

# Symbolic maximum likelihood estimation with Mathematica

Colin Rose

*Theoretical Research Institute, Sydney, Australia*

and Murray D. Smith

*University of Sydney, Australia*

[Received May 1998. Revised October 1999]

**Summary.** Mathematica is a symbolic programming language that empowers the user to undertake complicated algebraic tasks. One such task is the derivation of maximum likelihood estimators, demonstrably an important topic in statistics at both the research and the expository level. In this paper, a Mathematica package is provided that contains a function entitled `SuperLog`. This function utilizes pattern-matching code that enhances Mathematica's ability to simplify expressions involving the natural logarithm of a product of algebraic terms. This enhancement to Mathematica's functionality can be of particular benefit for maximum likelihood estimation.

**Keywords:** Computer algebra systems; Estimate; Estimator; Mathematica; Symbolic maximum likelihood; Teaching

## 1. Introduction

Although statistical software has long been used for maximum likelihood (ML) estimation, the focus of attention has almost always been on obtaining ML estimates (a *numerical* problem), rather than on deriving ML estimators (a *symbolic* problem). Even after the introduction of powerful computer algebra systems, in the context of ML estimation, symbolic engines have largely only been used to solve numerical problems (see, for example, Currie (1995) and Cook and Broemeling (1995)) or to derive large sample symbolic approximations to ML estimators (see Stafford and Andrews (1993) and Stafford *et al.* (1994)). By contrast, we use a computer algebra system for the first time to derive *exact* symbolic ML estimators from first principles.

This paper shows how Mathematica (see Wolfram (1996)) can be extended to handle symbolic ML estimation, which is demonstrably an important topic in statistics, at both the research and the expository level. The paper has five sections. Section 2 expands Mathematica's programming language to handle symbolic ML estimation. Section 3 illustrates the approach with three simple expository examples and is therefore well suited for teaching purposes. Section 4 extends the analysis to more challenging material. Section 5 concludes. Appendices follow giving a glossary of Mathematica terms and the package code.

*Address for correspondence:* Murray D. Smith, Econometrics and Business Statistics, University of Sydney, Sydney, NSW 2006, Australia.  
E-mail: Murray.Smith@econ.usyd.edu.au

## 2. Extending Mathematica: the package SMLE

Consider the following simple problem. Let  $(Y_1, \dots, Y_n)$  denote a random sample of size  $n$  collected on  $Y \sim \text{Poisson}(\theta)$ , where parameter  $\theta > 0$  is unknown—show that the mean of the random sample is the ML estimator of  $\theta$ .

We begin our answer in the usual way by inputting the likelihood function into Mathematica:

$$\text{In}[1]:= \quad L = \prod_{i=1}^n \frac{e^{-\theta} \theta^{Y_i}}{Y_i!}$$

If we try to evaluate the log-likelihood

$$\text{In}[2]:= \quad \text{Log}[L]$$

$$\text{Out}[2] = \quad \text{Log} \left[ \prod_{i=1}^n \frac{e^{-\theta} \theta^{Y_i}}{Y_i!} \right]$$

nothing happens, and for good reason! Mathematica will not normally ‘convert’ an expression such as  $\text{Log}[a \times b]$  into  $\text{Log}[a] + \text{Log}[b]$  because symbols  $a$  and  $b$  could represent anything. *Prima facie*, it may appear then that a student armed with a pencil and paper can achieve much more than Mathematica. Of course, when deriving the log-likelihood, the student makes use of information which Mathematica does not ‘know’; for example, the student knows that  $Y_i$  takes non-negative integer values, that  $n$  is a positive integer, that  $\theta$  is positive and real valued, that each term in the product is positive and real valued etc. Mathematica assumes nothing about the symbols that have been entered, so its inaction is perfectly reasonable.

