**\_\_max\_\_NumPoints** scalar, number of points to be computed for plots

**\_\_max\_\_EndPoints**  $K \times 2$  matrix, lower (in first column) and upper (in second column) endpoints of density. Default is minimum and maximum, respectively, of the parameter values. If  $1 \times 2$  matrix, endpoints are the same for all parameters.

**\_\_max\_\_Smoothing**  $K \times 1$  vector, smoothing coefficients for each plot. If scalar, smoothing coefficient is the same for each plot. If zero, smoothing coefficient is computed by **MAXDensity**. Default = 0.

**\_\_max\_\_Truncate**  $K \times 2$  matrix, lower (in first column) and upper (in second column) truncation limits for truncated normal kernel. If  $1 \times 2$  matrix, truncations limits are the same for all plots. Default is minimum and maximum, respectively.

**\_\_output** If nonzero,  $K$  density plots are printed to the screen, otherwise no plots are generated.

## ■ Source

`maxdens.src`

## ■ Purpose

Generates histograms and surface plots from **GAUSS** data sets

## ■ Library

maxlik

## ■ Format

{ *tab*, *cut* } = **MAXHist**(*dataset*,*vars*)

## ■ Input

*dataset*      string containing name of **GAUSS** data set  
                   – or –  
                   N×K matrix, data

*vars*             $K \times 1$  character vector, labels of variables selected for analysis  
                   – or –  
                    $K \times 1$  numeric vector, indices of variables selected for analysis.  
 If *dataset* is a matrix, *vars* may be a character vector containing either the standard labels created by **MAXHist** (i.e., either V1, V2,..., or V01, V02,...). See discussion of the global variable **\_\_vpad** below, or the user-provided labels in **\_\_altnam**).

## ■ Output

*tab*             **\_\_max\_NumCat** × K matrix, univariate distributions of bootstrapped parameters

*cut*             **\_\_max\_NumCat** × K matrix, cutting points

## ■ Globals

The **MAXLIK** procedure global variables are also applicable.

**\_\_max\_Center**  $K \times 1$  value of center category in histograms. Default is initial coefficient estimates.

**\_\_max\_CutPoint** **\_\_max\_NumCat** × 1 vector, output, cutting points for histograms

**\_\_max\_Increment**  $K \times 1$  vector, increments for cutting points of the histograms. Default is  $2 * \mathbf{__max\_Width} * \text{std dev} / \mathbf{__max\_NumCat}$ .

**\_\_max\_NumCat** scalar, number of categories in the histograms

**\_\_max\_Width** scalar, width of histograms, default = 2

**\_\_output** If nonzero,  $K$  density plots are printed to the screen, otherwise no plots are generated.

### ■ Remarks

If **\_\_output** is nonzero,  $K(K - 1)/2$  plots are printed to the screen displaying univariate histograms and bivariate surface plots of the bootstrapped parameter distributions in pairs.

The globals, **\_\_max\_Center**, **\_\_max\_Width**, and **\_\_max\_Increment** may be used to establish cutting points (which is stored in **\_\_max\_Increment**) for the tables of re-sampled coefficients in *tab*. The numbers in **\_\_max\_Center** fix the center categories, **\_\_max\_Width** is a factor which when multiplied times the standard deviation of the estimate determines the increments between categories. Alternatively, the increments between categories can be fixed directly by supplying them in **\_\_max\_Increment**.

### ■ Source

`maxhist.src`

## ■ Purpose

Computes profile t plots and likelihood profile traces for maximum likelihood models

## ■ Library

maxlik

## ■ Format

$\{ x, f, g, cov, retcode \} = \text{MAXProfile}(dataset, vars, \&fct, start)$

## ■ Input

*dataset* string containing name of **GAUSS** data set  
 – or –  
 $N \times NV$  matrix, data

*vars*  $NV \times 1$  character vector, labels of variables selected for analysis  
 – or –  
 $NV \times 1$  numeric vector, indices of variables selected for analysis.  
 If *dataset* is a matrix, *vars* may be a character vector containing either the standard labels created by **MAXProfile** (i.e., either V1, V2, ..., or V01, V02, .... See discussion of the global variable **\_\_\_vpad** below, or the user-provided labels in **\_\_\_altnam**).

*&fct* a pointer to a procedure that returns either the log-likelihood for one observation or a vector of log-likelihoods for a matrix of observations (see discussion of the global variable **\_\_\_row** in global variable section below).

*start*  $K \times 1$  vector, start values.

## ■ Output

*x*  $K \times 1$  vector, means of re-sampled parameters

*f* scalar, mean re-sampled function at minimum (the mean log-likelihood)

*g*  $K \times 1$  vector, means of re-sampled gradients evaluated at the estimates

*h*  $K \times K$  matrix, covariance matrix of the re-sampled parameters

*retcode* scalar, return code. If normal convergence is achieved, then *retcode* = 0, otherwise a positive integer is returned indicating the reason for the abnormal termination:

0 normal convergence



- 1 forced exit.
- 2 maximum iterations exceeded.
- 3 function calculation failed.
- 4 gradient calculation failed.
- 5 Hessian calculation failed.
- 6 line search failed.
- 7 function cannot be evaluated at initial parameter values.
- 8 error with gradient
- 9 gradient vector transposed
- 10 secant update failed
- 11 maximum time exceeded
- 12 error with weights
- 34 data set could not be opened.
- 99 termination condition unknown.

## ■ Globals

The **MAXLIK** procedure global variables are also relevant.

**\_\_max\_NumCat** scalar, number of categories in profile table. Default = 16.

**\_\_max\_Increment**  $K \times 1$  vector, increments for cutting points, default is  $2 * \text{__max_Width} * \text{std dev} / \text{__max\_NumCat}$ . If scalar zero, increments are computed by **MAXProfile**.

**\_\_max\_Center**  $K \times 1$  vector, value of center category in profile table. Default values are coefficient estimates.

**\_\_max\_Select** selection vector for selecting coefficients to be included in profiling, for example

```
__max_Select = { 1, 3, 4 };
```

selects the 1st, 3rd, and 4th parameters for profiling.

**\_\_max\_Width** scalar, width of profile table in units of the standard deviations of the parameters. Default = 2.

## ■ Remarks

For each pair of the selected parameters, three plots are printed to the screen. Two of the are the profile t trace plots that describe the univariate profiles of the parameters, and one of them is the profile likelihood trace describing the joint distribution of the

## MAXProfile

### 3. *MAXIMUM LIKELIHOOD REFERENCE*

two parameters. Ideally distributed parameters would have univariate profile t traces that are straight lines, and bivariate likelihood profile traces that are two straight lines intersecting at right angles. This ideal is generally not met by nonlinear models, however, large deviations from the ideal indicate serious problems with the usual statistical inference.

- **Source**

`maxprof.src`

## ■ Purpose

Computes profile likelihood confidence limits

## ■ Library

maxlik

## ■ Format

`cl = MAXPflLimits(b,f,dataset,vars,&fct)`

## ■ Input

- b*             $K \times 1$  vector, maximum likelihood estimates
- f*            scalar, function at minimum (mean log-likelihood)
- dataset*    string containing name of **GAUSS** data set  
               – or –  
                $N \times NV$  matrix, data
- vars*         $NV \times 1$  character vector, labels of variables selected for analysis  
               – or –  
                $NV \times 1$  numeric vector, indices of variables selected for analysis.  
               If *dataset* is a matrix, *vars* may be a character vector containing either the standard labels created by **MAXPflLimits** (i.e., either V1, V2,..., or V01, V02,...). See discussion of the global variable **\_\_vpad** below, or the user-provided labels in **\_\_altnam**).
- &*fct*        a pointer to a procedure that returns either the log-likelihood for one observation or a vector of log-likelihoods for a matrix of observations (see discussion of the global variable **\_\_row** in global variable section below).

## ■ Output

- cl*             $K \times 2$  vector, upper (first column) and lower (second column) confidence limits for the parameters in *b*

## ■ Globals

- \_\_max\_Alpha** (1-**\_\_max\_Alpha**)% confidence limits are computed. The default is .05
- \_\_max\_NumObs** scalar, number of observations. Must be set. If the call to **MaxPflLimits** comes after a call to **MAXLIK**, it will be set by **MAXLIK**.

## MAXPflimits

### 3. MAXIMUM LIKELIHOOD REFERENCE

**\_max\_Select** selection vector for selecting parameters for analysis. For example,

```
_max_Select = { 1, 3, 4 };
```

selects the 1st, 3rd, and 4th parameters for limits.

#### ■ Remarks

**MAXPflimits** computes profile likelihood confidence limits given a maximum likelihood estimation.  $b$  and  $f$  should be returns from a call to **MAXLIK**. This will also properly set up **\_max\_NumObs** for **MAXPflimits**.

**MAXPflimits** solves for the confidence limits as a parametric likelihood problem. Thus it itself calls **MAXLIK** several times for each confidence limit. The screen output is turned off for these runs. However, the computation can be time consuming, and if you wish to check on its progress, type O, or Alt-O, and revise the **\_\_OUTPUT** global. This will turn on the screen output for that run. The parameter number is printed on the title and this will tell you what parameter it is presently working on.

#### ■ Source

maxpflcl.src

## ■ Purpose

Formats and prints the output from a call to **MAXLIK**.

## ■ Library

maxlik

## ■ Format

`{ x,f,g,h,retcode } = MAXPrt(x,f,g,h,retcode);`

## ■ Input

|                |                                                      |
|----------------|------------------------------------------------------|
| <i>x</i>       | $K \times 1$ vector, parameter estimates             |
| <i>f</i>       | scalar, value of function at minimum                 |
| <i>g</i>       | $K \times 1$ vector, gradient evaluated at <i>x</i>  |
| <i>h</i>       | $K \times K$ matrix, covariance matrix of parameters |
| <i>retcode</i> | scalar, return code.                                 |

## ■ Output

The input arguments are returned unchanged.

## ■ Globals

**\_\_header** string. This is used by the printing procedure to display information about the date, time, version of module, etc. The string can contain one or more of the following characters:

|     |                                   |          |
|-----|-----------------------------------|----------|
| "t" | print title (see <b>__title</b> ) | Example: |
| "l" | bracket title with lines          |          |
| "d" | print date and time               |          |
| "v" | print version number of program   |          |
| "f" | print file name being analyzed    |          |

`__header = "tld";`

Default = "tldvf".

