

Discrete choice models with random parameters in R : The Rchoice Package

Mauricio Sarrias
Cornell University

Abstract

Rchoice is a package for R which enables the estimation of a variety of Binary, Count and Ordered models with unobserved and observed heterogeneity in the parameters for cross-section data. We implement simulated maximum likelihood methods for the estimation of the coefficients which can assume a variety of distributions such as the **mlogit** package does. This document is a general description of **Rchoice** and all functionalities are illustrated using real databases.

Keywords: discrete choice models, simulated maximum likelihood, R, econometric, random parameters, hierarchical models.

1. Introduction

A growing number of empirical studies involve the assessment of influences on a choice among binary, count or ordered discrete alternatives. Those models are well known. However, one of the traditional modeling shortcomings is their inability to control for possible unobserved heterogeneity that may exist across individuals. For instance, one might assume that some variable, such as income or education, does not affect equally the utility of individuals, therefore there may be a deviation from the mean of the respective coefficient. Binary, count and ordered choice models with random parameters extension allows heterogeneity among individuals assuming some distribution on the parameters.

In this document we present the package **Rchoice** for program R. **Rchoice** is a package for estimating a variety of Ordered and Binary Choice Models with observed and unobserved heterogeneity in the coefficients. The estimation procedure is based on Maximum Simulated Likelihood (MSL) which allows to control for observed and unobserved heterogeneity in a very flexible way. ¹ To our knowledge, only LIMDEP (Greene 2002) is able to estimate these type of models in a concise and flexible manner. Therefore, this package intends to make available these estimation methods to the general public and practitioners in a friendly and flexible way.

¹For multinomial discrete choice models with random parameters see **mlogit** package in R (Croissant *et al.* 2012)

2. Maximum Simulated Likelihood Estimation

One of the major capability of the **Rchoice** package is to estimate Binary, Count and Ordered models with random parameters. They are estimated by **Rchoice** using Simulated Maximum Likelihood (SML). In this section we briefly explain some basic ideas of SML procedure. For a more complete treatment of SML see for example [Gourieroux and Monfort \(1997\)](#); [Hajivassiliou and Ruud \(1986\)](#); [Train \(2009\)](#); [Cameron and Trivedi \(2005\)](#)

A random parameter model or random coefficient model permits regression parameter to vary across individuals according to some distribution. A fully parametric random parameter model specifies the dependent variable y_i conditional on regressors \mathbf{x}_i and given parameters β_i to have conditional density $f(y_i|\mathbf{x}_i, \beta_i)$, where β_i are iid with density $g(\beta_i|\theta)$. Inference is based on the density of y_i conditional on \mathbf{x}_i and given θ :

$$f(y|\mathbf{x}_i, \theta) = \int f(y|\mathbf{x}, \beta)g(\beta, \theta)d\beta$$

This integral will not have a closed-form solution except in some especial cases. For example, we can assume normally distributed random parameters, with $\beta_i \sim N(\mu, \Sigma)$. Then $\beta_i = \mu + \Sigma^{-1/2}\mathbf{v}_i$, where $\mathbf{v}_i \sim N(\mathbf{0}, \mathbf{I})$, thus:

$$f(y|\mathbf{x}, \theta) = \int_{-\infty}^{\infty} f(y|\mathbf{x}, \mu, \Sigma) \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\mathbf{v}^2\right) d\mathbf{v} \quad (1)$$

Note that (1) has no close-form solution, that is, it is difficult to integrate out the random parameter and hence it difficult to perform ML estimation. However ML estimation may still be possible if we instead use a good approximation $\hat{f}(y|\mathbf{x}, \theta)$ of $f(y|\mathbf{x}, \theta)$ to form a likelihood function.

But, how can we obtain $\hat{f}(y|\mathbf{x}, \theta)$? A good approximation can be obtained by **Monte Carlo integration**.² This procedure provides an alternative to deterministic numerical integration. Here we can *simulate* the integration using random draws from the distribution $g(\beta|\theta)$. For example, the researcher specifies the function form $g(\beta|\theta)$ and wants to estimate the parameter θ . The Monte Carlo approximation is:

$$\hat{f}(y|\mathbf{x}_i, \beta_{ir}, \theta) = \frac{1}{R} \sum_{r=1}^R \tilde{f}(y|\mathbf{x}, \beta_r, \theta),$$

where β_{ir} , for example, is the r th draw of β from $g(\beta_i|\theta)$ for individual i . Given independence over i , the MSL is the value θ that maximizes:

$$\hat{\theta}_{MSL} \equiv \arg \max_{\theta \in \Theta} \sum_{i=1}^N \log \hat{f}(y|\mathbf{x}_i, \beta_{ir}, \theta)$$

²Another numerical approximation is Gauss-Hermite quadrature. However, it has been documented that for models with more than 3 random parameters SML performs better.

The following proposition gives the asymptotic distribution of MSL estimator. For a complete derivation of the asymptotic properties of the MSL and a more comprehensive view see [Lee \(1992\)](#) or [Gourieroux and Monfort \(1997\)](#).

Proposition 1 (Distribution of MSL Estimator). *Assume the following:*

1. *The data are from a simple random sample from a dgp with likelihood function $f(y|\mathbf{x}, \boldsymbol{\theta}_0)$ that satisfies the regularity conditions so that the ML estimator is consistent and asymptotically normal with variance matrix $\mathbf{A}^{-1}(\boldsymbol{\theta}_0)$, where:*

$$\left[\frac{1}{N} \sum_{i=1}^N \frac{\partial \log f(y|\mathbf{x}_i, \boldsymbol{\theta})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \Big|_{\boldsymbol{\theta}_0} \right] \xrightarrow{p} \mathbf{A}(\boldsymbol{\theta}_0)$$

2. *The likelihood function f is estimated using the simulator \hat{f} with \tilde{f} unbiased for f .*

*Then the **maximum simulated likelihood** estimator is asymptotically equivalent to the ML estimator if $R \rightarrow \infty$, $N \rightarrow \infty$ and $\sqrt{N}/R \rightarrow \infty$, and it has a limit normal distribution with:*

$$\sqrt{N}(\hat{\boldsymbol{\theta}}_{MSL} - \boldsymbol{\theta}) \xrightarrow{d} N(0, \mathbf{A}^{-1}(\boldsymbol{\theta}_0))$$

3. Technical Aspects of Maximum Simulated Likelihood

In this section, we show the general technical aspect of the MSL. This will allow us to accommodate any index-type regression model such as probit, logit, ordered and count model. This section relies heavily on chapter 15 of [Greene \(2012\)](#) and chapter 10 of [Train \(2009\)](#).