Fortunately, ML problems have some common features which we can exploit: the contributions to the likelihood are all positive valued, the sample size is a positive integer and parameters and variables are real valued. These features can be added to Mathematica, not by providing explicit information about each symbol in turn, but rather by providing a pattern-matching function that assumes these common features. To do so, we load our small package entitled SMLE.m (see Appendix B) into Mathematica in the usual way:

$$\text{In}[3]:= \quad \ll \text{SMLE.m}$$

(The code given in the SMLE.m package should be stored as a text file on disc under the name SMLE.m. It should be placed in any one of the directories which Mathematica searches when attempting to find an external file—this list of directories can be determined by using the  $\$Path$  command.) This package is also available electronically by anonymous file transfer protocol at

<ftp://ftp.tri.org.au/SMLE.m>

We then ‘activate’ our SuperLog function as follows:

$$\begin{aligned} \text{In}[4]:= \quad & \text{SuperLog}[\text{On}] \\ & \text{SuperLog is now On.} \end{aligned}$$

If we now evaluate  $\text{Log}[L]$  again, we obtain a much more useful result:

$$\begin{aligned} \text{In}[5]:= \quad & \text{logL} = \text{Log}[L] \\ \text{Out}[5] = \quad & -n \theta - \sum_{i=1}^n \text{Log}[Y_i!] + \text{Log}[\theta] \sum_{i=1}^n Y_i \end{aligned}$$

In essence, SuperLog modifies Mathematica’s Log function to mimic what our student does:

it converts ‘logarithms of products’ into ‘sums of logarithms’. It simplifies log-likelihood expressions of the form  $\log(\prod_{i=1}^n f_i)$  where *we know* that  $f_i$  is defined on the positive real line. To do this, SuperLog exploits the power of the Mathematica programming language by using pattern-matching code (apart from reserved symbols, the function accepts any indexing symbol, any symbol for sample size and any functional form for the likelihood contributions). The SMLE.m package has been tested using both version 3 and version 4 of Mathematica on platforms including Macintosh, Windows 95, Windows 98, Windows NT, SGI Unix and DEC Unix. To return Mathematica’s Log function to its default status, simply enter SuperLog[Off].

### 3. Simple examples

In this section, we consider three expository examples.

#### 3.1. Example 1: Poisson distribution

We have already derived the log-likelihood for the Poisson distribution. The score is

$$\text{In}[6]:= \quad \text{score} = \partial_{\theta} \log L$$

$$\text{Out}[6] = \quad -n + \frac{\sum_{i=1}^n Y_i}{\theta}$$

where Mathematica’s internal partial differential operator  $\partial_{\theta}$  has been used. Setting the score to 0 defines the ML estimator  $\hat{\theta}$ . The resulting equation can be easily solved by using Mathematica’s Solve function. The ML estimator  $\hat{\theta}$  is given as a replacement rule  $\rightarrow$  for  $\theta$ :

$$\text{In}[7]:= \quad \hat{\theta} = \text{Solve}[\text{score} == 0, \theta]$$

$$\text{Out}[7] = \quad \left\{ \left\{ \theta \rightarrow \frac{\sum_{i=1}^n Y_i}{n} \right\} \right\}$$

where the second-order conditions confirm that  $\hat{\theta}$ , the mean of the random sample, is indeed the ML estimator.

$$\text{In}[8]:= \quad \partial_{\{\theta, 2\}} \log L$$

$$\text{Out}[8] = \quad -\frac{\sum_{i=1}^n Y_i}{\theta^2}$$

Finally, let us suppose that an observed random sample is (1, 0, 3, 4):

$$\text{In}[9]:= \quad \text{data} = \{1, 0, 3, 4\};$$

Then the ML estimate of  $\theta$  is obtained by substitution of these data into the ML estimator  $\hat{\theta}$ :

$$\text{In}[10]:= \quad \hat{\theta} /. \{n \rightarrow 4, Y_i :> \text{data}[[i]]\}$$

$$\text{Out}[10] = \quad \{\{\theta \rightarrow 2\}\}$$

(There are several ways to enter data into Mathematica objects. The method used here is based on one of Mathematica’s replacement rules  $:>$ . The notation  $\text{data}[[i]]$  is Mathematica notation for the  $i$ th element of the list  $\text{data}$ .)