**\_\_title** string, message printed at the top of the screen and printed out by **MAXPrt**. Default = "".

## ■ Remarks

The call to **MAXLIK** can be nested in the call to **MAXPrt**:

```
{ x,f,g,h,retcode } = MAXPrt(MAXLIK(dataset,vars,&fct,start));
```

## ■ Source

maxlik.src

## MAXSet

### 3. *MAXIMUM LIKELIHOOD REFERENCE*

- **Purpose**

Resets *MAXIMUM LIKELIHOOD* global variables to default values.

- **Library**

maxlik

- **Format**

MAXSet;

- **Input**

None

- **Output**

None

- **Remarks**

Putting this instruction at the top of all command files that invoke **MAXLIK** is generally good practice. This prevents globals from being inappropriately defined when a command file is run several times or when a command file is run after another command file has executed that calls **MAXLIK**.

- **Source**

maxlik.src

### ■ Purpose

computes Wald confidence limits

### ■ Library

maxlik

### ■ Format

$cl = \text{MAXTlimits}(b, cov)$

### ■ Input

$b$   $K \times 1$  vector, parameter estimates

$cov$   $K \times K$  matrix, covariance matrix of parameter estimates

### ■ Output

$cl$   $K \times 2$  matrix, lower (first column) and upper (second column) confidence limits of the selected parameters

### ■ Globals

**\_max\_Alpha** (1-**\_max\_Alpha**)% confidence limits are computed. The default is .05

**\_max\_NumObs** scalar, number of observations. Must be set.

**\_max\_Select** selection vector for selecting coefficients to be included in profiling, for example

```
_max_Select = { 1, 3, 4 };
```

selects the 1st, 3rd, and 4th parameters for profiling.

### ■ Remarks

**MAXTlimits** returns  $b[i] \pm t(\_max\_NumObs - K; \_max\_Alpha/2) \times \sqrt{cov[i, i]}$

The global **\_max\_NumObs** must be set. If **MAXTlimits** is called immediately after a call to **MAXLIK**, **\_max\_NumObs** will be set by **MAXLIK**.

### ■ Source

maxlik.src

**MAXTlimits**

3. *MAXIMUM LIKELIHOOD REFERENCE*



## Chapter 4

# Event Count and Duration Regression

by

Gary King

Department of Government

Harvard University

This module contains procedures for estimating statistical models of event count or duration data.

The programs included in this module implement maximum likelihood estimators for parametric statistical models of events data. Data based on events come in two forms: *event counts* and *durations* between events. Event counts are dependent variables that take on only nonnegative integer values, such as the number of wars in a year, the number of medical consultations in a month, the number of patents per firm, or even the frequency in the cell of a contingency table. Dependent variables that are measured as durations between events measure time and may take on any nonnegative real number; examples include the duration of parliamentary coalitions or time between coups d'état. Note that the *same* underlying phenomena may be represented as either event counts (e.g., number of wars) or durations (time between wars), and some of the programs included in the **COUNT** module enable you to estimate exactly the same parameters with either form of data.

A variety of statistical models have been proposed to analyze events data, and the programs here provide some that I have developed, along with others I have found

#### 4. EVENT COUNT AND DURATION REGRESSION

particularly useful in my research. I list here the specific programs included in this module, the models each program can estimate, and citations to the work for which I wrote each program. More complete references to the literature on event count and duration models appear at the end of this document.

|                 |                                                                                                                                                                                                                                                                      |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Poisson</b>  | Poisson regression (King, 1988, 1987), truncated Poisson regression (1989d: Section 5), and log-linear and log-proportion models for contingency tables (1989a: Chapter 6).                                                                                          |
| <b>Negbin</b>   | Negative binomial regression (1989b), truncated negative binomial regression (1989d: Section 5), truncated or untruncated variance function models (1989d: Section 5), overdispersed log-linear and log-proportion models for contingency tables (1989a: Chapter 6). |
| <b>Hurdlep</b>  | Hurdle Poisson regression model (1989d: Section 4).                                                                                                                                                                                                                  |
| <b>Supreme</b>  | Seemingly unrelated Poisson regression model (1989c).                                                                                                                                                                                                                |
| <b>Supreme2</b> | Poisson regression model with unobserved dependent variables (1989d: Section 6).                                                                                                                                                                                     |
| <b>Expon</b>    | Exponential duration model with or without censoring (King, Alt, Burns, and Laver, 1989).                                                                                                                                                                            |
| <b>Expgam</b>   | Exponential-Gamma duration model with or without censoring (King, Alt, Burns, and Laver, 1989).                                                                                                                                                                      |
| <b>Pareto</b>   | Pareto duration model with or without censoring (King, Alt, Burns, and Laver, 1989).                                                                                                                                                                                 |

### 4.1 Getting Started

**GAUSS 3.6.18+** is required to use these routines.

#### 4.1.1 README Files

The file **README.cn** contains any last minute information on this module. Please read it before using the procedures in this module.

#### 4.1.2 Setup

In order to use the procedures in the **COUNT** Module, the **COUNT** library must be active. This is done by including **count** in the **LIBRARY** statement at the top of your program or command file:

#### 4. EVENT COUNT AND DURATION REGRESSION

```
library count,quantal,pgraph;
```

This enables **GAUSS** to find the **COUNT** and required *MAXIMUM LIKELIHOOD* procedures. If you plan to make any right hand references to the global variables (which are described in a later section), you also need the statement:

```
#include count.ext;
```

To reset global variables in succeeding executions of the command file, the following instruction can be used:

```
countset;
```

This could be included with the above statements without harm and would insure the proper definition of the global variables for all executions of the command file.

The version number of each module is stored in a global variable. For the **COUNT** Module, this global is:

**\_\_cn\_ver**      3×1 matrix, the first element contains the major version number of the **COUNT** Module, the second element the minor version number, and the third element the revision number.

If you call for technical support, you may be asked for the version number of your copy of this module.

### 4.2 About the COUNT Procedures

The format of the programs included in this module are all very similar:

```
{ b,vc,llik } = Expon(dataset,dep,ind);
{ b,vc,llik } = Expgam(dataset,dep,ind);
{ b,vc,llik } = Pareto(dataset,dep,ind);
{ b,vc,llik } = Poisson(dataset,dep,ind);
{ b,vc,llik } = Negbin(dataset,dep,ind1,ind2);
{ b,vc,llik } = Hurdlep(dataset,dep,ind1,ind2);
{ b,vc,llik } = Supreme(dataset,dep1,dep2,ind1,ind2);
{ b,vc,llik } = Supreme2(dataset,dep1,dep2,ind1,ind2,ind3);
```

An example program file looks like this:

```
library count;
CountSet;
dep = { wars };
ind = { age, party, unem };
dataset = "sample";

call Poisson(dataset,dep,ind);
```

You may run these lines, or ones like them, from the **GAUSS** editor or interactively in command mode.

### 4.2.1 Inputs

The variable *dataset* is always the first argument. This may either be a matrix or a string containing the name of a **GAUSS** data set.

The dependent variable (or variables) is specified in each program by naming a symbol or a column number. For example,

```
dep = { durat };
```

or

```
dep = 7;
```

The independent variable vector (or vectors) is also specified in each program with variable names or column numbers. For example,

```
ind = { age, sex, race, height, size, iq };
```

or

```
ind = { 2, 4, 5, 6, 7 };
```

For each procedure, the data set and dependent variables must be specified. However, since constant terms are automatically included as part of independent variable vectors, you may occasionally wish to include no additional independent variables. You may do this easily by setting the relevant vector to zero. For example,  $ind = 0$ . For another example, you may wish to run the negative binomial regression model with a scalar dispersion parameter rather than a variance function:  $ind2 = 0$ .

### 4.2.2 Outputs

Printed output is controlled by the global **\_\_\_output**, described in the section below. This section describes the outputs *b*, *vc*, and *llik* on the left hand side of the expressions above.

*b*            vector, the maximum likelihood estimates for all the parameters. The mean vector comes first; the variance function, other mean vectors, and scalar dispersion parameters, if any, come next.

*vc*            matrix, the variance-covariance matrix evaluated at the maximum. The standard errors are **SQRT(DIAG(vc))**. If you choose the global **\_\_\_CovPar = 3**, *vc* contains heteroskedastic-consistent parameter estimates.. See Section 2.7 for more discussion of options for statistical inference in maximum likelihood models.

*llik*          scalar, the value of the log-likelihood function at the maximum.

## 4. EVENT COUNT AND DURATION REGRESSION

### 4.2.3 Global Control Variables

**\_\_cn\_\_Inference** scalar character. Determines the type of statistical inference.

**BOOT** generates bootstrapped estimates and covariance matrix of estimates

**MAXLIK** generates maximum likelihood estimates

Setting **\_\_cn\_\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_\_Censor** scalar, allows you to include a variable indicating which observations are censored. This is used by the exponential, exponential-gamma, and Pareto models of duration data. Alternatively, you may set it to a symbol **\_\_cn\_\_Censor** = “varname” if you are using a **GAUSS** data set, or a number (**\_\_cn\_\_Censor** = 11) if the data set is a matrix in memory. The censoring variable should be 0 for censored observations and 1 for others.

By default, no observations are censored.

**\_\_cn\_\_Fix** scalar, name of index number of logged variable among the regressors with coefficient fixed to 1.0. By default, no logged variables are included.

In some of the programs, you have the option of including the log of a variable and fixing its coefficient to 1.0. To include the variable (the program takes the log), set **\_\_cn\_\_Fix** to a variable name or number (**\_\_cn\_\_Fix** = “totals” or **\_\_cn\_\_Fix** = 12). The default (**\_\_cn\_\_Fix** = 0) includes no additional variable. In most event count data, the observation period is the same length for all  $i$  (a year, month, etc.). However, in others, the observation period varies. For example, suppose one observed the number of times a citizen was contacted by a candidate in the interval between two public opinion polls; since polls typically take some time to administer, the observation period would vary over the individuals. In still other situations, the observation period may be the same length but the population of potential events varies. For example, if one observed the number of suicides per state, one would need some way to include information on differing state sizes in the analysis. It turns out that both of these situations can be dealt with in the same way by including an additional variable in the stochastic portion of the model. But (as explained in King, 1989, Section 5.8), this procedure turns out to be mathematically equivalent to including the log of this additional variable in the regression component, and constraining its coefficient to

#### 4. EVENT COUNT AND DURATION REGRESSION

1.0. There is often little harm in just including the log of this variable and estimating its coefficient with all the others, but several of the programs allow one to make this constraint.

