



THE
POWER
TO KNOW.

SAS/STAT[®] 13.1 User's Guide

The BCHOICE Procedure

This document is an individual chapter from *SAS/STAT*[®] 13.1 *User's Guide*.

The correct bibliographic citation for the complete manual is as follows: SAS Institute Inc. 2013. *SAS/STAT*[®] 13.1 *User's Guide*. Cary, NC: SAS Institute Inc.

Copyright © 2013, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government License Rights; Restricted Rights: The Software and its documentation is commercial computer software developed at private expense and is provided with RESTRICTED RIGHTS to the United States Government. Use, duplication or disclosure of the Software by the United States Government is subject to the license terms of this Agreement pursuant to, as applicable, FAR 12.212, DFAR 227.7202-1(a), DFAR 227.7202-3(a) and DFAR 227.7202-4 and, to the extent required under U.S. federal law, the minimum restricted rights as set out in FAR 52.227-19 (DEC 2007). If FAR 52.227-19 is applicable, this provision serves as notice under clause (c) thereof and no other notice is required to be affixed to the Software or documentation. The Government's rights in Software and documentation shall be only those set forth in this Agreement.

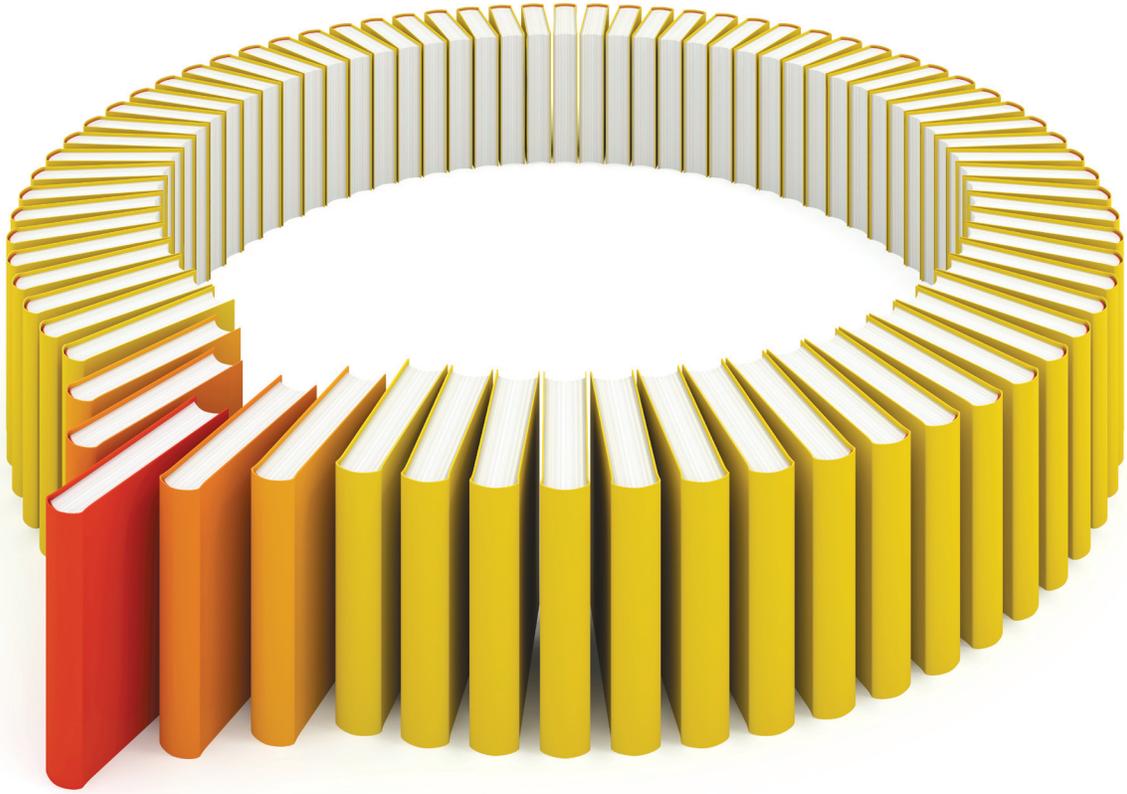
SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513-2414.

December 2013

SAS provides a complete selection of books and electronic products to help customers use SAS[®] software to its fullest potential. For more information about our offerings, visit support.sas.com/bookstore or call 1-800-727-3228.

SAS[®] and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.



Gain Greater Insight into Your SAS[®] Software with SAS Books.

Discover all that you need on your journey to knowledge and empowerment.

 support.sas.com/bookstore
for additional books and resources.


THE POWER TO KNOW.®

Chapter 27

The BCHOICE Procedure (Experimental)

Contents

Overview: BCHOICE Procedure	1002
PROC BCHOICE Compared with Other SAS Procedures	1003
Getting Started: BCHOICE Procedure	1004
A Simple Logit Model	1004
A Logit Model with Random Effects	1011
Syntax: BCHOICE Procedure	1016
PROC BCHOICE Statement	1016
BY Statement	1025
CLASS Statement	1026
MODEL Statement	1028
PREDDIST Statement	1031
RANDOM Statement	1031
Details: BCHOICE Procedure	1034
Discrete Choice Models	1034
Types of Choice Models	1035
Random Effects	1040
Identification and Specification	1043
Gamerman Algorithm	1046
Autocall Macros for Postprocessing	1048
Regenerating Diagnostics Plots	1049
Displayed Output	1050
ODS Table Names	1053
ODS Graphics	1054
Examples: BCHOICE Procedure	1054
Example 27.1: Alternative-Specific and Individual-Specific Effects	1054
Example 27.2: Nested Logit Modeling	1058
Example 27.3: Probit Modeling	1060
Example 27.4: A Random-Effects-Only Logit Model	1063
Example 27.5: Heterogeneity Affected by Individual Characteristics	1065
Example 27.6: Inference on Quantities of Interest	1070
Example 27.7: Predict the Choice Probabilities	1071
References	1072

Overview: BCHOICE Procedure

The BCHOICE (Bayesian choice) procedure performs Bayesian analysis for discrete choice models. Discrete choice models are used in marketing research to model decision makers' choices among alternative products and services. The decision maker might be people, households, companies and so on, and the alternatives might be products, services, actions, or any other options or items about which choices must be made (Train 2009). The collection of alternatives that are available to the decision makers is called a choice set.

Discrete choice models are derived under the assumption of utility-maximizing behavior by decision makers. When individuals are asked to make one choice among a set of alternatives, they usually determine the level of utility that each alternative offers. The utility that individual i obtains from alternative j among J alternatives is denoted as

$$u_{ij} = v_{ij} + \epsilon_{ij}, \quad i = 1, \dots, N \text{ and } j = 1, \dots, J$$

where the subscript i is an index for the individuals, the subscript j is an index for the alternatives in a choice set, v_{ij} is a nonstochastic utility function that relates observed factors to the utility, and ϵ_{ij} is the error component that captures the unobserved characteristics of the utility. In discrete choice models, the observed part of the utility function is assumed to be linear in the parameters,

$$v_{ij} = \mathbf{x}'_{ij} \boldsymbol{\beta}$$

where \mathbf{x}_{ij} is a p -dimensional design vector of observed attribute levels that relate to alternative j and $\boldsymbol{\beta}$ is the corresponding vector of fixed regression coefficients that indicate the utilities or part-worths of the attribute levels.

Decision makers choose the alternative that gives them the greatest utility. Let \mathbf{y}_i be the multinomial response vector for the i th individual. The value y_{ij} takes 1 if the j th component of $\mathbf{u}_i = (u_{i1}, \dots, u_{iJ})$ is the largest, and 0 otherwise:

$$\begin{aligned} u_{ij} &= \mathbf{x}'_{ij} \boldsymbol{\beta} + \epsilon_{ij} \\ y_{ij} &= \begin{cases} 1 & \text{if } u_{ij} \geq \max(\mathbf{u}_i) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Different specifications about the density of the error vector $\boldsymbol{\epsilon}_i = (\epsilon_{i1}, \dots, \epsilon_{iJ})$ result in different types of choice models: logit, nested logit, and probit, as detailed in the section “Types of Choice Models” on page 1035. Logit and nested logit models have closed-form likelihood, whereas a probit model does not.

The past 15 years have seen a dramatic increase in using a Bayesian approach to develop new methods of analysis and models of consumer behavior. The milestone breakthroughs are Albert and Chib (1993) and McCulloch and Rossi (1994) for choice probit models, and Allenby and Lenk (1994) and Allenby (1997) for logit models that have normally distributed random effects (these models are called mixed logit models). Train (2009) extends the Bayesian procedure for mixed logit models to include nonnormal distributions, such as lognormal, uniform, and triangular distributions.

An issue in choice models is that the quantity of relevant data at the individual level is very limited. Respondents frequently become fatigued after answering 15 to 20 questions in a survey. Rossi, McCulloch, and Allenby (1996) and Allenby and Rossi (1999) show how to obtain information about individual-level

parameters within a model by using random taste variation. The lack of data at the individual level, combined with the desire to account for individual differences instead of treating all respondents alike, presents challenges in marketing research. Bayesian methods are ideally suited for analyzing limited data.

Bayesian methods have several advantages. First, Bayesian methods do not require optimization of any function. For probit models and logit models that have random effects, optimization of the likelihood function can be numerically difficult. Different starting values might lead to different maximization results. The problem of a local maximum versus a global maximum is another issue, because convergence is not guaranteed to find the global maximum. Second, Bayesian procedures enable consistency and efficiency to be achieved under more relaxed conditions. When the likelihood function does not have a closed form or there are too many parameters (as in models with random effects), simulation can be used to estimate the likelihood function. Maximization that is based on such a simulated likelihood function is consistent only if the number of draws in simulation rises with the sample size, and it is efficient only if the number of draws in simulation increases faster than the square root of the sample size. On the other hand, Bayesian methods are consistent for a fixed number of draws and are efficient if the number of draws goes up at any rate with the sample size.

For a short introduction to Bayesian analysis and related basic concepts, see Chapter 7, “[Introduction to Bayesian Analysis Procedures](#).” Also see the section “[A Bayesian Reading List](#)” on page 158 for a guide to Bayesian textbooks of varying degrees of difficulty. It follows from Bayes’ theorem that a posterior distribution is proportional to the product of the likelihood function and the prior distribution of the parameter. When it is difficult to obtain the posterior distribution analytically, Bayesian methods often rely on simulations to generate samples from the posterior distribution, and they use the simulated draws to approximate the distribution and to make the inferences.

To use the BCHOICE procedure, you need to specify the type of model for the data. You can also supply a prior distribution for the parameters if you want something other than the default noninformative prior. PROC BCHOICE obtains samples from the corresponding posterior distributions, produces summary and diagnostic statistics, and saves the posterior samples in an output data set that can be used for further analysis. The procedure derives inferences from simulation rather than through analytic or numerical methods. You should expect slightly different answers from each run for the same problem, unless you use the same random number seed.

PROC BCHOICE Compared with Other SAS Procedures

The underlying structure of PROC BCHOICE is similar to the structure of PROC MCMC in that both procedures obtain samples from the posterior distributions and produce summary and diagnostic statistics when you specify the model or the priors or both. However, they differ in that PROC MCMC is a general-purpose Markov chain Monte Carlo (MCMC) simulation procedure that is designed to fit a wide range of Bayesian models, whereas PROC BCHOICE is designed specifically for discrete choice models. You can call PROC MCMC to analyze data that have any likelihood, prior, or hyperprior, as long as these functions can be programmed by using the SAS DATA step functions. For example, you can fit choice logit and nested logit models in PROC MCMC by using some SAS coding to specify the likelihood. PROC BCHOICE works only with choice models, but it is customized to fit special characteristics and features in a choice model. The syntax is quite different from PROC MCMC’s syntax. PROC BCHOICE provides a CLASS statement to handle categorical variables, and it requires less complicated SAS coding for choice models. The default sampling method for choice logit models when direct sampling is not available is the Metropolis-Hastings method, which is based on the Gamerman approach, which in turn has proved to be often more efficient than

the random walk Metropolis algorithm that PROC MCMC uses. In addition, it is difficult to fit choice probit models in PROC MCMC.

For a standard logit choice model, you can use the TIES=BRESLOW option in the PHREG procedure. The approach has the same likelihood as PROC BCHOICE of fitting the data. You use the STRATA statement in PROC PHREG to specify how to define the choice set. However, this is a frequentist approach, and it does not work for a choice model with random effects.

Getting Started: BCHOICE Procedure

The first example in this section is a simple logit conjoint analysis that illustrates some basic features of PROC BCHOICE. The second example discusses a mixed logit model with individual-level random effects, which has become increasingly popular.

A Simple Logit Model

This example uses a standard logit model in which the error component, ϵ_{ij} ($j = 1, \dots, J$), is independently and identically distributed (iid) with the Type I extreme-value distribution, $\exp(-\exp(-\epsilon_{ij}))$. This assumption provides a convenient form for the choice probability (McFadden 1974):

$$P(y_{ij} = 1) = \frac{\exp(\mathbf{x}'_{ij}\boldsymbol{\beta})}{\sum_{k=1}^J \exp(\mathbf{x}'_{ik}\boldsymbol{\beta})}, \quad i = 1, \dots, N \text{ and } j = 1, \dots, J$$

The likelihood is formed by the product of the N independent multinomial distributions:

$$p(\mathbf{Y}|\boldsymbol{\beta}) = \prod_{i=1}^N \prod_{j=1}^J P(y_{ij} = 1)^{y_{ij}}$$

Suppose you want a normal prior on $\boldsymbol{\beta}$

$$\pi(\boldsymbol{\beta}) = N(\mathbf{0}, c\mathbf{I})$$

where \mathbf{I} is the identity matrix and c is a scalar. c is often set to be large for a noninformative prior.

The posterior density of the parameter $\boldsymbol{\beta}$ is

$$p(\boldsymbol{\beta}|\mathbf{Y}) \propto p(\mathbf{Y}|\boldsymbol{\beta})\pi(\boldsymbol{\beta})$$

PROC BCHOICE obtains samples from the posterior distribution, produces summary and diagnostic statistics, and saves the posterior samples in an output data set that can be used for further analysis.

In this example (Kuhfeld 2010), each of 10 subjects is presented with eight different chocolate candies and asked to choose one. The eight candies consist of the 2^3 combinations of dark or milk chocolate, soft or chewy center, and nuts or no nuts. Each subject sees all eight alternatives and makes one choice. Experimental choice data such as these are usually analyzed by using a multinomial logit model.

The following statements read the data:

```

title 'Conjoint Analysis of Chocolate Candies';

data Chocs;
  input Subj Choice Dark Soft Nuts;
  datalines;
1 0 0 0 0
1 0 0 0 1
1 0 0 1 0
1 0 0 1 1
1 1 1 0 0
1 0 1 0 1

... more lines ...

10 0 1 0 0
10 1 1 0 1
10 0 1 1 0
10 0 1 1 1
;

proc print data=Chocs (obs=16);
  by Subj;
  id Subj;
run;

```

The data for the first two subjects are shown in [Figure 27.1](#).

Figure 27.1 Data for the First Two Subjects

Conjoint Analysis of Chocolate Candies				
Subj	Choice	Dark	Soft	Nuts
1	0	0	0	0
	0	0	0	1
	0	0	1	0
	0	0	1	1
	1	1	0	0
	0	1	0	1
	0	1	1	0
	0	1	1	1
2	0	0	0	0
	0	0	0	1
	0	0	1	0
	0	0	1	1
	0	1	0	0
	1	1	0	1
	0	1	1	0
	0	1	1	1

The data set contains 10 subjects and 80 observation lines. Each line of the data represents one alternative in the choice set for each subject. It is required that the response variable, which is Choice in this study, indicate the chosen alternative by the value 1 and the unchosen alternatives by the value 0. Dark is 1 for dark chocolate and 0 for milk chocolate; Soft is 1 for soft center and 0 for chewy center; Nuts is 1 if the candy contains nuts and 0 if it does not contain nuts. In this example, subject 1 chose the fifth alternative (Choice=1, which is Dark/Chewy/No Nuts) among the eight alternatives in the choice set. All the chosen and unchosen alternatives must appear in the data set. If you have choice data in which all alternatives appear on one line, then you must rearrange the data in the correct form (that is, the data must contain one observation line for each alternative of each choice set for each subject).

The following statements fit a multinomial logit model:

```
ods graphics on;
proc bchoice data=Chocs outpost=Bsamp nmc=10000 thin=2 diag=(AutoCorr
    ESS MCSE) seed=124;
    class Dark(ref='0') Soft(ref='0') Nuts(ref='0') Subj;
    model Choice = Dark Soft Nuts / choiceset=(Subj) cprior=normal(var=1000);
run;
```

The ODS GRAPHICS ON statement invokes the ODS Graphics environment and displays the diagnostic plots, such as the trace and autocorrelation function plots of the posterior samples. For more information about ODS, see Chapter 21, “Statistical Graphics Using ODS.”

The PROC BCHOICE statement invokes the procedure, and the DATA= option specifies the input data set Chocs. The OUTPOST= requests an output data set called Bsamp to contain all the posterior samples. The NMC= option specifies the number of posterior simulation iterations. The THIN= option controls the thinning of the Markov chain and specifies that one of every two samples be kept. Thinning is often used to reduce correlation among posterior sample draws. In this example, 5,000 simulated values are saved in the Bsamp data set. The DIAG=(AUTOCORR ESS MCSE) option requests that three convergence diagnostics be output to help determine whether the chain has converged: the autocorrelations, effective sample sizes, and Monte Carlo standard errors. The SEED= option specifies a seed for the random number generator, which guarantees the reproducibility of the random stream.

The CLASS statement names the classification variables to be used in the model. The CLASS statement must precede the MODEL statement (as it must in most other SAS procedures). The REF= option specifies the reference level.

The MODEL statement is required; it defines the dependent variable (y_{ij}) and independent variables (x_{ij}). To the left of the equal sign in the MODEL statement, you specify the dependent variable that indicates which alternatives are chosen and which ones are not chosen. The dependent variable Choice has values 1 (chosen) and 0 (unchosen). You specify the independent variables after the equal sign. The independent variables often indicate attributes or characteristics of the alternatives in the choice set, such as Dark, Soft, and Nuts in this example. The CHOICESET= option specifies how a choice set is defined. In this example, CHOICESET=(Subj) because there is one and only one choice set per subject. The variable that you specify in the CHOICESET= option must be a classification variable that appears in the CLASS statement. You should always use the CHOICESET= option to define the choice set.

The first table that PROC BCHOICE produces is the “Model Information” table, as shown in Figure 27.2. This table displays basic information about the analysis, such as the name of the input data set, response variable, model type, sampling algorithm, burn-in size, simulation size, thinning number, and random number seed. The random number seed initializes the random number generators. If you repeat the analysis and use the same seed, you get an identical stream of random numbers.

Figure 27.2 Model Information

The BCHOICE Procedure	
Model Information	
Data Set	WORK.CHOCS
Response Variable	Choice
Type of Model	Logit
Fixed Effects Included	Yes
Random Effects Included	No
Sampling Algorithm	Gamerman Metropolis
Burn-In Size	500
Simulation Size	10000
Thinning	2
Random Number Seed	124
Number of Threads	1

The “Choice Sets Summary” table is displayed by default and should be used to check the data entry. In this case, there are 10 choice sets, one for each subject. All 10 choice sets display the same pattern: they have eight total alternatives, one of which is chosen and seven of which are unchosen. PROC BCHOICE requires that all choice sets have a pattern of the same number of total alternatives, one of which is the chosen alternative. If there is more than one pattern among the choice sets (for example, some choice sets have nine alternatives or some subjects chose two alternatives), a warning message appears in the SAS log and those invalid choice sets are excluded from analysis.

Figure 27.3 Choice Sets Summary

Choice Sets Summary				
Pattern	Choice Sets	Total Alternatives	Chosen Alternatives	Not Chosen
1	10	8	1	7

The next table is the “Number of Observations” table, as shown in [Figure 27.4](#). This table lists the number of observations that are read from the `DATA=` data set and the number of nonmissing and valid observations that are used in the analysis. PROC BCHOICE does not impute missing values. If any missing value is encountered in a row of the `DATA=` data set, PROC BCHOICE skips that row and moves on to read the next row. If the missing value causes the corresponding choice set to have a pattern different from that of the rest of the choice sets, the entire choice set is discarded from analysis. In this example, all 80 observations are used.

Figure 27.4 Number of Observations

Number of Observations	
Number of Observations Read	80
Number of Observations Used	80

PROC BCHOICE reports posterior summary statistics (posterior means, standard deviations, and HPD intervals) for each parameter, as shown in Figure 27.5. For more information about posterior statistics, see the section “Summary Statistics” on page 155.

Figure 27.5 PROC BCHOICE Posterior Summary Statistics

Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Dark 1	5000	1.5308	0.7943	0.1848	3.2412
Soft 1	5000	-2.4312	0.9792	-4.4882	-0.8125
Nuts 1	5000	0.9671	0.7454	-0.3255	2.6675

Before you examine the posterior summary statistics and try to draw any conclusions, you might want to verify that the simulation has converged. Convergence diagnostics are essential to inferring from simulations that are based on Markov chains. If the Markov chain has not converged, all conclusions that are based on the samples might be misleading. PROC BCHOICE computes the effective sample size by default. There are a number of other convergence diagnostics to help you determine whether the chain has converged: the Monte Carlo standard errors, the autocorrelations at selected lags, and so on. These statistics are shown in Figure 27.6. For details and interpretations of these diagnostics, see the section “Assessing Markov Chain Convergence” on page 141.

The “Posterior Autocorrelations” table shows that the autocorrelations among posterior samples reduce quickly. The “Effective Sample Sizes” table reports the number of effective sample sizes of the Markov chain. The “Monte Carlo Standard Errors” table indicates that the standard errors of the mean estimates for each of the variables are relatively small with respect to the posterior standard deviations. The values in the MCSE/SD column (ratios of the standard errors and the standard deviations) are small. This means that only a fraction of the posterior variability is caused by the simulation.

Figure 27.6 PROC BCHOICE Convergence Diagnostics

Posterior Autocorrelations				
Parameter	Lag 1	Lag 5	Lag 10	Lag 50
Dark 1	0.4495	0.1427	0.0603	-0.0095
Soft 1	0.5389	0.1518	0.0730	0.0449
Nuts 1	0.4136	0.1473	0.0259	-0.0264

Effective Sample Sizes			
Parameter	ESS	Autocorrelation	
		Time	Efficiency
Dark 1	1078.4	4.6365	0.2157
Soft 1	904.9	5.5253	0.1810
Nuts 1	1161.6	4.3043	0.2323

Figure 27.6 *continued*

Monte Carlo Standard Errors			
Parameter	MCSE	Standard Deviation	MCSE/SD
Dark 1	0.0242	0.7943	0.0305
Soft 1	0.0326	0.9792	0.0332
Nuts 1	0.0219	0.7454	0.0293

PROC BCHOICE produces a number of graphs, shown in Figure 27.7, which also aid convergence diagnostic checks. The trace plots have two important aspects to examine. First, you want to check whether the mean of the Markov chain has stabilized and appears constant over the graph. Second, you want to check whether the chain has good mixing and is “dense,” in the sense that it quickly traverses the support of the distribution to explore both the tails and the mode areas efficiently. The plots show that the chains appear to have reached their stationary distributions.

Next, you want to examine the autocorrelation plots, which indicate the degree of autocorrelation for each of the posterior samples. High correlations usually imply slow mixing. Finally, the kernel density plots estimate the posterior marginal distributions for each parameter.