### 3.2. Example 2: exponential distribution

Suppose that the positive-valued continuous variable  $Y$  is such that  $Y \sim \text{Exp}(\theta)$ , where parameter  $\theta > 0$ . For a random sample of size  $n$  on  $Y$ , we have the log-likelihood function as

$$\begin{aligned} \text{In[11]} := & \quad \log L = \text{Log} \left[ \prod_{i=1}^n \frac{e^{-Y_i/\theta}}{\theta} \right] // \text{Apart} \\ \text{Out[11]} = & \quad -n \text{Log}[\theta] - \frac{\sum_{i=1}^n Y_i}{\theta} \end{aligned}$$

which has been further simplified by using the `Apart` function. The score is

$$\begin{aligned} \text{In[12]} := & \quad \text{score} = \partial_{\theta} \log L \\ \text{Out[12]} = & \quad -\frac{n}{\theta} + \frac{\sum_{i=1}^n Y_i}{\theta^2} \end{aligned}$$

The first-order conditions are given by

$$\begin{aligned} \text{In[13]} := & \quad \hat{\theta} = \text{Solve}[\text{score} == 0, \theta] \\ \text{Out[13]} = & \quad \left\{ \left\{ \theta \rightarrow \frac{\sum_{i=1}^n Y_i}{n} \right\} \right\} \end{aligned}$$

and the second-order conditions evaluated at  $\hat{\theta}$  are

$$\begin{aligned} \text{In[14]} := & \quad \partial_{\{\theta, 2\}} \log L /. \text{Flatten}[\hat{\theta}] \\ \text{Out[14]} = & \quad -\frac{n^3}{(\sum_{i=1}^n Y_i)^2} \end{aligned}$$

Hence, the Hessian is strictly negative. The ML estimator is thus  $\hat{\theta} = (1/n) \sum_{i=1}^n Y_i$ .

### 3.3. Example 3: beta distribution

Let the proportion  $Y$  be a random variable defined over the unit interval of the real line,  $0 < y < 1$ . Furthermore, let the probability density function (PDF) of  $Y$  be given by

$$\text{In[15]} := \quad f = \theta y^{\theta-1}$$

Thus  $Y \sim \text{beta}(\theta, 1)$ , where parameter  $\theta > 0$  is unknown. Clearly the distribution of  $Y$  is a special case of the standard two-parameter beta distribution. For a random sample of size  $n$  on  $Y$ , the log-likelihood for  $\theta$  is given by

$$\begin{aligned} \text{In[16]} := & \quad \log L = \text{Log} \left[ \prod_{i=1}^n (f /. Y \rightarrow Y_i) \right] \\ \text{Out[16]} = & \quad n \text{Log}[\theta] + (-1 + \theta) \sum_{i=1}^n \text{Log}[Y_i] \end{aligned}$$

The ML estimator of  $\theta$  is derived as

In[17]:=  $\hat{\theta} = \text{Solve}[\partial_{\theta} \log L == 0, \theta]$

Out[17] =  $\left\{ \left\{ \theta \rightarrow -\frac{n}{\sum_{i=1}^n \text{Log}[Y_i]} \right\} \right\}$

because on inspecting the Hessian

In[18]:=  $\partial_{\{\theta,2\}} \log L$

Out[18] =  $-\frac{n}{\theta^2}$

we see that it is negative for all  $\theta$ ; thus the log-likelihood is globally concave in  $\theta$  and the solution to the first-order condition corresponds to the unique maximum.

#### 4. Further illustrations

In this section, we consider three examples in which Mathematica is used to derive the ML estimator. All three are commonly used when teaching ML estimation. In addition, we show how to use Mathematica to tackle other important issues concerning ML estimation such as deriving the distribution of an ML estimator, proving concavity of a log-likelihood and concentration of a log-likelihood.

##### 4.1. Example 3 (continued): beta distribution

Continuing the  $\text{beta}(\theta, 1)$  model, consider the problem of determining the exact distribution of the ML estimator  $\hat{\theta} = -n/\sum_{i=1}^n \log(Y_i)$ . This can be attempted with Mathematica by using the moment-generating function (MGF) method (Mittelhammer (1996), section 3.5, provides a good description of this method). First, we derive the MGF of  $\log(Y)$ , i.e. we find  $E[\exp\{t \log(Y)\}] = E(Y^t)$ :