**\_\_cn\_\_Dispersion** scalar, set this to a value to change the starting value for only the dispersion parameter in the negative binomial (**Negbin**), generalized event count (**Hurdlep**), exponential-gamma (**Expgam**), Pareto (**Pareto**), and seemingly-unrelated Poisson models (**Supreme**, **Supreme2**). By default, a special starting value is not used for the dispersion parameter.

**\_\_cn\_\_Precision** scalar, the number of digits printed to the right of the decimal point on output. Default = 4.

**\_\_cn\_\_Start** scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.
- 1** uses a vector of user supplied start values stored in the global variable **\_\_cn\_\_StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_\_StartValue**  $L \times 1$  vector, start values if **\_\_cn\_\_Start** = 1.

**\_\_cn\_\_ZeroTruncate** scalar, specifies whether or not the model is a truncated model. For the Poisson and negative binomial models, **\_\_cn\_\_ZeroTruncate** = 0 estimates a truncated-at-zero version of the model. By default, the regular untruncated model is estimated.

**\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to a **COUNT** procedure. When a data matrix is passed to a **COUNT** procedure and the user is selecting from that matrix, the global variable **\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

**\_\_output** scalar, determines printing of intermediate results.

- 0** nothing is written.
- 1** serial ASCII output format suitable for disk files or printers.
- 2** (DOS only) output is suitable for screen only. ANSI.SYS must be active.

Default = 2.

#### 4. EVENT COUNT AND DURATION REGRESSION

**\_\_\_row** scalar, specifies how many rows of the data set are read per iteration of the read loop. By default, the number of rows to be read is calculated automatically.

**\_\_\_rowfac** scalar, “row factor”. If a **COUNT** procedure fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
___rowfac = 0.8;
```

causes **GAUSS** to read in 80% of the rows originally calculated.

This global has an affect only when **\_\_\_row** = 0.

Default = 1.

**\_\_\_title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = “”.

**\_\_\_vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:

**0** Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....

**1** Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.

#### 4.2.4 Statistical Inference

**MAXLIK** statistical inference features may be accessed through the **COUNT** global, **\_\_cn\_Inference**. **\_\_cn\_Inference** has the following settings:

|        |                              |
|--------|------------------------------|
| maxlik | maximum likelihood estimates |
| boot   | bootstrapped estimates       |

That is to generate bootstrapped estimates, set

```
__cn_Inference = "boot";
```

### Bootstrapping

In addition to the usual standard errors, you may generate bootstrap standard errors. Setting `_cn_Inference = BOOT` causes **COUNT** to call **MAXBoot**. This generates bootstrapped estimates and covariance matrices of the estimates.

The bootstrapped parameters are also stored in a **GAUSS** data set. The name of the data set can be determined by setting `_max_BootFname` to a file name, or by default it will be set to `BOOT#` where `#` is a four digit number incremented from 0001 until a name not in use is found. For further details about the bootstrap, see Section 2.7.4.

The data set thus generated can be used for computing confidence intervals of the coefficients using **MAXBlimits**. Also, density estimates and plots can be generated using **MAXDensity**, and histograms and surface plots of the coefficients can be produced using **MAXHist**. For further details about **MAXDensity**, see Section 2.7.4, and for further details about **MAXHist** see Section 2.7.4.

### 4.2.5 Problems with Convergence

All the programs use maximum likelihood estimation by numerically maximizing a different likelihood function. As with virtually all nonlinear iterative procedures, convergence works most of the time, but not every time. Problems to be aware of include the following:

1. The explanatory variables in each regression function should not be highly collinear among themselves.
2. The model should have more observations than parameters; indeed, the more observations, the better.
3. Starting values should not be too far from the optimal values.
4. The model specified should fit the data.
5. The Poisson hurdle model must have at least some observations with  $y_i = 0$  and should take on at least two other values greater than zero.
6. The truncated models should have no observations with zeros (if inadvertently included, a message appears and the program stops).
7. The models with scalar dispersion parameters and variance functions should have maximum likelihood estimates that are bounded so that, for example, in the negative binomial model  $\hat{\sigma}^2 > 1$



#### 4. EVENT COUNT AND DURATION REGRESSION

If you avoid the potential problems listed in the last paragraph, you should have little problem with convergence. Of course, avoiding these problems with difficult data sets is not always easy nor obvious. In these cases, problems may be indicated by the following situations:

1. iterations sending the parameters off in unreasonable directions or creating very large numbers.
2. the program actually bombing out.
3. a single iteration taking an extraordinarily long time.
4. the program taking more than 40 or 50 iterations with no convergence.

If one of these problems occur, you have several options. First, look over the list in the last paragraph. To verify that the problem does indeed exist, you might try running your data on the Poisson regression model if you have event count data, or the exponential regression model if you have duration data. Both are known to be globally concave and tend to converge very easily. If this model works, but another does not, you probably do have a problem.

In the case of problems, you must consider iteration a participatory process. When is iterating, you can press **Alt-H** to receive a list of options that may be changed during iteration. See *MAXLIK REFERENCE* for a full explanation of each. I find that the following practices tend to work well:

1. If the program has produced many iterations without much progress, try pressing **Alt-I** every few iterations to force the program to calculate the information matrix or switch Newton-Raphson iterations. Either of these may not work if the iterations are not far enough along.
2. The number of zeros to the right of the decimal point on the relative gradients (printed on the screen while the program is iterating) is the approximate precision of your final estimates. If the program is having trouble converging, but the gradients are small enough (i.e., you have sufficient precision for your substantive problem), press **Alt-C** to force the program to converge.
3. If the program bombs out very quickly, changing the starting values are your best bet (with the global **\_cn\_Start**). The default starting values created with least squares, **\_cn\_Start = 0**, usually works best. If that does not work, you can also try creating them yourself, by thinking about what the answer is likely to be or by running a simpler model. For example, the exponential-gamma model is sometimes problematic; however, the exponential model often provides good starting values for the effect parameters. Thus if the other methods do not work, you might try the following:

#### 4. EVENT COUNT AND DURATION REGRESSION

```
library count;
CountSet;
dep = { durat };
ind = { unem, infl, age };
dataset = "datafile";
{ b,vc,llik } = Expon(dataset,dep,ind);
_cn_StartValue = b;
_cn_Start = 1;
call Expgam(dataset,dep,ind);
```

You can also choose one of the other methods of creating starting values by changing the `_cn_Start` global (described above).

### 4.3 Annotated Bibliography

- Allison, Paul. 1984. *Event History Analysis*. Beverly Hills: Sage. [A simple overview of event history methods for duration data.]
- Bishop, Yvonne M.M.; Stephen E. Fienberg; and Paul W. Holland. 1975. *Discrete Multivariate Analysis* Cambridge, Mass.: M.I.T. Press. [Models for contingency tables.]
- Cameron, A. Colin and Pravin K. Trivedi. 1986. "Econometric Models Based on Count Data: Comparisons and Applications of Some Estimators and Tests," *Journal of Applied Econometrics* 1, 29–53. [Review of the econometric literature on event counts.]
- Grogger, Jeffrey T. and Richard T. Carson. 1988. "Models for Counts from Choice Based Samples," Discussion Paper 88-9, Department of Economics, University of California, San Diego. [Truncated event count models.]
- Gourieroux, C.; A. Monfort; and A. Trognon. 1984. "Pseudo Maximum Likelihood Methods: Applications to Poisson Models," *Econometrica* 52: 701–720. [A three-stage robust estimation method for count data.]
- Hall, Bronwyn H.; Zvi Griliches; and Jerry A. Hausman. 1986. "Patents and R and D: Is there a Lag?" *International Economic Review*. 27, 2 (June): 265–83. [Nice example of a applying a variety of different estimators to single equation count models.]
- Hausman, Jerry; Bronwyn H. Hall; and Zvi Griliches. 1984. "Econometrics Models for Count Data with An Application to the Patents-R&D Relationship," *Econometrica*. 52, 4 (July): 909-938. [Count models for time-series cross sectional panels.]

#### 4. EVENT COUNT AND DURATION REGRESSION

- Holden, Robert T. 1987. "Time Series Analysis of a Contagious Process," *Journal of the American Statistical Association*. 82, 400 (December): 1019–1026. [A time series model of count data applied to airline hijack attempts.]
- Jorgenson, Dale W. 1961. "Multiple Regression Analysis of a Poisson Process," *Journal of the American Statistical Association* 56,294 (June): 235–45. [The Poisson regression model.]
- Kalbfleisch, J.D. and R.L. Prentice. 1980. *The Statistical Analysis of Failure Time Data*. New York: Wiley. [Summary of research on many models of duration data.]
- King, Gary. 1989a. *Unifying Political Methodology: The Likelihood Theory of Statistical Inference*. New York: Cambridge University Press. [Introduction to likelihood, maximum likelihood, and a large variety of statistical models as special cases; Chapter 5 is discrete regression models.]
- . 1989b. "Variance Specification in Event Count Models: From Restrictive Assumptions to a Generalized Estimator," *American Journal of Political Science*, vol. 33, no. 3 (August):762-784. [Poisson-based models with over- and under-dispersion.]
- . 1989c. "A Seemingly Unrelated Poisson Regression Model," *Sociological Methods and Research*. 17, 3 (February, 1989): 235–255. [A model for simultaneously analyzing a pair of event count variables in a SURM framework.]
- . 1989d. "Event Count Models for International Relations: Generalizations and Applications," *International Studies Quarterly*, vol. 33, no. 3 (June):123-147. [Hurdle models, truncated models, and models with unobserved dependent variables, all for event count data.]
- . 1988. "Statistical Models for Political Science Event Counts: Bias in Conventional Procedures and Evidence for The Exponential Poisson Regression Model," *American Journal of Political Science*, 32, 3 (August): 838–863. [Introduction to count models; analytical and Monte Carlo comparisons of LS, logged-LS, and Poisson regression models.]
- . 1987. "Presidential Appointments to the Supreme Court: Adding Systematic Explanation to Probabilistic Description," *American Politics Quarterly*, 15, 3 (July): 373–386. [An application of the Poisson regression model.]
- King, Gary; James Alt; Nancy Burns; Michael Laver. 1990. "A Unified Model of Cabinet Duration in Parliamentary Democracies," *American Journal of Political Science*, vol. 34, no. 3 (August):846-871. [Exponential model of duration data with censoring.]

#### 4. EVENT COUNT AND DURATION REGRESSION

- McCullagh, P. And J.A. Nelder 1983. *Generalized Linear Models*. London: Chapman and Hall. [A unified approach to specifying and estimating this class of models. Some count and duration models are covered.]
- Mullahy, John. 1986. "Specification and Testing of Some Modified Count Data Models," *Journal of Econometrics*. 33: 341-65. [Several hurdle-type models of event count data.]
- Tuma, Nancy Brandon and Michael T. Hannan. 1984. *Social Dynamics*. New York: Academic Press.

# Chapter 5

## Count Reference

- **Purpose**

Formats and prints the output from calls to **COUNT** procedures with confidence limits

- **Library**

count

- **Format**

`{ b,cl,llik } = CountCLPrt(b,cl,llik);`

- **Input**

*b*  $(K+1) \times 1$  vector, maximum likelihood estimates of the effect parameters stacked on top of the dispersion parameter.

*cl*  $(K + 1) \times 2$  matrix, confidence limits

*llik* scalar, value of the log-likelihood function at the maximum.

- **Output**

The input arguments are returned unchanged.

- **Remarks**

Confidence limits computed by **MAXBLimits** may be passed in the fourth argument in the call to **CountCLPrt**:

```
_cn_Inference = { boot };
{ b,vc,llik } = Expgam(dataset,dep,ind);
cl = MAXBLimits(_max_BootFname);
call CountCLPrt(b,cl,llik);
```

- **Source**

count.src

## ■ Purpose

Formats and prints the output from calls to **COUNT** procedures.

## ■ Library

count

## ■ Format

```
{ b,vc,llik } = CountPrt(b,vc,llik);
```

## ■ Input

*b* (K+1)×1 vector, maximum likelihood estimates of the effect parameters stacked on top of the dispersion parameter.

*vc* (K+1)×(K+1) matrix, variance-covariance matrix of the estimated parameters

*llik* scalar, value of the log-likelihood function at the maximum.

## ■ Output

The input arguments are returned unchanged.

## ■ Remarks

The call to **COUNT** procedures can be nested in the call to the **CountPrt**:

```
{ b,vc,llik } = countprt(ExpGam(dataset,dep,ind));
```

## ■ Source

count.src

## CountSet

- **Purpose**

Resets **COUNT** global variables to default values.

- **Library**

count

- **Format**

**CountSet;**

- **Input**

None

- **Output**

None

- **Remarks**

Putting this instruction at the top of all command files that invoke **COUNT** procedures is generally good practice. This prevents globals from being inappropriately defined when a command file is run several times or when a command file is run after another command file has executed that calls a **COUNT** procedure.

**CountSet** calls **Set** which calls **GAUSSET**.

- **Source**

count.src



## ■ Purpose

Estimates an exponential-gamma regression model, for the analysis of duration data, with maximum likelihood.

## ■ Library

count

## ■ Format

$\{ b, vc, llik \} = \text{Expgam}(\text{dataset}, \text{dep}, \text{ind});$

## ■ Input

*dataset* string, name of **GAUSS** data set.  
 – or –  
 N×K matrix, data

*dep* string, the name of the dependent variable.  
 – or –  
 scalar, the index of the dependent variable.

*ind* K×1 character vector, names of the independent variables.  
 – or –  
 K×1 numeric vector, indices of independent variables.  
 Set to 0 to include only a constant term.

If *dataset* is a matrix, *dep* or *ind* may be a string or character variable containing either the standard labels created by (V1, V2, ..., or V01, V02, ..., depending on the value of `___vpad`), or the user-provided labels in `___altnam`.

## ■ Output

*b* (K+1)×1 vector, maximum likelihood estimates of the effect parameters stacked on top of the dispersion parameter.

*vc* (K+1)×(K+1) matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If you choose the global `___CovPar = 3`, *vc* contains heteroskedastic-consistent parameter estimates.

*llik* scalar, value of the log-likelihood function at the maximum.

## ■ Globals

**MAXLIK** globals are also relevant

**\_\_cn\_\_Inference** string, determines the type of statistical inference.

**boot** generates bootstrapped estimates and covariance matrix of estimates

**maxlik** generates maximum likelihood estimates (default)

Setting **\_\_cn\_\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_\_Censor** string, the name of the censor variable from *dataset*.

– or –

scalar, the index of the censor variable from *dataset*.

By default, no censoring is used.

**\_\_cn\_\_Start** scalar, selects method of calculating starting values. Possible values are:

**0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.

**1** uses a vector of user supplied start values stored in the global variable **\_\_cn\_\_StartValue**.

**2** uses a vector of zeros.

**3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_\_StartValue**  $(K+1) \times 1$  vector, start values if **\_\_cn\_\_Start** = 1.

**\_\_cn\_\_Precision** scalar, number of decimal points to print on output. Default = 4.

**\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to **Expgam**. When a data matrix is passed to **Expgam** and when the user is selecting from that matrix, the global variable **\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

**\_\_miss** scalar, determines how missing data will be handled.

**0** Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.

**1** Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.

Default = 0.

**\_\_output** scalar, determines printing of intermediate results.

- 0 nothing is written.
- 1 serial ASCII output format suitable for disk files or printers.
- 2 (DOS only) output is suitable for screen only. ANSI.SYS must be active.

Default = 2.

**\_\_row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.

**\_\_rowfac** scalar, “row factor”. If **Expgam** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
__rowfac = 0.8;
```

will cause **GAUSS** to read in 80% of the rows originally calculated.

This global has an affect only when **\_\_row** = 0.

Default = 1.

**\_\_title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = “”.

**\_\_vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:

- 0 Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
- 1 Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.

## ■ Remarks

Let the  $n$  duration observations (nonnegative real numbers) for the dependent variable be denoted as  $y_1, \dots, y_n$ . Assume that  $y_i$  follows a gamma distribution with expected value  $\mu_i$  and variance  $\mu_i^2 \sigma^2$ . Let the mean  $\mu_i$  be an exponential-linear function of a vector of explanatory variables,  $x_i$ :

$$E(y_i) \equiv \mu_i = \exp(x_i \beta) \quad (5.1)$$

The program includes a constant term as the first column of  $x_i$  and allows one to include any number of explanatory variables. Note that  $\mu_i$  from a duration model equals  $1/\lambda_i$  from an event count model; thus, one need only change the sign of the effect parameters to get estimates of the same parameters from these different kinds of data.

The dispersion  $\sigma^2$  is parametrized as follows:

$$\sigma_i^2 = \exp(\gamma) \quad (5.2)$$

**EXPGAM** reports estimates of  $\beta$  and  $\gamma$ .

For an introduction to the exponential gamma regression model see King, Alt, Burns, and Laver (1989) or Kalbfleisch and Prentice (1980).

## ■ Example

Exponential-Gamma Regression Model of Duration Data

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty, allianc };
{ b,vc,llik } = Expgam(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

A vector of effect parameters and a scalar dispersion parameter are estimated. The vector includes one element corresponding to each explanatory variable named in *ind* and a constant term. Five parameters are estimated in this example.

Censored Exponential-Gamma Regression Model of Duration Data

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty, allianc };
_Censor = { v12 };
{ b,vc,llik } = Expgam(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

## 5. COUNT REFERENCE

## Expgam

A vector of effect parameters and a scalar dispersion parameter are estimated. The vector includes one element corresponding to each explanatory variable named in *ind* and a constant term. Five parameters are estimated in this example.

### ■ Source

`expgam.src`

## ■ Purpose

Estimates an exponential regression model or censored exponential regression model with maximum likelihood.

## ■ Library

count

## ■ Format

$\{ b, vc, llik \} = \text{Expon}(\text{dataset}, \text{dep}, \text{ind});$

## ■ Input

*dataset*      string, name of **GAUSS** data set.  
                   – or –  
                   N×K matrix, data

*dep*            string, the name of the dependent variable  
                   – or –  
                   scalar, the index of the dependent variable

*ind*            K×1 character vector, names of the independent variables  
                   – or –  
                   K×1 numeric vector, indices of independent variables  
                   Set to 0 to include only a constant term.

If *dataset* is a matrix, *dep* or *ind* may be a string or character variable containing either the standard labels created by (V1, V2,..., or V01, V02,..., depending on the value of `___vpad`), or the user-provided labels in `___altnam`.

## ■ Output

*b*              K×1 vector, maximum likelihood estimates of the effect parameters.

*vc*            K×K matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If the global `___CovPar` is set to 3, *vc* will contain heteroskedastic-consistent parameter estimates.

*llik*          scalar, value of the log-likelihood function at the maximum.

## ■ Globals

**MAXLIK** globals are also relevant.

**\_\_cn\_\_Inference** string, determines the type of statistical inference.

**boot** generates bootstrapped estimates and covariance matrix of estimates

**maxlik** generates maximum likelihood estimates

Setting **\_\_cn\_\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_\_Censor** string, the name of the censor variable from *dataset*  
 – or –  
 scalar, the index of the censor variable from *dataset*  
 By default, no censoring is used.

**\_\_cn\_\_Start** scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.
- 1** will use a vector of user supplied start values stored in the global variable **\_\_cn\_\_StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_\_StartValue**  $K \times 1$  vector, start values if **\_\_cn\_\_Start** = 1.

**\_\_cn\_\_Precision** scalar, number of decimal points to print on output. Default = 4.

**\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to **Expon**. When a data matrix is passed to **Expon** and the user is selecting from that matrix, the global variable **\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

**\_\_miss** scalar, determines how missing data will be handled.

- 0** Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.
- 1** Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.

Default = 0.

- \_\_\_output** scalar, determines printing of intermediate results.
- 0** nothing is written.
  - 1** serial ASCII output format suitable for disk files or printers.
  - 2** (DOS only) output is suitable for screen only. ANSI.SYS must be active.
- Default = 2.
- \_\_\_row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.
- \_\_\_rowfac** scalar, “row factor”. If **EXPON** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting
- ```
___rowfac = 0.8;
```
- will cause **GAUSS** to read in 80% of the rows originally calculated.
- This global only has an affect when **___row** = 0.
- Default = 1.
- ___title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = “”.
- ___vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:
- 0** Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
 - 1** Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.
- Default = 1.