Figure 27.7 PROC BCHOICE Diagnostic Plots

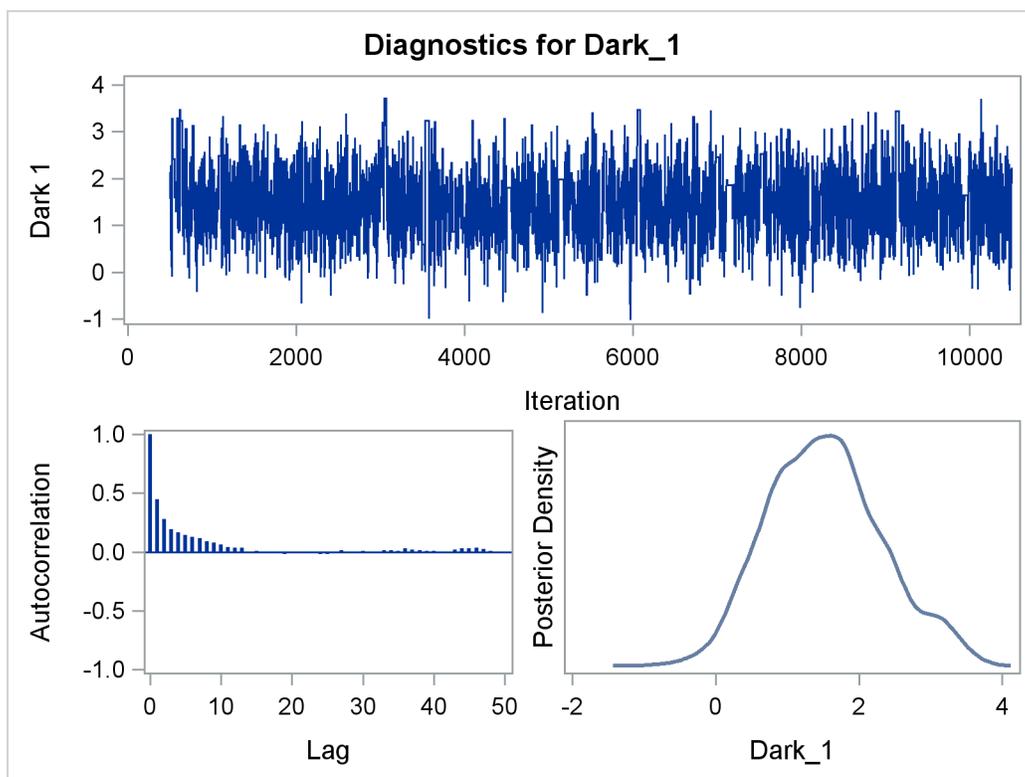
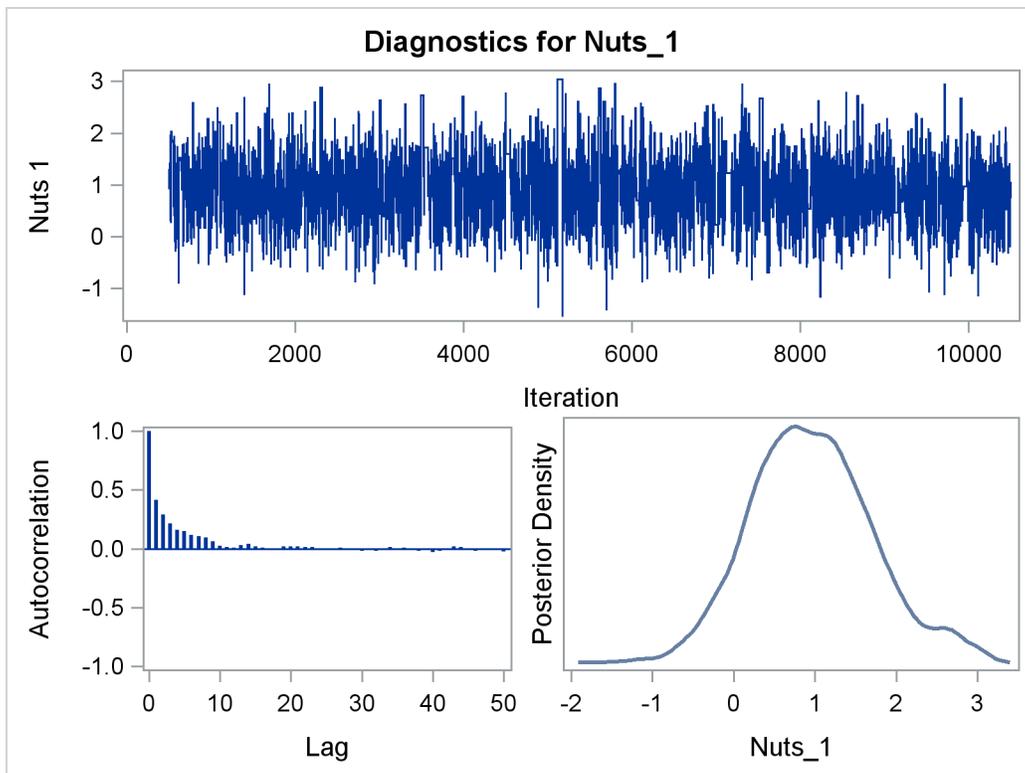
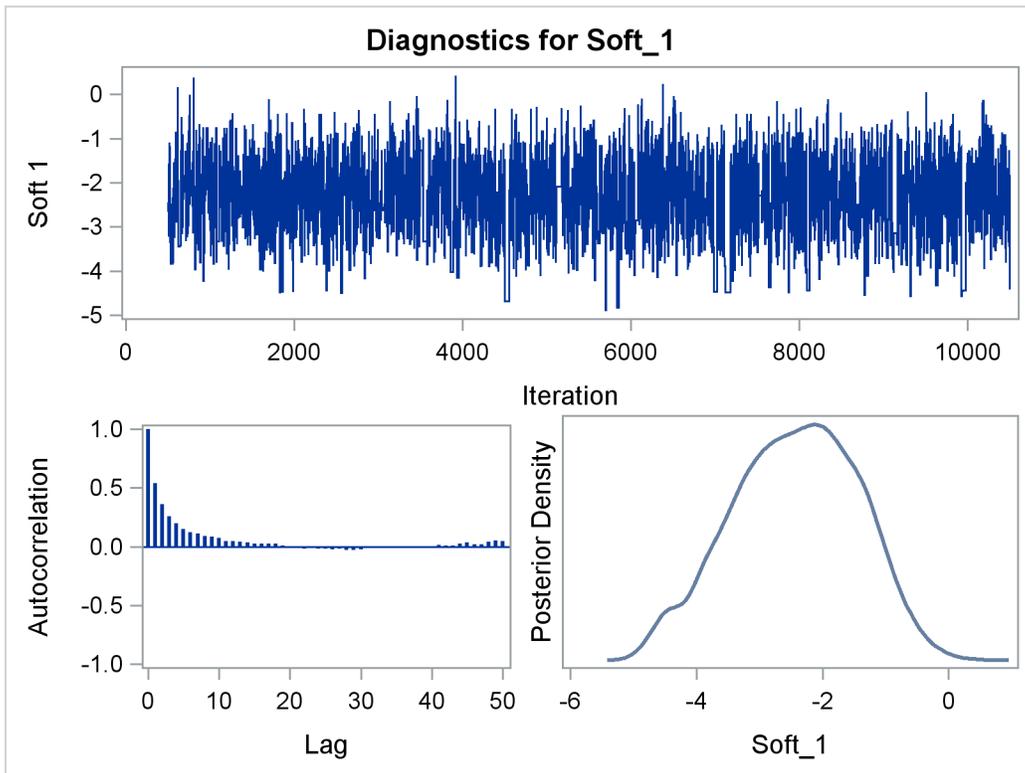


Figure 27.7 continued



Because the chains have converged, you can go back to the posterior summary table for results and conclusions. As seen in [Figure 27.5](#), the part-worth for dark chocolate is 1.5 and the part-worth for milk chocolate (base category) is structural 0; the part-worth for soft center is -2.4 and the part-worth for chewy center is structural 0; the part-worth for containing nuts is 1.0 and the part-worth for no nuts is structural 0. A positive part-worth implies being more favorable. Hence, dark chocolate is preferred over milk chocolate, soft centers are less popular than chewy centers, and candies with nuts are more popular than candies without nuts.

A Logit Model with Random Effects

Choice models that have random effects (or random coefficients) provide solutions to create individual-level or group-specific utilities. Because people have different preferences, it can be misleading to roll the whole sample together into a single set of utilities. The desire to account for individual differences, instead of treating all respondents alike, provides challenges in marketing research. For logit models that have random effects, using frequentist methods to optimize of the likelihood function can be numerically difficult. Bayesian methods are ideally suited for analysis with random effects.

Choice models that have random effects generalize the standard choice models to incorporate individual-level effects. Let the utility that individual i obtains from alternative j in choice situation t ($t = 1, \dots, T$) be

$$\begin{aligned} u_{ijt} &= \mathbf{x}'_{ijt}\boldsymbol{\beta} + \mathbf{z}'_{ijt}\boldsymbol{\gamma}_i + \epsilon_{ijt} \\ y_{ijt} &= 1 \text{ if } u_{ijt} \geq \max(u_{i1t}, u_{i2t}, \dots, u_{iJt}) \\ &= 0 \text{ otherwise} \end{aligned}$$

where y_{ijt} is the observed choice for individual i and alternative j in choice situation t ; \mathbf{x}_{ijt} is the fixed design vector for individual i and alternative j in choice situation t ; $\boldsymbol{\beta}$ are the fixed coefficients; \mathbf{z}_{ijt} is the random design vector for individual i and alternative j in choice situation t ; and $\boldsymbol{\gamma}_i$ are the random coefficients for individual i corresponding to \mathbf{z}_{ijt} .

It is assumed that each $\boldsymbol{\gamma}_i$ is drawn from a superpopulation and that this superpopulation is normal, $\boldsymbol{\gamma}_i \sim \text{iid } N(\mathbf{0}, \boldsymbol{\Omega}_\boldsymbol{\gamma})$. An additional stage is added to the model in which a prior for $\boldsymbol{\Omega}_\boldsymbol{\gamma}$ is specified:

$$\begin{aligned} \pi(\boldsymbol{\gamma}_i) &= N(\mathbf{0}, \boldsymbol{\Omega}_\boldsymbol{\gamma}) \\ \pi(\boldsymbol{\Omega}_\boldsymbol{\gamma}) &= \text{inverse Wishart}(\nu_0, \mathbf{V}_0) \end{aligned}$$

The covariance matrix $\boldsymbol{\Omega}_\boldsymbol{\gamma}$ characterizes the extent of unobserved heterogeneity among individuals. Large diagonal elements of $\boldsymbol{\Omega}_\boldsymbol{\gamma}$ indicate substantial heterogeneity in part-worths. Off-diagonal elements indicate patterns in the evaluation of attribute levels.

Consider a study that estimates the market demand for kitchen trash cans (Rossi 2013). There are four attributes, and each has two levels: touchless opening (Yes/No), material (Steel/Plastic), automatic trash bag replacement (Yes/No), and price (80/40). The number of all possible hypothetical types of trash cans is $2^4 = 16$. Including more attributes and more levels can easily become unmanageable. The study uses a fractional factorial design, in which the first three factors are set up to be a full factorial design and the fourth is generated as the product of the first three. This design confounds the three-way interaction with the effect of the fourth factor, shown in [Table 27.1](#).

Table 27.1 Design for the Trash Can Study

Obs	Touchless	Steel	AutoBag	Price80
1	-1	-1	-1	-1
2	-1	-1	1	1
3	-1	1	-1	1
4	-1	1	1	-1
5	1	-1	-1	1
6	1	-1	1	-1
7	1	1	-1	-1
8	1	1	1	1

In [Table 27.1](#), 1 means “Yes” and -1 means “No.” This is a balanced design, in which each level appears the same number of times. This study assigns only two alternatives to a choice set by randomly sampling two rows from the previous table and giving each individual 10 choice sets (or choice tasks) to pick from. For more information about how to design a choice model efficiently, see [Kuhfeld \(2010\)](#).

Data were obtained by enrolling 104 people and assigning 10 choice tasks to each of them: for each task, the participants stated their preference between two types of trash cans. The following steps read in the data:

```

data Trashcan;
  input ID Task Choice Index Touchless Steel AutoBag Price80 @@;
  datalines;
1 1 1 1 0 1 1 0 1 1 0 2 1 1 0 0 1 2 0 1 0 0 0 0 1 2 1 2 1 1 1 1 1 3 0 1
0 0 0 0 1 3 1 2 1 1 0 0 1 4 0 1 0 1 0 1 1 4 1 2 1 0 0 1 1 5 1 1 0 1 1 0
1 5 0 2 1 0 0 1 1 6 0 1 0 0 1 1 1 6 1 2 1 1 1 1 1 7 0 1 1 0 0 1 1 7 1 2
1 1 0 0 1 8 0 1 0 0 1 1 1 8 1 2 0 1 1 0 1 9 1 1 0 0 1 1 1 9 0 2 1 0 0 1
1 10 0 1 0 1 1 0 1 10 1 2 1 0 1 0 2 1 1 1 0 1 1 0 2 1 0 2 1 1 0 0 2 2 0
1 0 0 0 0 2 2 1 2 1 1 1 1 2 3 0 1 0 0 0 0 2 3 1 2 1 1 0 0 2 4 0 1 0 1 0
1 2 4 1 2 1 0 0 1 2 5 1 1 0 1 1 0 2 5 0 2 1 0 0 1 2 6 1 1 0 0 1 1 2 6 0
2 1 1 1 1 2 7 1 1 1 0 0 1 2 7 0 2 1 1 0 0 2 8 1 1 0 0 1 1 2 8 0 2 0 1 1
0 2 9 1 1 0 0 1 1 2 9 0 2 1 0 0 1 2 10 1 1 0 1 1 0 2 10 0 2 1 0 1 0 3 1

... more lines ...

2 1 2 1 1 1 1 104 3 0 1 0 0 0 0 104 3 1 2 1 1 0 0 104 4 0 1 0 1 0 1 104
4 1 2 1 0 0 1 104 5 1 1 0 1 1 0 104 5 0 2 1 0 0 1 104 6 0 1 0 0 1 1 104
6 1 2 1 1 1 1 104 7 0 1 1 0 0 1 104 7 1 2 1 1 0 0 104 8 0 1 0 0 1 1 104
8 1 2 0 1 1 0 104 9 0 1 0 0 1 1 104 9 1 2 1 0 0 1 104 10 0 1 0 1 1 0 104
10 1 2 1 0 1 0
;

proc print data=Trashcan (obs=20);
  by ID Task;
  id ID Task;
run;

```

The data for the first four choice tasks are shown in [Figure 27.8](#).

Figure 27.8 Data for the First Four Choice Tasks

ID	Task	Choice	Index	Touchless	Steel	Auto Bag	Price80
1	1	1	1	0	1	1	0
		0	2	1	1	0	0
1	2	0	1	0	0	0	0
		1	2	1	1	1	1
1	3	0	1	0	0	0	0
		1	2	1	1	0	0
1	4	0	1	0	1	0	1
		1	2	1	0	0	1

In the data, ID is the individual's ID number, and Task indexes the number of choice tasks. The response is Choice, which states each individual's choice for each choice task. Touchless, Steel, AutoBag, and Price80 are the attribute variables; for each of them, 1 means "Yes" and 0 means "No." In the data, 0 replaces the -1 values that are shown in the design matrix in Table 27.1.

The following statements fit a logit model with random effects:

```
proc bchoice data=Trashcan seed=123 nmc=20000 thin=4;
  class ID Task;
  model Choice = Touchless Steel AutoBag Price80 / choiceset=(ID Task);
  random Touchless Steel AutoBag Price80 / sub=ID monitor=(1 to 5) type=un;
run;
```

The choice set is specified by ID (which identifies the participants) and by Task (which identifies each of the 10 choice tasks that are assigned to each participant). The variables ID and Task are needed in the CLASS statement because they define the choice set in the MODEL statement.

In addition to the MODEL statement for fixed effects, the RANDOM statement is added for random effects. Note that Touchless, Steel, AutoBag, and Price80 are listed as both fixed and random effects, so that their average part-worth values in the population are estimated via fixed effects and the deviation from the overall mean for each individual is presented through random effects. The SUB=ID argument in the RANDOM statement defines ID as a subject index for the random effects grouping, so that each person with a different ID has his or her own random effects. The MONITOR option requests the production of the individual-level random-effects parameter estimates, and the MONITOR=(1 to 5) option requests the random-effects parameter estimates for the first five subjects, (By default, PROC BCHOICE does not output results for any individual-level random-effects parameters.) The TYPE=UN option in the RANDOM statement specifies an unstructured covariance matrix for the random effects. The unstructured type provides a mechanism for estimating the correlation between the random effects. The TYPE=VC (variance components) option, which is the default structure, models a different variance component for each random effect.

Summary statistics for the fixed coefficients (β), the covariance of the random coefficients (Ω_{γ}), and the random coefficients (γ_i) for the first five individuals are shown in Figure 27.9.

Figure 27.9 Posterior Summary Statistics

The BCHOICE Procedure						
Posterior Summaries and Intervals						
Parameter	Subject	N	Standard		95% HPD Interval	
			Mean	Deviation		
Touchless		5000	1.7432	0.2758	1.2225	2.2974
Steel		5000	1.0366	0.2749	0.5043	1.5951
AutoBag		5000	2.1776	0.3556	1.4799	2.8655
Price80		5000	-4.6574	0.6732	-5.9315	-3.3896
RECOV Touchless, Touchless		5000	3.2690	1.0960	1.4597	5.4938
RECOV Steel, Touchless		5000	-0.5786	0.9401	-2.6960	1.1540
RECOV Steel, Steel		5000	2.7386	0.9952	1.1160	4.7779
RECOV AutoBag, Touchless		5000	-0.6080	0.8920	-2.3911	1.2471
RECOV AutoBag, Steel		5000	-0.0456	0.7354	-1.5824	1.3500
RECOV AutoBag, AutoBag		5000	3.6331	1.5755	0.8716	6.5751
RECOV Price80, Touchless		5000	-1.4322	1.2425	-4.0486	0.6300
RECOV Price80, Steel		5000	-1.5099	1.2504	-4.1057	0.7940
RECOV Price80, AutoBag		5000	-2.3795	1.7869	-5.8962	0.5972
RECOV Price80, Price80		5000	7.9205	3.5695	2.3922	14.7155
Touchless	ID 1	5000	0.5492	1.0522	-1.5118	2.6518
Steel	ID 1	5000	-0.3691	1.1114	-2.7311	1.7119
AutoBag	ID 1	5000	1.6954	1.4487	-0.8981	4.5846
Price80	ID 1	5000	-0.7483	2.3010	-5.2892	3.6146
Touchless	ID 2	5000	-1.8146	0.9255	-3.6760	0.0196
Steel	ID 2	5000	-1.4133	0.8567	-3.0222	0.3404
AutoBag	ID 2	5000	0.4091	1.3343	-2.1911	2.9911
Price80	ID 2	5000	4.6933	1.5455	1.8588	7.8627
Touchless	ID 3	5000	-0.7099	1.0230	-2.7242	1.3434
Steel	ID 3	5000	0.6959	1.0397	-1.2993	2.7718
AutoBag	ID 3	5000	-1.9793	1.2506	-4.3813	0.5263
Price80	ID 3	5000	3.8323	1.7134	0.3076	7.0044
Touchless	ID 4	5000	-0.6107	0.9869	-2.5368	1.3259
Steel	ID 4	5000	0.3235	1.0106	-1.5620	2.3570
AutoBag	ID 4	5000	1.0191	1.4644	-1.5675	4.0437
Price80	ID 4	5000	-1.8942	2.3716	-6.8807	2.3238
Touchless	ID 5	5000	-1.3562	1.2145	-3.7166	0.9827
Steel	ID 5	5000	1.7447	1.2911	-0.7051	4.1791
AutoBag	ID 5	5000	1.1826	1.5781	-1.8454	4.3037
Price80	ID 5	5000	-0.3420	2.4111	-4.9867	4.3344

The fixed effects (Touchless, Steel, AutoBag, and Price80) are shown in the first four rows. Across all the respondents in the data, the average part-worths for touchless opening, steel material, and automatic trash bag replacement are all positive, indicating that most people favor those features; the average part-worth for having to pay USD80 for a trash can instead of USD40 is negative (−4.7), which is very intuitive, because spending more money is usually unfavorable.

The covariance estimate of the random coefficients ($\hat{\Omega}_y$) is displayed by the parameters whose label begins with “RECOV”:

$$\hat{\Omega}_y = \begin{pmatrix} 3.3 & \cdot & \cdot & \cdot \\ -0.6 & 2.7 & \cdot & \cdot \\ -0.6 & -0.1 & 3.6 & \cdot \\ -1.4 & -1.5 & -2.4 & 7.9 \end{pmatrix}$$

where the dots refer to the corresponding elements in the lower part of the symmetric covariance matrix. The covariance estimate of the random coefficients ($\hat{\Omega}_y$) characterizes the variability of part-worths across respondents. Some of the diagonal elements of the matrix are large. For example, the variance for price (labeled “RECOV Price80, Price80”) is quite large, indicating substantial unexplained difference in response to price. Off-diagonal elements of the matrix illustrate attribute levels that tend to be evaluated similarly (positive covariance) or differently (negative covariance) across all the respondents. The covariances between each of the attributes (Touchless, Steel, and AutoBag) and Price80 are all negative, implying that the respondents who prefer some of the new features are those who are also unwilling to pay a higher price for the trash can. Therefore, offering a discounted price might be a particularly effective method of introducing the new features to customers.

The next set of parameters that are displayed are the estimates for the individual-level random effects for the first five respondents (see [Figure 27.9](#)). These estimates are the deviation from the overall means (which are estimated via the fixed effects). The part-worth for touchless opening for the first respondent (who is labeled “ID 1” in the Subject column) is $1.7 + 0.5 = 2.2$.

Allenby and Rossi (1999) and Rossi, Allenby, and McCulloch (2005) propose a hierarchical Bayesian random-effects model that is set up in a different way such that there are no fixed effects but only random effects. For more information about this type of model, see the section “[Random Effects](#)” on page 1040 and a follow-up example in “[Example 27.4: A Random-Effects-Only Logit Model](#)” on page 1063.

Syntax: BCHOICE Procedure

The following statements are available in the BCHOICE procedure. Items within `< >` are optional.

```
PROC BCHOICE < options > ;
  BY variables ;
  CLASS variable < (options) > < ... variable < (options) > > < / options > ;
  MODEL response < (response-options) > = < fixed-effects > < / model-options > ;
  RANDOM random-effects < / options > ;
  PREDDIST < 'label' > OUTPRED=SAS-data-set < options > ;
```

The PROC BCHOICE and MODEL statements are required. The CLASS statement, if present, must precede the MODEL statement.

The rest of this section provides detailed syntax information for each statement, beginning with the PROC BCHOICE statement. The remaining statements are presented in alphabetical order.

PROC BCHOICE Statement

```
PROC BCHOICE < options > ;
```

The PROC BCHOICE statement invokes the BCHOICE procedure.

Table 27.2 summarizes the *options* available in the PROC BCHOICE statement.

Table 27.2 PROC BCHOICE Statement Options

Option	Description
Basic Options	
DATA=	Names the input data set
NBI=	Specifies the number of burn-in iterations
NMC=	Specifies the number of iterations, excluding the burn-in iterations
NTHREADS=	Specifies the number of threads for the computation
NTU=	Specifies the number of tuning iterations for the random walk sampler
OUTPOST=	Names the output data set to contain posterior samples of parameters
SEED=	Specifies the random seed for simulation
THIN=	Specifies the thinning rate
Sampling, Summary, Diagnostics, and Plotting options	
ALGORITHM=	Specifies the algorithm to use to sample the posterior distribution
DIAGNOSTICS=	Controls the convergence diagnostics
DIC	Computes the deviance information criterion (DIC)
HITPROB	Outputs the average of estimated probabilities of chosen alternatives in the input data
PLOTS=	Controls plotting
STATISTICS=	Controls posterior statistics
Other Options	
ACCEPTTOL=	Specifies the acceptance rate tolerance for the random walk sampler

Table 27.2 (continued)

Option	Description
INF=	Specifies the machine numerical limit for infinity
LOGPOST	Calculates the logarithm of the posterior density and likelihood
MAXTUNE=	Specifies the maximum number of tuning loops for the random walk sampler
MINTUNE=	Specifies the minimum number of tuning loops for the random walk sampler
NOCLPRINT	Suppresses the “Class Level Information” table completely or partially
TARGACCEPT=	Specifies the target acceptance rate for the random walk sampler

You can specify the following *options*.

ACCEPTTOL=*n*

specifies a tolerance for acceptance probabilities for the random walk sampler. By default, ACCEPTTOL=0.075. You can specify this option for logit models.

ALGORITHM=GAMERMAN | RWM | LATENT

ALG=GAMERMAN | RWM | LATENT

specifies the algorithm to use to sample the posterior distribution for the regression coefficients. You can specify the following algorithms:

GAMERMAN

uses the Metropolis-Hastings approach of Gamerman (1997). This is the default for logit models.

RWM

uses the random walk Metropolis algorithm along with the normal proposal, as suggested in Rossi, Allenby, and McCulloch (2005).

LATENT

uses the latent variables via the data augmentation method. This is the default for probit models. This option is ignored for logit models.

When possible, PROC BCHOICE samples directly from the full conditional distribution. Otherwise, the default sampling algorithm is the Gamerman algorithm for logit models and the latent variables via the data augmentation method for probit models.

DATA=SAS-*data-set*

specifies the input data set.

DIAGNOSTICS=NONE | (*keyword-list*)

DIAG=NONE | (*keyword-list*)

specifies options for convergence diagnostics. By default, PROC BCHOICE computes the effective sample sizes. The sample autocorrelations, Monte Carlo errors, Geweke test, Raftery-Lewis test, and Heidelberger-Welch test are also available. For more information about convergence diagnostics, see the section “[Assessing Markov Chain Convergence](#)” on page 141. You can request all the diagnostic tests by specifying DIAGNOSTICS=ALL. You can suppress all the diagnostic tests by specifying DIAGNOSTICS=NONE.

You can specify one or more of the following *keyword-list* options:

ALL

computes all diagnostic tests and statistics. You can combine the option ALL with any other specific tests to modify test options. For example, `DIAGNOSTICS=(ALL AUTOCORR(LAGS=(1 5 35)))` computes all tests by using default settings and autocorrelations at lags 1, 5, and 35.

AUTOCORR < (*autocorrelation-options*) >**AC** < (*autocorrelation-options*) >

computes default autocorrelations at lags 1, 5, 10, and 50 for each variable. You can choose other lags by using the following *autocorrelation-option*:

LAGS=(*numeric-list*)

specifies autocorrelation lags. The *numeric-list* must take positive integer values.