In[19]:= `SetOptions[Integrate, GenerateConditions -> False];`

In[20]:=  $\text{mgf} = \int_0^1 y^t f \, dy$

Out[20] =  $\frac{\theta}{t + \theta}$

where, for simplicity, we have turned off the `Integrate` option which yields conditional output. Next, because  $(Y_1, \dots, Y_n)$  is a random sample of size  $n$ , the MGF of

$$\overline{\log(Y)} = -\frac{1}{n} \sum_{i=1}^n \log(Y_i) = \frac{1}{\theta}$$

is, by the MGF theorem, equal to

$$E \left[ \exp \left\{ -\frac{t}{n} \log(Y) \right\} \right]^n = E(Y^{-t/n})^n.$$

Given our previous output, we only need to use a replacement rule to determine the MGF of  $\overline{\log(Y)}$ :

$$\text{In}[21] := \left( \text{mgf} /. t \rightarrow -\frac{t}{n} \right)^n // \text{Simplify}$$

$$\text{Out}[21] = \left( \frac{n \theta}{-t + n \theta} \right)^n$$

This expression can be recognized as the MGF of a random variable  $W \sim \text{gamma}(n, 1/n\theta)$ , as we verify by comparing it with

$$\text{In}[22] := g = \frac{w^{a-1} e^{-w/b}}{\text{Gamma}[a] b^a} /. \left\{ a \rightarrow n, b \rightarrow \frac{1}{n \theta} \right\};$$

$$\text{In}[23] := \int_0^\infty e^{tw} g dw // \text{PowerExpand}$$

$$\text{Out}[23] = n^n \theta^n (-t + n \theta)^{-n}$$

Given that  $\overline{\log(Y)}$  is gamma distributed, and that  $\hat{\theta} = 1/\overline{\log(Y)}$ , it follows that the ML estimator  $\hat{\theta}$  has an inverted gamma distribution with parameters  $n$  and  $1/n\theta$ . The PDF of  $\hat{\theta} = q > 0$  is easily derived by transformation:

$$\text{In}[24] := w = \frac{1}{q}; \quad \text{pdf} = \text{Abs}[\partial_q w] g$$

$$\text{Out}[24] = \frac{e^{-n\theta/q} \left(\frac{1}{q}\right)^{-1+n} \left(\frac{1}{n \theta}\right)^{-n}}{\text{Abs}[q]^2 \text{Gamma}[n]}$$

If desired, this expression can be further simplified by using a replacement rule, since the argument of the absolute value is always positive. The mean of the ML estimator is easily derived:

$$\text{In}[25] := \int_0^\infty q \text{pdf} dq // \text{FullSimplify}$$

$$\text{Out}[25] = \frac{n \theta}{-1 + n}$$

It is easy to see from this output that the ML estimator  $\hat{\theta}$  is biased upwards.

#### 4.2. Example 4: normal linear regression model

A statistical model of considerable practical importance is the normal linear regression model. For illustration, we shall consider a simple case of this model, namely a regression model with a constant dummy and one regressor variable  $X$ . For a given value of  $X = x$ , the conditional distribution of the dependent variable  $Y$  is assumed to be

$$Y|(X = x) \sim N(\beta_1 + \beta_2 x, \gamma),$$

where parameter  $\theta = (\beta_1, \beta_2, \gamma)$  (we use  $\gamma$  to denote the variance parameter). Denote a random sample of  $T$  pairs on  $(Y, X)$  by  $((Y_1, X_1), \dots, (Y_T, X_T))$ . We assume, conditional on each  $X_k = x_k$ , that  $Y_i$  is independent of  $Y_j$  for all  $i \neq j$  ( $i, j, k = 1, \dots, T$ ). Under these assumptions, the log-likelihood is given by

$$\text{In}[26] := \log L = \text{Log} \left[ \prod_{k=1}^T \frac{e^{-(Y_k - \mu)^2 / (2\gamma)}}{\sqrt{2\pi\gamma}} \ /. \ \mu \rightarrow \beta_1 + x_k \beta_2 \right]$$

$$\begin{aligned} \text{Out}[26] = & -\frac{1}{2\gamma} \left( T \gamma \text{Log}[2] + T \gamma \text{Log}[\pi] + T \gamma \text{Log}[\gamma] \right. \\ & \left. + T \beta_1^2 + \beta_2^2 \sum_{k=1}^T x_k^2 + 2\beta_1 \left( \beta_2 \sum_{k=1}^T x_k - \sum_{k=1}^T Y_k \right) - 2\beta_2 \sum_{k=1}^T x_k Y_k + \sum_{k=1}^T Y_k^2 \right) \end{aligned}$$