■ Remarks

Let y_i ($i = 1, \dots, n$) take on any non-negative real number representing a duration. Often y_i is only measured as an integer, such as the number of days or months. Even so, if your dependent variable is a measure of time, duration models, and not event count models, are called for. Let y_i be distributed exponentially with mean μ_i . Also let $E(y_i) \equiv \mu_i = \exp(x_i\beta)$. Note that μ_i from a duration model equals $1/\lambda_i$ from an event

5. COUNT REFERENCE

count model; thus, one need only change the sign of the effect parameters to get estimates of the same parameters from these different kinds of data.

For an introduction to the exponential regression model and the censored exponential regression model see Kalbfleisch and Prentice (1980) and King, Alt, Burns, and Laver (1989).

■ Example

Exponential Regression Model

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty, allianc };
{ b,vc,llik } = Expon(dataset,dep,ind);
output file = count.out on;
call CountPrt(b,vc,llik);
output off;
```

A single vector of effect parameters are estimated. This vector includes one element corresponding to each explanatory variable named in *ind* and a constant term.

Censored Exponential Regression Model

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty, allianc };
_cn_Censor = { notseen };
{ b,vc,llik } = Expon(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

A single vector of effect parameters are estimated. This vector includes one element corresponding to each explanatory variable named in *ind* and a constant term.

■ Source

expon.src

■ Purpose

Estimates a hurdle Poisson regression model, for the analysis of event counts, with maximum likelihood.

■ Library

count

■ Format

$\{ b, vc, llik \} = \text{Hurdlep}(dataset, dep, ind);$

■ Input

dataset string, name of **GAUSS** data set.
 – or –
 $N \times K$ matrix, data

dep string, the name of the dependent variable
 – or –
 scalar, the index of the dependent variable

ind1 $K \times 1$ character vector, names of first event independent variables
 – or –
 $K \times 1$ numeric vector, indices of first event independent variables

ind2 $K \times 1$ character vector, names of second event independent variables
 – or –
 $K \times 1$ numeric vector, indices of second event independent variables

If *dataset* is a matrix, *dep*, *ind1*, or *ind2* may be a string or character variable containing either the standard labels created by (V1, V2, ..., or V01, V02, ..., depending on the value of **___vpad**), or the user-provided labels in **___altnam**.

■ Output

b $(K+L) \times 1$ vector, maximum likelihood estimates of the effect parameters stacked on top of the dispersion parameter.

vc $(K+L) \times (K+L)$ matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If you choose the global **___CovPar** = 3, *vc* will contain heteroskedastic-consistent parameter estimates.

llik scalar, value of the log-likelihood function at the maximum.

■ Globals

MAXLIK globals are also relevant

__cn__Inference string, determines the type of statistical inference.

boot generates bootstrapped estimates and covariance matrix of estimates

maxlik generates maximum likelihood estimates

Setting **__cn__Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **__max__BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

__cn__Start scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing $\ln(y + 0.5)$ on the explanatory variables.
- 1** will use a vector of user supplied start values stored in the global variable **__cn__StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval $[-\frac{1}{2}, \frac{1}{2}]$.

Default = 0.

__cn__StartValue $(K+L) \times 1$ vector, start values if **__cn__Start** = 1.

__cn__Precision scalar, number of decimal points to print on output. Default = 4.

__altnam $K \times 1$ vector, alternate names for variables when a matrix is passed to **Hurdlep**. When a data matrix is passed to **Hurdlep** and the user is selecting from that matrix, the global variable **__altnam**, if it is used, must contain names for the columns of the original matrix.

__miss scalar, determines how missing data will be handled.

- 0** Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.
- 1** Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.

Default = 0.

___output scalar, determines printing of intermediate results.

0 nothing is written.

1 serial ASCII output format suitable for disk files or printers.

2 (DOS only) output is suitable for screen only. ANSI.SYS must be active.

Default = 2.

___row scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.

___rowfac scalar, “row factor”. If **Hurdlep** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **___rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
___rowfac = 0.8;
```

will cause **GAUSS** to read in 80% of the rows originally calculated.

This global only has an affect when **___row** = 0.

Default = 1.

___title string, message printed at the top of the screen and printed out by **CountPrt**. Default = “”.

___vpad scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:

- 0** Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
- 1** Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.

■ Remarks

Let the n event count observations (nonnegative integers) for the dependent variable be denoted as y_1, \dots, y_n . y_i is then a random dependent variable representing the number of events that have occurred during observation period i . Let λ_{0i} be the rate of the first event occurrence and λ_{+i} be the rate for all additional events after the first. If these are

the expected values of two separate Poisson processes, we have the hurdle Poisson regression model. These means are parametrized as usual:

$$\lambda_{0i} = \exp(x_i\beta) \quad (5.3)$$

and

$$\lambda_{+i} = \exp(z_i\gamma) \quad (5.4)$$

where x_i and z_i are (possibly) different vectors of explanatory variables. The program produces estimates of β and γ . If $\beta = \gamma$ and $x = z$, this model reduces to the Poisson.

For an introduction to the Hurdle Poisson regression model see Mullahy (1986) and King (1989d).

■ Example

Hurdle Poisson Regression Model:

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind1 = { unem, poverty, allianc };
ind2 = { race, sex, age, partyid, x4, v5 };
{ b,vc,llik } = Hurdlep(dataset,dep,ind1,ind2);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

Two vectors of effect parameters are estimated. Each includes one element corresponding to each explanatory variable plus a constant term (in the example, four parameters appear in the first regression function and seven in the second).

■ Source

hurdlep.src

■ Purpose

Estimates a negative binomial regression model or truncated-at-zero negative binomial regression model with maximum likelihood.

■ Library

count

■ Format

{ *b,vc,llik* } = **Negbin**(*dataset,dep,ind1,ind2*);

■ Input

dataset string, name of **GAUSS** data set.
 – or –
 N×K matrix, data

dep string, the name of the dependent variable
 – or –
 scalar, the index of the dependent variable

ind1 K×1 character vector, names of first event independent variables
 – or –
 K×1 numeric vector, indices of first event independent variables
 Set to 0 to include only a constant term.

ind2 K×1 character vector, names of second event independent variables
 – or –
 K×1 numeric vector, indices of second event independent variables
 Set to 0 for a scalar dispersion parameter.

If *dataset* is a matrix, *dep*, *ind1*, or *ind2* may be a string or character variable containing either the standard labels created by (V1, V2,..., or V01, V02,..., depending on the value of **___vpad**), or the user-provided labels in **___altnam**.

■ Output

b (K+1)×1 or (K+L)×1 vector, maximum likelihood estimates of the effect parameters stacked on top of either the dispersion parameter or the coefficients of the variance function.

vc (K+1)×(K+1) or (K+L)×(K+L) matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If you choose the global **___CovPar** = 3, *vc* will contain heteroskedastic-consistent parameter estimates.

llik scalar, value of the log-likelihood function at the maximum.

■ Globals

MAXLIK globals are also relevant.

__cn__Inference string, determines the type of statistical inference.

boot generates bootstrapped estimates and covariance matrix of estimates

maxlik generates maximum likelihood estimates

Setting **__cn__Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **__max__BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

__cn__Fix scalar, name of index number of logged variable among the regressors with coefficient constrained to 1.0 By default, no logged variables are included.

__cn__Start scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing $\ln(y + 0.5)$ on the explanatory variables.
- 1** will use a vector of user supplied start values stored in the global variable **__cn__StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval $[-\frac{1}{2}, \frac{1}{2}]$.

Default = 0.

__cn__StartValue $(K+1) \times 1$ or $(K+L) \times 1$ vector, start values if **__cn__Start** = 1.

__cn__Dispersion scalar, start value for scalar dispersion parameter. Default = 3.

__cn__Precision scalar, number of decimal points to print on output. Default = 4.

__cn__ZeroTruncate scalar, specifies which model is used:

- 0** truncated-at-zero negative binomial model
- 1** negative binomial model is used.

- __altnam** $K \times 1$ vector, alternate names for variables when a matrix is passed to **Negbin**. When a data matrix is passed to **Negbin** and the user is selecting from that matrix, the global variable **__altnam**, if it is used, must contain names for the columns of the original matrix.
- __miss** scalar, determines how missing data will be handled.
- 0** Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.
 - 1** Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.
- Default = 0.
- __output** scalar, determines printing of intermediate results.
- 0** nothing is written.
 - 1** serial ASCII output format suitable for disk files or printers.
 - 2** (DOS only) output is suitable for screen only. ANSI.SYS must be active.
- Default = 2.
- __row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.
- __rowfac** scalar, "row factor". If **Negbin** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **__rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting
- ```
__rowfac = 0.8;
```
- will cause **GAUSS** to read in 80% of the rows originally calculated. This global only has an affect when **\_\_row** = 0.
- Default = 1.
- \_\_title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = "".
- \_\_vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:
- 0** Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....



- 1 Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.

## ■ Remarks

Let  $y_i$  be a random dependent variable representing the number of events that have occurred during observation period  $i$  ( $i = 1, \dots, n$ ). Assume that  $y_i$  follows a negative binomial distribution with expected value  $\lambda_i$  and variance  $\lambda_i\sigma^2$ . Let the mean  $\lambda_i$  (the rate of event occurrence, which must be greater than zero) be an exponential-linear function of a vector of explanatory variables,  $x_i$ :

$$E(y_i) \equiv \lambda_i = \exp(x_i\beta) \quad (5.5)$$

The program includes a constant term as the first column of  $x_i$  and allows one to include any number of explanatory variables.

$\sigma^2$  is parametrized as follows:

$$\sigma_i^2 = 1 + \exp(z_i\gamma) \quad (5.6)$$

where  $z_i = 1$ , if estimating a scalar dispersion parameter, or a vector of explanatory variables, if estimating a variance function. The program calculates estimates of  $\beta$  and  $\gamma$ .