ESS

computes the effective sample sizes (Kass et al. 1998) of the posterior samples of each parameter. It also computes the correlation time and the efficiency of the chain for each parameter. Small values of ESS might indicate a lack of convergence. For more information, see the section “Effective Sample Size” on page 154.

GEWEKE < (*Geweke-options*) >

computes the Geweke spectral density diagnostics; this is a two-sample *t*-test between the first f_1 portion (as specified by the `FRAC1=` option) and the last f_2 portion (as specified by the `FRAC2=` option) of the chain. For more information, see the section “Geweke Diagnostics” on page 148. By default, `FRAC1=0.1` and `FRAC2=0.5`, but you can choose other fractions by using the following *Geweke-options*:

FRAC1=*value***F1**=*value*

specifies the beginning proportion of the Markov chain. By default, `FRAC1=0.1`.

FRAC2=*value***F2**=*value*

specifies the end proportion of the Markov chain. By default, `FRAC2=0.5`.

HEIDELBERGER < (*Heidel-options*) >**HEIDEL** < (*Heidel-options*) >

computes the Heidelberg-Welch diagnostic (which consists of a stationarity test and a half-width test) for each variable. The stationary diagnostic test tests the null hypothesis that the posterior samples are generated from a stationary process. If the stationarity test is passed, a half-width test is then carried out. For more information, see the section “Heidelberg and Welch Diagnostics” on page 150.

You can also specify suboptions, such as `DIAGNOSTICS=HEIDELBERGER(EPS=0.05)`, as follows:

EPS=*value*

specifies a small positive number ϵ such that if the half-width is less than ϵ times the sample mean of the retaining iterations, the half-width test is passed. By default, `EPS=0.1`.

HALPHA=*value*

specifies the α level ($0 < \alpha < 1$) for the half-width test. By default, HALPHA=0.05.

SALPHA=*value*

specifies the α level ($0 < \alpha < 1$) for the stationarity test. By default, SALPHA=0.05.

MAXLAG=*n*

specifies the maximum number of autocorrelation lags to use to compute the effective sample size; for more information, see the section “Effective Sample Size” on page 154. The value of n is also used in the calculation of the Monte Carlo standard error; see the section “Standard Error of the Mean Estimate” on page 155. By default, MAXLAG=MIN(500, MCsample/4), where MCsample is the Markov chain sample size that is kept after thinning—that is, MCsample = $\left\lceil \frac{NMC}{NTHIN} \right\rceil$. If n is too low, you might observe significant lags and the effective sample size cannot be calculated accurately. A warning message appears in the SAS log, and you can increase either the MAXLAG= option or the NMC= option, accordingly. Specifying this option implies the ESS and MCSE options.

MCSE**MCERROR**

computes the Monte Carlo standard error for the posterior samples of each parameter.

NONE

suppresses all the diagnostic tests and statistics. This option is not recommended.

RAFTERY <(Raftery-options)>**RL** <(Raftery-options)>

computes the Raftery-Lewis diagnostic, which evaluates the accuracy of the estimated quantile ($\hat{\theta}_Q$ for a given $Q \in (0, 1)$) of a chain. $\hat{\theta}_Q$ can achieve any degree of accuracy when the chain is allowed to run for a long time. The algorithm stops when the estimated probability $\hat{P}_Q = \Pr(\theta \leq \hat{\theta}_Q)$ reaches within $\pm R$ of the value Q with probability S ; that is, $\Pr(Q - R \leq \hat{P}_Q \leq Q + R) = S$. For more information, see the section “Raftery and Lewis Diagnostics” on page 151.

You can specify Q , R , S , and a precision level ϵ for a stationary test by specifying the following *Raftery-options*—for example, DIAGNOSTICS=RAFTERY(QUANTILE=0.05):

ACCURACY=*value***R=***value*

specifies a small positive number as the margin of error for measuring the accuracy of estimation of the quantile. By default, ACCURACY=0.005.

EPS=*value*

specifies the tolerance level (a small positive number) for the stationary test. By default, EPS=0.001.

PROB=*value***S=***value*

specifies the probability of attaining the accuracy of the estimation of the quantile. By default, PROB=0.95.

QUANTILE=*value*

Q=*value*

specifies the order (a value between 0 and 1) of the quantile of interest. By default, QUANTILE=0.025.

DIC

computes the deviance information criterion (DIC). DIC is calculated by using the posterior mean estimates of the parameters. For more information, see the section “[Deviance Information Criterion \(DIC\)](#)” on page 157.

HITPROB

HITPROBABILITY

calculates the average of estimated probabilities of all chosen alternatives in the input data set. You can use this average as a measure of goodness of fit.

INF=*value*

specifies the numerical definition of infinity in PROC BCHOICE. By default, INF=1E15. For example, PROC BCHOICE considers 1E16 to be outside the support of the normal distribution and assigns a missing value to the log density evaluation. The minimum *value* that is allowed is 1E10, and the maximum *value* that is allowed is 1E27.

LOGPOST

calculates the logarithm of the posterior density of the parameters and the likelihood at each iteration. As a result, the OUTPOST= data set will contain the LOGLIKE and LOGPOST variables. You can specify this option for logit models and nested logit models, but not probit models.

MAXTUNE=*n*

specifies an upper limit for the number of proposal tuning loops for the random walk sampler. By default, MAXTUNE=6. You can specify this option for logit models.

MINTUNE=*n*

specifies a lower limit for the number of proposal tuning loops for the random walk sampler. By default, MINTUNE=2. You can specify this option for logit models.

NBI=*n*

specifies the number of burn-in iterations to perform before beginning to save parameter estimate chains. By default, NBI=500. For more information, see the section “[Burn-in, Thinning, and Markov Chain Samples](#)” on page 140.

NMC=*n*

specifies the number of iterations in the main simulation loop. This is the MCMC sample size if THIN=1. By default, NMC=5000.

NOCLPRINT<=*number*>

suppresses the display of the “Class Level Information” table if you do not specify *number*. If you specify *number*, the values of the classification variables are displayed for only those variables whose number of levels is less than *number*. Specifying a *number* helps reduce the size of the “Class Level Information” table if some classification variables have a large number of levels.

NTHREADS=*n***NTHREAD=*n***

specifies the number of threads for analytic computations and overrides the SAS system option `THREADS` | `NOTHREADS`. If you do not specify the `NTHREADS=` option or if you specify `NTHREADS=0`, the default number of threads is 1.

NTU=*n*

specifies the number of iterations to use in each proposal tuning phase for the random walk sampler. By default, `NTU=500`. You can specify this option for logit models.

OUTPOST=*SAS-data-set***OUT=*SAS-data-set***

specifies an output data set to contain the posterior samples of all parameters and the iteration numbers. It contains the log of the posterior density (`LOGPOST`) and the log likelihood (`LOGLIKE`) if you specify the `LOGPOST` option. By default, no `OUTPOST=` data set is created.

PLOTS < (*global-plot-options*) > <= *plot-request* < (*options*) > >**PLOTS** < (*global-plot-options*) > <= (*plot-request* < (*options*) > <... *plot-request* < (*options*) > > >

controls the display of diagnostic plots. You can request three types of plots: trace plots, autocorrelation function plots, and kernel density plots. By default, the plots are displayed in panels unless you specify the *global-plot-option* `UNPACK`. Also, when you specify more than one type of plot, the plots are grouped by parameter unless you specify the *global-plot-option* `GROUPBY=TYPE`. When you specify only one *plot-request*, you can omit the parentheses around it, as shown in the following example:

```
plots=none
plots(unpack)=trace
plots=(trace density)
```

If ODS Graphics is enabled but you do not specify the `PLOTS=` option, then PROC BCHOICE produces, for each parameter, a panel that contains the trace plot, the autocorrelation function plot, and the density plot. This is equivalent to specifying `PLOTS=(TRACE AUTOCORR DENSITY)`.

You can specify the following *global-plot-options*:

FRINGE

adds a fringe plot to the horizontal axis of the density plot.

GROUPBY=PARAMETER | TYPE**GROUP=PARAMETER | TYPE**

specifies how the plots are grouped when there is more than one type of plot. By default, `GROUPBY=PARAMETER`. You can specify the following values:

TYPE

groups the plots by type.

PARAMETER

groups the plots by parameter.

LAGS=*n*

specifies the number of autocorrelation lags that are used in plotting the ACF graph. By default, LAGS=50.

SMOOTH

smooths the trace plot by using a fitted penalized B-spline curve (Eilers and Marx 1996).

UNPACKPANEL**UNPACK**

unpacks all paneled plots, so that each plot in a panel is displayed separately.

You can specify the following *plot-requests*:

ALL

requests all types of plots. PLOTS=ALL is equivalent to specifying PLOTS=(TRACE AUTO-CORR DENSITY).

AUTOCORR**ACF**

displays the autocorrelation function plots for the parameters.

DENSITY**D****KERNEL****K**

displays the kernel density plots for the parameters.

NONE

suppresses the display of all plots.

TRACE**T**

displays the trace plots for the parameters.

Consider a model that has four parameters, X1–X4. The following list shows which plots are produced for various option settings:

- PLOTS=(TRACE AUTOCORR) displays the trace and autocorrelation plots for each parameter side by side, with two parameters per panel:

Display 1	Trace(X1)	Autocorr(X1)
	Trace(X2)	Autocorr(X2)
Display 2	Trace(X3)	Autocorr(X3)
	Trace(X4)	Autocorr(X4)

- PLOTS(GROUPBY=TYPE)=(TRACE AUTOCORR) displays all the paneled trace plots, followed by panels of autocorrelation plots:

Display 1	Trace(X1)	
	Trace(X2)	
Display 2	Trace(X3)	
	Trace(X4)	
Display 3	Autocorr(X1)	Autocorr(X2)
	Autocorr(X3)	Autocorr(X4)

- PLOTS(UNPACK)=(TRACE AUTOCORR) displays a separate trace plot and a separate correlation plot, parameter by parameter:

Display 1	Trace(X1)
Display 2	Autocorr(X1)
Display 3	Trace(X2)
Display 4	Autocorr(X2)
Display 5	Trace(X3)
Display 6	Autocorr(X3)
Display 7	Trace(X4)
Display 8	Autocorr(X4)

- PLOTS(UNPACK GROUPBY=TYPE)=(TRACE AUTOCORR) displays all the separate trace plots, followed by the separate autocorrelation plots:

Display 1	Trace(X1)
Display 2	Trace(X2)
Display 3	Trace(X3)
Display 4	Trace(X4)
Display 5	Autocorr(X1)
Display 6	Autocorr(X2)
Display 7	Autocorr(X3)
Display 8	Autocorr(X4)

SEED=*n*

specifies the random number seed. By default, SEED=0, and PROC BCHOICE gets a random number seed from the system clock. Negative seed values are treated as the default. The largest possible value for the seed is $2^{31} - 1$. Analyses that use the same nonzero seed are reproducible. The seed value is reported in the “Model Information” table.

STATISTICS<(global-stats-options)> = NONE | ALL |stats-request**STATS<(global-stats-options)> = NONE | ALL |stats-request**

specifies options for posterior statistics. By default, PROC BCHOICE computes the posterior mean, standard deviation, quantiles, and two 95% credible intervals: equal-tail and highest posterior density (HPD). Other available statistics include the posterior correlation and covariance. For more information, see the section “[Summary Statistics](#)” on page 155. You can request all the posterior statistics by specifying STATS=ALL. You can suppress all the calculations by specifying STATS=NONE.

You can specify the following *global-stats-options*:

ALPHA=numeric-list

specifies the α level for the equal-tail and HPD intervals. The value of α must be between 0 and 0.5. By default, ALPHA=0.05.

PERCENT=numeric-list**PERCENTAGE=numeric-list**

calculates the posterior percentages. The *numeric-list* contains values between 0 and 100. By default, PERCENTAGE=(25 50 75).

You can specify the following *stats-requests*:

ALL

computes all posterior statistics. You can combine the ALL option with any other options. For example, STATS(ALPHA=(0.02 0.05 0.1))=ALL computes all statistics by using the default settings and intervals at α levels of 0.02, 0.05, and 0.1.

BRIEF

computes the posterior means, standard deviations, and $100(1 - \alpha)\%$ HPD credible interval for each variable. By default, ALPHA=0.05, but you can use the global ALPHA= option to request other values. This is the default output for posterior statistics.

CORR

computes the posterior correlation matrix.

COV

computes the posterior covariance matrix.

INTERVAL**INT**

computes the $100(1 - \alpha)\%$ equal-tail and HPD credible intervals for each variable. For more information, see the sections “[Equal-Tail Credible Interval](#)” on page 156 and “[Highest Posterior Density \(HPD\) Interval](#)” on page 156. By default, ALPHA=0.05, but you can use the global ALPHA= option to request other intervals of any probabilities.

NONE

suppresses all the statistics.

SUMMARY**SUM**

computes the posterior means, standard deviations, and percentile points for each variable. By default, the 25th, 50th, and 75th percentile points are produced, but you can use the global PERCENT= option to request specific percentile points.

TARGACCEPT=value

specifies the target acceptance rate for the random walk–based Metropolis algorithm for logit models. For more information, see the section “[Metropolis and Metropolis-Hastings Algorithms](#)” on page 136 in Chapter 7, “[Introduction to Bayesian Analysis Procedures](#).” The numeric *value* must be between 0.01 and 0.99. By default, TARGACCEPT=0.45 for models with one parameter; TARGACCEPT=0.35 for models with two, three, or four parameters; and TARGACCEPT=0.234 for models with more than four parameters (Roberts, Gelman, and Gilks 1997; Roberts and Rosenthal 2001).

THIN=n**NTHIN=n**

controls the thinning rate of the simulation. PROC BCHOICE keeps every *n*th simulation sample and discards the rest. All posterior statistics and diagnostics are calculated by using the thinned samples. By default, THIN=1. For more information, see the section “[Burn-in, Thinning, and Markov Chain Samples](#)” on page 140.

BY Statement

BY variables ;

You can specify a BY statement with PROC BCHOICE to obtain separate analyses of observations in groups that are defined by the BY variables. When a BY statement appears, the procedure expects the input data set to be sorted in order of the BY variables. If you specify more than one BY statement, only the last one specified is used.

If your input data set is not sorted in ascending order, use one of the following alternatives:

- Sort the data by using the SORT procedure with a similar BY statement.
- Specify the NOTSORTED or DESCENDING option in the BY statement for the BCHOICE procedure. The NOTSORTED option does not mean that the data are unsorted but rather that the data are arranged in groups (according to values of the BY variables) and that these groups are not necessarily in alphabetical or increasing numeric order.
- Create an index on the BY variables by using the DATASETS procedure (in Base SAS software).

For more information about BY-group processing, see the discussion in *SAS Language Reference: Concepts*. For more information about the DATASETS procedure, see the discussion in the *Base SAS Procedures Guide*.

CLASS Statement

```
CLASS variable <(options)> ... <variable <(options)>> </ global-options > ;
```

The CLASS statement names the classification variables to be used as explanatory variables in the analysis.

The CLASS statement must precede the MODEL statement. Most options can be specified either as individual variable *options* or as *global-options*. You can specify *options* for each variable by enclosing the options in parentheses after the variable name. You can also specify *global-options* for the CLASS statement by placing them after a slash (/). *Global-options* are applied to all the variables specified in the CLASS statement. If you specify more than one CLASS statement, the *global-options* specified in any one CLASS statement apply to all CLASS statements. However, individual CLASS variable *options* override the *global-options*. You can specify the following values for either an *option* or a *global-option*:

CPREFIX=*n*

specifies that, at most, the first *n* characters of a CLASS variable name be used in creating names for the corresponding design variables. The default is $32 - \min(32, \max(2, f))$, where *f* is the formatted length of the CLASS variable.

DESCENDING

DESC

reverses the sort order of the classification variable. If both the DESCENDING and ORDER= options are specified, PROC BCHOICE orders the categories according to the ORDER= option and then reverses that order.

LPREFIX=*n*

specifies that, at most, the first *n* characters of a CLASS variable label be used in creating labels for the corresponding design variables. The default is $256 - \min(256, \max(2, f))$, where *f* is the formatted length of the CLASS variable.

MISSING

treats missing values (., _, .A, ..., .Z for numeric variables and blanks for character variables) as valid values for the CLASS variable.

ORDER=DATA | FORMATTED | FREQ | INTERNAL

specifies the sort order for the levels of classification variables. This ordering determines which parameters in the model correspond to each level in the data, so the ORDER= option can be useful when you use the CONTRAST statement. By default, ORDER=FORMATTED. For ORDER=FORMATTED and ORDER=INTERNAL, the sort order is machine-dependent. When ORDER=FORMATTED is in effect for numeric variables for which you have supplied no explicit format, the levels are ordered by their internal values.

The following table shows how PROC BCHOICE interprets values of the ORDER= option.

Value of ORDER=	Levels Sorted By
DATA	Order of appearance in the input data set
FORMATTED	External formatted values, except for numeric variables with no explicit format, which are sorted by their unformatted (internal) values
FREQ	Descending frequency count; levels with more observations come earlier in the order
INTERNAL	Unformatted value

For more information about sort order, see the chapter on the SORT procedure in the *Base SAS Procedures Guide* and the discussion of BY-group processing in *SAS Language Reference: Concepts*.

PARAM=*keyword*

specifies the parameterization method for the classification variable or variables. You can specify any of the *keywords* shown in the following table; the default is PARAM=REF. Design matrix columns are created from CLASS variables according to the corresponding coding schemes:

Value of PARAM=	Coding
EFFECT	Effect coding
GLM	Less-than-full-rank reference cell coding (this <i>keyword</i> can be used only in a global option)
REFERENCE REF	Reference cell coding

All parameterizations are full rank, except for the GLM parameterization. The REF= option in the CLASS statement determines the reference level for EFFECT and REFERENCE coding and for their orthogonal parameterizations. It also indirectly determines the reference level for a singular GLM parameterization through the order of levels.

REF='*level*' | *keyword*

specifies the reference level for PARAM=EFFECT, PARAM=REFERENCE, and their orthogonalizations. For PARAM=GLM, the REF= option specifies a level of the classification variable to be put at the end of the list of levels. This level thus corresponds to the reference level in the usual interpretation of the linear estimates with a singular parameterization.

For an individual variable REF= option (but not for a global REF= option), you can specify the *level* of the variable to use as the reference level. Specify the formatted value of the variable if a format is assigned. For a global or individual variable REF= option, you can use one of the following *keywords*. The default is REF=LAST.

FIRST designates the first ordered level as reference.

LAST designates the last ordered level as reference.

TRUNCATE<=*n*>

specifies the length *n* of CLASS variable values to use in determining CLASS variable levels. The default is to use the full formatted length of the CLASS variable. If you specify TRUNCATE without the length *n*, the first 16 characters of the formatted values are used. When formatted values are longer

than 16 characters, you can use this option to revert to the levels as determined in releases before SAS 9. The TRUNCATE option is available only as a global option.

MODEL Statement

MODEL *response* <(response-options)> = <fixed-effects> </model-options> ;

The MODEL statement is required; it defines the dependent variable and the fixed effects. The *fixed-effects* determine the **X** matrix of the model. The specification of effects is the same as in the GLIMMIX and MIXED procedures, where you do not specify random effects in the MODEL statement.

Table 27.3 summarizes the options available in the MODEL statement. These are subsequently discussed in detail in alphabetical order by option category.

Table 27.3 MODEL Statement Options

Option	Description
CHOICESET=	Specifies the variables for defining a choice set
COEFFPRIOR=	Specifies the prior of the regression coefficients
COVPRIOR=	Specifies the prior of the covariance parameter for a probit model
COVTYPE=	Specifies the structure of the covariance matrix of the error difference for a probit model
INIT=	Controls the generation of initial values of the regression coefficients
LAMBDA PRIOR=	Specifies the prior of the log-sum coefficients for a nested logit model
NEST=	Defines the nonoverlapping nests for a nested logit model
SAMELAMBDA	Constrains the log-sum coefficients to be the same for all the nests in a nested logit model
TYPE=	Specifies the type of the model

CHOICESET=(variables)

specifies one or more variables for defining the choice sets. You must specify how the choice sets are constructed, and you can use more than one variable. PROC BCHOICE does not sort by the values of the choice set variable; rather, it considers the data to be from a new choice set whenever the value of the choice set variable changes from the previous observation.

COEFFPRIOR=NORMAL <(options)>

CPRIOR=NORMAL <(options)>

specifies the prior distribution for the regression coefficients. The default is the normal prior $N(\mathbf{0}, 10^2\mathbf{I})$, where **I** is the identity matrix. You can specify the following *options*, enclosed in parentheses:

INPUT=SAS-data-set

specifies a SAS data set that contains the mean and covariance information of the normal prior. The data set must have a `_TYPE_` variable to represent the type of each observation and a variable for each regression coefficient. If the data set also contains a `_NAME_` variable, the values of this variable are used to identify the covariances for the `_TYPE_='COV'` observations; otherwise, the `_TYPE_='COV'` observations are assumed to be in the same order as the explanatory variables in the MODEL statement. PROC BCHOICE reads the mean vector from

the observation for which `_TYPE_='MEAN'` and reads the covariance matrix from observations for which `_TYPE_='COV'`. For an independent normal prior, the variances can be specified with `_TYPE_='VAR'`; alternatively, the precisions (inverse of the variances) can be specified with `_TYPE_='PRECISION'`.

VAR<=c>

specifies the normal prior $N(\mathbf{0}, c\mathbf{I})$, where \mathbf{I} is the identity matrix and c is a scalar.

COVPRIOR=IWISHART<(options)>

specifies an inverse Wishart prior distribution, **IWISHART**(a,b), for the covariance matrix for the vector of error differences. For models that do not have a covariance matrix for the error differences (the logit and nested logit models), this option is ignored.

You can specify the following *options*, enclosed in parentheses:

DF=a

specifies the degrees of freedom of the inverse Wishart distribution. The default is the number of alternatives in the choice set plus 2, which is equivalent to the dimension of the covariance matrix of the error differences plus 3.

SCALE=b

specifies $b\mathbf{I}$ for the scale parameter of the inverse Wishart distribution, where \mathbf{I} is the identity matrix. The default is the number of alternatives in the choice set plus 2.

COVTYPE=UN | VC

specifies the covariance structure of the error difference vector for a probit model. Although a variety of structures are available, most applications call for either **COVTYPE=VC** or **COVTYPE=UN** for the error difference vector. The **COVTYPE=VC** (variance components) models a different variance component for each error term. The **TYPE=UN** (unstructured) specifies a full structured covariance matrix. The unstructured form accommodates any pattern of correlation in addition to fitting a different variance component for each error difference term.

INIT=keyword-list | (numeric-list)

INITIAL=keyword-list | (numeric-list)

specifies options for generating the initial values for the coefficients parameters that are specified as *fixed-effects* in the MODEL statement. By default, **INIT=POSTMODE** for logit models and **INIT=PRIORMODE** for probit models. You can specify the following *keywords*:

LIST=numeric-list

assigns the numbers to be the initial values of the fixed effects in the corresponding list order. The length of the list must be the same as the number of fixed effects. For example, the following statement assigns the values 1, 2, and 3 to the first, second, and third coefficients in the model and prints the table of initial values:

```
model y = x / choicest=(ID Index) init=(list=(1 2 3) pinit);
```

If the length of the list is less than the number of fixed effects, the initial value of each remaining parameter will be replaced by the corresponding default initial value. For example, the corresponding mode of the posterior density is used for a logit model. If the length of the list is greater than the number of fixed effects, the extra ones are ignored.

PINIT

tabulates initial values for the fixed effects. (By default, PROC BCHOICE does not display the initial values.)

POSTMODE

uses the mode of the posterior density as the initial value of the parameter, if you do not provide one. If the mode does not exist or if it is on the boundary of the support of the density, the mean value is used. If you specify POSTMODE for a probit model, where the posterior density is difficult to obtain, PROC BCHOICE resets it to PRIORMODE.

PRIORMODE

uses the mode of the prior density as the initial value of the parameter.

LAMBDA PRIOR=SEMIFLAT < options >**LPRIOR=SEMIFLAT** < options >

specifies a semi-flat prior distribution (Lahiri and Gao 2002) for the log-sum coefficient, λ , for each nest in a nested logit model. For models that are not nested logit, this option is ignored.

You can specify the following *option*, enclosed in parentheses:

PHI=a

specifies the parameter ϕ of the semi-flat prior. By default, $\phi = 0.8$.

NEST=(numeric-list)

defines the nonoverlapping nests for a nested logit model. For a nested logit model, you must specify the nests for all the alternatives in the choice set. Otherwise, the standard logit model is assumed. The number of values in the list should match the number of alternatives in the choice sets, and each of the actual values represents the nest that the particular alternative goes to. For example, NEST=(1 2 1 1 2) arranges the first, third, and fourth alternatives in the first nest and the second and fifth alternatives in the second nest. Currently, this option can accommodate only two-level nested logit models.

SAMELAMBDA

constrains the log-sum coefficients to be the same for all the nests in a nested logit model.

TYPE=keyword

specifies the type of the model. You can specify the following *keywords*:

LOGIT

specifies a standard logit model.

NLOGIT

specifies a nested logit model. If you do not also specify the NEST= option to define the nests, this option is ignored, and a standard logit model is fit.