The score vector is given by

$$\text{In}[27] := \text{score} = \{\partial_{\beta_1} \log L, \partial_{\beta_2} \log L, \partial_{\gamma} \log L\} // \text{Simplify}$$

$$\begin{aligned} \text{Out}[27] = & \left\{ \frac{-T \beta_1 - \beta_2 \sum_{k=1}^T x_k + \sum_{k=1}^T Y_k}{\gamma}, \frac{-\beta_1 \sum_{k=1}^T x_k - \beta_2 \sum_{k=1}^T x_k^2 + \sum_{k=1}^T x_k Y_k}{\gamma}, \right. \\ & \left. \frac{1}{2\gamma^2} \left( -T \gamma + T \beta_1^2 + \beta_2^2 \sum_{k=1}^T x_k^2 + 2\beta_1 \left( \beta_2 \sum_{k=1}^T x_k - \sum_{k=1}^T Y_k \right) \right. \right. \\ & \left. \left. - 2\beta_2 \sum_{k=1}^T x_k Y_k + \sum_{k=1}^T Y_k^2 \right) \right\} \end{aligned}$$

The ML estimators of  $\theta$  are derived by Mathematica as

$$\text{In}[28] := \hat{\theta} = \text{Solve}[\text{score} == \{0, 0, 0\}, \{\beta_1, \beta_2, \gamma\}]$$

$$\begin{aligned} \text{Out}[28] = & \left\{ \left\{ \gamma \rightarrow \left( -\sum_{k=1}^T x_k^2 \left( \sum_{k=1}^T Y_k \right)^2 + 2 \left( \sum_{k=1}^T x_k \right) \left( \sum_{k=1}^T Y_k \right) \sum_{k=1}^T x_k Y_k - T \left( \sum_{k=1}^T x_k Y_k \right)^2 \right. \right. \right. \\ & \left. \left. - \left( \sum_{k=1}^T x_k \right)^2 \sum_{k=1}^T Y_k^2 + T \left( \sum_{k=1}^T x_k^2 \right) \sum_{k=1}^T Y_k^2 \right) / \left( T \left( -\left( \sum_{k=1}^T x_k \right)^2 + T \sum_{k=1}^T x_k^2 \right) \right), \right. \\ & \left. \beta_1 \rightarrow \frac{\left( \sum_{k=1}^T x_k^2 \right) \sum_{k=1}^T Y_k - \left( \sum_{k=1}^T x_k \right) \sum_{k=1}^T x_k Y_k}{-\left( \sum_{k=1}^T x_k \right)^2 + T \sum_{k=1}^T x_k^2}, \beta_2 \rightarrow \frac{\left( \sum_{k=1}^T x_k \right) \sum_{k=1}^T Y_k - T \sum_{k=1}^T x_k Y_k}{\left( \sum_{k=1}^T x_k \right)^2 - T \sum_{k=1}^T x_k^2} \right\} \end{aligned}$$

The functional form given for the estimator is unfamiliar and imposing. However, if we iteratively solve the first-order conditions

In[29]:=  $\tilde{\beta}_1 = \text{Solve}[\text{score}[[1]] == 0, \beta_1] // \text{Simplify}$

$$\text{Out[29]} = \left\{ \left\{ \beta_1 \rightarrow \frac{-\beta_2 \sum_{k=1}^T x_k + \sum_{k=1}^T Y_k}{T} \right\} \right\}$$

In[30]:=  $\tilde{\beta}_2 = \text{Solve}[(\text{score}[[2]]) /. \tilde{\beta}_1 == 0, \beta_2] // \text{Simplify}$

$$\text{Out[30]} = \left\{ \left\{ \beta_2 \rightarrow \frac{\left( \sum_{k=1}^T x_k \right) \sum_{k=1}^T Y_k - T \sum_{k=1}^T x_k Y_k}{\left( \sum_{k=1}^T x_k \right)^2 - T \sum_{k=1}^T x_k^2} \right\} \right\}$$

it is fairly easy to see that these results are just the well-known formulae for the estimators that are presented in most elementary texts:

$$\hat{\beta}_1 = \bar{Y} - \hat{\beta}_2 \bar{x}$$

and

$$\hat{\beta}_2 = \frac{\sum_{k=1}^T (Y_k - \bar{Y})(x_k - \bar{x})}{\sum_{k=1}^T (x_k - \bar{x})^2}.$$

Further analysis might include the verification of the equivalence between functional forms and a formal check of the second-order conditions. Both exercises involve a mixture of work with Mathematica and with a pencil and paper.