For an introduction to the negative binomial regression model, see Hausman, Hall, and Griliches (1984) and King (1989b); for information on the truncated negative binomial model, see Grogger and Carson (1988), and on the variance function model with or without truncation see King (1989d: Section 5)

## ■ Example

Negative Binomial Regression Model

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind1 = { unem, poverty, allianc };
{ b,vc,llik } = Negbin(dataset,dep,ind1,0);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

A single vector of effect parameters and one scalar dispersion parameter are estimated. The vector of effect parameters includes one element corresponding to each explanatory variable and a constant term. In the example, five parameters are estimated.

#### Negative Binomial Variance Function Regression Model

```

library count;
#include count.ext;
Countset;
dataset = "wars";
dep1 = { wars };
ind1 = { unem, poverty, allianc };
ind2 = { partyid, x4 };
{ b,vc,llik } = Negbin(dataset,dep,ind1,ind2);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;

```

Two vectors of effect parameters are estimated, one for the mean *ind1* and one for the variance function *ind2*. Each vector includes a constant term and one element corresponding to each explanatory variable. The example estimates seven parameters.

#### Truncated-at-zero Negative Binomial Regression Model

```

library count;
#include count.ext;
Countset;
dataset = "wars";
dep1 = { wars };
ind1 = { unem, poverty, allianc };
_cn_ZeroTruncate = 0;
{ b,vc,llik } = Negbin(dataset,dep,ind1,0);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;

```

A single vector of effect parameters and one scalar dispersion parameter are estimated. The vector of effect parameters includes one element corresponding to each explanatory variable and a constant term. In the example, five parameters are estimated.

#### Truncated-at-zero Negative Binomial Variance Function Regression Model

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep1 = { wars };
ind1 = { unem, poverty, allianc };
ind2 = { partyid, x4 };
_cn_ZeroTruncate = 0;
{ b,vc,llik } = Negbin(dataset,dep,ind1,0);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

Two vectors of effect parameters are estimated, one for the mean and one for the variance function. Each vector includes a constant term and one element corresponding to each explanatory variable. In the example, the variables specified in *ind1* pertain to the expected value and *ind2* to the variance. Seven parameters are estimated.

#### ■ Source

negbin.src

## ■ Purpose

Estimates a Pareto regression model, for the analysis of duration data, with maximum likelihood.

## ■ Library

count

## ■ Format

$\{ b, vc, llik \} = \text{Pareto}(\text{dataset}, \text{dep}, \text{ind});$

## ■ Input

*dataset*      string, name of **GAUSS** data set.  
                   – or –  
                    $N \times K$  matrix, data

*dep*            string, the name of the dependent variable  
                   – or –  
                   scalar, the index of the dependent variable

*ind*             $K \times 1$  character vector, names of the independent variables  
                   – or –  
                    $K \times 1$  numeric vector, indices of independent variables  
                   Set to 0 to include only a constant term.

If *dataset* is a matrix, *dep* and *ind* may be a string or character variable containing either the standard labels created by (V1, V2, ..., or V01, V02, ..., depending on the value of **\_\_\_vpad**), or the user-provided labels in **\_\_\_altnam**.

## ■ Output

*b*               $(K+1) \times 1$  vector, maximum likelihood estimates of the effect parameters stacked on top of the dispersion parameter.

*vc*             $(K+1) \times (K+1)$  matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If the global **\_\_\_CovPar** is set to 3, *vc* will contain heteroskedastic-consistent parameter estimates.

*llik*          scalar, value of the log-likelihood function at the maximum.

## ■ Globals

**MAXLIK** globals are also relevant.

**\_\_cn\_\_Inference** string, determines the type of statistical inference.

**boot** generates bootstrapped estimates and covariance matrix of estimates

**maxlik** generates maximum likelihood estimates

Setting **\_\_cn\_\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_\_Censor** string, the name of the censor variable from *dataset*  
– or –  
scalar, the index of the censor variable from *dataset*

Each element of censor variable is 0 if censored, or 1 if not.

By default, no censoring is used.

**\_\_cn\_\_Start** scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.
- 1** will use a vector of user supplied start values stored in the global variable **\_\_cn\_\_StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_\_StartValue**  $(K+1) \times 1$  vector, start values if **\_\_cn\_\_Start** = 1.

**\_\_cn\_\_Dispersion** scalar, start value for scalar dispersion parameter. Default = 3.

**\_\_cn\_\_Precision** scalar, number of decimal points to print on output. Default = 4.

**\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to **Pareto**. When a data matrix is passed to **Pareto** and the user is selecting from that matrix, the global variable **\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

**\_\_miss** scalar, determines how missing data will be handled.

- 0** Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.

- 1 Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.

Default = 0.

**\_\_\_output** scalar, determines printing of intermediate results.

- 0 nothing is written.
- 1 serial ASCII output format suitable for disk files or printers.
- 2 (DOS only) output is suitable for screen only. ANSI.SYS must be active.

Default = 2.

**\_\_\_row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.

**\_\_\_rowfac** scalar, "row factor". If **Pareto** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
___rowfac = 0.8;
```

will cause **GAUSS** to read in 80% of the rows originally calculated.

This global only has an affect when **\_\_\_row** = 0.

Default = 1.

**\_\_\_title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = "".

**\_\_\_vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:

- 0 Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
- 1 Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.

## ■ Remarks

Let the  $n$  duration observations (non-negative real numbers) for the dependent variable be denoted as  $y_1, \dots, y_n$ . Assume that  $y_i$  follows a Pareto distribution with expected value  $\mu_i$  and variance  $\mu_i\sigma^2 + \mu_i^2$ . Let the mean  $\mu_i$  be an exponential-linear function of a vector of explanatory variables,  $x_i$ :

$$E(y_i) \equiv \mu_i = \exp(x_i\beta) \quad (5.7)$$

The program includes a constant term as the first column of  $x_i$  and allows one to include any number of explanatory variables. Note that  $\mu_i$  from a duration model equals  $1/\lambda_i$  from an event count model; thus, one need only change the sign of the effect parameters to get estimates of the same parameters from these different kinds of data.

The dispersion  $\sigma^2$  is parametrized as follows:

$$\sigma_i^2 = \exp(\gamma) \quad (5.8)$$

The program gives estimates of  $\beta$  and  $\gamma$ .

For an introduction to the Pareto regression model see Hannan and Tuma (1984) and King, Alt, Burns, and Laver (1989).

## ■ Example

Pareto Regression Model of Duration Data

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty, allianc };
{ b,vc,llik } = Pareto(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

A vector of effect parameters and a scalar dispersion parameter are estimated. The vector includes one element corresponding to each explanatory variable named in *ind* and a constant term. Five parameters are estimated in this example.

Censored Pareto Regression Model of Duration Data

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty };
_cn_Censor = { cvar };
{ b,vc,llik } = Pareto(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

A vector of effect parameters and a scalar dispersion parameter are estimated. The vector includes one element corresponding to each explanatory variable named in *ind* and a constant term. Five parameters are estimated in this example.

#### ■ Source

pareto.src



## ■ Purpose

Estimates a Poisson regression model or truncated-at-zero Poisson regression model with maximum likelihood.

## ■ Library

count

## ■ Format

$\{ b, vc, llik \} = \text{Poisson}(dataset, dep, ind);$

## ■ Input

*dataset*      string, name of **GAUSS** data set.  
                   – or –  
                    $N \times K$  matrix, data

*dep*            string, the name of the dependent variable  
                   – or –  
                   scalar, the index of the dependent variable

*ind*             $K \times 1$  character vector, names of the independent variables  
                   – or –  
                    $K \times 1$  numeric vector, indices of independent variables  
                   Set to 0 to include only a constant term.

If *dataset* is a matrix, *dep* and *ind* may be a string or character variable containing either the standard labels created by (V1, V2,..., or V01, V02,..., depending on the value of **\_\_vpad**), or the user-provided labels in **\_\_altnam**.

## ■ Output

*b*               $K \times 1$  vector, maximum likelihood estimates of the effect parameters.

*vc*             $K \times K$  matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If you choose the global **\_\_CovPar** = 3, *vc* will contain heteroskedastic-consistent parameter estimates.

*llik*          scalar, value of the log-likelihood function at the maximum.

## ■ Globals

**MAXLIK** globals are also relevant.

**\_\_cn\_\_Inference** string, determines the type of statistical inference.

**boot** generates bootstrapped estimates and covariance matrix of estimates

**maxlik** generates maximum likelihood estimates

Setting **\_\_cn\_\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_\_Fix** scalar, name of index number of logged variable among the regressors with coefficient constrained to 1.0 By default, no logged variables are included.

**\_\_cn\_\_Start** scalar, selects method of calculating starting values. Possible values are:

**0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.

**1** will use a vector of user supplied start values stored in the global variable **\_\_cn\_\_StartValue**.

**2** uses a vector of zeros.

**3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_\_StartValue**  $K \times 1$  vector, start values if **\_\_cn\_\_Start** = 1.

**\_\_cn\_\_Precision** scalar, number of decimal points to print on output. Default = 4.

**\_\_cn\_\_ZeroTruncate** scalar, specifies which model is used:

**0** truncated-at-zero negative binomial model

**1** negative binomial model is used.

Default = 1.

**\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to **Poisson**. When a data matrix is passed to **Poisson** and the user is selecting from that matrix, the global variable **\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

**\_\_miss** scalar, determines how missing data will be handled.

- 0 Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.
- 1 Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.

Default = 0.

**\_\_\_output** scalar, determines printing of intermediate results.

- 0 nothing is written.
- 1 serial ASCII output format suitable for disk files or printers.
- 2 (DOS only) output is suitable for screen only. ANSI.SYS must be active.

Default = 2.

**\_\_\_row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.

**\_\_\_rowfac** scalar, “row factor”. If **POISSON** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
__rowfac = 0.8;
```

will cause **GAUSS** to read in 80% of the rows originally calculated.

**\_\_\_title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = “”.

**\_\_\_vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:

- 0 Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
- 1 Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.