3.1. Simulated Maximum Likelihood

In order to develop a general set of results, it is convenient to write each single density in the simulated function as:

$$f(y_i|\mathbf{x}_i, \boldsymbol{\beta}_{ir}, \boldsymbol{\theta}) = P_{ir}(\boldsymbol{\theta}) = P_{ir}$$

where $\boldsymbol{\theta}$ is the vector that collects all the parameters. The simulated log-likelihood is :

$$\log L_s = \sum_{i=1}^N \log \left(\frac{1}{R} \sum_{r=1}^R P_{ir}(\boldsymbol{\theta}) \right) \quad (2)$$

If we define:

$$P_i = P_i(\boldsymbol{\theta}) = \frac{1}{R} \sum_{r=1}^R P_{ir},$$

then, the maximum likelihood can be written as:

$$\log L_s = \sum_{i=1}^N \log P_i(\boldsymbol{\theta}) \quad (3)$$

With this notation, we will be able to accommodate richer specifications of the index function and discrete choice models by simply changing the specification of P_{ir} . As typical, the index model represents a latent process of the form:

$$U_{ir}^*(\boldsymbol{\theta}) = \mathbf{z}'_i \boldsymbol{\delta} + \mathbf{x}'_i \boldsymbol{\beta}_{ir} + \epsilon_i, \quad (4)$$

where \mathbf{z}_i is a vector of variables with fixed parameters $\boldsymbol{\delta}$; \mathbf{x}_i is a vector of variables with random coefficients $\boldsymbol{\beta}_{ir}$; and ϵ_i is the error term. For simplicity, assume that $\boldsymbol{\beta}_{ir} \sim N(\boldsymbol{\beta}, \boldsymbol{\Sigma})$, then the random vector of coefficients can be written as:

$$\boldsymbol{\beta}_{ir} = \boldsymbol{\beta} + \mathbf{L}\boldsymbol{\omega}_{ir}$$

where $\boldsymbol{\omega}_{ir}$ is a vector of random draws from normal standard distribution. If the random parameters are correlated normal, then \mathbf{L} is a lower triangular which produces the covariance matrix of the random parameters, $\mathbf{L}\mathbf{L}' = \boldsymbol{\Sigma}$; otherwise, the matrix \mathbf{L} is a diagonal matrix of standard deviations. A hierarchical model is obtained by allowing the parameter heterogeneity to be partly systematic in terms of observed variables:

$$\boldsymbol{\beta}_{ir} = \boldsymbol{\beta} + \mathbf{I}\mathbf{s}_i + \mathbf{L}\boldsymbol{\omega}_{ir},$$

where \mathbf{I} is a matrix of parameters and \mathbf{s}_i is a vector of covariates. Then, $E(\boldsymbol{\beta}_i) = \boldsymbol{\beta} + \boldsymbol{\pi}\mathbf{s}_i + \mathbf{L}E(\boldsymbol{\omega}) = \boldsymbol{\beta} + \boldsymbol{\pi}\mathbf{s}_i$ and its covariance is $\text{Var}(\boldsymbol{\beta}_i) = E(\mathbf{L}\boldsymbol{\omega}(\boldsymbol{\omega}')') = \mathbf{L}E(\boldsymbol{\omega}\boldsymbol{\omega}')\mathbf{L} = \mathbf{L}\mathbf{L} = \mathbf{L}\mathbf{L}' = \boldsymbol{\Sigma}$.

Example 1 (Representation of Correlated Random Parameters). Suppose two correlated random parameters, β_1 and β_2 , whose mean depend upon variables S, B and C . Then:

$$\begin{aligned} \beta_{1,ir} &= \bar{\beta}_1 + \pi_{1,1}S_i + \pi_{1,2}B_i + \pi_{1,3}C_i + s_{11}\omega_{1,ir} \\ \beta_{2,ir} &= \bar{\beta}_2 + \pi_{2,1}S_i + \pi_{2,2}B_i + \pi_{2,3}C_i + s_{21}\omega_{1,ir} + s_{22}\omega_{2,ir} \end{aligned}$$

or in vector form:

$$\begin{pmatrix} \beta_{1,ir} \\ \beta_{2,ir} \end{pmatrix} = \begin{pmatrix} \bar{\beta}_1 \\ \bar{\beta}_2 \end{pmatrix} + \begin{pmatrix} \pi_{1,1} & \pi_{1,2} & \pi_{1,3} \\ \pi_{2,1} & \pi_{2,2} & \pi_{2,3} \end{pmatrix} \begin{pmatrix} S_i \\ B_i \\ C_i \end{pmatrix} + \begin{pmatrix} s_{11} & 0 \\ s_{21} & s_{22} \end{pmatrix} \begin{pmatrix} \omega_{1,ir} \\ \omega_{2,ir} \end{pmatrix}$$

In this case, the variance-covariance matrix of the random parameters is:

$$\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}' = \begin{pmatrix} s_{11} & 0 \\ s_{21} & s_{22} \end{pmatrix} \begin{pmatrix} s_{11} & s_{21} \\ 0 & s_{22} \end{pmatrix} = \begin{pmatrix} s_{11}^2 & s_{11}s_{22} \\ s_{21}s_{22} & s_{21}^2 + s_{22}^2 \end{pmatrix},$$

and the conditional mean vector is:

$$E(\boldsymbol{\beta}_i|\mathbf{s}_i) = \boldsymbol{\beta} + \mathbf{I}\mathbf{s}_i$$

Rchoice allows to specify different distribution for the random parameters. See section 4.

Finally, depending in the nature of the dependent variable and the distribution of the error term, the probability for each individual can be specified. It is well known that, if the dependent variable is binary, then the probability for each individual in each draw is:

$$P_{ir}(\boldsymbol{\theta}) = F(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) \quad (5)$$

where $q_i = 2y_i - 1$ ³. Furthermore, if the model is probit, then:

$$\begin{aligned} F(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) &= \Phi(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) \\ f(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) &= \phi(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) \end{aligned}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the CDF and the pdf for the standard normal distribution. Likewise, if the model is logit, then:

$$\begin{aligned} F(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) &= \Lambda(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) = \frac{\exp(q_i \cdot U_{ir}^*(\boldsymbol{\theta}))}{1 + \exp(q_i \cdot U_{ir}^*(\boldsymbol{\theta}))} \\ f(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) &= \Lambda(q_i \cdot U_{ir}^*(\boldsymbol{\theta})) [1 - \Lambda(q_i \cdot U_{ir}^*(\boldsymbol{\theta}))] \end{aligned}$$

For the Poisson model, the probability for individual i for the r draw is:

$$P_{ir}(\boldsymbol{\theta}) = \frac{\exp(-\exp(U_{ir}^*(\boldsymbol{\theta}))) \exp(U_{ir}^*(\boldsymbol{\theta}))^{y_i}}{y_i!} \quad (6)$$

and for the ordered model, we have:

$$P_{ir}(\boldsymbol{\theta}) = F(\kappa_j - U_{ir}^*(\boldsymbol{\theta})) - F(\kappa_{j-1} - U_{ir}^*(\boldsymbol{\theta})) \quad (7)$$

where:

$$\kappa_j = \kappa_{j-1} + \exp(\alpha_j)$$

This last reparametrization ensure the ordering of the thresholds.