For data collected on the variables, it is straightforward to compute ML estimates. For example, if the data are

In[31]:=  $\text{depY} = \{1, 0, 3, 4\}; \quad \text{regX} = \{1, 2, 3, 4\};$

the ML estimates are

In[32]:=  $\hat{\theta} /. \{\text{T} \rightarrow 4, Y_k :> \text{depY}[[k]], x_k :> \text{regX}[[k]]\}$

$$\text{Out[32]} = \left\{ \left\{ \gamma \rightarrow \frac{7}{10}, \beta_1 \rightarrow -1, \beta_2 \rightarrow \frac{6}{5} \right\} \right\}$$

### 4.3. Example 5: gamma distribution

Let  $(Y_1, \dots, Y_T)$  denote a random sample of size  $T$  collected on a random variable  $Y = y > 0$  which is gamma distributed; i.e.  $Y \sim \text{gamma}(\alpha, \beta)$ , where parameter  $\theta = (\alpha, \beta)$ , with  $\alpha > 0$  and  $\beta > 0$ . We wish to determine the ML estimator of  $\theta$ . The log-likelihood is

In[33]:=  $\log L = \text{Log} \left[ \prod_{i=1}^T \frac{Y_i^{\alpha-1} e^{-Y_i/\beta}}{\text{Gamma}[\alpha] \beta^\alpha} \right] // \text{Apart}$

$$\text{Out[33]} = -T \alpha \text{Log}[\beta] - T \text{Log}[\text{Gamma}[\alpha]] - \sum_{i=1}^T \text{Log}[Y_i] + \alpha \sum_{i=1}^T \text{Log}[Y_i] - \frac{\sum_{i=1}^T Y_i}{\beta}$$

For this problem, the ML estimator of  $\theta$  admits, only in part, a closed form solution. To see this, consider the score:



$$\begin{aligned} \text{In}[34] := & \quad \text{score} = \{\partial_\alpha \log L, \partial_\beta \log L\} \\ \text{Out}[34] = & \quad \left\{ -T \text{Log}[\beta] - T \text{PolyGamma}[0, \alpha] + \sum_{i=1}^T \text{Log}[Y_i], -\frac{T \alpha}{\beta} + \frac{\sum_{i=1}^T Y_i}{\beta^2} \right\} \end{aligned}$$

where the `PolyGamma[0,  $\alpha$ ]` term is Mathematica's notation for the digamma function (the first derivative of the natural logarithm of the gamma function). For Mathematica to derive the ML estimator in closed form, it must successfully execute `Solve[score == {0, 0}, { $\alpha$ ,  $\beta$ }]`. Unfortunately, this task cannot be performed in this case because of the presence of the `PolyGamma` function. Nevertheless, we can make progress by concentrating the log-likelihood in  $\beta$ . This is because

$$\begin{aligned} \text{In}[35] := & \quad \hat{\beta} = \text{Solve}[\text{score}[[2]] == 0, \beta] \\ \text{Out}[35] = & \quad \left\{ \left\{ \beta \rightarrow \frac{\sum_{i=1}^T Y_i}{T \alpha} \right\} \right\} \end{aligned}$$

is in a closed form. Mathematica has given us  $\hat{\beta}$  as a function of  $\alpha$ , i.e.  $\hat{\beta} = \hat{\beta}(\alpha)$ . Hence, the ML estimator of  $\beta$  is

$$\hat{\beta}(\hat{\alpha}) = \frac{1}{\hat{\alpha} T} \sum_{i=1}^T Y_i,$$

where  $\hat{\alpha}$  denotes the ML estimator of  $\alpha$ . Replacing  $\beta$  with  $\hat{\beta}(\alpha)$  in the log-likelihood yields the concentrated log-likelihood:

$$\begin{aligned} \text{In}[36] := & \quad \text{ConlogL} = \log L /. \text{Flatten}[\hat{\beta}] \\ \text{Out}[36] = & \quad -T \alpha - T \text{Log}[\text{Gamma}[\alpha]] - T \alpha \text{Log} \left[ \frac{\sum_{i=1}^T Y_i}{T \alpha} \right] - \sum_{i=1}^T \text{Log}[Y_i] \\ & \quad + \alpha \sum_{i=1}^T \text{Log}[Y_i] \end{aligned}$$