### ■ Remarks

Let the  $n$  event count observations (non-negative integers) for the dependent variable be denoted as  $y_1, \dots, y_n$ .  $y_i$  is then a random dependent variable representing the number of events that have occurred during observation period  $i$ . By assuming that the events occurring within each period are independent and have constant rates of occurrence,  $y_i$  can be shown to follow a Poisson distribution:

$$f_p(y_i|\lambda_i) = \begin{cases} \frac{e^{-\lambda_i}(\lambda_i)^{y_i}}{y_i!} & \text{for } \lambda_i > 0 \text{ and } y_i = 0, 1, \dots \\ 0 & \text{otherwise} \end{cases} \quad (5.9)$$

with expected value and variance  $\lambda_i$ . Under the Poisson regression model,  $\lambda_i$  (the rate of event occurrence, which must be greater than zero) is assumed to be an exponential-linear function of a vector of explanatory variables,  $x_i$ :

$$E(y_i) \equiv \lambda_i = \exp(x_i\beta) \quad (5.10)$$

The program includes a constant term as the first element of  $x_i$  and allows one to include any number of explanatory variables.

For an introduction to the Poisson regression model see King (1988); on the truncated model, see Grogger and Carson (1988) and King (1989d).

### ■ Example

Poisson Regression Model

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep = { wars };
ind = { unem, poverty, allianc };
{ b,vc,llik } = Poisson(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

Truncated-at-zero Poisson Regression Model

```
library count;
#include count.ext;
Countset;
dataset = "wars";
```

```
dep = { wars };
ind = { unem, poverty, allianc };
_cn_ZeroTruncate = 0;
{ b,vc,llik } = Poisson(dataset,dep,ind);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

■ **Source**

poisson.src

## ■ Purpose

Estimates a seemingly unrelated Poisson regression model, for the analysis of two event **COUNT** variables, with maximum likelihood.

## ■ Library

count

## ■ Format

$\{ b, vc, llik \} = \text{Supreme}(\text{dataset}, \text{dep1}, \text{dep2}, \text{ind1}, \text{ind2});$

## ■ Input

*dataset* string, name of **GAUSS** data set.  
 – or –  
 N×K matrix, data

*dep1* string, name of the first dependent variable  
 – or –  
 scalar, index of the first dependent variable

*dep2* string, name of the second dependent variable  
 – or –  
 scalar, index of the second dependent variable

*ind1* K×1 character vector, names of first event independent variables  
 – or –  
 K×1 numeric vector, indices of first event independent variables  
 Set to 0 to include only a constant term.

*ind2* K×1 character vector, names of second event independent variables  
 – or –  
 K×1 numeric vector, indices of second event independent variables  
 Set to 0 to include only a constant term.

If *dataset* is a matrix, *dep1*, *dep2*, *ind1* and *ind2* may be a string or character variable containing either the standard labels created by (V1, V2, ..., or V01, V02, ..., depending on the value of **\_\_vpad**), or the user-provided labels in **\_\_altnam**.

## ■ Output

*b* (K+L+2)×1 vector, maximum likelihood estimates of the effect parameters of  $\beta$  and  $\gamma$  stacked on top of the covariance parameter  $\xi$ .

- vc*  $(K+L+2) \times (K+L+2)$  matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If you choose the global `___CovPar = 3`, *vc* will contain heteroskedastic-consistent parameter estimates.
- llik* scalar, value of the log-likelihood function at the maximum.

## ■ Globals

**MAXLIK** globals are also relevant.

**\_\_cn\_\_Inference** string, determines the type of statistical inference.

**boot** generates bootstrapped estimates and covariance matrix of estimates

**maxlik** generates maximum likelihood estimates

Setting **\_\_cn\_\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_\_Start** scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.
- 1** will use a vector of user supplied start values stored in the global variable **\_\_cn\_\_StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_\_StartValue**  $(K+L+2) \times 1$  vector, start values if **\_\_cn\_\_Start** = 1.

**\_\_cn\_\_Precision** scalar, number of decimal points to print on output. Default = 4.

**\_\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to **Supreme**. When a data matrix is passed to **Supreme** and the user is selecting from that matrix, the global variable **\_\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

**\_\_\_miss** scalar, determines how missing data will be handled.

- 0 Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.
- 1 Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.

Default = 0.

**\_\_\_output** scalar, determines printing of intermediate results.

- 0 nothing is written.
- 1 serial ASCII output format suitable for disk files or printers.
- 2 (DOS only) output is suitable for screen only. ANSI.SYS must be active.

Default = 2.

**\_\_\_row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.

**\_\_\_rowfac** scalar, “row factor”. If **Supreme** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting

```
___rowfac = 0.8;
```

will cause **GAUSS** to read in 80% of the rows originally calculated.

This global only has an affect when **\_\_\_row** = 0.

Default = 1.

**\_\_\_title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = “”.

**\_\_\_vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:

- 0 Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
- 1 Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.

Default = 1.



## ■ Remarks

Suppose we observe two event count dependent variables  $y_{1i}$  and  $y_{2i}$  for  $n$  observations. Let these variables be distributed as a bivariate Poisson with  $E(y_{1i}) = \lambda_{1i}$  and  $E(y_{2i}) = \lambda_{2i}$ . These means are parametrized as follows:

$$\lambda_{0i} = \exp(x_i\beta) \quad (5.11)$$

and

$$\lambda_{+i} = \exp(z_i\gamma) \quad (5.12)$$

where  $x_i$  and  $z_i$  are (possibly) different vectors of explanatory variables. The covariance parameter is  $\xi$ .

If you have convergence problems, you might try **Supreme2** with argument  $ind3 = 0$  instead.

For details about this model, see King (1989c).

## ■ Example

Seemingly Unrelated Poisson Regression Model (Supreme)

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep1 = { wars };
ind1 = { unem, poverty, allianc };
dep2 = { coups };
ind2 = { unem, age, sex, race };
{ b,vc,llik } = Supreme(dataset,dep1,dep2,ind1,ind2);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

Two vectors of effect parameters and one scalar covariance parameter are estimated. The vectors of effect parameters each include one element corresponding to each explanatory variable and a constant term. In the example, ten parameters are estimated.

## ■ Source

supreme.src

## ■ Purpose

Estimates a Poisson regression model with unobserved dependent variables, for the analysis of two observed (and three unobserved) event count variables, with maximum likelihood.

## ■ Library

count

## ■ Format

{ *b,vc,llik* } = **Supreme2**(*dataset,dep1,dep2,ind1,ind2,ind3*);

## ■ Input

|                |                                                                                                                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>dataset</i> | string, name of <b>GAUSS</b> data set.<br>– or –<br>N×K matrix, data                                                                                                                          |
| <i>dep1</i>    | string, name of the first dependent variable<br>– or –<br>scalar, index of the first dependent variable                                                                                       |
| <i>dep2</i>    | string, name of the second dependent variable<br>– or –<br>scalar, index of the second dependent variable                                                                                     |
| <i>ind1</i>    | K×1 character vector, names of first event independent variables<br>– or –<br>K×1 numeric vector, indices of first event independent variables<br>Set to 0 to include only a constant term.   |
| <i>ind2</i>    | L×1 character vector, names of second event independent variables<br>– or –<br>L×1 numeric vector, indices of second event independent variables<br>Set to 0 to include only a constant term. |
| <i>ind3</i>    | M×1 character vector, names of second event independent variables<br>– or –<br>M×1 numeric vector, indices of second event independent variables<br>Set to 0 to include only a constant term. |

If *dataset* is a matrix, *dep1*, *dep2*, *ind1*, *ind2*, or *ind3* may be a string or character variable containing either the standard labels created by (V1, V2, ..., or V01, V02, ..., depending on the value of `___vpad`), or the user-provided labels in `___altnam`.

## ■ Output

|             |                                                                                                                                                                                                                                                |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>b</i>    | $(K+L+M) \times 1$ vector, maximum likelihood estimates of the effect parameters of $\beta$ and $\gamma$ stacked on top of the covariance parameter $\xi$ .                                                                                    |
| <i>vc</i>   | $(K+L+M) \times (K+L+M)$ matrix, variance-covariance matrix of the estimated parameters evaluated at the maximum. If you choose the global <code>___CovPar = 3</code> , <i>vc</i> will contain heteroskedastic-consistent parameter estimates. |
| <i>llik</i> | scalar, value of the log-likelihood function at the maximum.                                                                                                                                                                                   |

## ■ Globals

**MAXLIK** globals are also relevant.

**\_\_cn\_Inference** string, determines the type of statistical inference.

**boot** generates bootstrapped estimates and covariance matrix of estimates

**maxlik** generates maximum likelihood estimates

Setting **\_\_cn\_Inference** to **BOOT** generates a **GAUSS** data set containing the bootstrapped parameters. The file name of this data set is either a temporary name, or the name in the **MAXLIK** global variable, **\_\_max\_BootFname**. This data set can be used with **MAXBlimits** for generating confidence limits, with **MAXDensity** for generating density estimates and plots of the bootstrapped parameters, or with **MAXHist** for generating histogram and surface plots.

**\_\_cn\_Start** scalar, selects method of calculating starting values. Possible values are:

- 0** calculates them by regressing  $\ln(y + 0.5)$  on the explanatory variables.
- 1** will use a vector of user supplied start values stored in the global variable **\_\_cn\_StartValue**.
- 2** uses a vector of zeros.
- 3** uses random uniform numbers on the interval  $[-\frac{1}{2}, \frac{1}{2}]$ .

Default = 0.

**\_\_cn\_StartValue**  $(K+L+M) \times 1$  vector, start values if **\_\_cn\_Start** = 1.

**\_\_cn\_Precision** scalar, number of decimal points to print on output. Default = 4.

**\_\_\_altnam**  $K \times 1$  vector, alternate names for variables when a matrix is passed to **Supreme2**. When a data matrix is passed to **Supreme2** and the user is selecting from that matrix, the global variable **\_\_\_altnam**, if it is used, must contain names for the columns of the original matrix.

- \_\_miss** scalar, determines how missing data will be handled.
- 0** Missing values will not be checked for, and so the data set must not have any missings. This is the fastest option.
  - 1** Listwise deletion. Removes from computation any observation with a missing value on any variable included in the analysis.
- Default = 0.
- \_\_output** scalar, determines printing of intermediate results.
- 0** nothing is written.
  - 1** serial ASCII output format suitable for disk files or printers.
  - 2** (DOS only) output is suitable for screen only. ANSI.SYS must be active.
- Default = 2.
- \_\_row** scalar, specifies how many rows of the data set will be read per iteration of the read loop. By default, the number of rows to be read will be calculated automatically.
- \_\_rowfac** scalar, "row factor". If **Supreme2** fails due to insufficient memory while attempting to read a **GAUSS** data set, then **\_\_rowfac** may be set to some value between 0 and 1 to read a *proportion* of the original number of rows of the **GAUSS** data set. For example, setting
- ```
__rowfac = 0.8;
```
- will cause **GAUSS** to read in 80% of the rows originally calculated. This global only has an affect when **__row** = 0.
- Default = 1.
- __title** string, message printed at the top of the screen and printed out by **CountPrt**. Default = "".
- __vpad** scalar, if *dataset* is a matrix in memory, the variable names are automatically created by . Two types of names can be created:
- 0** Variable names automatically created by are not padded to give them equal length. For example, V1, V2,...V10, V11,....
 - 1** Variable names created by the procedure are padded with zeros to give them an equal number of characters. For example, V01, V02, ..., V10, V11,.... This is useful if you want the variable names to sort properly.
- Default = 1.

■ Remarks

This model assumes the existence of three independent unobserved variables, y_{1i}^* , y_{2i}^* , and y_{3i}^* , with means $E(y_{ji}^*) = \lambda_{ji}$, for $j = 1, 2, 3$. Although these are not observed, we do observe y_{1i} and y_{2i} , which are functions of these three variables:

$$\begin{aligned} y_{1i} &= y_{1i}^* + y_{3i}^* \\ y_{2i} &= y_{2i}^* + y_{3i}^* \end{aligned}$$

The procedure estimates three separate regression functions, one for the expected value of each of the unobserved variables:

$$\lambda_{1i} = \exp(x_{1i}\beta_1) \tag{5.13}$$

$$\lambda_{2i} = \exp(x_{2i}\beta_2)$$

$$\lambda_{3i} = \exp(x_{3i}\beta_3)$$

where x_{1i} , x_{2i} and x_{3i} are (possibly) different sets of explanatory variables and β_1 , β_2 , and β_3 are separate parameter vectors. This option produces maximum likelihood estimates for these three parameter vectors.

■ Example

Poisson Regression Model with Unobserved Dependent Variables