PROBIT

specifies a probit model.

PREDDIST Statement

PREDDIST <'label'> **OUTPRED**=SAS-data-set <**COVARIATES**=SAS-data-set> ;

The PREDDIST statement creates a new SAS data set that contains random samples from the posterior predictive distribution of the choice probabilities. The posterior predictive distribution is the distribution of unobserved observations (prediction) conditional on the observed data. Let \mathbf{Y} be the observed data, \mathbf{X} be the covariates, θ be the parameter, and \mathbf{Y}_{pred} be the unobserved data. The posterior predictive distribution is defined as follows:

$$\begin{aligned} p(\mathbf{Y}_{\text{pred}}|\mathbf{Y}, \mathbf{X}) &= \int p(\mathbf{Y}_{\text{pred}}, \theta | \mathbf{Y}, \mathbf{X}) d\theta \\ &= \int p(\mathbf{Y}_{\text{pred}}|\theta, \mathbf{Y}, \mathbf{X}) p(\theta | \mathbf{Y}, \mathbf{X}) d\theta \end{aligned}$$

Assuming that the observed and unobserved data are conditionally independent given θ , the posterior predictive distribution can be further simplified as follows:

$$p(\mathbf{Y}_{\text{pred}}|\mathbf{Y}, \mathbf{X}) = \int p(\mathbf{Y}_{\text{pred}}|\theta) p(\theta | \mathbf{Y}, \mathbf{X}) d\theta$$

The posterior predictive distribution is an integral of the likelihood function $p(\mathbf{Y}_{\text{pred}}|\theta)$ with respect to the posterior distribution $p(\theta | \mathbf{Y})$. The PREDDIST statement generates samples from a posterior predictive distribution based on draws from the posterior distribution of θ .

You can specify the following options:

COVARIATES=SAS-data-set

names the SAS data set to contain the sets of explanatory variable values for which the predictions are established. This data set must contain data that has the same variables used in the model. If you omit the COVARIATES= option, the DATA= data set that is specified in the PROC BCHOICE statement is used instead.

NALTER=n

NALTERNATIVE=n

specifies the number of alternatives in a choice set in the COVARIATES= data set. All choice sets in the data must have the same number of alternatives.

OUTPRED=SAS-data-set

creates an output data set to contain the samples from the posterior predictive distribution.

RANDOM Statement

RANDOM random-effects </ options > ;

The RANDOM statement defines the random effects in the choice model. You can use this statement to specify traditional variance component models and to specify random coefficients, which can be classification

or continuous. You can specify only one RANDOM statement. Include the **TYPE=** option to define the covariance structure. PROC BCHOICE does not include the intercept in the RANDOM statement.

Table 27.4 summarizes the *options* available in the RANDOM statement. All *options* are subsequently discussed in alphabetical order.

Table 27.4 Summary of RANDOM Statement Options

Option	Description
COVPRIOR=	Specifies the prior distribution for the covariance of the random effects
MONITOR	Displays solutions of the random effects
NOOUTPOST	Suppresses the output of the posterior samples of random effects to the OUTPOST= data set
REMEAN	Models the prior mean of the random effects
SUBJECT=	Identifies the subjects in the model
TYPE=	Specifies the covariance structure

You can specify the following *options* in the RANDOM statement after a slash (/).

COVPRIOR=IWISHART<(options)>

specifies an inverse Wishart prior distribution, **IWISHART**(*a*,*b*), for the covariance matrix of the random effects.

You can specify the following *options*, enclosed in parentheses:

DF=*a*

specifies the degrees of freedom of the inverse Wishart distribution. The default is the dimension of the covariance matrix of the random effects plus 3.

SCALE=*b*

specifies *b***I** for the scale parameter of the inverse Wishart distribution, where **I** is the identity matrix. The default is the dimension of the covariance matrix of the random effects plus 3.

MONITOR

MONITOR=(*numeric-list*)

MONITOR=**RANDOM** (*number*)

outputs results for the individual-level random-effects parameters. (By default, PROC BCHOICE does not output results for individual-level random-effects parameters.) You can monitor a subset of the random-effects parameters. You can provide a numeric list of the SUBJECT indexes, or PROC BCHOICE can randomly choose a subset of all subjects for you.

To monitor a list of random-effects parameters for certain subjects, you can provide their indexes as follows:

```
random x / subject=index monitor=(1 to 5 by 2 23 57);
```

PROC BCHOICE outputs results of random effects for subjects 1, 3, 5, 23, and 57. PROC BCHOICE can also randomly choose a subset of all the subjects to monitor, if you submit a statement such as the following:

```
random x / subject=index monitor=(random(12));
```

PROC BCHOICE outputs results of random effects for 12 randomly selected subjects. You control the sequence of the random indexes by specifying the **SEED=** option in the PROC BCHOICE statement.

When you specify the **MONITOR** option, it uses the specification of the **STATISTICS=** and **PLOTS=** options in the PROC BCHOICE statement. By default, PROC BCHOICE outputs all the posterior samples of all random-effects parameters to the **OUTPOST=** output data set. You can use the **NOOUTPOST** option to suppress the saving of the random-effects parameters.

NOOUTPOST

suppresses the output of the posterior samples of random-effects parameters to the **OUTPOST=** data set. In models that have a large number of random-effects parameters (for example, tens of thousands), PROC BCHOICE can run faster if it does not save the posterior samples of the random-effects parameters.

When you specify both the **NOOUTPOST** option and the **MONITOR** option, PROC BCHOICE outputs the list of variables that are monitored.

The maximum number of variables that can be saved to an **OUTPOST=** data set is 32,767. If you run a large-scale random-effects model in which the number of parameters exceeds this limit, the **NOOUTPOST** option is invoked automatically and PROC BCHOICE does not save the random-effects draws to the posterior output data set. You can use the **MONITOR** option to select a subset of the parameters to store in the **OUTPOST=** data set.

REMEAN

models the mean of the random effects, so that $\bar{\gamma}$ is estimated. You can also model the prior mean of the random effects as a function of individual demographic variables. Rossi, McCulloch, and Allenby (1996) and Rossi, Allenby, and McCulloch (2005) propose adding another layer of flexibility to the random-effects-only model by allowing heterogeneity that is driven by observable (demographic) characteristics of the individuals. The following **REMEAN=(AGE GENDER)** option in the **RANDOM** statement estimates the mean of the random effect, X , and models the mean as a function of Age and Gender:

```
random X / subject=Index remean=(Age Gender);
```

SUBJECT=*effect*

SUB=*effect*

identifies the subjects in the model. PROC BCHOICE assumes complete independence across subjects; thus, for the **RANDOM** statement, the **SUBJECT=** option produces a block-diagonal structure that has identical blocks. Specifying a subject effect is equivalent to nesting all other effects in the **RANDOM** statement within the subject effect.

The *effect* can be continuous variables. PROC BCHOICE does not sort by the values of the continuous variable; rather, it considers the data to be from a new subject or group whenever the value of the continuous variable changes from the previous observation.

TYPE=UN | VC

specifies the covariance structure. Although a variety of structures are available, most applications call for either TYPE=VC or TYPE=UN. The TYPE=VC (variance components) option, which is the default structure, models a different variance component for each random effect. The TYPE=UN (unstructured) specifies a full structured covariance matrix for the random effects. The unstructured form accommodates any pattern of correlation between the random effects in addition to fitting a different variance component for each random effect.

Details: BCHOICE Procedure

Discrete Choice Models

Discrete choice models are used in marketing research to model decision makers' choices among alternative products and services. The decision makers might be people, households, companies and so on, and the alternatives might be products, services, actions, or any other options or items about which choices must be made (Train 2009). The collection of alternatives that are available to the decision makers is called a choice set.

Discrete choice models are derived under the assumption of utility-maximizing behavior by the decision maker. When individuals are asked to make one choice among a set of alternatives, they usually determine the level of utility that each alternative offers. The utility that individual i obtains from alternative j among J alternatives is denoted as

$$u_{ij} = v_{ij} + \epsilon_{ij}, \quad i = 1, \dots, N, \quad \text{and} \quad j = 1, \dots, J$$

where the subscript i is an index for the individuals, the subscript j is an index for the alternatives in a choice set, v_{ij} is a nonstochastic utility function that relates those observed factors to the utility, and ϵ_{ij} is the error component that captures the unobserved characteristics of the utility. In discrete choice models, the observed part of the utility function is assumed to be linear in the parameters,

$$v_{ij} = \mathbf{x}'_{ij} \boldsymbol{\beta}$$

where \mathbf{x}_{ij} is a p -dimensional design vector of observed attribute levels that relate to alternative j and $\boldsymbol{\beta}$ is the corresponding vector of fixed regression coefficients that indicate the utilities or part-worths of the attribute levels.

Decision makers choose the alternative that gives them the greatest utility. Let \mathbf{y}_i be the multinomial response vector for the i th individual. The value y_{ij} takes 1 if the j th component of $\mathbf{u}_i = (u_{i1}, \dots, u_{iJ})$ is the largest, and 0 otherwise:

$$\begin{aligned} u_{ij} &= \mathbf{x}'_{ij} \boldsymbol{\beta} + \epsilon_{ij} \\ y_{ij} &= \begin{cases} 1 & \text{if } u_{ij} \geq \max(\mathbf{u}_i) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

The probability that the individual i chooses alternative j is

$$\begin{aligned}
 P(y_{ij} = 1) &= \Pr(u_{ij} > u_{ik} \text{ for all } k \neq j) \\
 &= \Pr(v_{ij} + \epsilon_{ij} > v_{ik} + \epsilon_{ik} \text{ for all } k \neq j) \\
 &= \Pr(\epsilon_{ik} - \epsilon_{ij} < v_{ij} - v_{ik} \text{ for all } k \neq j) \\
 &= \int_{\epsilon} \mathbf{I}(\epsilon_{ik} - \epsilon_{ij} < v_{ij} - v_{ik} \text{ for all } k \neq j) f(\epsilon_i) d\epsilon_i
 \end{aligned}$$

where $\mathbf{I}(\cdot)$ is the indicator function and $f(\epsilon_i)$ denotes the joint density of the error vector $\epsilon_i = (\epsilon_{i1}, \dots, \epsilon_{iJ})$. Different specifications about the density result in different types of choice models, as detailed in the next section. Logit and nested logit models have a closed form for this integral, whereas a probit model does not.

Types of Choice Models

Logit

In the multinomial logit (MNL) model, each ϵ_{ij} ($j = 1, \dots, J$) is independently and identically distributed (iid) with the Type I extreme-value distribution, $\exp(-\exp(-\epsilon_{ij}))$, which is also known as the Gumbel distribution. The essential part of the model is that the unobserved parts of utility are uncorrelated among all alternatives, in addition to having the same variance. This assumption provides a convenient form for the choice probability (McFadden 1974):

$$P(y_{ij} = 1) = \frac{\exp(\mathbf{x}'_{ij}\boldsymbol{\beta})}{\sum_{k=1}^J \exp(\mathbf{x}'_{ik}\boldsymbol{\beta})}, \quad i = 1, \dots, N \text{ and } j = 1, \dots, J$$

The MNL likelihood is formed by the product of the N independent multinomial distributions:

$$p(\mathbf{Y}|\boldsymbol{\beta}) = \prod_{i=1}^N \prod_{j=1}^J P(y_{ij} = 1)^{y_{ij}}$$

The MNL model is by far the easiest and most widely used discrete choice model. It is popular because the formula for the choice probability has a closed form and is readily interpretable. In his Nobel lecture, McFadden (2001) tells a history of the development of this path-breaking model.

Bayesian analysis for the MNL model requires the specification of a prior over the coefficient parameters $\boldsymbol{\beta}$ and computation of the posterior density. It is convenient to use a normal prior on $\boldsymbol{\beta}$, $\pi(\boldsymbol{\beta}) = \mathbf{N}(\bar{\boldsymbol{\beta}}, \boldsymbol{\Omega}_{\boldsymbol{\beta}})$. When the investigator has no strong prior beliefs about the location of the parameters, it is recommended that diffuse, but proper, priors be used.

The posterior density of the parameter $\boldsymbol{\beta}$ is

$$p(\boldsymbol{\beta}|\mathbf{Y}) \propto p(\mathbf{Y}|\boldsymbol{\beta})\pi(\boldsymbol{\beta})$$

Because discrete choice logit models fall under the framework of a generalized linear model (GLM) with a logit link, you can use Metropolis-Hastings sampling with the Gamerman approach for the sampling procedure. This approach extends the usual iterative weighted least squares method to include a sampling step based on the Metropolis-Hastings algorithm for GLMs. For more information about the Gamerman approach,

see the section “[Gamerman Algorithm](#)” on page 1046. You can also incorporate this methodology for a hierarchical structure that has random effects. For more information, see the section “[Logit with Random Effects](#)” on page 1041. The Gamerman approach is the default sampling method when a direct sampling that uses a closed-form posterior distribution cannot be attained. The random walk Metropolis is also available when you specify `ALGORITHM=RWM` in the PROC BCHOICE statement.

The iid assumption about the error terms of the utilities of the different alternatives imposes the independence from irrelevant alternatives (IIA) property, which implies proportional substitution across alternatives. Sometimes this property might seem restrictive or unrealistic, prompting the need for more flexible models to overcome it.

Nested Logit

You derive the nested logit model by allowing the error components to be identical but nonindependent. Instead of being independent Type I extreme-value errors, the errors are assumed to have a generalized extreme-value (GEV) distribution. This model generalizes the standard logit model to allow for particular patterns of correlation in unobserved utility (McFadden 1978).

In a nested logit model, all the alternatives in a choice set can be partitioned into nests in such a way that the following conditions are true:

- The probability ratio of any two alternatives that are in the same nest is independent of the existence of all other alternatives. Hence, the IIA assumption holds within each nest.
- The probability ratio of any two alternatives that are in different nests is not independent of the existence of other alternatives in the two nests. Hence, the IIA assumption does not hold between different nests.

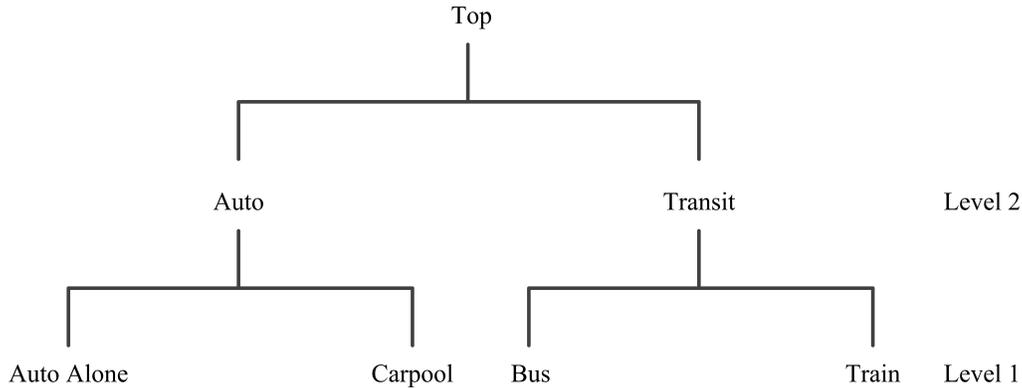
A nested logit example from Train (2009) best explains how a nested logit model works. Suppose the alternatives for commuting to work are driving an auto alone, carpooling, taking the bus, and taking the train. If you remove any alternative, the probabilities of the other alternatives increase. But by what percentage would each probability increase? Suppose the changes in probabilities occur as shown in [Table 27.5](#).

Table 27.5 Example of IIA Holding within Nests of Alternatives

Alternative	Original Prob.	Probabilities with One Alternative Removed			
		Auto Alone	Carpool	Bus	Train
Auto alone	0.40	–	0.45(+12.5%)	0.52(+30.0%)	0.48(+20.0%)
Carpool	0.10	0.20(+100%)	–	0.13(+30.0%)	0.12(+20.0%)
Bus	0.30	0.48(+60.0%)	0.33(+10.0%)	–	0.40(+33.0%)
Train	0.20	0.32(+60.0%)	0.22(+10.0%)	0.35(+70.0%)	–

In [Table 27.5](#), the probabilities for bus and train increase by the same proportion whenever you remove one of the other alternatives. Therefore, IIA is valid between bus and train, so they can be placed in the same nest called “transit.” Similarly, auto alone and carpool are placed in one nest called “auto.” IIA does not hold between these two nests.

A convenient way to represent the model and substitution patterns is to draw a tree diagram, where each branch denotes a subset of alternatives within which IIA holds and where each leaf on each branch denotes an alternative. The tree diagram for the commuting-to-work example is shown in [Figure 27.10](#).

Figure 27.10 Two-Level Decision Tree

Although decision trees in nested logit analyses, such as the one shown in Figure 27.10, are often interpreted as implying that the higher-level decisions are made first, followed by decisions at lower levels, no such temporal ordering is necessarily required. A good way to think about nested logit models is that they are appropriate when there are groups of alternatives that are similar to each other in unobserved ways. In other words, nested logit models are appropriate when there is correlation between the alternatives in each nest but no correlation between the alternatives in different nests. Currently, PROC BCHOICE considers only two-level nested logit models.

In mathematical notation, let the set of alternatives be partitioned into K nonoverlapping subsets (nests) that are denoted by S_1, S_2, \dots, S_K . The nested logit model is derived by assuming that the vector of the unobserved part of the utility, $\epsilon_i = (\epsilon_{i1}, \dots, \epsilon_{iJ})$, has a cumulative distribution

$$\exp\left(-\sum_{k=1}^K \left(\sum_{j \in S_k} \exp(-\epsilon_{ij}/\lambda_k)\right)^{\lambda_k}\right)$$

This is a type of generalized extreme value (GEV) distribution. For a standard logit model, each ϵ_{ij} is iid with an extreme-value distribution, whereas for a nested logit model, the marginal distribution of ϵ_{ij} is an extreme-value distribution. The ϵ_{ij} are correlated within nests. If alternatives j and m belong to the same nest, then ϵ_{ij} is correlated with ϵ_{im} . But if any two alternatives are in different nests, their unobserved part of utility is still independent. The parameter λ_k measures the degree of independence among alternatives in nest k . The higher the value of λ_k is, the less correlation there is. However, the correlation is actually more complicated than the parameter λ_k . The equation $\lambda_k = 1$ represents no correlation in nest k . If $\lambda_k = 1$ for all nests, the nested logit model reduces to the standard logit model.

The choice probability for alternative $j \in S_k$ has a closed form:

$$P(y_{ij} = 1) = \frac{\exp(\mathbf{x}'_{ij}\boldsymbol{\beta}/\lambda_k) \left(\sum_{m \in S_k} \exp(\mathbf{x}'_{im}\boldsymbol{\beta}/\lambda_k)\right)^{\lambda_k - 1}}{\sum_{l=1}^K \left(\sum_{m \in S_l} \exp(\mathbf{x}'_{im}\boldsymbol{\beta}/\lambda_l)\right)^{\lambda_l}}$$

From this form, the nested logit likelihood is derived as the product of the N multinomial distributions:

$$p(\mathbf{Y}|\boldsymbol{\beta}, \boldsymbol{\lambda}) = \prod_{i=1}^N \prod_{j=1}^J P(y_{ij} = 1)^{y_{ij}}$$

A prior for $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_K)$ is needed in addition to a prior for the coefficient parameters $\boldsymbol{\beta}$. The λ_k in nest k is often called the log-sum coefficient. The value of λ_k must be positive for the model to be consistent with

utility-maximizing behavior. If $\lambda_k \in [0, 1]$ for all k , the model is consistent with utility maximization for all possible values of the explanatory covariates, whereas the model is consistent only for some range of the covariates but not for all values if $\lambda_k > 1$. Kling and Herriges (1995) and Herriges and Kling (1996) suggest tests of consistency of nested logit with utility maximization when $\lambda_k > 1$. Train, McFadden, and Ben-Akiva (1987) show an example of models for which $\lambda_k > 1$. A negative value of λ_k should be avoided because the model will be inconsistent and imply that consumers are choosing the alternative to minimize their utilities.

For the nested logit model, noninformative priors are not ideal for λ . Flat priors on different versions of the parameter space can yield different posterior distributions.

Lahiri and Gao (2002) suggest the following semi-flat priors by using the parameter ϕ for each λ :

$$\pi(\lambda) = \begin{cases} 0 & \text{if } \lambda \leq 0 \\ \phi & \text{if } 0 < \lambda < 1 \\ \phi \exp\left[\frac{\phi}{1-\phi}(1-\lambda)\right] & \text{if } \lambda \geq 1 \end{cases}$$

PROC BCHOICE uses this prior with a default value of 0.8 for ϕ . You can specify other values for ϕ . Sims' priors are also attractive with various s values as follows:

$$\pi(\lambda) = \begin{cases} 0 & \text{if } \lambda \leq 0 \\ s\lambda^{s-1} \exp(-\lambda^s) & \text{if } \lambda > 0 \end{cases}$$

In addition, a gamma or beta distribution is a good choice for the prior of λ . However, PROC BCHOICE supports only semi-flat priors in this release.

When the priors for the parameters (β, λ) are specified, the posterior density is as follows:

$$p(\beta, \lambda | \mathbf{Y}) \propto p(\mathbf{Y} | \beta, \lambda) \pi(\beta) \pi(\lambda)$$

Another generalization of the conditional logit model, the heteroscedastic extreme-value (HEV) model, is obtained by allowing independent but nonidentical errors that are distributed with a Type I extreme-value distribution (Bhat 1995). The HEV model permits different variances on the error components of utility across the alternatives. Currently, PROC BCHOICE does not consider this type of model.

Probit

The multinomial probit (MNP) model is derived when the errors, $\epsilon'_i = (\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{iJ})$, have a multivariate normal (MVN) distribution with a mean vector of $\mathbf{0}$ and a covariance matrix Σ .

$$\begin{aligned} u_{ij} &= \mathbf{x}'_{ij} \beta + \epsilon_{ij}, \quad \epsilon_i \sim N(\mathbf{0}, \Sigma), \quad i = 1, \dots, N \quad \text{and} \quad j = 1, \dots, J \\ y_{ij} &= \begin{cases} 1 & \text{if } u_{ij} \geq \max(\mathbf{u}_i) \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

where Σ is the $J \times J$ covariance matrix of the error vector. For a full covariance matrix, any pattern of correlation and heteroscedasticity can be accommodated. Thus, this model fits a very general error structure.

However, the choice probability, from which the likelihood is derived, does not have a closed form. McCulloch and Rossi (1994) propose an algorithm that is a multivariate version of the probit regression algorithm of Albert and Chib (1993) and construct a Gibbs sampler that is based on a Markov chain to draw directly from the exact posteriors of the MNP model. This approach avoids direct evaluation of the likelihood. Hence,

it avoids problems that are associated with calculating choice probabilities that affect both the standard likelihood and the method of simulated moments approaches.

To avoid the parameter identification problems, a usual practice is to take differences against one particular alternative, because only differences in utility matter. Suppose you take differences against the last alternative in the choice set and you define $w_{ij} = u_{ij} - u_{iJ}$, $\tilde{\mathbf{x}}_{ij} = \mathbf{x}_{ij} - \mathbf{x}_{iJ}$, $\tilde{\epsilon}_{ij} = \epsilon_{ij} - \epsilon_{iJ}$, and $\tilde{\boldsymbol{\epsilon}}'_i = (\tilde{\epsilon}_{i1}, \tilde{\epsilon}_{i2}, \dots, \tilde{\epsilon}_{iJ-1})$.

$$w_{ij} = \tilde{\mathbf{x}}'_{ij}\boldsymbol{\beta} + \tilde{\epsilon}_{ij}, \quad \tilde{\boldsymbol{\epsilon}}_i \sim N(0, \tilde{\boldsymbol{\Sigma}}), \quad i = 1, \dots, N \quad \text{and} \quad j = 1, \dots, J - 1$$

$$y_{ij} = \begin{cases} 1 & \text{if } w_{ij} \geq \max(0, \mathbf{w}_{i,-j}) \\ 0 & \text{otherwise} \end{cases}$$

where $\tilde{\boldsymbol{\Sigma}}$ is the $(J - 1) \times (J - 1)$ covariance matrix of the vector of error differences, $\tilde{\boldsymbol{\epsilon}}_i$, and $\mathbf{w}_{i,-j} = (w_{i1}, \dots, w_{i(j-1)}, w_{i(j+1)}, \dots, w_{i(J-1)})$. The dimension has been reduced to $J - 1$. McCulloch and Rossi (1994) call w_{ij} the latent variable. The introduction of the latent variables in what is known as the data augmentation step makes the application of a Gibbs sampler feasible.