As in the original problem, maximization of the concentrated log-likelihood with respect to  $\alpha$  does not yield a closed form solution. Hence, the ML estimator is defined implicitly by

$$\hat{\alpha} = \arg \max_{\alpha > 0} (\text{ConlogL}).$$


For a specific set of data, numerical techniques are required to determine the ML estimate  $\hat{\alpha}$ .

With actual data, the issue of determining a suitable numerical algorithm to maximize the observed concentrated log-likelihood becomes important. Mathematica's in-built optimizer `FindMinimum` has a suite of the more popular gradient method algorithms available to the user, including the Newton–Raphson algorithm. Of course, which of these algorithms, if any, is appropriate depends very much on the particular situation at hand. In the present case, further algebra with Mathematica enables us to uncover concavity in  $\alpha$ :

$$\begin{aligned} \text{In}[37] := & \quad \partial_{\{a,2\}} \text{ConlogL} // \text{Factor} \\ \text{Out}[37] = & \quad -\frac{T(-1 + \alpha \text{PolyGamma}[1, \alpha])}{\alpha} \end{aligned}$$

The sign of the derivative obviously depends on the quantity within round brackets. A plot against various  $\alpha > 0$  is a simple, informal device for determining its sign; for example,

```
In[38]:= Plot[-1 +  $\alpha$  PolyGamma[1,  $\alpha$ ], { $\alpha$ , 0.1, 5}]
```



As the plot is everywhere positive, the Hessian must be negative for all  $\alpha > 0$ . Thus the concentrated log-likelihood is concave in  $\alpha$ , and we may conclude that the Newton–Raphson algorithm is well suited for estimating  $\alpha$ .

## 5. Conclusion

Although computer software has long been used for ML estimation, the focus has almost always been on numerical rather than on symbolic problems. In this paper, we have shown how the Mathematica computer algebra system can be used to derive *symbolic* ML estimators. This was done by modifying Mathematica’s `Log` command using our `SuperLog` function. This is especially useful when constructing the log-likelihood for a random sample consisting of independent random variables. We have shown how this new function enhances Mathematica’s functionality and have demonstrated its particular benefits for the teaching of ML estimation.

## Acknowledgements

Revisions to this paper were undertaken while Smith was visiting the Sonderforschungsbereich 386, Institut für Statistik, Ludwig-Maximilians-Universität, München, Germany. The content of the paper benefitted from presentations at the Worldwide Mathematica Conference (Chicago, 1998), the International Association for Mathematics and Computers in Simulation Conference on Applications of Computer Algebra (Prague, 1998) and the Royal Statistical Society international conference (Glasgow, 1998) and from discussions with seminar participants at the University of Sydney. We also thank the Joint Editor and referee for helpful comments. Smith acknowledges gratefully the support of the Alexander von Humboldt-Stiftung.

## Appendix A: Glossary

For completeness, we provide a short summary of the use and syntax of some of the Mathematica functions that appear in the paper. Further details of these, and all other functions, can be found in Wolfram (1996); see also the on-line help provided with the Mathematica software.

### A.1. Simplification rules

Functions `Apart`, `Factor`, `FullSimplify`, `PowerExpand`, and `Simplify` represent a portion of a suite of commands designed to perform various algebraic manipulations to a specified expression. The syntax of each of these commands is, for example, `Simplify[expr]` or `expr // Simplify`. The latter is the form which we generally prefer.

### A.2. Brackets

Brackets in Mathematica have distinct interpretations and are not interchangeable. Square brackets `[ ]` are reserved for use with Mathematica function names; for example the logarithm function is `Log[ ]`. A list of elements is collected between braces `{ }`; vectors and matrices are built from lists. Round brackets `( )` are the only parentheses which have the usual term-collective interpretation.