```
library count;
#include count.ext;
Countset;
dataset = "wars";
dep1 = { wars };
ind1 = { unem, poverty, allianc };
dep2 = { coups };
ind2 = { unem, age, sex, race };
ind3 = { us, sov };
{ b,vc,llik } = Supreme2(dataset,dep1,dep2,ind1,ind2,ind3);
output file = count.out reset;
call CountPrt(b,vc,llik);
output off;
```

Three vectors of effect parameters are estimated. Each includes one element corresponding to each explanatory variable plus a constant term. In the example, twelve parameters are estimated.

■ Source

supreme2.src

Supreme2

5. *COUNT REFERENCE*

Index

active parameters, 11
algorithm, 34

Alt-1, 34

Alt-3, 34

Alt-2, 34

Alt-5, 34

Alt-4, 34

Alt-6, 34

Alt-A, 34

Alt-H, 33

___altnam, 95, 96, 108, 113, 117, 122,
127, 132, 137, 141

B _____

Bayesian inference, 23, 48, 70

BFGS, 8, 34, 42, 61

BHHH, 9, 34, 42, 61

BHHHStep, 10

bootstrap, 29, 95, 98

BRENT, 10

C _____

__cn_Censor, 95

__cn_Dispersion, 95

__cn_Fix, 95

__cn_Inference, 95, 97

__cn_Precision, 95

__cn_Start, 95

__cn_StartValue, 96, 108, 113, 117,
121, 127, 132, 137, 141

__cn_ZeroTruncate, 95

condition of Hessian, 13

conjugate gradient, 9

convergence, 44, 64

`count.src`, 104, 105, 106

CountCLPrt, 104

CountPrt, 105

CountSet, 106

covariance matrix, parameters, 23, 24,
28, 42, 61

___CovPar, 94

cubic step, 44, 64

D _____

density, 95

derivatives, 7, 18, 42, 61, 68

DFP, 8, 34, 42, 61

diagnosis, 17

direction, 7

E _____

Expgam, 107

`expgam.src`, 111

Expon, 112

`expon.src`, 115

F _____

FASTBayes, 48

FASTBAYES, 22

`fastbayes.src`, 50

FASTBoot, 51

FASTBOOT, 22

`fastboot.src`, 53

FASTMAX, 22, 40

`fastmax.src`, 47

`fastpflcl.src`, 55

FASTPflClimits, 54

`fastprof.src`, 58

FASTPROFILE, 22**FASTProfile**, 56**G** _____

global variables, 33

gradient, 40, 48, 51, 56, 59, 70, 73, 76,
82, 87

gradient procedure, 18, 43, 46, 62, 68

GRADRE, 66**H** _____

HALF, 10

Hessian, 8, 13, 34

Hessian procedure, 19, 21, 47, 69

Hurdlep, 116

hurdlep.src, 119

I _____

inactive parameters, 11

Installation, 1

K _____

Kiss-Monster, 30

L _____

likelihood profile trace, 28

line search, 7, 9, 34

linear congruential, 30

log-likelihood function, 6, 21, 40, 46,
47, 48, 51, 59, 67, 68, 70, 73,
82, 85

log-linear, 92

M _____**__max_Active**, 11, 41, 42, 60, 61**__max_Algorithm**, 41, 42, 60, 61**__max_Alpha**, 54, 75, 85, 89**__max_BayesAlpha**, 31, 49, 71**__max_BootFname**, 49, 50, 52, 53, 71,
72, 74**__max_Center**, 57, 80, 83**__max_CovPar**, 24, 40, 41, 42, 59, 60,
61**__max_CutPoint**, 80**__max_Delta**, 41, 42, 60, 61**__max_Diagnostic**, 61**__max_Diagnostic****__max_Diagnostic**, 17, 60**__max_Extrap**, 41, 42, 60, 62**__max_FinalHess**, 25, 41, 42, 60, 62**__max_GradCheckTol**, 19, 47, 60, 62,
69**__max_GradMethod**, 34, 41, 42, 60, 62**__max_GradProc**, 41, 43, 46, 47, 60,
62, 68**__max_GradStep**, 41, 43, 60, 63**__max_GradTol**, 34, 41, 42, 60, 62**__max_HessCov**, 24, 41, 43, 60, 63**__max_HessProc**, 19, 41, 43, 47, 60,
63, 69**__max_Increment**, 57, 80, 83**__max_Interp**, 41, 43, 60, 63**__max_IterData**, 41, 44, 60, 63**__max_Kernel**, 30, 78**__max_Lag**, 60, 63**__max_Lagrange**, 24**__max_LineSearch**, 41, 44, 60, 64**__max_MaxIters**, 41, 44, 60, 64**__max_MaxTime**, 41, 44, 49, 52, 60,
64, 71, 74**__max_Maxtry**, 60**__max_MaxTry**, 44**__max_Maxtry**, 41**__max_MaxTry**, 34, 64**__max_NumCat**, 57, 80, 83**__max_NumObs**, 29, 41, 50, 52, 60, 71,
74**__max_NumPoints**, 78**__max_NumSample**, 29, 44, 49, 52, 57,
64, 71, 74, 83**__max_Options**, 41, 44, 60, 64**__max_ParNames**, 41, 45, 60, 65**__max_PriorProc**, 49, 71**__max_RandRadius**, 10, 41, 44, 45, 60,
64, 65

INDEX

__max__RandType, 30, 41, 45, 49, 52, 60, 65
__max__Select, 57, 75, 83, 89
__max__Smoothing, 30, 78
__max__State, 41, 45, 50, 52, 60, 65
__max__Switch, 21, 41, 45, 60, 65
__max__Truncate, 78
__max__UserHess, 21
__max__UserNumGrad, 41, 60, 66
__max__UserNumHess, 60, 66
__max__UserSearch, 60, 67
__max__Width, 57, 80, 83
__max__XprodCov, 24, 41, 46, 60, 67
MAXBayes, 70
maxbayes.src, 72
maxblim.src, 75
MAXBlimits, 75
MAXBoot, 29, 73
maxboot.src, 74
MAXCLPrt, 76
maxdens.src, 79
MAXDensity, 30, 78, 98
MAXHist, 30, 80, 95, 98
maxhist.src, 81
maximum likelihood, 6, 40, 51, 59, 73, 91
MAXLIK, 59
maxlik.src, 69, 77, 87, 88, 89
maxpflcl.src, 86
MAXPflClimits, 85
maxprof.src, 84
MAXProfile, 82
MAXPrt, 87
MAXSet, 88
MAXTlimits, 89
__miss, 108, 113, 117, 122, 127, 132, 137, 142

N _____

Negbin, 120
negbin.src, 125
NEWTON, 8, 34, 42, 47, 61, 69
NR, 34

O _____

__output, 34, 65, 78, 80, 95, 96, 109, 114, 118, 122, 128, 133, 138, 142

P _____

Pareto, 126
pareto.src, 130
Poisson, 131
poisson.src, 135
PRCG, 9, 34, 42, 61
profile likelihood, 23, 25, 54, 85
profile t plot, 28
pseudo-random numbers, 30

Q _____

quadratic step, 44, 64
quasi-maximum likelihood covariance matrix, 24, 42, 61
quasi-Newton, 8

R _____

random search, 11
regression, Hurdle Poisson, 92
regression, negative binomial, 92
regression, seemingly unrelated Poisson, 92
regression, truncated negative binomial, 92
regression, truncated Poisson, 92
resampling, 29
Richardson Extrapolation, 66
__row, 6, 40, 59, 60, 63, 66, 67, 68, 70, 73, 82, 85
__rowfac, 60, 66, 97, 109, 114, 118, 122, 128, 133, 138, 142
run-time switches, 33

S _____

scaling, 13
Shift-1, 34
Shift-2, 34
Shift-4, 34
Shift-3, 34

Shift-5, 34
starting point, 17
statistical inference, 23, 97
STEEP, 34, 42, 61
step length, 9, 34, 44, 64
STEPBT, 10
Supreme, 136
supreme.src, 139
Supreme2, 140
supreme2.src, 143
switching algorithms, 21

T _____

___title, 41, 46, 60, 66

U _____

UNIX, 1, 3

V _____

vput, 17
vread, 17

W _____

Wald, 23, 89
Wald inference, 23
___weight, 11, 41, 46, 60, 67
weighted maximum likelihood, 11
Windows/NT/2000, 3