A normal prior is used for $\boldsymbol{\beta}$ and an inverse Wishart prior on $\tilde{\boldsymbol{\Sigma}}$:

$$\pi(\boldsymbol{\beta}) = N(\bar{\boldsymbol{\beta}}, \boldsymbol{\Omega}_\beta)$$

$$\pi(\tilde{\boldsymbol{\Sigma}}) = \text{inverse Wishart}(\nu, \mathbf{V})$$

Then sampling is carried out consecutively from the following three groups of conditional posterior distributions:

- (1) $(w_{ij} | \mathbf{w}_{i,-j}, \boldsymbol{\beta}, \tilde{\boldsymbol{\Sigma}}, \mathbf{Y}), \quad i = 1, \dots, N \quad \text{and} \quad j = 1, \dots, J - 1$
- (2) $(\boldsymbol{\beta} | \mathbf{W}, \tilde{\boldsymbol{\Sigma}}, \mathbf{Y})$
- (3) $(\tilde{\boldsymbol{\Sigma}} | \mathbf{W}, \boldsymbol{\beta}, \mathbf{Y})$

where \mathbf{W} is obtained by stacking all \mathbf{w}_i . All three groups of conditional distributions have closed forms that are easily drawn from. Conditional (1) is truncated normals, (2) is a regular normal, and (3) is an inverse Wishart distribution. This sampling avoids the tuning stage that is required by most Metropolis-Hastings algorithms, and its analytically tractable complete data likelihood also makes it easy to embed probit models into more elaborate hierarchical structures, such as random-effects models. For more information, see McCulloch and Rossi (1994).

Originally, the multinomial probit model required burdensome computation compared to a family of multinomial choice models that are derived from the Gumbel distributed utility function, because the computation involves multidimensional integration (with dimension $J - 1$) in the estimation process. In addition, the multinomial probit model requires more parameters than other multinomial choice models. However, it fully relaxes the restrictive IIA assumption that is imposed in logit models and allows any pattern of substitution among alternatives. Because of the breakthrough Bayesian methods that Albert and Chib (1993) and McCulloch and Rossi (1994) introduced, the probit model can be estimated without the need to compute the choice probabilities. This advantage has greatly increased the popularity of the probit model.

Random Effects

Choice models that have random effects (or random coefficients) provide solutions to create individual-level or group-specific utilities. They are also referred as “mixed models” or “hybrid models.” One of the greatest challenges in marketing research is to acknowledge the diversity of preferences and sensitivities in the marketplace. Heterogeneity in individual preferences is the reason for differentiated product programs and market segmentation. As individual preferences become more diverse, it becomes less appropriate to consider the analyses in an aggregated way. Individual utilities are useful because they make segmentation easy and provide a way to detect groups. Because people have different preferences, it can be misleading to roll the whole sample together into a single set of utilities.

For example, imagine studying the popularity of a new brand. Some participants in the study who are interviewed love the new brand, whereas others dislike it. If you simply aggregate the data and look at the average, the conclusion is that the sample is ambivalent toward the new brand. This would be the least helpful conclusion that could be drawn, because it does not fit for anyone.

Choice models that have random effects generalize the standard choice models to incorporate individual-level effects. Let the utility that individual i obtains from alternative j in choice situation t ($t = 1, \dots, T$) be

$$u_{ijt} = \mathbf{x}'_{ijt}\boldsymbol{\beta} + \mathbf{z}'_{ijt}\boldsymbol{\gamma}_i + \epsilon_{ijt}$$

$$y_{ijt} = \begin{cases} 1 & \text{if } u_{ijt} \geq \max(u_{i1t}, u_{i2t}, \dots, u_{iJt}) \\ 0 & \text{otherwise} \end{cases}$$

where y_{ijt} is the observed choice for individual i and alternative j in choice situation t ; \mathbf{x}_{ijt} is the fixed design vector for individual i and alternative j in choice situation t ; $\boldsymbol{\beta}$ is the vector of fixed coefficients; \mathbf{z}_{ijt} is the random design vector for individual i and alternative j in choice situation t ; and $\boldsymbol{\gamma}_i$ is the vector of random coefficients for individual i that correspond to \mathbf{z}_{ijt} . Sometimes the random coefficients might be at a level different from the individual level. For example, it is common to assume that there are random coefficients at the household level or group level of participants. For the convenience of notation, random effects are assumed to be at the individual level.

In the random-effects model, it is assumed that each $\boldsymbol{\gamma}_i$ is drawn from a superpopulation and this superpopulation is normal, $\boldsymbol{\gamma}_i \sim \text{iid } N(\mathbf{0}, \boldsymbol{\Omega}_{\boldsymbol{\gamma}})$. An additional stage is added to the model where a prior for $\boldsymbol{\Omega}_{\boldsymbol{\gamma}}$ is specified:

$$\pi(\boldsymbol{\gamma}_i) = N(\mathbf{0}, \boldsymbol{\Omega}_{\boldsymbol{\gamma}})$$

$$\pi(\boldsymbol{\Omega}_{\boldsymbol{\gamma}}) = \text{inverse Wishart}(\nu_0, \mathbf{V}_0)$$

The covariance matrix $\boldsymbol{\Omega}_{\boldsymbol{\gamma}}$ characterizes the extent of unobserved heterogeneity among individuals. Large diagonal elements of $\boldsymbol{\Omega}_{\boldsymbol{\gamma}}$ indicate substantial heterogeneity in part-worths. Off-diagonal elements indicate patterns in the evaluation of attribute levels. For example, positive covariances specify pairs of attribute levels that tend to be evaluated similarly across respondents. Product offerings that consist of these attribute levels are more strongly preferred or disliked by certain individuals.

In this setup, the prior mean is 0 for the random effects, meaning that the random effects either are truly around 0 or have been centered by the fixed effects. For random effects whose mean is not around 0, you can follow the usual practice of specifying them in the fixed effects. For example, if one random effect is price and you do not think that the population mean of price is around 0, then you should add price as a fixed effect as follows:

```

proc bchoice data=randeffdata;
  class subj set;
  model y = price / choicest=(subj set);
  random price / subject=subj;
run;

```

Thus, you obtain the estimate for the population mean of the price effect through the fixed effect, and you obtain the deviation from the population mean for each individual through random effects.

Allenby and Rossi (1999) and Rossi, Allenby, and McCulloch (2005) propose a hierarchical Bayesian random-effects model that is set up in a different way. In their model, there are no fixed effects but only random effects. This model, which is referred to as the random-effects-only model in the rest of this chapter, is as follows:

$$\begin{aligned}
 u_{ijt} &= \mathbf{z}'_{ijt} \boldsymbol{\gamma}_i + \epsilon_{ijt} \\
 y_{ijt} &= \begin{cases} 1 & \text{if } u_{ijt} \geq \max(u_{i1t}, u_{i2t}, \dots, u_{iJt}) \\ 0 & \text{otherwise} \end{cases} \\
 \pi(\boldsymbol{\gamma}_i) &= N(\bar{\boldsymbol{\gamma}}, \boldsymbol{\Omega}_{\boldsymbol{\gamma}}) \\
 \pi(\boldsymbol{\Omega}_{\boldsymbol{\gamma}}) &= \text{inverse Wishart}(\nu_0, \mathbf{V}_0)
 \end{aligned}$$

where $\bar{\boldsymbol{\gamma}}$ is a mean vector of regression coefficients, which models the central location of distribution of the random coefficients; $\bar{\boldsymbol{\gamma}}$ represents the average part-worths across the respondents in the data. If you want to use this setup, specify the **REMEAN** option in the **RANDOM** statement to request estimation on $\bar{\boldsymbol{\gamma}}$ and do not specify any fixed effects in the **MODEL** statement.

Rossi, McCulloch, and Allenby (1996) and Rossi, Allenby, and McCulloch (2005) add another layer of flexibility to the random-effects-only model by allowing heterogeneity that is driven by observable (demographic) characteristics of the individuals. They model the prior mean of the random coefficients as a function of the individual's demographic variables (such as age and gender),

$$\begin{aligned}
 \pi(\boldsymbol{\gamma}_i) &= N(\boldsymbol{\Gamma} d_i, \boldsymbol{\Omega}_{\boldsymbol{\gamma}}) \\
 \pi(\boldsymbol{\Omega}_{\boldsymbol{\gamma}}) &= \text{inverse Wishart}(\nu_0, \mathbf{V}_0)
 \end{aligned}$$

where d_i is a vector that consists of an intercept and some observable demographic variables. $\boldsymbol{\Gamma}$ is a matrix of regression coefficients, which affects the location of distribution of the random coefficients. $\boldsymbol{\Gamma}$ should be useful for identifying respondents who have part-worths that are different from those in the rest of the sample if some individual-level characteristics are included in d_i . This specification allows the preferences or intercepts to vary by both demographic variables and the slopes. If d_i consists of only an intercept, this model reduces to the previous one. For more information about how to sample $\boldsymbol{\Gamma}$, see Rossi, McCulloch, and Allenby (1996), Rossi, Allenby, and McCulloch (2005) (Section 2.12 in Chapter 2), and Rossi (2012).

Logit with Random Effects

The logit model with random effects consists of the fixed-coefficients parameters $\boldsymbol{\beta}$, the random-coefficients parameters $\boldsymbol{\gamma}_i$, and the covariance parameters for the random coefficients $\boldsymbol{\Omega}_{\boldsymbol{\gamma}}$. You can use the Metropolis-Hastings sampling with Gamerman approach to draw samples through the following three conditional posterior distributions:

- (1) $(\boldsymbol{\beta} | \boldsymbol{\gamma}_i, \mathbf{Y})$
- (2) $(\boldsymbol{\gamma}_i | \boldsymbol{\beta}, \boldsymbol{\Omega}_{\boldsymbol{\gamma}}, \mathbf{Y}) \quad i = 1, \dots, N$
- (3) $(\boldsymbol{\Omega}_{\boldsymbol{\gamma}} | \boldsymbol{\gamma}_i, \mathbf{Y})$

All chains are initialized with random effects that are set to 0 and a covariance matrix that is set to an identity matrix. Updating is done first for the fixed effects, β , as a block to position the chain in the correct region of the parameter space. Then the random effects are updated, and finally the covariance of the random effects is updated. For more information, see Gamerman (1997) and the section “Gamerman Algorithm” on page 1046.

The hierarchical Bayesian random-effects-only model as proposed in Allenby and Rossi (1999) and Rossi, Allenby, and McCulloch (2005) contains the random-coefficients parameters γ_i , the population mean of the random-coefficients parameters $\bar{\gamma}$, and the covariance parameters for the random coefficients Ω_γ . The sampling can be carried out by the following conditional posteriors:

- (1) $(\gamma_i | \bar{\gamma}, \Omega_\gamma, \mathbf{Y}) \quad i = 1, \dots, N$
- (2) $(\bar{\gamma} | \gamma_i, \Omega_\gamma)$
- (3) $(\Omega_\gamma | \gamma_i, \bar{\gamma})$

The second and third conditional posteriors are easy to draw because they have direct sampling distributions: the second has a normal distribution with a mean of $\sum_{i=1}^N \gamma_i / N$ and a covariance of Ω_γ / N ; the third is an inverse Wishart($\nu_0 + N, \mathbf{V}_0 + S$), where $S = \sum_{i=1}^N (\gamma_i - \bar{\gamma})(\gamma_i - \bar{\gamma})' / N$. There is no closed form for the first conditional posterior. The Metropolis-Hastings sampling with Gamerman approach is the default sampling algorithm for it.

You can also use random walk Metropolis sampling when direct sampling is not an option, such as for the fixed-coefficients parameters β and random-coefficients parameters γ_i . You can specify `ALGORITHM=RWM` to choose random walk Metropolis sampling. For the random-coefficients parameters γ_i , Rossi, McCulloch, and Allenby (1996) and Rossi, Allenby, and McCulloch (2005) suggest random walk Metropolis sampling in which increments have covariance $s^2 \Omega_\gamma^t$, where s is a scaling constant whose value is usually set at $\frac{2.93}{\sqrt{\dim(\gamma_i)}}$ and Ω_γ^t is the current draw of Ω_γ .

The sampling for the random-effects-only setup is often faster, because the first conditional posterior for each i involves only the data for individual i and because the second and third conditional posteriors depend not on the data directly but only on the draws of γ_i .

Probit with Random Effects

The probit model with random effects has the following parameters: the fixed-coefficients parameters β , the covariance parameters for the error differences $\tilde{\Sigma}$, the random-coefficients parameters γ_i , and the covariance parameters for the random coefficients Ω_γ . It has extra parameters $(\gamma_i, \Omega_\gamma)$ in addition to $(\beta, \tilde{\Sigma})$ in a fixed-effects-only model. You can conveniently adapt the Gibbs sampler proposed in McCulloch and Rossi (1994) to handle the model by appending the new parameters to the set of parameters that would be drawn for a probit model without random coefficients. The sampling can be carried out by the following conditional posteriors:

- (1) $(w_{ij} | \mathbf{w}_{i,-j}, \beta, \tilde{\Sigma}, \gamma_i, \mathbf{Y}) \quad i = 1, \dots, N \text{ and } j = 1, \dots, J - 1$
- (2) $(\beta | \mathbf{W}, \tilde{\Sigma}, \gamma_i, \mathbf{Y})$
- (3) $(\gamma_i | \mathbf{W}, \beta, \tilde{\Sigma}, \Omega_\gamma, \mathbf{Y}) \quad i = 1, \dots, N$
- (4) $(\tilde{\Sigma} | \mathbf{W}, \beta, \mathbf{Y})$
- (5) $(\Omega_\gamma | \mathbf{W}, \beta, \tilde{\Sigma}, \gamma_i, \mathbf{Y})$

All the groups of conditional distributions have closed forms that are easily drawn from. Conditional (1) is truncated normal, (2) and (3) are normal, and (4) and (5) are inverse Wishart distribution. For more information, see McCulloch and Rossi (1994).

Identification and Specification

Two aspects of identification in choice models play an important role in model specification: location shift and scale shift. For more detailed information, see Chapter 2 in Train (2009).

Location Shift

Decision makers choose the alternative that gives them the greatest utility. If a constant is added to the utilities of all alternatives, the location of the utilities will change, but the decision maker will choose the same alternative.

The probability that the individual i chooses alternative j is

$$\begin{aligned} P(y_{ij} = 1) &= \Pr(u_{ij} > u_{ik} \text{ for all } k \neq j) \\ &= \Pr(u_{ij} - u_{ik} > 0 \text{ for all } k \neq j) \end{aligned}$$

Only the difference in utilities matters, not the absolute values (Train 2009). This explains why there is no overall intercept term in a choice model.

Alternative-Specific Main Effects

People often want to observe the main effect that is related to each of the alternatives. For example,

$$u_{ij} = \alpha_j + \mathbf{x}'_{ij}\boldsymbol{\beta} + \epsilon_{ij} \text{ for all } j$$

where α_j is the main effect that is related to alternative j , \mathbf{x}_{ij} is the design variable vector that is specific to alternative j , and $\boldsymbol{\beta}$ are the corresponding coefficients of \mathbf{x}_{ij} . The main effect α_j implies the average effect for alternative j on utility of all factors that are not included in the model.

Consider a person who can take either a car or a bus to work. The form of utilities for taking a car or bus can be specified as follows:

$$\begin{aligned} u_{\text{car}} &= \alpha_{\text{car}}^1 + \mathbf{x}'_{\text{car}}\boldsymbol{\beta} + \epsilon_{\text{car}} \\ u_{\text{bus}} &= \alpha_{\text{bus}}^1 + \mathbf{x}'_{\text{bus}}\boldsymbol{\beta} + \epsilon_{\text{bus}} \end{aligned}$$

The preceding form is equivalent to the form

$$\begin{aligned} u_{\text{car}} &= \alpha_{\text{car}}^2 + \mathbf{x}'_{\text{car}}\boldsymbol{\beta} + \epsilon_{\text{car}} \\ u_{\text{bus}} &= \alpha_{\text{bus}}^2 + \mathbf{x}'_{\text{bus}}\boldsymbol{\beta} + \epsilon_{\text{bus}} \end{aligned}$$

as long as $\alpha_{\text{car}}^1 - \alpha_{\text{bus}}^1 = \alpha_{\text{car}}^2 - \alpha_{\text{bus}}^2$. There are infinitely many values that will result in the same choice probabilities. The two main effects are not estimable.

To avoid this problem, you must normalize one of the main effects to 0. If you set α_{car} to 0, then α_{bus} is interpreted as the average effect on the utility of bus relative to car. PROC BCHOICE automatically normalizes for you by setting one of the main effects to 0 if you specify the alternative specific main effects in the CLASS statement. If you want to designate which alternative to set to 0, use the REF= option in the CLASS statement.

Individual-Specific Effects

Attributes of the alternatives vary, so they are also called alternative-specific effects. However, demographic variables, such as age, gender, and race, do not change among alternatives for an individual. These are individual-specific effects. They can enter the model only if they are specified in ways that create differences in utilities among alternatives.

Suppose you want to examine the effect of an individual's age on choosing car or bus:

$$\begin{aligned} u_{\text{car}} &= \mathbf{x}'_{\text{car}}\boldsymbol{\beta} + \theta_{\text{car}}\text{Age} + \epsilon_{\text{car}} \\ u_{\text{bus}} &= \alpha_{\text{bus}} + \mathbf{x}'_{\text{bus}}\boldsymbol{\beta} + \theta_{\text{bus}}\text{Age} + \epsilon_{\text{bus}} \end{aligned}$$

Because only differences in utility matter, you cannot estimate both θ_{car} and θ_{bus} . You need to set one of them to 0. You can do this by interacting Age with the alternative-specific main effects, one of which has been normalized to 0:

$$\begin{aligned} u_{\text{car}} &= \mathbf{x}'_{\text{car}}\boldsymbol{\beta} + \epsilon_{\text{car}} \\ u_{\text{bus}} &= \alpha_{\text{bus}} + \mathbf{x}'_{\text{bus}}\boldsymbol{\beta} + \theta_{\text{bus}}\text{Age} + \epsilon_{\text{bus}} \end{aligned}$$

where θ_{bus} is interpreted as the differential effect of age on the utility of bus versus car.

You need to make sure that individual-specific effects have interacted with some alternative-specific effects, so that there is no identification problem. PROC BCHOICE checks whether there is any variable that does not change among alternatives; it issues a warning message if any such variable is found.

Scale Shift

The scale of utility is also irrelevant for choice models. The alternative that has the largest utility is the largest, regardless of how the utilities are scaled. Multiplying the utilities by a constant does not change the choice probabilities.

Scale Normalization for Logit Models

Consider the utility function

$$u_{ij}^1 = \mathbf{x}'_{ij}\boldsymbol{\beta} + \epsilon_{ij}^1 \text{ for all } j$$

where $\text{Var}(\epsilon_{ij}^1) = 1$. There is another utility function

$$u_{ij}^2 = \mathbf{x}'_{ij}(\boldsymbol{\beta} * \sigma) + \epsilon_{ij}^2 \text{ for all } j$$

where $\text{Var}(\epsilon_{ij}^2) = \sigma^2$. These two utility functions are equivalent. The coefficients $(\boldsymbol{\beta} * \sigma)$ reflect the effect of the observed variables relative to the standard deviation of the unobserved factors (Train 2009).

In both standard and nested logit models, the error variance is equal to $\pi^2/6$, which is around 1.6. So the coefficients are larger by a factor of $\sqrt{1.6}$ than the coefficients of a model that has an error variance of 1.

Scale Normalization for Probit Models

Parameters in logit models are automatically normalized because of the specific variance assumption on the error term. In probit models, the covariance matrix of the error is estimated together with the β coefficients; therefore, the normalization is not done automatically. The researcher has to be careful in interpreting the results.

For notational convenience, use a four-alternative model. The error vector is $\epsilon' = (\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4)$, which is assumed to have a normal distribution with 0 mean and a covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} & \sigma_{14} \\ \cdot & \sigma_{22} & \sigma_{23} & \sigma_{24} \\ \cdot & \cdot & \sigma_{33} & \sigma_{34} \\ \cdot & \cdot & \cdot & \sigma_{44} \end{pmatrix}$$

where the dots refer to the corresponding elements in the upper part of the symmetric matrix. There are $J(J+1)/2$ elements in the covariance matrix, but not all of them are estimable, because only the difference in the utilities matters.

A typical way of normalizing the covariance matrix is to reduce it by differencing the utilities with respect to the last alternative; some prefer the first alternative (Train 2009). You can choose which alternative to subtract by rearranging the data in the right order. Define the error differences as $(\tilde{\epsilon}_{14}, \tilde{\epsilon}_{24}, \tilde{\epsilon}_{34})$; the subscripts mean that the error difference is taken against the fourth one. The covariance matrix for the transformed new vector of error differences is of the form

$$\tilde{\Sigma} = \begin{pmatrix} \tilde{\sigma}_{11} & \tilde{\sigma}_{12} & \tilde{\sigma}_{13} \\ \cdot & \tilde{\sigma}_{22} & \tilde{\sigma}_{23} \\ \cdot & \cdot & \tilde{\sigma}_{33} \end{pmatrix}$$

where

$$\begin{aligned} \tilde{\sigma}_{11} &= \sigma_{11} + \sigma_{44} - 2\sigma_{14} \\ \tilde{\sigma}_{22} &= \sigma_{22} + \sigma_{44} - 2\sigma_{24} \\ \tilde{\sigma}_{33} &= \sigma_{33} + \sigma_{44} - 2\sigma_{34} \\ \tilde{\sigma}_{12} &= \sigma_{12} + \sigma_{44} - \sigma_{14} - \sigma_{24} \\ \tilde{\sigma}_{13} &= \sigma_{13} + \sigma_{44} - \sigma_{14} - \sigma_{34} \\ \tilde{\sigma}_{23} &= \sigma_{23} + \sigma_{44} - \sigma_{24} - \sigma_{34} \end{aligned}$$

Furthermore, the top left entry on the diagonal is set to 1:

$$\tilde{\tilde{\Sigma}} = \begin{pmatrix} 1 & \tilde{\sigma}_{12}/\tilde{\sigma}_{11} & \tilde{\sigma}_{13}/\tilde{\sigma}_{11} \\ \cdot & \tilde{\sigma}_{22}/\tilde{\sigma}_{11} & \tilde{\sigma}_{23}/\tilde{\sigma}_{11} \\ \cdot & \cdot & \tilde{\sigma}_{33}/\tilde{\sigma}_{11} \end{pmatrix}$$

PROC BCHOICE outputs estimates of $\tilde{\tilde{\Sigma}}$.

In the multinomial logit (MNL) model, each ϵ_{ij} is independently and identically distributed (iid) with the Type I extreme-value distribution, and the covariance matrix is $\pi^2/6I$. The normalized covariance matrix of the vector of error differences would be

$$\begin{pmatrix} 1 & 0.5 & .5 \\ \cdot & 1 & .5 \\ \cdot & \cdot & 1 \end{pmatrix}$$

This provides a method of model selection: you might want to fit a probit model first, check the normalized covariance matrix estimation to see how close it is to the preceding matrix, and then decide whether a simpler logit model would be appropriate.

Gamerman Algorithm

Discrete choice logit models fall in the framework of a generalized linear model (GLM) with a logit link. The Metropolis-Hastings sampling approach of Gamerman (1997) is well suited to this type of model.

In the GLM setting, the data are assumed to be independent with exponential family density

$$f(y_i|\theta_i) = \exp[(y_i\theta_i - b(\theta_i))/\phi_i]c(y_i, \phi_i)$$

The means that $\mu_i = E(y_i|\theta_i)$ are related to the canonical parameters θ_i via $\mu_i = b'(\theta_i)$ and to the regression coefficients via the link function

$$g(\mu_i) = \eta_i = \mathbf{X}_i\boldsymbol{\beta}$$

The maximum likelihood (ML) estimator in a GLM and the asymptotic variance are obtained by iterative application of weighted least squares (IWLS) to transformed observations. Following McCullagh and Nelder (1989), define the transformed response as

$$\tilde{y}_i(\boldsymbol{\beta}) = \eta_i + (y_i - \mu_i)g'(\mu_i)$$

and define the corresponding weights as

$$\mathbf{W}_i^{-1}(\boldsymbol{\beta}) = b''(\theta_i)[g'(\mu_i)]^2$$

Suppose a normal prior is specified on $\boldsymbol{\beta}$, $\pi(\boldsymbol{\beta}) = \mathbf{N}(\bar{\boldsymbol{\beta}}, \boldsymbol{\Omega}_\beta)$. The posterior density is as follows:

$$\begin{aligned} p(\boldsymbol{\beta}|\mathbf{Y}) &\propto p(\mathbf{Y}|\boldsymbol{\beta})\pi(\boldsymbol{\beta}) \\ &\propto \exp\left\{\sum_{i=1}^N \frac{y_i\theta_i - b(\theta_i)}{\phi_i} - \frac{1}{2}(\boldsymbol{\beta} - \bar{\boldsymbol{\beta}})' \boldsymbol{\Omega}_\beta^{-1}(\boldsymbol{\beta} - \bar{\boldsymbol{\beta}})\right\} \end{aligned}$$

Gamerman (1997) proposes that Metropolis-Hastings sampling be combined with iterative weighted least squares as follows:

1. Start with $\boldsymbol{\beta}^{(0)}$ and $t = 1$.
2. Sample $\boldsymbol{\beta}^*$ from the proposal density $\mathbf{N}(\mathbf{m}^{(t)}, \mathbf{C}^{(t)})$, where

$$\begin{aligned} \mathbf{m}^{(t)} &= \{\boldsymbol{\Omega}_\beta^{-1} + \mathbf{X}'\mathbf{W}(\boldsymbol{\beta}^{(t-1)})\mathbf{X}\}^{-1} \{\boldsymbol{\Omega}_\beta^{-1}\bar{\boldsymbol{\beta}} + \mathbf{X}'\mathbf{W}(\boldsymbol{\beta}^{(t-1)})\tilde{\mathbf{Y}}(\boldsymbol{\beta}^{(t-1)})\} \\ \mathbf{C}^{(t)} &= \{\boldsymbol{\Omega}_\beta^{-1} + \mathbf{X}'\mathbf{W}(\boldsymbol{\beta}^{(t-1)})\mathbf{X}\}^{-1} \end{aligned}$$

3. Accept β^* with probability

$$\alpha(\beta^{(t-1)}, \beta^*) = \min\left[1, \frac{p(\beta^*|\mathbf{Y})q(\beta^*, \beta^{(t-1)})}{p(\beta^{(t-1)}|\mathbf{Y})q(\beta^{(t-1)}, \beta^*)}\right]$$

where $p(\beta|\mathbf{Y})$ is the posterior density and $q(\beta^*, \beta^{(t-1)})$ and $q(\beta^{(t-1)}, \beta^*)$ are the transitional probabilities that are based on the proposal density $\mathbf{N}(\mathbf{m}^{(\cdot)}, \mathbf{C}^{(\cdot)})$. More specifically, $q(\beta^*, \beta^{(t-1)})$ is an $\mathbf{N}(\mathbf{m}^*, \mathbf{C}^*)$ density that is evaluated at $\beta^{(t-1)}$, whereas \mathbf{m}^* and \mathbf{C}^* have the same expression as $\mathbf{m}^{(t)}$ and $\mathbf{C}^{(t)}$ but depend on β^* instead of $\beta^{(t-1)}$. If β^* is not accepted, the chain stays with $\beta^{(t-1)}$.