3.2. Gradient of the Simulated Maximum Likelihood

SML procedure is very time consuming. Providing the gradient to the maximization procedure can considerable reduce the time to achieve convergence. **Rchoice** package provides the

³As explained by Greene (2012), if the distribution is symmetric, as the normal and logistic are, then $1 - F(\mathbf{x}'\boldsymbol{\beta}) = F(\mathbf{x}'\boldsymbol{\beta})$. Then, $\log L = \sum_i F(q_i \mathbf{x}_i \boldsymbol{\beta})$

gradient for all random parameter models. Next, the formulas used by **Rchoice** to obtain the gradient are given. To obtain the derivatives, we begin with:

$$\frac{\partial \log L_s}{\partial \boldsymbol{\theta}} = \sum_{i=1} \left(\frac{\frac{1}{R} \sum_{r=1}^R \left(\frac{\partial P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)}{\frac{1}{R} \sum_{r=1}^R P_{ir}(\boldsymbol{\theta})} \right) \quad (8)$$

For the derivative term,

$$\begin{aligned} \frac{\partial P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} &= P_{ir} \frac{\partial \log P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \\ &= P_{ir}(\boldsymbol{\theta}) \mathbf{g}_{ir}(\boldsymbol{\theta}). \end{aligned} \quad (9)$$

where we use the fact that $\frac{\partial \log p}{\partial \theta} = \frac{1}{p} \frac{\partial p}{\partial \theta}$. Now, inserting 9 into 8 we get:

$$\frac{\partial \log L_s}{\partial \boldsymbol{\theta}} = \sum_{i=1} \left(\frac{\sum_{r=1}^R P_{ir}(\boldsymbol{\theta}) \mathbf{g}_{ir}(\boldsymbol{\theta})}{\sum_{r=1}^R P_{ir}(\boldsymbol{\theta})} \right). \quad (10)$$

Define the weight $Q_{ir}(\boldsymbol{\theta}) = P_{ir}(\boldsymbol{\theta}) / \sum_{r=1}^R P_{ir}(\boldsymbol{\theta})$ so that $0 < Q_{ir}(\boldsymbol{\theta}) < 1$ and $\sum_{r=1}^R Q_{ir}(\boldsymbol{\theta}) = 1$. Then,

$$\frac{\partial \log L_s}{\partial \boldsymbol{\theta}} = \sum_{i=1}^N \sum_{r=1}^R Q_{ir}(\boldsymbol{\theta}) \mathbf{g}_{ir}(\boldsymbol{\theta}) = \sum_{i=1}^N \bar{\mathbf{g}}_i(\boldsymbol{\theta}). \quad (11)$$

Example 2 (Gradients formulas used by Rchoice). For binary models, taking first derivatives on the log of equation 5 we get:

$$\frac{\partial \log P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{g}_{ir} = \lambda_{ir}(\boldsymbol{\theta}) \left(\frac{\partial U_{ir}^*(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right), \quad (12)$$

where:

$$\lambda_{ir} = \frac{q_i f(q_i U_{ir}^*(\boldsymbol{\theta}))}{F(q_i U_{ir}^*(\boldsymbol{\theta}))},$$

and $f = \frac{\partial F(\cdot)}{\partial \theta}$. If the model is the Poisson regression model, then \mathbf{g}_{ir} is given by:

$$\frac{\partial \log P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \mathbf{g}_{ir} = (y_i - \exp(U_{ir}^*(\boldsymbol{\theta}))) \left(\frac{\partial U_{ir}^*(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right).$$

For the ordered model, let $\boldsymbol{\theta}$ the vector collecting all the parameters except for the thresholds parameters. Then, \mathbf{g}_{ir} is given by:

$$\frac{\partial \log P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = - \frac{f_{ir,j}(\boldsymbol{\theta}) - f_{ir,j-1}(\boldsymbol{\theta})}{F_{ir,j}(\boldsymbol{\theta}) - F_{ir,j-1}(\boldsymbol{\theta})} \left(\frac{\partial U_{ir}^*(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right),$$

and by:

$$\frac{\partial \log P_{ir}(\boldsymbol{\theta})}{\partial \boldsymbol{\alpha}_k} = \frac{d_{j,k} f_{ir,j}(\boldsymbol{\theta}) - d_{j-1,k} f_{ir,j-1}(\boldsymbol{\theta})}{F_{ir,j}(\boldsymbol{\theta}) - F_{ir,j-1}(\boldsymbol{\theta})} \left(\frac{\partial \kappa_j}{\partial \boldsymbol{\alpha}_k} - \frac{\partial \kappa_{j-1}}{\partial \boldsymbol{\alpha}_k} \right)$$

with $\delta_{j,k} = 1$ if $j = k$ and 0 otherwise. Finally, if parameters are uncorrelated, then:

$$\frac{\partial U_{itr}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \begin{pmatrix} \mathbf{z}_{it} \\ \mathbf{x}_{it} \\ \mathbf{s}_i \otimes \mathbf{x}_{it} \\ \boldsymbol{\omega}_{it} \bullet \mathbf{x}_{it} \end{pmatrix}.$$

Rchoice uses this formulas to compute the gradient and uses the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (as default) to iteratively solve the MSL.

4. Drawing from Densities

SML procedure requires to draw pseudo-random numbers from the specified distribution for simulation. A good performance of SML requires very large number of draws. The main drawback to this approach is that with large samples and complex models, the maximization of $\log L_s$ can be very time consuming. Researchers have gained speed with no degradation in simulation performance through the use of small number of Halton draws (Bhat 2001; Train 2000). The idea is that, instead of taking independent random draws, simulation can potentially be improved by selecting evaluation points more systematically and with better coverage (Sándor and Train 2004). In this section, we detail how draws are computed by **Rchoice**.⁴

⁴In terms of programming, we have modified the functions from **mlogit** (Croissant *et al.* 2012) to allow variation in the mean by observed covariates (Hierarchical model).

Suppose that there are K_2 random parameters. Then, the K_2 elements of ω_{ir} are drawn as follows. We begin with a K_2 random vector ω_{ir} that is:

- K_2 independent draws from the standard uniform $(0, 1)$ distribution or
- K_2 independent draws from the m th Halton sequence, where m is the m th prime number in the sequence of K_2 prime numbers beginning with 2.

An important attribute of the Halton values is that they are also distributed in the $(0, 1)$ interval. Then, the primitive draw (Pseudo or Halton draws) is then transformed to the distribution specified by the user as follows:

- $u_{k,ir} \sim U(0, 1)$: primitive draw from halton or pseudo-random number generator
- $w_{k,ir} = \Phi^{-1}(u_{k,ir}) \sim N(0, 1)$

Using these two primitive draws, **Rchoice** creates the random parameters as follows:

1. Normal Parameter:

$$\begin{aligned}\beta_{k,ir} &= \beta_k + \sigma_k w_{k,ir} \\ w_{k,ir} &\sim N(0, 1)\end{aligned}$$

where β_k and σ_k are estimated. Then, $\beta_{k,i} \sim N(\beta_k, \sigma_k^2)$

2. Truncated normal Parameter:

$$\begin{aligned}\beta_{k,ir} &= \begin{cases} \beta_k + \sigma_k w_{k,ir} & \text{if } \beta_k + \sigma_k w_{k,ir} > 0 \\ 0 & \text{otherwise} \end{cases} \\ w_{k,ir} &\sim N(0, 1)\end{aligned}$$

where β_k and σ_k are estimated. Then, $\beta_{k,i} \sim N(\beta_k, \sigma_k^2)$ with the share below zero massed at zero

3. Log-Normal Distribution:

$$\begin{aligned}\beta_{k,ir} &= \exp(\beta_k + \sigma_k w_{k,ir}) \\ w_{k,ir} &\sim N(0, 1)\end{aligned}$$

where β_k and σ_k are estimated. Then, $\beta_{k,i} \sim \log N(\beta_k, \sigma_k^2)$

4. Uniform:

$$\begin{aligned}\beta_{k,ir} &= \beta_k - \sigma_k + 2\sigma_k u_{k,ir} \\ u_{k,ir} &\sim U(0,1)\end{aligned}$$

where β_k and σ_k are estimated.