### A.3. Partial differentiation

$\partial_\theta$  is Mathematica's partial differential operator. Its syntax, for the partial derivative of an expression `expr` with respect to a symbol  $\theta$ , is  `$\partial_\theta$  expr`. The syntax for the second-order partial derivative is  `$\partial_{\{\theta,2\}}$  expr`.

### A.4. Replacement rules

The ability to replace, in a given expression, a symbol with other symbols or numeric quantities is a vital component of the Mathematica programming language. For instance, `expr /.  $\theta \rightarrow \tau$` , acts on `expr` by replacing symbol  $\theta$  with  $\tau$ , wherever the former occurs in `expr`. By contrast, `expr /.  $\theta :> \tau$` , acts on `expr` by replacing symbol  $\theta$  with  $\tau$ , but delaying the evaluation of  $\tau$  until after the replacement has been performed. This form of transformation is especially useful when `expr` is a symbolic sum, and we wish to evaluate it for a specific set of data.

### A.5. Search functions

`Solve` searches for solutions to an equation or set of equations. The command `Solve[expr == 0,  $\theta$ ]` attempts to find the values of  $\theta$  for which `expr == 0`. If `sol` denotes the solution to an equation, the output from `Solve` is reported in the form of a list of replacement rules, e.g. `{ $\theta \rightarrow$  sol}`. In the case of polynomial equations, Mathematica attempts to output all solutions, including complex-valued solutions. Sets of equations must be enclosed within a list, e.g. `{expr1 == 0, expr2 == 0}`. Similarly, if the solutions are sought in two or more variables, then those variable symbols must be placed in a list also, e.g. `{ $\theta_1, \theta_2$ }`; hence, `Solve[{expr1 == 0, expr2 == 0}, { $\theta_1, \theta_2$ }]`.

## Appendix B: Code

```
(*:Name:          SMLE *)
(*:Authors:      Colin Rose and Murray D. Smith *)
(*:Version:      Mathematica v3, or v4, or later required *)
(*:Legal:        Copyright 1999 *)
(*:Summary:      Symbolic Maximum Likelihood Estimation *)

BeginPackage["SMLE`"]

SuperLog::usage =
  "SuperLog[On] activates the enhanced Log operator, so that
  Log[Product[_ _]] objects get converted into sums of logs.
  SuperLog[Off] switches the enhancement off."

Begin["`Private`"]

SuperLog[Q_]:=Module[{erk, iii, nnn},
```

```

Product[iii, {iii, nnn}]; (* pre-load Product *)
Which[
  Q === On,
    Unprotect[Log]; Clear[Log];
    Log[Product[x_, {k_, a_, b_}]] :=
      Log[Product[Times[erk, x], {k, a, b}]] /. erk -> 1;
    Log[Product[HoldPattern[Times[x_]], {k_, a_, b_}]] := Simplify[
      Map[Sum[#, {k, a, b}]&, Plus@@Map[Expand[PowerExpand[Log[#]]]&,
        List[x]]] // .
      Sum[u. w_, {kk_, aa_, bb_}] :> u Sum[w, {kk, aa, bb}] /; FreeQ[u, kk] ==
      True];
    Protect[Log]; Print["SuperLog is now On."],

  Q === Off,
    Unprotect[Log]; Clear[Log]; Protect[Log]; Print["SuperLog is now Off."],

  True,
    Print["Carumbah! Please use SuperLog[On] or SuperLog[Off]."]
  ]

End[]
Protect[ SuperLog ];
EndPackage[]

```

## References

- Cook, P. and Broemeling, L. (1995) Bayesian statistics using Mathematica. *Am. Statistn*, **49**, 70–76.
- Currie, I. D. (1995) Maximum likelihood estimation with Mathematica. *Appl. Statist.*, **44**, 379–394.
- Mittelhammer, R. (1996) *Mathematical Statistics for Economics and Business*. New York: Springer.
- Stafford, J. and Andrews, D. (1993) A symbolic algorithm for studying adjustments to the profile likelihood. *Biometrika*, **80**, 715–730.
- Stafford, J., Andrews, D. and Wang, Y. (1994) Symbolic computation: a unified approach to studying likelihood. *Statist. Comput.*, **4**, 235–245.
- Wolfram, S. (1996) *The Mathematica Book*, 3rd edn. Cambridge: Cambridge University Press.