4. Set $t = t + 1$ and return to step 1.

You can extend this methodology to logit models that have random effects. If there are random effects, the link function is extended to

$$g(\mu_i) = \eta_i = \mathbf{X}_i\beta + \mathbf{Z}_i\gamma_i$$

where the random effects are assumed to have a normal distribution, $\pi(\gamma_i) = \mathbf{N}(\mathbf{0}, \mathbf{\Omega}_\gamma)$, and $\pi(\mathbf{\Omega}_\gamma) =$ inverse Wishart(ν_0, \mathbf{V}_0). The posterior density is

$$p(\beta, \gamma_1, \dots, \gamma_N, \mathbf{\Omega}_\gamma | \mathbf{Y}) \propto p(\mathbf{Y} | \beta, \gamma_1, \dots, \gamma_N, \mathbf{\Omega}_\gamma) \pi(\beta) \prod_{i=1}^N \pi(\gamma_i) \pi(\mathbf{\Omega}_\gamma)$$

The parameters are divided into blocks, $\beta, \gamma_1, \dots, \gamma_N$, and $\mathbf{\Omega}_\gamma$. For the fixed-effects β block, the conditional posterior has the same form, but the link changes to include $\mathbf{Z}_i\gamma_i, i = 1, \dots, N$, which are taken as known constants (offsets) at each iteration. The only change that is needed is to replace the transformed response $\tilde{y}_i(\beta^{(t-1)})$ with $\tilde{y}_i(\beta^{(t-1)}) - \mathbf{Z}_i\gamma_i$ in step 2 of the previous Gamerman procedure.

For the random-effects γ_i block, the same Metropolis-Hastings sampling with the least square proposal can apply. The conditional posterior is

$$p(\gamma_i | \mathbf{Y}, \beta, \mathbf{\Omega}_\gamma) \propto \exp\left\{\frac{y_i\theta_i - b(\theta_i)}{\phi_i} - \frac{1}{2}\gamma_i'\mathbf{\Omega}_\gamma^{-1}\gamma_i\right\}$$

The transformed response is now $\tilde{y}_i(\gamma_i^{(t-1)})$, and the proposal density is $\mathbf{N}(\mathbf{m}_i^{(t)}, \mathbf{C}_i^{(t)})$, where

$$\begin{aligned} \mathbf{m}_i^{(t)} &= \{\mathbf{\Omega}_\gamma^{-1} + \mathbf{Z}'\mathbf{W}(\gamma_i^{(t-1)})\mathbf{Z}\}^{-1}\mathbf{Z}'\mathbf{W}(\gamma_i^{(t-1)})\{\tilde{\mathbf{Y}}(\gamma_i^{(t-1)}) - \mathbf{X}_i\beta\} \\ \mathbf{C}_i^{(t)} &= \{\mathbf{\Omega}_\gamma^{-1} + \mathbf{Z}'\mathbf{W}(\gamma_i^{(t-1)})\mathbf{X}_i\}^{-1} \end{aligned}$$

Finally, for the covariance matrix $\mathbf{\Omega}_\gamma$ block, direct sampling from an inverse Wishart($\nu_0 + N, \mathbf{V}_0 + S$) is used, where $S = \sum_{i=1}^N \gamma_i\gamma_i'/N$.

The chain is initialized with random effects set to 0 and the covariance set to the identity matrix. Updating is done first for the fixed effects, β , as a block to position the chain in the correct region of the parameter space. Then the random effects are updated, and finally the covariance of the random effects is updated. For more information about this algorithm, see Gamerman (1997).

Autocall Macros for Postprocessing

Although PROC BCHOICE provides a number of convergence diagnostic tests and posterior summary statistics, it performs the calculations only for the default tests and statistics or only if you specify the options in advance. If you want to analyze the posterior draws of unmonitored parameters or functions of the parameters that are calculated in later DATA step calls, you can use the autocall macros that are listed in Table 27.6.

Table 27.6 Postprocessing Autocall Macros

Macro	Description
%ESS	Effective sample sizes
%GEWEKE*	Geweke diagnostic
%HEIDEL*	Heidelberger-Welch diagnostic
%MCSE	Monte Carlo standard errors
%POSTACF	Autocorrelation
%POSTCOR	Correlation matrix
%POSTCOV	Covariance matrix
%POSTINT	Equal-tail and HPD intervals
%POSTSUM	Summary statistics
%RAFTERY	Raftery diagnostic

*The %GEWEKE and %HEIDEL macros use a different optimization routine than that used in PROC BCHOICE. As a result, there might be numerical differences in some cases, especially when the sample size is small.

Table 27.7 lists options that are shared by all postprocessing autocall macros. For macro-specific options, see Table 27.8.

Table 27.7 Shared Options

Option	Description
DATA=SAS-data-set	Names the input data set that contains posterior samples
OUT=SAS-data-set	Specifies a name for the output SAS data set to contain the results
PRINT=YES NO	Displays the results (the default is YES)
VAR=variable-list	Specifies the variables on which to perform the calculation

Suppose the data set that contains posterior samples is called `Post` and the variables of interest are defined in the macro variable `&PARMS`. The following statements call the %ESS macro and calculate the effective sample sizes for each variable:

```
%ESS(data=Post, var=Alpha Beta U_1-U_17)
```

By default, the ESS estimates are displayed. You can choose not to display the result and instead use the following statement to save the output to a data set:

```
%ESS(data=Post, var=&parms, print=NO, out=eout)
```

Some of the macros can take additional options, which are listed in Table 27.8.

Table 27.8 Macro-Specific Options

Macro	Option	Description
%ESS	AUTOCORLAG= <i>numeric</i>	Specifies the maximum number of autocorrelation lags used in computing the ESS estimates. By default, AUTOCORLAG=MIN(500, NOBS/4), where NOBS is the sample size of the input data set.
	HIST=YES NO	Displays a histogram of all ESS estimates. By default, HIST=NO.
%GEWEKE	FRAC1= <i>numeric</i>	Specifies the earlier portion of the Markov chain used in the test. By default, FRAC1=0.1.
	FRAC2= <i>numeric</i>	Specifies the latter portion of the Markov chain used in the test. By default, FRAC2=0.5.
%HEIDEL	SALPHA= <i>numeric</i>	Specifies the α level for the stationarity test. By default, SALPHA=0.05.
	HALPHA= <i>numeric</i>	Specifies the α level for the halfwidth test. By default, HALPHA=0.05.
	EPS= <i>numeric</i>	Specifies a small positive number ϵ such that if the halfwidth is less than ϵ times the sample mean of the remaining iterations, the halfwidth test is passed. By default, EPS=0.1.
%MCSE	AUTOCORLAG= <i>numeric</i>	Specifies the maximum number of autocorrelation lags used in computing the Monte Carlo standard error estimates. By default, AUTOCORLAG=MIN(500, NOBS/4), where NOBS is the sample size of the input data set.
%POSTACF	LAGS=%str(<i>numeric-list</i>)	Specifies which autocorrelation lags to calculate. The default values are 1, 5, 10, and 50.
%POSTINT	ALPHA= <i>value</i>	Specifies the α level ($0 < \alpha < 1$) for the interval estimates. By default, ALPHA=0.05.
%RAFTERY	Q= <i>numeric</i>	Specifies the order of the quantile of interest. By default, Q=0.025.
	R= <i>numeric</i>	Specifies the margin of error for measuring the accuracy of estimation of the quantile. By default, R=0.005.
	S= <i>numeric</i>	Specifies the probability of attaining the accuracy of the estimation of the quantile. By default, S=0.95.
	EPS= <i>numeric</i>	Specifies the tolerance level for the stationary test. By default, EPS=0.001.

For example, the following statement calculates and displays autocorrelation at lags 1, 6, 11, 50, and 100. Note that the lags in the *numeric-list* must be separated by commas.

```
%PostACF(data=Post, var=&parms, lags=%str(1 to 15 by 5, 50, 100))
```

Regenerating Diagnostics Plots

By default, PROC BCHOICE generates three plots: the trace plot, the autocorrelation plot, and the kernel density plot. Unless ODS Graphics is enabled before the procedure is called, it is hard to generate the

same graph afterward. Directly using the `Stat.BCHOICE.Graphics.TraceAutocorrDensity` template is not feasible. The easiest way to regenerate the same graph is to use the `%TADPlot` autocall macro. The `%TADPlot` macro requires you to specify an input data set (which is the output data set from a previous PROC BCHOICE call) and a list of variables that you want to plot.

Suppose that the output data set `Postsamp` contains posterior draws for the regression coefficients of `Mode1`, `Mode2`, and `Mode3`. If you want to examine these parameters graphically, you can use the following statements to regenerate the graphs:

```
ods graphics on;
%tadplot(data=Postsamp, var=Mode1 Mode2 Mode3)
ods graphics off;
```

Displayed Output

This section describes the displayed output from PROC BCHOICE. For a quick reference of all ODS table names, see the section “ODS Table Names” on page 1053. ODS tables are arranged in four groups, listed in the following sections: “Model- and Data-Related ODS Tables” on page 1050, “Sampling-Related ODS Tables” on page 1051, “ODS Tables Related to Posterior Statistics” on page 1051, and “ODS Tables Related to Convergence Diagnostics” on page 1052.

Model- and Data-Related ODS Tables

Model Information

The “Model Information” table (ODS table name `ModelInfo`) displays basic information about the model, such as the name of the input data set, response variable, model type, sampling algorithm, burn-in size, simulation size, thinning, and random number seed. The “Model Information” table is produced irrespective of the sampling technique.

The “Model Information” table also displays the random number seed that is used to initialize the random number generators. If you repeat the analysis and specify this seed value in the `SEED=` option in the PROC BCHOICE statement, you get an identical stream of random numbers. This table is displayed by default.

Class Level Information

The “Class Level Information” table (ODS table name `ClassLevels`) lists the levels of every variable that is specified in both the `CLASS` statement and the `MODEL` statement. You should check this information to make sure that the data are correct. You can adjust the order of the `CLASS` variable levels by specifying the `ORDER=` option in the `CLASS` statement. You can suppress the “Class Level Information” table completely or partially by specifying the `NOCLPRINT` option in the PROC BCHOICE statement.

Choice Sets Summary

The “Choice Sets Summary” table (ODS table name `ChoiceSummary`) is displayed by default and should be used to check the data entry. It shows the numbers of choice sets, the number of alternatives, and the number of chosen and unchosen alternatives. You should always check this summary table to ensure that the data are arrayed correctly. This table is displayed by default.

Number of Observations

The “NObs” table (ODS table name NOBS) shows the number of observations in the data set and the number of observations that are used in the analysis. By default, observations that have missing values are not used. This table is displayed by default.

Sampling-Related ODS Tables**Parameters Initial Value**

The “Parameters Initial” table (ODS table name ParametersInit) shows the initial value of each fixed regression coefficient. This table is not displayed by default, but you can request it by specifying the option `INIT=PINIT` in the MODEL statement.

Prior for Fixed Effects

The “Prior for Fixed Effects” table (ODS table name CoeffPrior) shows the prior distribution for the regression coefficients. The table is displayed by default for a model with fixed effects.

ODS Tables Related to Posterior Statistics

PROC BCHOICE calculates some essential posterior statistics and outputs them to a number of ODS tables. Some of the ODS tables are produced by default, and you can request others by specifying an option in the PROC BCHOICE statement. For more information about the calculations, see the section “[Summary Statistics](#)” on page 155.

Summary and Interval Statistics

The “Posterior Summaries and Intervals” table (ODS table name PostSumInt) contains basic summary and interval statistics for each parameter. The table lists the number of posterior samples, the posterior mean and standard deviation estimates, and the highest posterior density (HPD) interval estimates. This table is displayed by default.

Summary Statistics

The “Posterior Summaries” table (ODS table name PostSummaries) contains basic statistics for each parameter. The table lists the number of posterior samples, the posterior mean and standard deviation estimates, and the percentile estimates. This table is not displayed by default, but you can request it by specifying the option `STATISTICS=SUM`.

Correlation Matrix

The “Posterior Correlation Matrix” table (ODS table name Corr) contains the posterior correlation of model parameters. This table is not displayed by default, but you can request it by specifying the option `STATISTICS=CORR`.

Covariance Matrix

The “Posterior Covariance Matrix” table (ODS table name Cov) contains the posterior covariance of model parameters. This table is not displayed by default, but you can request it by specifying the option `STATISTICS=COV`.

Deviance Information Criterion

The “Deviance Information Criterion” table (ODS table name DIC) contains the deviance information criterion (DIC) of the model. This table is not displayed by default, but you can request it by specifying the option `DIC`. For more information about the calculations, see the section “[Deviance Information Criterion \(DIC\)](#)” on page 157.

Interval Statistics

The “Posterior Intervals” table (ODS table name PostIntervals) contains the equal-tail and highest posterior density (HPD) interval estimates for each parameter. The default α value is 0.05, and you can change it to other levels by using the `STATISTICS=` option. This table is not displayed by default, but you can request it by specifying the option `STATISTICS=INT`.

ODS Tables Related to Convergence Diagnostics

PROC BCHOICE provides convergence diagnostic tests that check for Markov chain convergence. The procedure produces a number of ODS tables that you can request and save individually. For information about calculation, see the section “[Statistical Diagnostic Tests](#)” on page 146.

Autocorrelation

The “Autocorrelations” table (ODS table name AUTOCORR) contains the first-order autocorrelations of the posterior samples for each parameter. The Parameter column states the name of the parameter. By default, PROC BCHOICE displays lag 1, 5, 10, and 50 estimates of the autocorrelations. You can request different autocorrelations by using the `DIAGNOSTICS=AUTOCORR(LAGS=)` option. This table is displayed by default.

Effective Sample Size

The “Effective Sample Sizes” table (ODS table name ESS) calculates the effective sample size of each parameter. For more information, see the section “[Effective Sample Size](#)” on page 154. This table is displayed by default.

Monte Carlo Standard Errors

The “Monte Carlo Standard Errors” table (ODS table name MCSE) calculates the standard errors of the posterior mean estimate. For more information, see the section “[Standard Error of the Mean Estimate](#)” on page 155. This table is not displayed by default, but you can request it by specifying the option `DIAGNOSTICS=MCSE`.

Geweke Diagnostics

The “Geweke Diagnostics” table (ODS table name Geweke) lists the results of the Geweke diagnostic test. For more information, see the section “[Geweke Diagnostics](#)” on page 148. This table is not displayed by default, but you can request it by specifying `DIAGNOSTICS=GEWEKE`.

Heidelberger-Welch Diagnostics

The “Heidelberger-Welch Diagnostics” table (ODS table name Heidelberger) lists the results of the Heidelberger-Welch diagnostic test. The test consists of two parts: a stationary test and a half-width test. For more information, see the section “[Heidelberger and Welch Diagnostics](#)” on page 150. This table is not displayed by default, but you can request it by specifying `DIAGNOSTICS=HEIDEL`.

Raftery-Lewis Diagnostics

The “Raftery-Lewis Diagnostics” table (ODS table name Raftery) lists the results of the Raftery-Lewis diagnostic test. For more information, see the section “[Raftery and Lewis Diagnostics](#)” on page 151. This table is not displayed by default, but you can request it by specifying `DIAGNOSTICS=RAFTERY`.

ODS Table Names

PROC BCHOICE assigns a name to each table that it creates. You can use these names to refer to the tables when you use the output delivery system (ODS) to select tables and create output data sets. These names are listed in [Table 27.9](#). For more information about ODS, see Chapter 21, “[Statistical Graphics Using ODS](#).”

Table 27.9 ODS Tables Produced in PROC BCHOICE

ODS Table Name	Description	Statement or Option
AutoCorr	Autocorrelation statistics for each parameter	DIAG=AUTOCORR
ChoiceSummary	Choice sets summary	Default
ClassLevels	Class level information	Default
CoeffPrior	Prior information for fixed effects	Default
Corr	Correlation matrix of the posterior samples	STATS=CORR
Cov	Covariance matrix of the posterior samples	STATS=COV
DIC	Deviance information criterion	DIC
ESS	Effective sample size for each parameter	Default
Geweke	Geweke diagnostics for each parameter	DIAG=GEWEKE
Heidelberger	Heidelberger-Welch diagnostics for each parameter	DIAG=HEIDEL
MCSE	Monte Carlo standard error for each parameter	STATS=MCSE
ModelInfo	Model information	Default
NObs	Number of observations	Default
ParametersInit	Parameter initial values	INIT=PINIT
PostSumInt	Brief posterior statistics for each parameter, including sample size, mean, standard deviation, and HPD intervals	Default
PostIntervals	Equal-tail and HPD intervals for each parameter	STATS=INT
PostSummaries	Basic posterior statistics for each parameter, including sample size, mean, standard deviation, and percentiles	STATS=SUM
Raftery	Raftery-Lewis diagnostics for each parameter	DIAG=RAFTERY

ODS Graphics

You can refer by name to every graph that is produced through ODS Graphics. The names of the graphs that PROC BCHOICE generates are listed in Table 27.10.

Table 27.10 Graphs Produced by PROC BCHOICE

ODS Graph Name	Plot Description	Statement and Option
ADPanel	Autocorrelation function and density panel	PLOTS=(AUTOCORR DENSITY)
AutocorrPanel	Autocorrelation function panel	PLOTS=AUTOCORR
AutocorrPlot	Autocorrelation function plot	PLOTS(UNPACK)=AUTOCORR
DensityPanel	Density panel	PLOTS=DENSITY
DensityPlot	Density plot	PLOTS(UNPACK)=DENSITY
TAPanel	Trace and autocorrelation function panel	PLOTS=(TRACE AUTOCORR)
TADPanel	Trace, density, and autocorrelation function panel	PLOTS=(TRACE AUTOCORR DENSITY)
TDPanel	Trace and density panel	PLOTS=(TRACE DENSITY)
TracePanel	Trace panel	PLOTS=TRACE
TracePlot	Trace plot	PLOTS(UNPACK)=TRACE

Examples: BCHOICE Procedure

Example 27.1: Alternative-Specific and Individual-Specific Effects

In many situations, a choice model that includes characteristics of both the alternatives and the individuals is needed for investigating consumer choice.

Consider an example of travel demand. People are asked to choose among travel by auto, plane, or public transit (bus or train). The following SAS statements create the data set `Travel`. The variables `AutoTime`, `PlanTime`, and `TranTime` represent the total travel time that is required to get to a destination by using auto, plane, or public transit, respectively. The variable `Age` represents the age of each individual who is surveyed, and the variable `Chosen` contains each individual's choice of travel mode.

```
data Travel;
  input AutoTime PlanTime TranTime Age Chosen $;
  AgeCtr=Age-34;
  datalines;
10.0 4.5 10.5 32 Plane
5.5 4.0 7.5 13 Auto
```

```

4.5 6.0 5.5 41 Transit
3.5 2.0 5.0 41 Transit
1.5 4.5 4.0 47 Auto
10.5 3.0 10.5 24 Plane
7.0 3.0 9.0 27 Auto
9.0 3.5 9.0 21 Plane
4.0 5.0 5.5 23 Auto
22.0 4.5 22.5 30 Plane
7.5 5.5 10.0 58 Plane
11.5 3.5 11.5 36 Transit
3.5 4.5 4.5 43 Auto
12.0 3.0 11.0 33 Plane
18.0 5.5 20.0 30 Plane
23.0 5.5 21.5 28 Plane
4.0 3.0 4.5 44 Plane
5.0 2.5 7.0 37 Transit
3.5 2.0 7.0 45 Auto
12.5 3.5 15.5 35 Plane
1.5 4.0 2.0 22 Auto
;

```

In this example, the AutoTime, PlanTime, and TranTime variables apply to the alternatives, whereas Age is a characteristic of the individuals. AgeCtr, a centered version of Age, is created by subtracting the sample's mean age from each individual's age. To study how the choice depends on both the travel time and age of the individuals, you need to incorporate both types of variables.

Before you invoke PROC BCHOICE to fit a choice logit model, you must arrange your data in such a way that there is one observation for each combination of individual and alternative. In this example, let Subject identify the individuals, let TravTime represent the travel time for each mode of transportation, and let Choice have the value 1 if the alternative is chosen and 0 otherwise. The following SAS statements rearrange the data set Travel into a new data set, Travel2, and display the first nine observations:

```

data Travel2(keep=Subject Mode TravTime Age AgeCtr Choice);
  array Times[3] AutoTime PlanTime TranTime;
  array Allmodes[3] $ _temporary_ ('Auto' 'Plane' 'Transit');
  set Travel;
  Subject = _n_;
  do i = 1 to 3;
    Mode = Allmodes[i];
    TravTime = Times[i];
    Choice = (Chosen eq Mode);
    output;
  end;
run;
proc print data=Travel2 (obs=20);
  by Subject;
  id Subject;
run;

```

The data for the first nine observations is shown in [Output 27.1.1](#).

Output 27.1.1 Data for the First Nine Observations

Subject	Age	Age Ctr	Mode	Trav Time	Choice
1	32	-2	Auto	10.0	0
	32	-2	Plane	4.5	1
	32	-2	Transit	10.5	0
Subject	Age	Age Ctr	Mode	Trav Time	Choice
2	13	-21	Auto	5.5	1
	13	-21	Plane	4.0	0
	13	-21	Transit	7.5	0
Subject	Age	Age Ctr	Mode	Trav Time	Choice
3	41	7	Auto	4.5	0
	41	7	Plane	6.0	0
	41	7	Transit	5.5	1

Notice that each subject in the data set `Travel` corresponds to a block of three observations in the data set `Travel2`, one for each travel alternative. The response variable `Choice` indicates the chosen alternative by the value 1 and the unchosen alternative by the value 0; exactly one alternative is chosen. The following SAS statements invoke PROC BCHOICE to fit the choice logit model:

```
proc bchoice data=Travel2 seed=124;
  class Mode Subject / param=ref order=data;
  model Choice = Mode TravTime / choiceset=(Subject);
run;
```

The “Choice Sets Summary” table shows that there are 21 choice sets and that each consists of three alternatives and one chosen alternative (each subject chooses one out of the three travel modes). It seems that the data are arrayed correctly.

Output 27.1.2 Choice Sets Summary

The BCHOICE Procedure				
Choice Sets Summary				
Pattern	Choice Sets	Total Alternatives	Chosen Alternatives	Not Chosen
1	21	3	1	2

Summary statistics are shown in [Output 27.1.3](#).

Output 27.1.3 PROC BCHOICE Posterior Summary Statistics

Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Mode Auto	5000	-0.1678	0.7440	-1.7017	1.2396
Mode Plane	5000	-1.8794	1.2683	-4.6055	0.3801
TravTime	5000	-0.5695	0.2047	-0.9943	-0.2328

When Transit is the reference mode (normalized to 0), the part-worth (posterior mean) of Auto, which is negative, might reflect that driving is more inconvenient than traveling by bus or train, and the negative part-worth of Plane might reflect that traveling by plane is more expensive than traveling by bus or train. However, neither is statistically significant, because both the 95% HPD intervals have 0 in them. The posterior mean of TravTime is negative, and the estimate is significant. These results make sense because having to spend more time en route is often unfavorable.