5. Triangular distribution:

$$\begin{aligned}\beta_{k,ir} &= \beta_k + \sigma_k v_{k,ir} \\ v_{k,ir} &\sim 1(u_{k,ir} < 0.5) (\sqrt{2u_{k,ir}} - 1) + 1(u_{k,ir} \geq 0.5) \left(1 - \sqrt{2(1 - u_{k,ir})}\right)\end{aligned}$$

where β_k and σ_k are estimated.

Rchoice allows to the user to specify two type of random draws by the argument `haltons`: pseudo-random draws (`haltons = NULL`) and Halton draws (`haltons = NA`) as default. If `haltons = NULL`, the seed is set to `set.seed(123)`. The user can change this by the `seed` argument. For the Halton draws, the default is to use the first K_2 primes numbers starting with 3. Within each series, the first 100 draws are discarded, as the first draws tend to be highly correlated across different draw. The user can also change the prime number and the element dropped for each serie. For example, if $K_2 = 2$, and the user wants to use the primes numbers 5 and 31 along with dropping the first 10 draws, he could specify `haltons = list("prime" = c(5,31), "drop" = c(10,10))`.

Note that log-normal and truncated normal give positive coefficients only. If the user wants a variable to have only negative coefficients, he should create the negative of the variable.

5. Poisson Regression Model Examples

5.1. Standard Poisson Model

In R there exist several package to estimate binary, count and ordered models. `glm` function allows to estimate different kind of discrete choice models such as Poisson and binary models. The function `probit` from the package **micEcon** allows to estimate probit model. Moreover, the function `polr` from the package **MASS** allows to estimate ordered models (Venables and Ripley 2002). The advantage of **Rchoice** is that allows more flexibility in the optimization routines which improves the convergence speed. **Rchoice** uses the function `maxLik` in order to maximize the log-likelihood function, which permits to estimate models by the Newton-Raphson (NR), BGFs and Berndt-Hall-Hall-Hausman (BHHH) procedures (see Henningsen and Toomet 2011).

In this section, we show the capabilities of **Rchoice** to estimate Poisson regression model with and without random parameters.

Rchoice is loaded by typing:

```
library("Rchoice")
```

In order to show how to estimate Poisson regression models using **Rchoice**, we will use data on scientific productivity (Long 1990, 1997). We load the data using

```
data("Articles")
head(Articles, 3)

##   art fem mar kid5  phd ment
## 1   0  0  1   0 2.52   7
## 2   0  1  0   0 2.05   6
## 3   0  1  0   0 3.75   6
```

To see more information about the data, one can use:

```
help(Articles)
```

The work by Long (1990) suggest that gender, marital status, number of young children, prestige of the graduate program, and the number of articles written by a scientist's mentor could affect a scientist's level of publication. In order to see this, we estimate a Poisson regression model use the **Rchoice** function specifying `link = "poisson"`:

```
poisson <- Rchoice(art ~ fem + mar + kid5 + phd + ment, data = Articles,
                  family = poisson)

##
## Starting values Parameters:
## (Intercept)      fem      mar      kid5      phd      ment
## 1.33425 -0.38089 0.26320 -0.29144 -0.01137 0.06149

summary(poisson)

##
## Model: poisson
## Model estimated on: Tue Apr 01 12:41:11 2014
##
## Call:
## Rchoice(formula = art ~ fem + mar + kid5 + phd + ment, data = Articles,
##         family = poisson, method = "nr")
##
```

```
## The estimation took: 0h:0m:0s
##
## Coefficients:
##           Estimate Std. Error t-value Pr(>|t|)
## (Intercept)  0.30462    0.10298    2.96  0.0031 **
## fem          -0.22459    0.05461   -4.11  3.9e-05 ***
## mar           0.15524    0.06137    2.53  0.0114 *
## kid5         -0.18488    0.04013   -4.61  4.1e-06 ***
## phd           0.01282    0.02640    0.49  0.6271
## ment         0.02554    0.00201   12.73 < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Optimization of log-likelihood by Newton-Raphson maximisation
## Log Likelihood: -1650
## Number of observations: 915
## Number of iterations: 7
## Exit of MLE: gradient close to zero
```

The output shows that the log-likelihood function is estimated using NR algorithm in 7 iterations. If the user wants to estimate the model using another algorithm he should type `method = "bfgs"` for the BGFS method or `method = "bhhh"` for BHHH method.

In terms of interpretation, we can say that, being a female scientist decreases the expected number of articles by a factor of 0.8 ($= \exp(-.225)$), holding all other variables constant. Or equivalently, being a female scientist decreases the expected number of articles by 20% ($= 100[\exp(-.225) - 1]$), holding all other variables constant. Prestige of PhD department is not important for productivity.

Another capability of **Rchoice** is its interaction with other packages in R. For example, we can compute the robust standard error by using the package **sandwich**:

```
library(sandwich)
library(lmtest)

coeftest(poisson, vcov = sandwich)

##
## t test of coefficients:
##
##           Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.30462    0.14652    2.08  0.03790 *
## fem          -0.22459    0.07166   -3.13  0.00178 **
## mar           0.15524    0.08193    1.89  0.05843 .
## kid5         -0.18488    0.05596   -3.30  0.00099 ***
```

```
## phd          0.01282    0.04196    0.31  0.76001
## ment        0.02554    0.00382    6.69  3.9e-11 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

In order to get the same robust standard errors as STATA (StataCorp 2011), we need to make a small sample correction:

```
vcov.stata <- vcovHC(poisson, type = "HC0") * nObs(poisson)/(nObs(poisson) -
  1)
coefTest(poisson, vcov = vcov.stata)

##
## t test of coefficients:
##
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.30462    0.14660    2.08  0.0380 *
## fem          -0.22459    0.07170   -3.13  0.0018 **
## mar           0.15524    0.08197    1.89  0.0586 .
## kid5         -0.18488    0.05599   -3.30  0.0010 ***
## phd           0.01282    0.04199    0.31  0.7601
## ment         0.02554    0.00382    6.69  4e-11 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

where the correction is $n/(n - 1)$.