To study the relationship between the choice of transportation and the age of people who make the choice, you need to create an interaction between AgeCtr and Mode. AgeCtr is not estimable by itself, because it is the same throughout a choice set for an individual. The following statements request the interaction between AgeCtr and Mode:

```
proc bchoice data=Travel2 seed=124;
  class Mode Subject / param=ref order=data;
  model Choice = Mode Mode*AgeCtr TravTime / choiceset=(Subject);
run;
```

Output 27.1.4 PROC BCHOICE Posterior Summary Statistics

The BCHOICE Procedure					
Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Mode Auto	5000	-0.2634	0.7883	-1.8072	1.1395
Mode Plane	5000	-2.8210	1.5370	-5.9228	-0.0686
AgeCtr*Mode Auto	5000	-0.0986	0.0678	-0.2182	0.0350
AgeCtr*Mode Plane	5000	0.0251	0.0775	-0.1268	0.1618
TravTime	5000	-0.7608	0.2564	-1.2943	-0.3473

The parameter estimate for Mode Auto reflects the part-worth of Auto for an individual of mean age (34 years old), whereas the parameter estimate for Mode Plane is the part-worth of Plane for an individual of mean age. There are two interaction effects: the first corresponds to the effect of a one-unit change in age on the probability of choosing Auto over Transit, and the second corresponds to the effect of a one-unit change in age on the probability of choosing Plane over Transit.

Example 27.2: Nested Logit Modeling

The standard logit model imposes the restriction of the independence from irrelevant alternatives (IIA) property, which implies proportional substitution across alternatives. When the IIA assumption does not hold, models with more flexibility are needed.

One of the most widely used models is called the nested logit. In a nested logit model, all the alternatives in a choice set can be partitioned into nests in such a way that the following conditions are true:

- The ratio of any two alternatives that are in the same nest is independent of the existence of all other alternatives. Hence, the IIA assumption holds within each nest.
- The ratio of any two alternatives that are in different nests is not independent of the existence of other alternatives in the two nests. Hence, the IIA assumption does not hold between different nests.

For more information about nested logit models, see the section “[Nested Logit](#)” on page 1036.

In the previous example of travel demand data, people are asked to choose among travel by auto, plane, or public transit. It seems that auto and public transit are more similar to each other than either of them are to plane, because the probability of choosing auto and public transit might rise by about the same proportion whenever the option of taking a plane is unavailable. A nested logit model that places auto and public transit in one nest and plane in another nest might seem more reasonable than the standard logit model.

You specify a nested logit model by using the `TYPE=NLOGIT` option, and you allocate the nests by using the `NEST=` option. In the following code, `NEST=(1 2 1)` implies that travel alternatives 1 (auto) and 3 (public transit) are in the first nest and travel alternative 2 (plane) is in the second nest. There are two nests in total.

The deviance information criterion (DIC) is used to select the model. The “Deviance Information Criterion” table (ODS table name DIC) contains the DIC of the model. The table is not displayed by default, but you can request it by specifying the `DIC` option in the PROC BCHOICE statement. For more information about the calculations, see the section “[Deviance Information Criterion \(DIC\)](#)” on page 157. The `DIC` option requests the calculation of DIC.

```
proc bchoice data=Travel2 seed=124 dic;
  class Mode Subject / param=ref order=data;
  model Choice = Mode TravTime / choiceset=(Subject) type=nlogit nest=(1 2 1);
run;
```

Output 27.2.1 PROC BCHOICE Posterior Summary Statistics and DIC

The BCHOICE Procedure					
Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Mode Auto	5000	-0.2283	0.6911	-1.6005	1.1456
Mode Plane	5000	-2.0001	1.1697	-4.1169	0.3449
TravTime	5000	-0.5568	0.1979	-0.9243	-0.1872
Lambda 1	5000	0.8867	0.3557	0.2407	1.5001

Output 27.2.1 *continued*

Deviance Information Criterion	
Dbar (Posterior Mean of Deviance)	32.925
Dmean (Deviance Evaluated at Posterior Mean)	30.289
pD (Effective Number of Parameters)	2.636
DIC (Smaller is Better)	35.561

There are two alternatives in the first nest and one alternative in the second nest. A nest that has only one alternative is degenerate; therefore its λ is not estimable. This explains why there is only one λ estimate in the output.

The following statements revisit the standard logit model (the default type) that is discussed in the previous example and request that the calculation of DIC be displayed:

```
proc bchoice data=Travel2 seed=124 dic;
  class Mode Subject / param=ref order=data;
  model Choice = Mode TravTime / choiceset=(Subject);
run;
```

Output 27.2.2 PROC BCHOICE Posterior Summary Statistics and DIC

The BCHOICE Procedure					
Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Mode Auto	5000	-0.1678	0.7440	-1.7017	1.2396
Mode Plane	5000	-1.8794	1.2683	-4.6055	0.3801
TravTime	5000	-0.5695	0.2047	-0.9943	-0.2328

Deviance Information Criterion	
Dbar (Posterior Mean of Deviance)	33.217
Dmean (Deviance Evaluated at Posterior Mean)	30.440
pD (Effective Number of Parameters)	2.777
DIC (Smaller is Better)	35.994

The DIC values are very close. A smaller DIC value indicates a better fit to the data; hence, the nested logit model might be a little better.

Example 27.3: Probit Modeling

A more general model is the probit model, which is derived under the assumption of jointly normal random utility components, where the error vector has a multivariate normal distribution with a mean vector of $\mathbf{0}$ and a covariance matrix Σ . For a full covariance matrix, PROC BCHOICE can accommodate any pattern of correlation and heteroscedasticity. Thus, this model fits a very general error structure. For more information about probit models, see the section “Probit” on page 1038.

In Train (2009), a project team collected the following data from commuters about four available travel modes for their trips to work: car alone, carpool, bus, and railway. The time and cost of travel for each mode were determined for each commuter, based on the location of the commuter’s home and workplace.

```
data Commuter;
  input Subject Mode Choice Cost Time @@;
  datalines;
1 1 1 1.51 18.5 1 2 0 2.34 26.34 1 3 0 1.8 20.87 1 4 0 2.36 30.03 2 1 0
6.06 31.31 2 2 0 2.9 34.26 2 3 0 2.24 67.18 2 4 1 1.86 60.29 3 1 1 5.79
22.55 3 2 0 2.14 23.26 3 3 0 2.58 63.31 3 4 0 2.75 49.17 4 1 1 1.87
26.09 4 2 0 2.57 29.9 4 3 0 1.9 19.75 4 4 0 2.27 13.47 5 1 1 2.5 4.7 5 2
0 1.72 12.41 5 3 0 2.69 43.09 5 4 0 2.97 39.74 6 1 1 4.73 3.07 6 2 0
0.62 9.22 6 3 0 1.85 12.83 6 4 0 2.31 43.54 7 1 1 4.73 13.14 7 2 0 0.6
17.77 7 3 0 2.43 54.09 7 4 0 2 42.22 8 1 1 5.35 52.9 8 2 0 2.91 48.78 8
3 0 2.61 69.16 8 4 0 2.78 53.25 9 1 0 4.41 61.06 9 2 0 1.59 62.13 9 3 1

... more lines ...

450 1 1 4.59 29.44 450 2 0 2.89 33.73 450 3 0 1.9 66.12 450 4 0 1.79
39.84 451 1 1 3.24 16.35 451 2 0 1.21 18.98 451 3 0 1.75 23.39 451 4 0
2.02 43.3 452 1 0 6.93 65.42 452 2 0 1.17 60.48 452 3 1 2.46 52.4 452 4
0 2.61 48.37 453 1 0 6.53 59.57 453 2 1 1.41 55.14 453 3 0 2.21 67.82
453 4 0 1.86 73.45
;
proc print data=Commuter(obs=20);
run;
```

The variable Mode has the value 1 for car alone, 2 for carpool, 3 for bus, and 4 for railway. The variable Choice is the response variable that represents the decision among the four travel modes for each commuter. The data for the first five commuters are shown in [Output 27.3.1](#).

Output 27.3.1 Data for the First Five Commuters

Obs	Subject	Mode	Choice	Cost	Time
1	1	1	1	1.51	18.50
2	1	2	0	2.34	26.34
3	1	3	0	1.80	20.87
4	1	4	0	2.36	30.03
5	2	1	0	6.06	31.31
6	2	2	0	2.90	34.26
7	2	3	0	2.24	67.18
8	2	4	1	1.86	60.29
9	3	1	1	5.79	22.55
10	3	2	0	2.14	23.26
11	3	3	0	2.58	63.31
12	3	4	0	2.75	49.17
13	4	1	1	1.87	26.09
14	4	2	0	2.57	29.90
15	4	3	0	1.90	19.75
16	4	4	0	2.27	13.47
17	5	1	1	2.50	4.70
18	5	2	0	1.72	12.41
19	5	3	0	2.69	43.09
20	5	4	0	2.97	39.74

The following statements fit a probit model by specifying `TYPE=PROBIT`. The `BCHOICE` procedure's implementation of the Gibbs sampler for the probit model exhibits a higher autocorrelation than that for the logit model. High autocorrelation is created by introducing the latent variable via data augmentation because of the dependence between the latent variable and the regression parameters. You might want to control the thinning rate of the simulation. For example, `THIN=10` keeps every 10th sample in the simulation and discards the rest.

```
proc bchoice data=Commuter outpost=Commupostsamp thin=10 nmc=50000 seed=123;
  class Mode(ref='1') Subject;
  model Choice = Cost Time Mode / choiceset=(Subject) type=probit;
run;
```

Output 27.3.2 shows the summary statistics for the part-worth (β) of each of the attributes (Cost, Time, Mode 2, Mode 3, and Mode 4) and the covariance of the error difference vector ($\tilde{\Sigma}$), which is displayed by parameters labeled “Sigma 1 1,” “Sigma 2 1,” and so on.

Output 27.3.2 Posterior Summary Statistics

The BCHOICE Procedure					
Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Cost	5000	-0.4641	0.0736	-0.6152	-0.3299
Time	5000	-0.0494	0.00567	-0.0603	-0.0377
Mode 2	5000	-3.3847	0.6746	-4.8022	-2.1983
Mode 3	5000	-2.0056	0.2853	-2.5709	-1.4590
Mode 4	5000	-1.6277	0.2085	-2.0489	-1.2382
Sigma 1 1	5000	1.0000	0	1.0000	1.0000
Sigma 2 1	5000	1.1695	0.5505	0.0179	2.2003
Sigma 2 2	5000	4.3809	1.9423	1.3685	8.3868
Sigma 3 1	5000	0.5700	0.2203	0.1429	1.0094
Sigma 3 2	5000	1.6328	0.9052	-0.00171	3.6863
Sigma 3 3	5000	1.3917	0.4795	0.5699	2.3242

It is well known that an identification problem exists in probit models, because location and scale transformations do not change the choices that are made. The solution to the location shift is differencing with respect to the last alternative in each choice set. (See the section “Probit” on page 1038.) After that, a scale shift problem remains, because the parameters $(c\beta, c^2\tilde{\Sigma})$ for any constant $c > 0$ are equivalent to $(\beta, \tilde{\Sigma})$. A solution to the scaling problem is to normalize the parameters with respect to one of the diagonal elements of the covariance of the error difference vector, $\tilde{\Sigma}$. PROC BCHOICE reports $(\beta/\sqrt{\sigma_{11}}, \tilde{\Sigma}/\sigma_{11})$ at each draw, where σ_{11} is the first diagonal entry of $\tilde{\Sigma}$. This explains why “Sigma 1 1” is always 1 in [Output 27.3.2](#).

By the IIA property in a logit model, it is assumed that all alternatives are independent and have the same variance. Therefore, the normalized covariance matrix after differencing with respect to one of the alternatives is of the form

$$\tilde{\Sigma} = \begin{pmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{pmatrix}$$

Obviously, this matrix is quite different from the estimated normalized covariance matrix for this data set. Fitting a standard logit model would be inappropriate.

Example 27.4: A Random-Effects-Only Logit Model

This example revisits the trash can study that is described earlier in this chapter in the getting-started section “A Logit Model with Random Effects” on page 1011.

If you want to create a random-effects-only model using the random walk Metropolis sampling as suggested in Rossi, Allenby, and McCulloch (2005), you can add the `ALG=RWM` option to the `PROC BCHOICE` statement to specify the random walk Metropolis sampling algorithm, add the `REMEAN` option to the `RANDOM` statement to request estimation of the nonzero mean of the random effects, and not specify any fixed effects in the `MODEL` statement as follows:

```
proc bchoice data=Trashcan outpost=Postsamp seed=123 nmc=40000 thin=4
      alg=rwm nthreads=4;
  class ID Task;
  model Choice = / choicest=(ID Task);
  random Touchless Steel AutoBag Price80 / sub=ID remean
      monitor=(1 to 5) type=un;
run;
```

The `NTHREADS` option in the `PROC BCHOICE` statement specifies the number of threads to be used for analytic computations and sampling. Using four threads at the same time enhances the efficiency and reduces the run time by about four times. If you do not specify the `NTHREADS` option, the default number is 1. The maximum number of threads should not exceed the total number of CPUs on the host where the analytic computations execute.

Summary statistics for the mean of the random coefficients ($\bar{\gamma}$), the covariance of the random coefficients (Ω_{γ}), and the random coefficients (γ_i) for the first five individuals are shown in [Output 27.4.1](#).

Output 27.4.1 Posterior Summary Statistics

The BCHOICE Procedure						
Posterior Summaries and Intervals						
Parameter	Subject	N	Standard		95% HPD Interval	
			Mean	Deviation		
REMean Touchless		10000	1.7124	0.2660	1.1859	2.2359
REMean Steel		10000	1.0504	0.2610	0.5677	1.5814
REMean AutoBag		10000	2.1839	0.3497	1.5420	2.9100
REMean Price80		10000	-4.6499	0.6400	-5.9539	-3.4282
RECov Touchless, Touchless		10000	3.0571	1.1503	1.3122	5.2392
RECov Steel, Touchless		10000	-0.4706	1.0031	-2.2026	1.1375
RECov Steel, Steel		10000	2.6142	1.0170	1.0026	4.5185
RECov AutoBag, Touchless		10000	-0.7271	0.9559	-2.6179	1.0821
RECov AutoBag, Steel		10000	-0.00084	0.8108	-1.5522	1.4319
RECov AutoBag, AutoBag		10000	3.6211	1.5239	1.2185	6.7463
RECov Price80, Touchless		10000	-1.2188	1.1945	-3.6437	0.8218
RECov Price80, Steel		10000	-1.5669	1.2639	-4.1530	0.7019
RECov Price80, AutoBag		10000	-2.2400	1.6576	-5.5705	0.6001
RECov Price80, Price80		10000	7.6739	3.1808	2.3981	14.0942
Touchless	ID 1	10000	2.1940	1.0468	0.2013	4.2928
Steel	ID 1	10000	0.7208	1.0666	-1.2843	2.8607
AutoBag	ID 1	10000	3.7719	1.4263	1.1515	6.6288
Price80	ID 1	10000	-5.1558	2.2854	-9.8578	-1.0303
Touchless	ID 2	10000	-0.0412	0.8572	-1.6932	1.6535
Steel	ID 2	10000	-0.3953	0.8200	-2.0387	1.1833
AutoBag	ID 2	10000	2.7178	1.3321	0.1872	5.3265
Price80	ID 2	10000	0.00135	1.3107	-2.5173	2.6351
Touchless	ID 3	10000	1.1149	0.9607	-0.6416	3.1861
Steel	ID 3	10000	1.6568	0.9606	-0.1567	3.6094
AutoBag	ID 3	10000	0.3119	1.1598	-2.0512	2.4132
Price80	ID 3	10000	-0.9705	1.5496	-4.0250	2.0511
Touchless	ID 4	10000	1.1545	0.9697	-0.7475	3.0413
Steel	ID 4	10000	1.3730	0.9891	-0.5385	3.3589
AutoBag	ID 4	10000	3.1780	1.4497	0.5477	6.1348
Price80	ID 4	10000	-6.5924	2.4704	-11.4607	-1.8296
Touchless	ID 5	10000	0.4705	1.1716	-1.7910	2.7956
Steel	ID 5	10000	2.6725	1.2193	0.5348	5.1762
AutoBag	ID 5	10000	3.4376	1.6172	0.5508	6.7088
Price80	ID 5	10000	-5.0509	2.5350	-10.0811	-0.1987

The average part-worths (\bar{y}) for touchless opening, steel material, automatic trash bag replacement, and price, which are shown as “REMean Touchless,” “REMean Steel,” “REMean AutoBag,” and “REMean Price80,” are very similar to the estimates of the fixed effects in the previous model as shown in [Figure 27.9](#), indicating the equivalence of these two setups. The covariance estimates of the random coefficients (Ω_{γ}) are also very similar.

The next set of parameters show the estimates for the random effects for the first five respondents (see [Output 27.4.1](#)). However, these estimates are no longer the deviation from the overall means but are their own effects. The part-worth for touchless opening for the first respondent (ID 1) is 2.2, which is close to the part-worth that is estimated in the model as shown in [Figure 27.9](#).

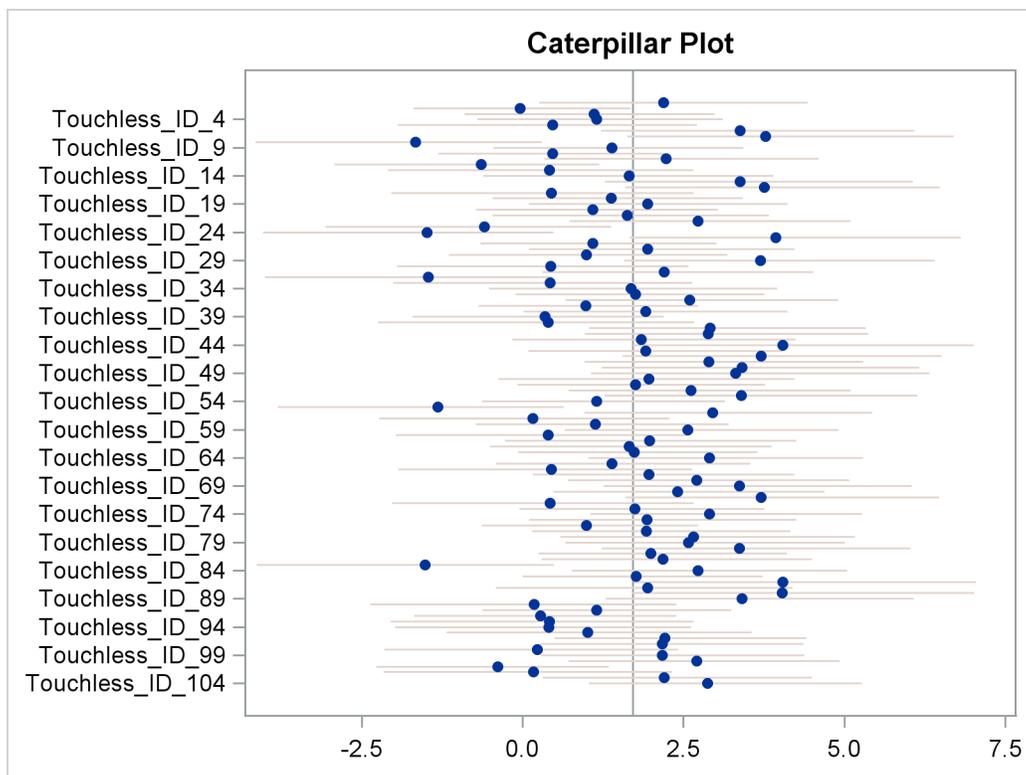
A caterpillar plot is a side-by-side bar plot of the 95% intervals for multiple variables. A caterpillar plot is usually used to visualize and compare random-effects parameters, which can exist in large numbers in certain models. You can use the %CATER autocall macro to create a caterpillar plot. The %CATER macro requires you to specify an input data set and a list of variables that you want to plot.

In this example, the output data set Postsamp contains posterior draws for all the random-effects parameters. The following statements generate a caterpillar plot for Touchless:

```
ods graphics on;
%CATER(data=Postsamp, var=Touchless_ID_);
```

Output 27.4.2 is a caterpillar plot of the random effects of touchless opening for the 104 participants in this study. It displays the heterogeneity in preferences for touchless opening among the sample. As shown, the individuals have diverse preferences for the feature of touchless opening for a trash can. Heterogeneity of preferences is an important aspect of the model, and ignoring its presence can lead to incorrect inferences.

Output 27.4.2 Caterpillar Plot of the Random-Effects Parameters



If you want to change the display of the caterpillar plot (for example, to use a different line pattern, color, or size of the markers), you must first modify the `Stat.MCMC.Graphics.Caterpillar` template and then call the %CATER macro again.

Example 27.5: Heterogeneity Affected by Individual Characteristics

Rossi, Allenby, and McCulloch (2005) studied a scanner panel data about purchases of margarine. The data were first analyzed in Allenby and Rossi (1991) and are about purchases of ten brands of margarine. This

example considers a subset of data about six margarine brands: Parkay stick, Blue Bonnet stick, Fleischmann's stick, a house brand stick, a generic stick, and Shedd's Spread tub. There are 313 households, which made a total of 3,405 purchases. Information about a few demographic characteristics of these households (income and family size) is expected to have effects on the central location of the distribution of heterogeneity.

The data set, which is called Sashelp.Margarin, comes from the SASHELP library.

```
proc print data=Sashelp.Margarin (obs=24);
  by HouseID Set;
  id HouseID Set;
run;
```

The data for the first four choice sets are shown in Figure 27.5.1.

Output 27.5.1 Data for the First Four Choice Sets

HouseID	Set	Choice	Brand	LogPrice	LogInc	Fam Size
2100016	1	1	PPk	-0.41552	3.48124	2
		0	PBB	-0.40048	3.48124	2
		0	PF1	0.08618	3.48124	2
		0	PHse	-0.56212	3.48124	2
		0	PGen	-1.02165	3.48124	2
		0	PSS	-0.16252	3.48124	2
2100016	2	1	PPk	-0.46204	3.48124	2
		0	PBB	-0.40048	3.48124	2
		0	PF1	-0.01005	3.48124	2
		0	PHse	-0.56212	3.48124	2
		0	PGen	-1.02165	3.48124	2
		0	PSS	-0.16252	3.48124	2
2100016	3	1	PPk	-1.23787	3.48124	2
		0	PBB	-0.69315	3.48124	2
		0	PF1	-0.01005	3.48124	2
		0	PHse	-0.56212	3.48124	2
		0	PGen	-1.02165	3.48124	2
		0	PSS	-0.23572	3.48124	2
2100016	4	1	PPk	-0.47804	3.48124	2
		0	PBB	-0.49430	3.48124	2
		0	PF1	-0.01005	3.48124	2
		0	PHse	-0.56212	3.48124	2
		0	PGen	-1.02165	3.48124	2
		0	PSS	-0.16252	3.48124	2

The variable `HouseID` represents the household ID, and each household made at least five purchases, which are defined by `Set`. The variable `Choice` represents the choice made among the six margarine brands for each purchase or choice set. The variable `Brand` has the value `PPk` for Parkay stick, `PBB` for Blue Bonnet stick, `PFL` for Fleischmann's stick, `PHse` for the house brand stick, `PGen` for the generic stick, and `PSS` for Shedd's Spread tub. The variable `LogPrice` is the logarithm of the product price. The variables `LogInc` and variable `FamSize` provide information about household income and family size, respectively.

For the purpose of comparison, the following statements fit the random-effects-only logit model by using random walk Metropolis sampling as suggested in Rossi, Allenby, and McCulloch (2005):

```
proc bchoice data=Sashelp.Margarin seed=123 nmc=20000 thin=4 alg=rwm nthreads=4;
  class Brand(ref='PPk') HouseID Set;
  model Choice = / choicest=(HouseID Set);
  random Brand LogPrice / subject=HouseID remean=(LogInc FamSize)
    type=un monitor=(1);
run;
```

The `ALG=RWM` option in the `PROC BCHOICE` statement requests the random walk Metropolis sampling algorithm, and the `REMEAN=(LOGINC FAMSIZ)` option in the `RANDOM` statement requests estimation of the nonzero mean of the random effects, which is a function of household income and family size. No fixed effects are specified in the `MODEL` statement. The `NMC=20000` option in the `PROC BCHOICE` statement runs the chain for 20,000 iterations, and the `THIN=4` option keeps one of every four samples. Summary statistics for the mean matrix of the random coefficients (Γ), the covariance of the random coefficients (Ω_{γ}), and the random coefficients (γ_i) for the first household are shown in [Output 27.5.2](#).