Rchoice also interacts with `linearHypothesis` and `deltaMethod` functions from **car** (Fox, Bates, Firth, Friendly, Gorjanc, Graves, Heiberger, Monette, Nilsson, Ogle *et al.* 2009) and the `lrtest` and `waldtest` functions from **lmtest** package (Zeileis and Hothorn 2002). For example, we can test $H_0 : \text{phd}/\text{ment} = 0$ by:

```
library(car)
```

```
deltaMethod(poisson, "phd/ment")
```

```
##           Estimate      SE
## phd/ment    0.502 1.043
```

5.2. Random Parameter Poisson Model

Now, we estimate a Poisson regression model with random parameters. In this case, we will assume that the effect of `kid5`, `phd` and `ment` are not fixed, but rather heterogeneous

among the population. Specifically, we will assume that the coefficients for those variables are independent normally distributed, that is, we will not allow correlation among them:

$$\begin{aligned}\beta_{\text{kid5},i} &= \beta_{\text{kid5}} + \sigma_{\text{kid5}}\omega_{\text{kid5},ir} \\ \beta_{\text{phd},i} &= \beta_{\text{phd}} + \sigma_{\text{phd}}\omega_{\text{phd},ir} \\ \beta_{\text{ment},i} &= \beta_{\text{ment}} + \sigma_{\text{ment}}\omega_{\text{ment},ir}\end{aligned}$$

Then, in order to estimate this model, we can write:

```
poisson.ran <- Rchoice(art ~ fem + mar + kid5 + phd + ment,
                      data = Articles, ranp = c(kid5 = "n", phd = "n", ment = "n"),
                      family = poisson,
                      R = 10)

##
## Starting values Parameters:
## (Intercept)      fem      mar  mean.kid5  mean.phd  mean.ment
##    0.30462    -0.22459    0.15524   -0.18488    0.01282    0.02554
##    sd.kid5      sd.phd      sd.ment
##    0.00000      0.00000      0.00000
```

It is important to discuss the arguments for the `Rchoice` function. First, the argument `ranp` indicates which variables are random in the formula and their distributions. In this case, we have specified that all of them are normal distributed using "n". The number of draws are not specified. Therefore, `Rchoice` will set `R = 40` as default. The user can change this by changing the `R` argument. The type of draws are Halton draws as a default, but if the user wants pseudo-random draws he can specify `haltons = NULL`. As explained before, the default maximization algorithm for SML is BGFS.

```
summary(poisson.ran)

##
## Model: poisson
## Model estimated on: Tue Apr 01 12:41:57 2014
##
## Call:
## Rchoice(formula = art ~ fem + mar + kid5 + phd + ment, data = Articles,
##         family = poisson, ranp = c(kid5 = "n", phd = "n", ment = "n"),
##         R = 10, method = "bfgs", iterlim = 2000)
##
## The estimation took: 0h:0m:45s
##
## Coefficients:
##              Estimate Std. Error t-value Pr(>|t|)
## (Intercept)  0.24215     0.12327    1.96  0.0495 *
```

```
## fem          -0.21443    0.06560   -3.27    0.0011 **
## mar           0.18122    0.07337    2.47    0.0135 *
## mean.kid5    -0.26565    0.06292   -4.22   2.4e-05 ***
## mean.phd     -0.01816    0.03510   -0.52    0.6048
## mean.ment     0.03040    0.00354    8.59   < 2e-16 ***
## sd.kid5       0.32641    0.07691    4.24   2.2e-05 ***
## sd.phd        0.12990    0.01675    7.75   8.9e-15 ***
## sd.ment       0.01572    0.00295    5.32   1.0e-07 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Optimization of log-likelihood by BFGS maximisation
## Log Likelihood: -1590
## Number of observations: 915
## Number of iterations: 115
## Exit of MLE: successful convergence
## Simulation based on 10 Halton draws
```

The result shows that indeed the three coefficients are random in the sample. We can also say that :

```
pnorm(coef(poisson.ran)["mean.kid5"]/coef(poisson.ran)["sd.kid5"])
## mean.kid5
##      0.2079
```

a 21% of the individuals have a positive coefficient for kid5. Note also that the mean coefficient for phd is 0 (not significant). This is due to the fact that the unobserved heterogeneity among scientists in the sample cancel out positive and negative effects. These observations are not possible with a Poisson regression with fixed effect.

Suppose that now we want to test if $H_0 = \sigma_{\text{kid5}} = \sigma_{\text{phd}} = \sigma_{\text{ment}} = 0$. This can be done by using the function `waldtest` or `lrtest` from package **lmtest**:

```
waldtest(poisson.ran, poisson)
## Wald test
##
## Model 1: art ~ fem + mar + kid5 + phd + ment
## Model 2: art ~ fem + mar + kid5 + phd + ment
##   Res.Df Df Chisq Pr(>Chisq)
## 1      906
## 2      909 -3    202    <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
lrtest(poisson.ran, poisson)

## Likelihood ratio test
##
## Model 1: art ~ fem + mar + kid5 + phd + ment
## Model 2: art ~ fem + mar + kid5 + phd + ment
##   #Df LogLik Df Chisq Pr(>Chisq)
## 1   9 -1590
## 2   6 -1651 -3   122   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Both test reject the null hypothesis. We can also specify different distribution of the parameters by using the S3 method `update`:

```
poisson.ran2 <- update(poisson.ran, ranp = c(kid5 = "u", phd = "t", ment = "cn"))
summary(poisson.ran2)
```

Now, we estimate the model `poisson.ran`, but assuming that the random parameters are correlated:

```
poissonc.ran <- Rchoice(art ~ fem + mar + kid5 + phd + ment, data = Articles,
                       ranp = c(kid5 = "n", phd = "n", ment = "n"),
                       family = poisson, correlation = TRUE,
                       R = 10)

##
## Starting values Parameters:
## (Intercept)          fem          mar    mean.kid5    mean.phd
##   0.30462    -0.22459    0.15524    -0.18488    0.01282
##   mean.ment sd.kid5.kid5  sd.kid5.phd sd.kid5.ment  sd.phd.phd
##   0.02554    0.00000    0.00000    0.00000    0.00000
##   sd.phd.ment sd.ment.ment
##   0.00000    0.00000

summary(poissonc.ran)

##
## Model: poisson
## Model estimated on: Tue Apr 01 12:43:39 2014
##
## Call:
## Rchoice(formula = art ~ fem + mar + kid5 + phd + ment, data = Articles,
##         family = poisson, ranp = c(kid5 = "n", phd = "n", ment = "n"),
##         R = 10, correlation = TRUE, method = "bfgs", iterlim = 2000)
```

```
##
## The estimation took: 0h:1m:42s
##
## Coefficients:
##           Estimate Std. Error t-value Pr(>|t|)
## (Intercept)  0.22781    0.12539   1.82 0.06926 .
## fem         -0.21160    0.06736  -3.14 0.00168 **
## mar          0.13569    0.07590   1.79 0.07381 .
## mean.kid5   -0.21418    0.06100  -3.51 0.00045 ***
## mean.phd    -0.01125    0.03628  -0.31 0.75651
## mean.ment   0.02858    0.00400   7.15 8.5e-13 ***
## sd.kid5.kid5 0.30468    0.09019   3.38 0.00073 ***
## sd.kid5.phd  0.12742    0.04497   2.83 0.00460 **
## sd.kid5.ment -0.02534    0.00395  -6.42 1.4e-10 ***
## sd.phd.phd   0.11332    0.03060   3.70 0.00021 ***
## sd.phd.ment  0.00280    0.00356   0.79 0.43052
## sd.ment.ment 0.00335    0.00469   0.71 0.47519
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Optimization of log-likelihood by BFGS maximisation
## Log Likelihood: -1570
## Number of observations: 915
## Number of iterations: 178
## Exit of MLE: successful convergence
## Simulation based on 10 Halton draws
```

We can extract the $\Sigma = \mathbf{LL}'$ matrix of variance-covariance matrix and the correlation matrix of the random parameters using `cov.Rchoice` and `cor.Rchoice`:

```
cov.Rchoice(poissonc.ran)

##           kid5           phd           ment
## kid5  0.092830  0.038823 -0.0077220
## phd   0.038823  0.029077 -0.0029117
## ment -0.007722 -0.002912  0.0006614

cor.Rchoice(poissonc.ran)

##           kid5           phd           ment
## kid5  1.0000  0.7473 -0.9855
## phd   0.7473  1.0000 -0.6639
## ment -0.9855 -0.6639  1.0000
```

5.3. Hierarchical Poisson Random Parameter Model

In this section we show how to estimate a Hierarchical Poisson Random Parameter Model. In this case, we assume that there exist not only unobserved heterogeneity in the coefficients for `kid5`, `phd` and `ment`, but also observed heterogeneity in the mean. Specifically, we assume that:

$$\begin{aligned}\beta_{\text{kid5},i} &= \beta_{\text{kid5}} + \pi_1 \text{fem} + \sigma_{\text{kid5},\text{kid5}} \omega_{\text{kid5},ir} \\ \beta_{\text{phd},i} &= \beta_{\text{phd}} + \pi_2 \text{fem} + \sigma_{\text{phd},\text{kid5}} \omega_{\text{kid5},ir} + \sigma_{\text{phd},\text{phd}} \omega_{\text{phd},ir} \\ \beta_{\text{ment},i} &= \beta_{\text{ment}} + \pi_3 \text{fem} + \sigma_{\text{ment},\text{kid5}} \omega_{\text{kid5},ir} + \sigma_{\text{ment},\text{phd}} \omega_{\text{phd},ir} + \sigma_{\text{ment},\text{ment}} \omega_{\text{ment},ir}\end{aligned}$$

The formulation above implies that those three coefficients (or marginal effect on latent productivity) varies also if the individual is female or male. **Rchoice** manages the variables in the hierarchical model by the `formula` object: all the hierarchical variables are included after the `|` symbol. For example, we can estimate this model by typing:

```
poissonH.ran <- Rchoice(art ~ fem + mar + kid5 + phd + ment | fem,
  data = Articles,
  ranp = c(kid5="n", phd = "n", ment = "n"),
  family = poisson,
  correlation = TRUE,
  R = 10)

## Warning: Model assumes no constant in S variables...updating formula

##
## Starting values Parameters:
## (Intercept)      fem      mar  mean.kid5  mean.phd
## 0.30462 -0.22459 0.15524 -0.18488 0.01282
## mean.ment  kid5.fem  phd.fem  ment.fem sd.kid5.kid5
## 0.02554 0.00000 0.00000 0.00000 0.00000
## sd.kid5.phd sd.kid5.ment sd.phd.phd sd.phd.ment sd.ment.ment
## 0.00000 0.00000 0.00000 0.00000 0.00000

summary(poissonH.ran)

##
## Model: poisson
## Model estimated on: Tue Apr 01 12:46:24 2014
##
## Call:
## Rchoice(formula = art ~ fem + mar + kid5 + phd + ment | fem,
## data = Articles, family = poisson, ranp = c(kid5 = "n", phd = "n",
## ment = "n"), R = 10, correlation = TRUE, method = "bfgs",
## iterlim = 2000)
##
```

```
## The estimation took: 0h:2m:45s
##
## Coefficients:
##           Estimate Std. Error t-value Pr(>|t|)
## (Intercept)  0.39314    0.15924    2.47 0.01356 *
## fem          -0.59283    0.22640   -2.62 0.00883 **
## mar           0.15045    0.07728    1.95 0.05156 .
## mean.kid5    -0.17707    0.06545   -2.71 0.00682 **
## mean.phd     -0.08937    0.04821   -1.85 0.06376 .
## mean.ment     0.03265    0.00457    7.14 9.4e-13 ***
## kid5.fem     -0.06172    0.11432   -0.54 0.58930
## phd.fem       0.15122    0.07117    2.12 0.03359 *
## ment.fem     -0.00495    0.00774   -0.64 0.52181
## sd.kid5.kid5  0.13224    0.06771    1.95 0.05081 .
## sd.kid5.phd  -0.10946    0.04239   -2.58 0.00982 **
## sd.kid5.ment -0.00641    0.00538   -1.19 0.23369
## sd.phd.phd   0.13512    0.04100    3.30 0.00098 ***
## sd.phd.ment  0.02127    0.00371    5.73 9.9e-09 ***
## sd.ment.ment 0.00728    0.00451    1.61 0.10666
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Optimization of log-likelihood by BFGS maximisation
## Log Likelihood: -1580
## Number of observations: 915
## Number of iterations: 185
## Exit of MLE: successful convergence
## Simulation based on 10 Halton draws
```

The estimated parameters indicates that gender matters only for phd mean coefficient. We can test if the interaction variables are jointly significant by using `lrtest`:

```
lrtest(poissonH.ran, poissonc.ran)

## Likelihood ratio test
##
## Model 1: art ~ fem + mar + kid5 + phd + ment | fem - 1
## Model 2: art ~ fem + mar + kid5 + phd + ment
##   #Df LogLik Df Chisq Pr(>Chisq)
## 1  15  -1578
## 2  12  -1573 -3  10.2      0.017 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

5.4. Plotting Conditional Means

It is important to note that the estimates of the model parameters provide the unconditional estimates of the parameter vector, but we can form a person specific conditional estimator (see Train 2009; Greene 2012). The estimator of the conditional mean of the distribution of the random parameters, conditioned on the person specific data, is:

$$\widehat{\beta}_i = \widehat{E}(\beta_i | \text{data}_i) = \sum_{r=1}^R \widehat{Q}_{ir} \widehat{\beta}_{ir}$$

where:

$$\widehat{\beta}_{ir} = \widehat{\beta} + \widehat{\Pi} s_i + \widehat{\mathbf{L}} \omega_{ir}$$

We can also estimate the standard deviation of this distribution by estimating:

$$\widehat{E}(\beta_i^2 | \text{data}_i) = \sum_{r=1}^R \widehat{Q}_{ir} \widehat{\beta}_{ir}^2,$$