Output 27.5.2 Posterior Summary Statistics

The BCHOICE Procedure

Posterior Summaries and Intervals

Parameter	Subject	N	Mean	Standard Deviation	95% HPD Interval	
REMean Brand PBB		5000	-1.1809	0.6122	-2.3775	-0.0144
REMean Brand PFl		5000	-3.5116	2.0188	-7.5975	0.2097
REMean Brand PGen		5000	-4.9821	1.1709	-7.2338	-2.6884
REMean Brand PHse		5000	-3.2291	0.8966	-5.0406	-1.5415
REMean Brand PSS		5000	-0.0157	1.2214	-2.4096	2.3238
REMean LogPrice		5000	-3.3194	0.8545	-5.0431	-1.6453
REMean Brand PBB LogInc		5000	0.0550	0.2043	-0.3367	0.4451
REMean Brand PFl LogInc		5000	0.8263	0.6562	-0.4800	2.0819
REMean Brand PGen LogInc		5000	-0.5483	0.3930	-1.3391	0.2101
REMean Brand PHse LogInc		5000	0.0272	0.3023	-0.5640	0.6158
REMean Brand PSS LogInc		5000	-0.5861	0.4168	-1.3908	0.2276
REMean LogPrice LogInc		5000	-0.3175	0.2990	-0.8955	0.2697
REMean Brand PBB FamSize		5000	-0.0336	0.0966	-0.2275	0.1437
REMean Brand PFl FamSize		5000	-0.7625	0.3231	-1.4176	-0.1522
REMean Brand PGen FamSize		5000	0.5773	0.1836	0.2248	0.9402
REMean Brand PHse FamSize		5000	0.2346	0.1352	-0.0467	0.4808
REMean Brand PSS FamSize		5000	0.0406	0.1997	-0.3719	0.4186
REMean LogPrice FamSize		5000	0.1085	0.1331	-0.1504	0.3580
RECOV Brand PBB, Brand PBB		5000	2.1824	0.3916	1.4411	2.9386
RECOV Brand PFl, Brand PBB		5000	2.0848	0.8517	0.4737	3.7633
RECOV Brand PFl, Brand PFl		5000	13.1762	3.9078	6.3684	21.0989
RECOV Brand PGen, Brand PBB		5000	1.9526	0.5760	0.8471	3.0718
RECOV Brand PGen, Brand PFl		5000	1.5669	1.7949	-1.6053	5.2591
RECOV Brand PGen, Brand PGen		5000	8.3864	1.3307	6.0342	11.1880
RECOV Brand PHse, Brand PBB		5000	1.5108	0.4554	0.6690	2.4392
RECOV Brand PHse, Brand PFl		5000	2.4807	1.3482	-0.2733	4.9664
RECOV Brand PHse, Brand PGen		5000	5.7498	0.9319	4.0263	7.6664
RECOV Brand PHse, Brand PHse		5000	5.4675	0.8456	3.8925	7.1852
RECOV Brand PSS, Brand PBB		5000	1.1780	0.6307	-0.0181	2.4864
RECOV Brand PSS, Brand PFl		5000	0.8534	2.1434	-3.5777	5.0734
RECOV Brand PSS, Brand PGen		5000	4.9601	1.1765	2.7987	7.4216
RECOV Brand PSS, Brand PHse		5000	3.4612	0.8947	1.8634	5.3448
RECOV Brand PSS, Brand PSS		5000	8.8466	1.7773	5.5686	12.3094
RECOV LogPrice, Brand PBB		5000	-0.2652	0.3179	-0.8778	0.3718
RECOV LogPrice, Brand PFl		5000	2.1725	0.8906	0.4017	3.8500
RECOV LogPrice, Brand PGen		5000	-1.1063	0.5708	-2.2724	-0.0595
RECOV LogPrice, Brand PHse		5000	-0.4487	0.4543	-1.2981	0.4822
RECOV LogPrice, Brand PSS		5000	0.1428	0.6313	-1.0706	1.3744
RECOV LogPrice, LogPrice		5000	2.0727	0.4839	1.1196	3.0223
Brand PBB	HouseID 2100016	5000	-2.2892	1.0105	-4.3643	-0.4129
Brand PFl	HouseID 2100016	5000	-4.1063	2.7737	-9.7447	0.5782
Brand PGen	HouseID 2100016	5000	-6.5999	1.6118	-9.7086	-3.5667
Brand PHse	HouseID 2100016	5000	-2.9469	1.1851	-5.3047	-0.7302
Brand PSS	HouseID 2100016	5000	-3.3378	2.1715	-7.5906	0.7075
LogPrice	HouseID 2100016	5000	-4.4983	1.1717	-6.8224	-2.2033

Table 27.11 collects the posterior means and standard deviations of $\bar{\Gamma}$ that are shown in Output 27.5.2. The first column corresponds to the parameters that are specified in the model, namely the brands and price. The second column shows the average part-worths of each brand (versus the brand, Parkay stick) and the price at $\text{LogInc}=0$ and $\text{FamSize}=0$. The LogInc and Age columns list the modifying effects on the preference for each brand and price by household income and family size, respectively. Larger families show more interest in the generic and house brands and tend to stay away from the Fleischmann’s brand. For example, consider the part-worth estimates for Fleischmann’s. The posterior mean for $\text{REMean Brand PFI FamSize}$ (the Fleischmann’s row and the Famsize column) is -0.76 with a standard deviation of 0.32 , meaning that an additional unit increase in family size is associated with a reduction of 0.76 in the estimated part-worth for Fleischmann’s. In general, the demographics of households are only weakly associated with preference for brand and price. These results are in good agreement with those of Rossi, Allenby, and McCulloch (2005).

Table 27.11 Posterior Mean and Standard Deviation of Γ

Parameter		Intercept	LogInc	FamSize
Blue Bonnet	Name	REMean Brand PBB	REMean Brand PBB LogInc	REMean Brand PBB FamSize
	Mean	-1.18	0.06	-0.03
	Std	0.61	0.20	0.10
Fleischmann’s	Name	REMean Brand PFI	REMean Brand PFI LogInc	REMean Brand PFI FamSize
	Mean	-3.51	0.83	-0.76
	Std	2.02	0.66	0.32
Generic	Name	REMean Brand PGen	REMean Brand PGen LogInc	REMean Brand PGen FamSize
	Mean	-4.98	-0.55	0.58
	Std	1.17	0.39	0.18
House	Name	REMean Brand PHse	REMean Brand PHse LogInc	REMean Brand PHse FamSize
	Mean	-3.23	0.03	0.23
	Std	0.90	0.30	0.14
Shedd’s Spread	Name	REMean Brand PSS	REMean Brand PSS LogInc	REMean Brand PSS FamSize
	Mean	-0.02	-0.59	-0.04
	Std	1.22	0.42	0.20
LogPrice	Name	REMean LogPrice	REMean LogPrice LogInc	REMean LogPrice FamSize
	Mean	-3.32	-0.32	0.11
	Std	0.85	0.30	0.13

Because the demographic variables are not zero-centered, the Intercept column shows the average part-worths of each brand and price for households with $\text{LogInc}=0$ and $\text{FamSize}=0$, which are not very meaningful. It is better to center demographic variables by their means, so that the posterior means listed in the Intercept column can be interpreted as the part-worths of a household that has an average income and average size.

Nevertheless, you can obtain the utilities of households that have any income levels and sizes. For example, the average part-worth of the Fleischmann’s brand for a household with average income ($\text{LogInc}=3.1$) and family size ($\text{FamSize}=3$) would be as follows, because the estimated LogInc coefficient is 0.83 and the estimated FamSize coefficient is -0.76 for Fleischmann’s:

$$-3.51 + 0.83 \times 3.1 - 0.76 \times 3 = -3.22$$

You can obtain part-worths for all other brands and compare their popularity among average households.

The posterior means and standard deviations of the covariance matrix of the random coefficients (Ω_{γ}) are displayed by parameters that are labeled “RECOV Brand PBB, Brand PBB,” “RECOV Brand PFI, Brand PBB,”

and so on. Some of the diagonal terms are fairly large, indicating that there is quite a bit of heterogeneity between households in margarine brand preference and price sensitivity. The covariance between the generic and house brands, “RECOV Brand PHse, Brand PGen,” is fairly large, suggesting that household preferences for these two brands are highly correlated.

The next set of parameters, which are displayed in [Output 27.5.2](#), contain the estimates for the random effects for the first household.

Example 27.6: Inference on Quantities of Interest

It is easy to give a model an odds interpretation in terms of a probability ratio. The odds of choosing alternative a instead of alternative b while holding all other attributes or conditions the same are

$$\text{ODDS}_{ab} = \frac{P_a}{P_b} = \frac{\exp(\beta_a)}{\exp(\beta_b)} = \exp(\beta_a - \beta_b)$$

where P_a and P_b are the probabilities of choosing alternative a and b , respectively.

In “[Example 27.1: Alternative-Specific and Individual-Specific Effects](#)” on page 1054, a study for travel demand was considered. Summary statistics of that study are displayed again in [Output 27.6.1](#) for the purpose of illustration.

Output 27.6.1 PROC BCHOICE Posterior Summary Statistics

Posterior Summaries and Intervals					
Parameter	N	Mean	Standard Deviation	95% HPD Interval	
Mode Auto	5000	-0.1678	0.7440	-1.7017	1.2396
Mode Plane	5000	-1.8794	1.2683	-4.6055	0.3801
TravTime	5000	-0.5695	0.2047	-0.9943	-0.2328

The overall probability ratio of choosing airplane as opposed to public transit (which is the reference category) simplifies to

$$\text{ODDS}_{(\text{Plane vs. Transit})} = \frac{P(\text{Plane})}{P(\text{Transit})} = \frac{\exp(-1.88)}{\exp(0)} = 0.15$$

This ratio indicates that the likelihood of choosing planes as means of travel is 0.15 times that of choosing public transit. Public transit is a lot more favorable than planes.

The overall probability ratio of choosing a plane instead of an automobile is

$$\text{ODDS}_{(\text{Plane vs. Auto})} = \frac{P(\text{Plane})}{P(\text{Transit})} = \frac{\exp(-1.88)}{\exp(-0.17)} = 0.18$$

You can derive the probabilities for any alternatives and any combinations of attributes, even imaginary ones, after you obtain the parameter estimates.

Although it is easy to obtain the point estimates, sometimes you might want to estimate other quantities, such as the standard deviations or various quantiles. If you are interested in making inference based on any quantities that are transformations of the random variables, you can do it either directly in PROC BCHOICE or by using the DATA step after you run the simulation.

Suppose you want to set the reference category as Plane so that you can directly compare the other two modes with it. You can specify REF='PLANE' in the CLASS statement and rerun the PROC BCHOICE simulation. You can also use the DATA step to calculate the quantities of interest. The following DATA step uses the simulated values of Mode_Auto and Mode_Plane from the output data set Bsamp to create a new series of posterior samples for Auto_Plane, a new variable that is created to directly compare airplane with automobile:

```
data Transout;
  set Bsamp;
  Auto_Plane=Mode_Auto -Mode_Plane;
run;
```

Then you can use some autocall macros to analyze the posterior samples of Auto_Plane. For example, the %POSTSUM macro provides summary statistics, and the %POSTINT macro provides equal-tail and HPD intervals as described in the section “Autocall Macros for Postprocessing” on page 1048.

```
%postsum(data=Transout, var=Auto_Plane)
%postint(data=Transout, var=Auto_Plane)
```

You can also generate some ODS graphs of those posterior samples for diagnostic purposes, as described in the section “Regenerating Diagnostics Plots” on page 1049.

Example 27.7: Predict the Choice Probabilities

This example shows how to obtain the posterior predictive distribution of the choice probability that each alternative is chosen from a choice set. The posterior predictive distribution enables you to get the expected choice probabilities of all the alternatives in the data, or even to predict market share for simulated or hypothetical products or marketplaces that do not directly reflect the choice set in the data.

Suppose you have a data set that contains all the attribute variables (the design matrix) for all the alternatives in a choice set. For example, in the candy study earlier in the chapter, in the section “A Simple Logit Model” on page 1004, you can use the same eight alternatives: Dark is 1 for dark chocolate and 0 for soft chocolate; Soft is 1 for soft center and 0 for chewy center; Nuts is 1 if the candy contains nuts and 0 if it contains no nuts. The following data set contains the eight alternatives:

```
data DesignMatrix;
  input Dark Soft Nuts;
datalines;
0 0 0
0 0 1
0 1 0
0 1 1
1 0 0
1 0 1
1 1 0
1 1 1
;
```

You can use the **PREDDIST** statement, which obtains samples from the posterior predictive distribution of each of the choice probabilities by using the posterior samples of parameters in the model:

```
proc bchoice data=Chocs outpost=Bsamp nmc=10000 thin=2 seed=124;
  class Dark(ref='0') Soft(ref='0') Nuts(ref='0') Subj;
  model Choice = Dark Soft Nuts / choiceset=(Subj);
  preddist covariates=DesignMatrix nalter=8 outpred=Predout;
run;

%POSTSUM(data=Predout, var=Prob_1_.);
```

In the **PREDDIST** statement, the **COVARIATES=** option names the data set to contain the explanatory variable values for which the predictions are established. This data set must contain data that have the same variables that are used in the model. The **NALTER=** option specifies the number of alternatives in each choice set in the **COVARIATES=** data set. All choice sets in the data must have the same number of alternatives. If you omit the **COVARIATES=** option, the **DATA=** data set that you specify in the **PROC BCHOICE** statement is used instead. The **OUTPRED=** option creates an output data set to contain the samples from the posterior predictive distribution of the choice probabilities. Then you can use SAS autocall macros to analyze the posterior samples. For example, the **%POSTSUM** macro provides summary statistics.

You can predict the choice probabilities by using the means of the posterior distributions. The results from using the **%POSTSUM** macro are shown in [Output 27.7.1](#). There is only one choice set for choice probability prediction, in which there are a total of eight alternatives. This explains the parameter names in the first column of the output, where the first number indexes the choice sets and the second number indexes the alternatives in each choice set. The most preferred chocolate candy is the sixth one, Dark/Chewy/Nuts, which takes about half the market.

Output 27.7.1 Choice of Chocolate Candies

Summary Statistics						
Parameter	N	Mean	StdDev	P25	P50	P75
Prob_1_1	5000	0.05541	0.04252	0.02396	0.04529	0.07282
Prob_1_2	5000	0.13093	0.08009	0.06949	0.11487	0.17965
Prob_1_3	5000	0.00686	0.00896	0.00137	0.00366	0.00853
Prob_1_4	5000	0.01578	0.01786	0.00389	0.00966	0.02061
Prob_1_5	5000	0.21016	0.10328	0.13431	0.19301	0.27497
Prob_1_6	5000	0.49462	0.13152	0.40465	0.49734	0.58797
Prob_1_7	5000	0.02617	0.02793	0.00717	0.01683	0.03550
Prob_1_8	5000	0.06008	0.05261	0.02032	0.04571	0.08514

References

- Albert, J. H. and Chib, S. (1993), “Bayesian Analysis of Binary and Polychotomous Response Data,” *Journal of the American Statistical Association*, 88, 669–679.
- Allenby, G. M. (1997), “An Introduction to Hierarchical Bayesian Modeling,” American Marketing Association, Advanced Research Techniques Forum, Tutorial Notes.

- Allenby, G. M. and Lenk, P. (1994), "Modeling Household Purchase Behavior with Logistic Normal Regression," *Journal of the American Statistical Association*, 89, 1218–1231.
- Allenby, G. M. and Rossi, P. E. (1991), "Quality Perceptions and Asymmetric Switching between Brands," *Marketing Science*, 10, 185–205.
- Allenby, G. M. and Rossi, P. E. (1999), "Marketing Models of Consumer Heterogeneity," *Journal of Econometrics*, 89, 57–78.
- Bhat, C. R. (1995), "A Heteroscedastic Extreme Value Model of Intercity Travel Mode Choice," *Transportation Research*, 29, 471–483.
- Eilers, P. H. C. and Marx, B. D. (1996), "Flexible Smoothing with *B*-Splines and Penalties," *Statistical Science*, 11, 89–121, with discussion.
- Gamerman, D. (1997), "Sampling from the Posterior Distribution in Generalized Linear Models," *Statistics and Computing*, 7, 57–68.
- Herriges, J. A. and Kling, C. L. (1996), "Testing the Consistency of Nested Logit Models with Utility Maximization," *Economics Letters*, 50, 33–39.
- Kass, R. E., Carlin, B. P., Gelman, A., and Neal, R. M. (1998), "Markov Chain Monte Carlo in Practice: A Roundtable Discussion," *American Statistician*, 52, 93–100.
- Kling, C. L. and Herriges, J. A. (1995), "An Empirical Investigation of the Consistency of Nested Logit Models with Utility Maximization," *American Journal of Agricultural Economics*, 77, 875–884.
- Kuhfeld, W. F. (2010), *Marketing Research Methods in SAS*, Technical report, SAS Institute Inc., http://support.sas.com/resources/papers/tnote/tnote_marketresearch.html.
- Lahiri, K. and Gao, J. (2002), "Bayesian Analysis of Nested Logit Model by Markov Chain Monte Carlo," *Journal of Econometrics*, 11, 103–133.
- McCullagh, P. and Nelder, J. A. (1989), *Generalized Linear Models*, 2nd Edition, London: Chapman & Hall.
- McCulloch, R. and Rossi, P. E. (1994), "An Exact Likelihood Analysis of the Multinomial Probit Model," *Journal of Econometrics*, 64, 207–240.
- McFadden, D. (1974), "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press.
- McFadden, D. (1978), "Modelling the Choice of Residential Location," in A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull, eds., *Spatial Interaction Theory and Planning Models*, Amsterdam: North-Holland.
- McFadden, D. (2001), "Economic Choices," *American Economic Review*, 91, 351–378.
- Roberts, G. O., Gelman, A., and Gilks, W. R. (1997), "Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms," *Annals of Applied Probability*, 7, 110–120.
- Roberts, G. O. and Rosenthal, J. S. (2001), "Optimal Scaling for Various Metropolis-Hastings Algorithms," *Statistical Science*, 16, 351–367.

Rossi, P. E. (2012), "Bayesm Package: Bayesian Inference for Marketing/Micro-econometrics," <http://www.perossi.org/home/bsm-1>.

Rossi, P. E. (2013), personal communication.

Rossi, P. E., Allenby, G. M., and McCulloch, R. (2005), *Bayesian Statistics and Marketing*, Chichester, UK: John Wiley & Sons.

Rossi, P. E., McCulloch, R., and Allenby, G. M. (1996), "The Value of Purchase History Data in Target Marketing," *Marketing Science*, 15, 321–340.

Train, K. E. (2009), *Discrete Choice Methods with Simulation*, Cambridge: Cambridge University Press.

Train, K. E., McFadden, D., and Ben-Akiva, M. (1987), "The Demand for Local Telephone Service: A Fully Discrete Model of Residential Calling Patterns and Service Choice," *Rand Journal of Economics*, 18, 109–123.

Subject Index

- BCHOICE procedure, 1001
 - class levels, 1020
 - compared with other SAS procedures, 1003
 - continuous effects, 1033
 - examples, see also examples, BCHOICE, 1054
 - MODEL statement options, 1028
 - ODS Graphics, 1021
 - output ODS Graphics table names, 1054
 - output table names, 1053
 - posterior predictive distribution, 1031
 - posterior samples data set, 1021
 - PROC BCHOICE statement options, 1016
 - RANDOM statement options, 1032
 - random-effects parameters, 1032
 - SUBJECT= option in RANDOM statement, 1033
- BCHOICE procedure, MODEL statement
 - initial values, 1029
- choice set, 1034
- class levels
 - BCHOICE procedure, 1020
- compared with other SAS procedures
 - BCHOICE procedure, 1003
- continuous effects
 - BCHOICE procedure, 1033
- discrete choice models, 1034
- examples, BCHOICE
 - a random-effects-only logit model, 1063
 - alternative-specific and individual-specific effects, 1054
 - getting started, 1004
 - heterogeneity affected by individual characteristics, 1065
 - inference on quantities of interest, 1070
 - logit model with random effects, 1011
 - nested logit modeling, 1057
 - predict the choice probabilities, 1071
 - probit modeling, 1060
 - simple logit model, 1004
- initial values
 - BCHOICE procedure, MODEL statement, 1029
- MODEL statement options
 - BCHOICE procedure, 1028
- ODS Graphics
 - BCHOICE procedure, 1021
 - output ODS Graphics table names
 - BCHOICE procedure, 1054
 - output table names
 - BCHOICE procedure, 1053
 - posterior predictive distribution
 - BCHOICE procedure, 1031
 - posterior samples data set
 - BCHOICE procedure, 1021
 - random-effects parameters
 - BCHOICE procedure, 1032
 - SUBJECT= option in RANDOM statement
 - BCHOICE procedure, 1033

Syntax Index

- ACCEPTTOL= option
 - PROC BCHOICE statement, 1017
- BCHOICE procedure
 - CLASS statement, 1025
 - MODEL statement, 1028
 - PREDDIST statement, 1031
 - RANDOM statement, 1031
 - syntax, 1016
- BCHOICE procedure, BY statement, 1025
- BCHOICE procedure, MODEL statement, 1028
 - CHOICESET= option, 1028
 - COEFFPRIOR= option, 1028
 - COVPRIOR= option, 1029
 - COVTYPE= option, 1029
 - INIT= option, 1029
 - LAMBDA PRIOR= option, 1030
 - NEST= option, 1030
 - SAMELAMBDA option, 1030
 - TYPE= option, 1030
- BCHOICE procedure, PREDDIST statement, 1031
 - COVARIATES= option, 1031
 - NALTER= option, 1031
 - NALTERNATIVE= option, 1031
 - OUTPRED= option, 1031
- BCHOICE procedure, PROC BCHOICE statement
 - ACCEPTTOL= option, 1017
 - DATA= option, 1017
 - DIAG= option, 1017
 - DIAGNOSTICS= option, 1017
 - DIC option, 1020
 - HITPROB option, 1020
 - INF= option, 1020
 - LOGPOST option, 1020
 - MAXTUNE= option, 1020
 - MINTUNE= option, 1020
 - NBI= option, 1020
 - NMC= option, 1020
 - NOCLPRINT option, 1020
 - NTHREADS= option, 1021
 - NTU= option, 1021
 - OUTPOST= option, 1021
 - PLOTS= option, 1021
 - SEED= option, 1024
 - STATISTICS= option, 1024
 - STATS= option, 1024
 - TARGACCEPT= option, 1025
 - THIN= option, 1025
- BCHOICE procedure, RANDOM statement, 1031
 - COVPRIOR= option, 1032
 - MONITOR option, 1032
 - NOOUTPOST option, 1033
 - REMEAN option, 1033
 - SUBJECT= option, 1033
 - TYPE= option, 1034
- BCHOICE procedure, CLASS statement, 1026
 - CPREFIX= option, 1026
 - DESCENDING option, 1026
 - LPREFIX= option, 1026
 - MISSING option, 1026
 - ORDER= option, 1026
 - PARAM= option, 1027
 - REF= option, 1027
 - TRUNCATE option, 1027
- BY statement
 - BCHOICE procedure, 1025
- CHOICESET= option
 - MODEL statement (BCHOICE), 1028
- CLASS statement
 - BCHOICE procedure, 1026
- COEFFPRIOR= option
 - MODEL statement (BCHOICE), 1028
- COVARIATES= option
 - PREDDIST statement (BCHOICE), 1031
- COVPRIOR= option
 - MODEL statement (BCHOICE), 1029
 - RANDOM statement (BCHOICE), 1032
- COVTYPE= option
 - MODEL statement (BCHOICE), 1029
- CPREFIX= option
 - CLASS statement (BCHOICE), 1026
- DATA= option
 - PROC BCHOICE statement, 1017
- DESCENDING option
 - CLASS statement (BCHOICE), 1026
- DIAG= option
 - PROC BCHOICE statement, 1017
- DIAGNOSTICS= option
 - PROC BCHOICE statement, 1017
- DIC option
 - PROC BCHOICE statement, 1020
- HITPROB option
 - PROC BCHOICE statement, 1020

INF= option
 PROC BCHOICE statement, 1020

INIT= option
 MODEL statement (BCHOICE), 1029

LAMBDA PRIOR= option
 MODEL statement (BCHOICE), 1030

LOGPOST option
 PROC BCHOICE statement, 1020

LPREFIX= option
 CLASS statement (BCHOICE), 1026

MAXTUNE= option
 PROC BCHOICE statement, 1020

MINTUNE= option
 PROC BCHOICE statement, 1020

MISSING option
 CLASS statement (BCHOICE), 1026

MODEL statement
 BCHOICE procedure, 1028

MONITOR option
 RANDOM statement (BCHOICE), 1032

NALTER= option
 PREDDIST statement (BCHOICE), 1031

NALTERNATIVE= option
 PREDDIST statement (BCHOICE), 1031

NBI= option
 PROC BCHOICE statement, 1020

NEST= option
 MODEL statement (BCHOICE), 1030

NMC= option
 PROC BCHOICE statement, 1020

NOCLPRINT option
 PROC BCHOICE statement, 1020

NOOUTPOST option
 RANDOM statement (BCHOICE), 1033

NTHREADS= option
 PROC BCHOICE statement, 1021

NTU= option
 PROC BCHOICE statement, 1021

ORDER= option
 CLASS statement (BCHOICE), 1026

OUTPOST= option
 PROC BCHOICE statement, 1021

OUTPRED= option
 PREDDIST statement (BCHOICE), 1031

PARAM= option
 CLASS statement (BCHOICE), 1027

PLOTS= option
 PROC BCHOICE statement, 1021

PREDDIST statement
 BCHOICE procedure, 1031

RANDOM statement
 BCHOICE procedure, 1031

REF= option
 CLASS statement (BCHOICE), 1027

REMEAN option
 RANDOM statement (BCHOICE), 1033

SAMELAMBDA option
 MODEL statement (BCHOICE), 1030

SEED= option
 PROC BCHOICE statement, 1024

STATISTICS= option
 PROC BCHOICE statement, 1024

STATS= option
 PROC BCHOICE statement, 1024

SUBJECT= option
 RANDOM statement (BCHOICE), 1033

TARGACCEPT= option
 PROC BCHOICE statement, 1025

THIN= option
 PROC BCHOICE statement, 1025

TRUNCATE option
 CLASS statement (BCHOICE), 1027

TYPE= option
 MODEL statement (BCHOICE), 1030
 RANDOM statement (BCHOICE), 1034