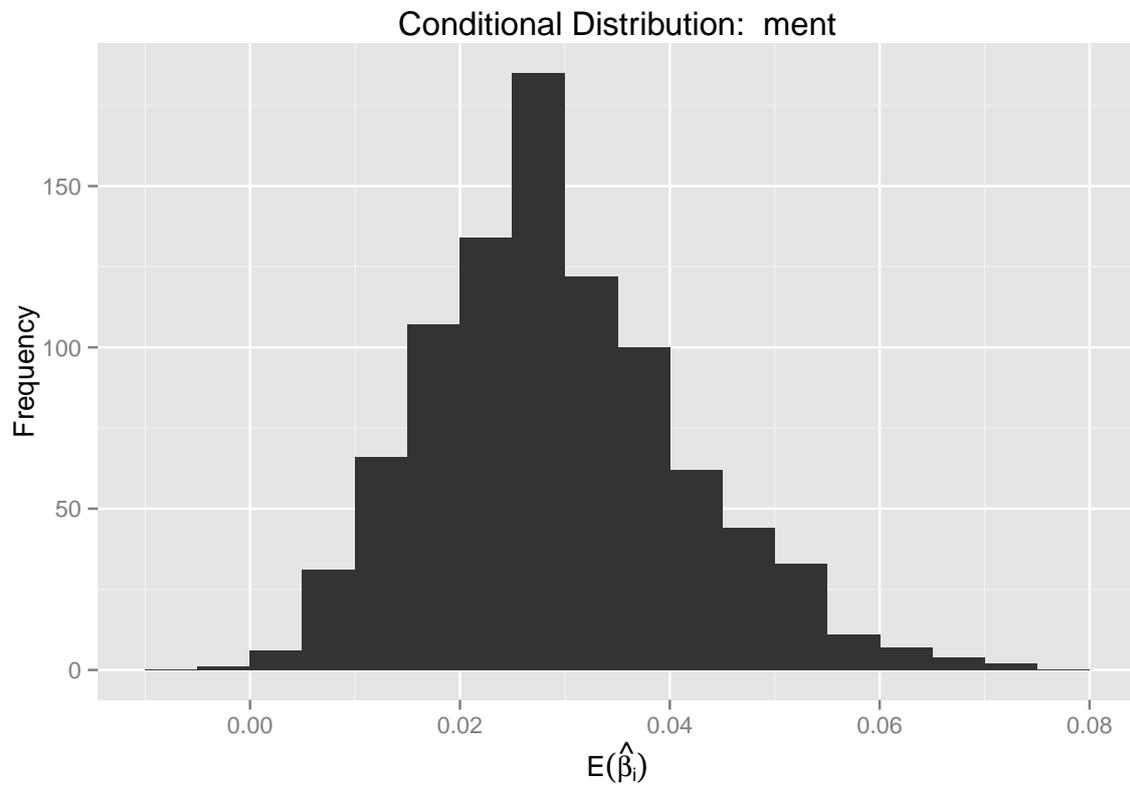
then computing the square root of the estimated variance,

$$\sqrt{\widehat{E}(\beta_i^2 | \text{data}_i) - \widehat{E}(\beta_i | \text{data}_i)^2}$$

With the estimates of the conditional mean and conditional variance, we can then compute the limits of an interval that resembles a confidence interval as the mean plus and minus two estimated standard deviation. This will construct an interval that contains at least 95 percent of the conditional distribution of β_i (Greene 2012).

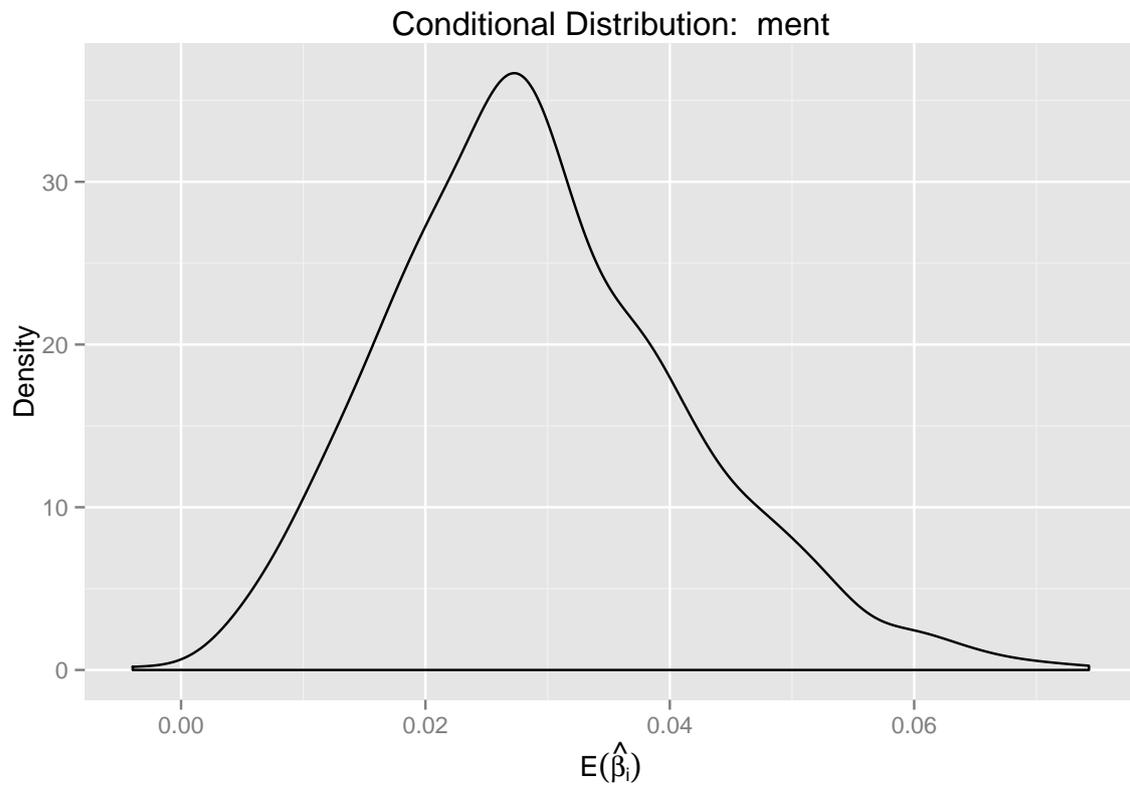
Rchoice allows to plot the histogram and kernel density of conditional means of random parameters using the function `plot`. For the histogram of the conditional mean of $\beta_{\text{ment},i}$, we can write:

```
plot(poissonH.ran, par = "ment", type = "histogram", bin = 0.005)
```



and for the kernel density:

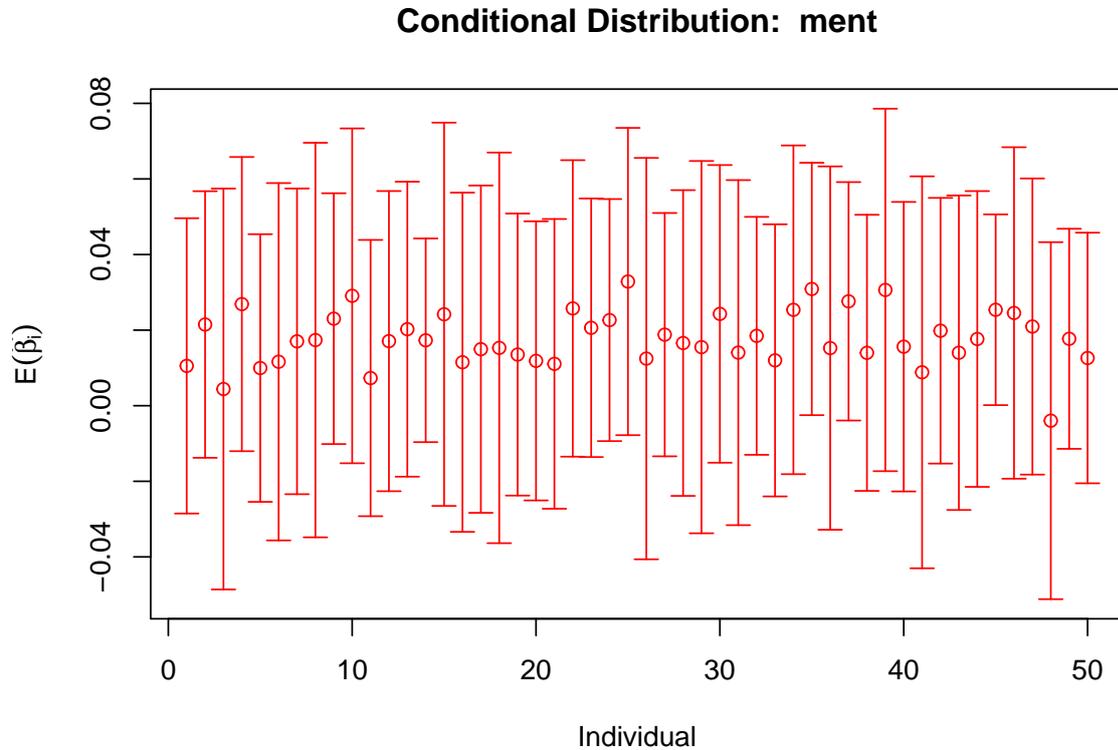
```
plot(poissonH.ran, par = "ment")
```



As [Greene \(2012\)](#) points out, even if the analysis departs from normal marginal distributions β_i , the sample distribution of the n estimated conditional means is not necessarily normal. Therefore, the kernel estimator based on the n estimators can have a variety of shapes.

We may also plot the individual confident interval for the conditional means for the first 50 individuals:

```
plot(poissonH.ran, par = "ment", ind = TRUE, id = seq(1, 50, 1))
```



6. Binary and Ordered Model Examples

In this section, we show how to estimate Binary and Ordered (Probit or Logit) models. Since the main characteristics of **Rchoice** package to estimate random parameter models were shown in the previous section, in this section we briefly show how to use **Rchoice** for Binary and Ordered models with random parameters.⁵

The argument `family` allows to specify between ordinal and binomial models. We estimate a Probit model with random parameters using the `Workmroz` data base:

```
data("Workmroz")
probit.ran <- Rchoice(lfp ~ k5 + k618 + age + wc + hc + lwg + inc,
  ranp = c(k5 = "n", hc = "n"),
  family = binomial('probit'),
  data = Workmroz,
  R = 100)
```

⁵MSL procedure is very complex. Sometimes, the model may not converge, or it may converge to a local maximum, if the initial values are not chosen well enough. For this reason, we recommend users to always use the argument `codeprint.level` in the estimation. If the model converges, but the Hessian is singular, we recommend using a vector of zeros as initial values, use different prime numbers, or use a different method of optimization. As default, the initial values correspond to those estimated by standard models with zeros for the parameters of **L** and **II** matrices.

```
summary(probit.ran)
```

In order to analyze ordered models, we use the "Health" database:

```
data("Health")
```

We estimate a Random Parameter Ordered Probit model for the variable `newhsat` for year 1988 (see `help(Health)`)

```
oprobit.ran <- Rchoice(newhsat ~ age + educ + hhinc + married + hhkids,  
  data = Health, family = ordinal('probit'),  
  subset = year == 1988,  
  ranp = c(age = "n", hhinc = "n"),  
  start = rep(0, 11))  
summary(oprobit.ran)
```

References

- Bhat CR (2001). “Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model.” *Transportation Research Part B: Methodological*, **35**(7), 677–693.
- Cameron AC, Trivedi PK (2005). *Microeconometrics: methods and applications*. Cambridge university press.
- Croissant Y, *et al.* (2012). “Estimation of multinomial logit models in R: The mlogit Packages.” *R package version 0.2-2*. URL: <http://cran.r-project.org/web/packages/mlogit/vignettes/mlogit.pdf>.
- Fox J, Bates D, Firth D, Friendly M, Gorjanc G, Graves S, Heiberger R, Monette G, Nilsson H, Ogle D, *et al.* (2009). “CAR: Companion to applied regression, R Package version 1.2-16.” Online at <http://cran.r-project.org/web/packages/car/index.html> (accessed on August 2012).
- Gourieroux C, Monfort A (1997). *Simulation-based econometric methods*. Oxford University Press.
- Greene WH (2002). *LIMDEP: Version 8.0: Econometric Modeling Guide*. Econometric Software.
- Greene WH (2012). *Econometric analysis*. 7 edition. Prentice Hall. ISBN 0131395386.
- Hajivassiliou VA, Ruud PA (1986). “Classical estimation methods for LDV models using simulation.” In RF Engle, D McFadden (eds.), *Handbook of Econometrics*, volume 4 of *Handbook of Econometrics*, chapter 40, pp. 2383–2441. Elsevier.
- Henningsen A, Toomet O (2011). “maxLik: A package for maximum likelihood estimation in R.” *Computational Statistics*, **26**(3), 443–458.
- Lee LF (1992). “On efficiency of methods of simulated moments and maximum simulated likelihood estimation of discrete response models.” *Econometric Theory*, **8**(04), 518–552.
- Long JS (1990). “The origins of sex differences in science.” *Social Forces*, **68**(4), 1297–1316.
- Long JS (1997). *Regression models for categorical and limited dependent variables*, volume 7. Sage.
- Sándor Z, Train K (2004). “Quasi-random simulation of discrete choice models.” *Transportation Research Part B: Methodological*, **38**(4), 313–327.
- StataCorp (2011). “Stata Statistical Software: Release 12.” *College Station, TX: StataCorp LP*.
- Train K (2000). “Halton sequences for mixed logit.” *Department of Economics, UCB*.
- Train K (2009). *Discrete choice methods with simulation*. 2 edition. Cambridge university press.

Venables WN, Ripley BD (2002). *Modern Applied Statistics with S*. Fourth edition. Springer, New York. ISBN 0-387-95457-0, URL <http://www.stats.ox.ac.uk/pub/MASS4>.

Zeileis A, Hothorn T (2002). “Diagnostic Checking in Regression Relationships.” *R News*, 2(3), 7–10. URL <http://CRAN.R-project.org/doc/Rnews/>.

Affiliation:

Mauricio Sarrias

Cornell University

E-mail: msarrias86@gmail.com