SAS/STAT® 9.2 User's Guide
The MCMC Procedure
(Experimental)
(Book Excerpt)

Chapter 52

The MCMC Procedure (Experimental)

Contents

Overview: MCMC Procedure

The MCMC procedure is a general purpose Markov chain Monte Carlo (MCMC) simulation procedure that is designed to fit Bayesian models. Bayesian statistics is different from traditional statistical methods such as frequentist or classical methods. 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 173 for a guide to Bayesian textbooks of varying degrees of difficulty.

In essence, Bayesian statistics treats parameters as unknown random variables, and it makes inferences based on the posterior distributions of the parameters. There are several advantages associated with this approach to statistical inference. Some of the advantages include its ability to use prior information and to directly answer specific scientific questions that can be easily understood. For further discussions of the relative advantages and disadvantages of Bayesian analysis, see the section "Bayesian Analysis: Advantages and Disadvantages" on page 149.

It follows from Bayes' theorem that a posterior distribution is the product of the likelihood function and the prior distribution of the parameter. In all but the simplest cases, it is very difficult to obtain the posterior distribution directly and analytically. Often, Bayesian methods rely on simulations to

generate sample from the desired posterior distribution and use the simulated draws to approximate the distribution and to make all of the inferences.

PROC MCMC is a flexible simulation-based procedure that is suitable for fitting a wide range of Bayesian models. To use the procedure, you need to specify a likelihood function for the data and a prior distribution for the parameters. You might also need to specify hyperprior distributions if you are fitting hierarchical models. PROC MCMC then 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. You can analyze data that have any likelihood, prior, or hyperprior with PROC MCMC, as long as these functions are programmable using the SAS DATA step functions. The parameters can enter the model linearly or in any nonlinear functional form. The default algorithm that PROC MCMC uses is an adaptive blocked random walk Metropolis algorithm that uses a normal proposal distribution.

PROC MCMC Compared with Other SAS Procedures

PROC MCMC is unlike most other SAS/STAT procedures in that the nature of the statistical inference is Bayesian. You specify prior distributions for the parameters with PRIOR statements and the likelihood function for the data with MODEL statements. 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 the same random number seed is used. The model specification is similar to PROC NLIN, and PROC MCMC shares much of the syntax of PROC NLMIXED.

Note that you can also carry out a Bayesian analysis with the GENMOD, PHREG, and LIFEREG procedures for generalized linear models, accelerated life failure models, Cox regression models, and piecewise constant baseline hazard models (also known as piecewise exponential models). See Chapter 37, "The GENMOD Procedure," Chapter 64, "The PHREG Procedure," and Chapter 48, "The LIFEREG Procedure."

Getting Started: MCMC Procedure

There are three examples in this "Getting Started" section: a simple linear regression, the Behrens-Fisher estimation problem, and a random effects model. The regression model is chosen for its simplicity; the Behrens-Fisher problem illustrates some advantages of the Bayesian approach; and the random effects model is one of the most prevalently used models.

Keep in mind that PARMS statements declare the parameters in the model, PRIOR statements declare the prior distributions, and MODEL statements declare the likelihood for the data. In most cases, you do not need to supply initial values. The procedure advises you if it is unable to generate starting values for the Markov chain.

Simple Linear Regression

This section illustrates some basic features of PROC MCMC by using a linear regression model. The model is as follows:

$$Y_i = \beta_0 + \beta_1 X_i + \epsilon_i$$

for the observations $i = 1, 2, \ldots, n$.

The following statements create a SAS data set with measurements of Height and Weight for a group of children:

```
title 'Simple Linear Regression';

data Class;
   input Name $ Height Weight @@;
   datalines;
Alfred   69.0 112.5    Alice   56.5   84.0    Barbara 65.3   98.0
Carol    62.8 102.5    Henry   63.5 102.5    James   57.3   83.0
Jane     59.8  84.5    Janet   62.5 112.5    Jeffrey 62.5   84.0
John     59.0  99.5    Joyce   51.3  50.5    Judy    64.3   90.0
Louise   56.3  77.0    Mary    66.5 112.0    Philip  72.0 150.0
Robert   64.8 128.0    Ronald 67.0 133.0    Thomas  57.5   85.0
William 66.5 112.0
;
```

The equation of interest is as follows:

$$\text{Weight}_i = \beta_0 + \beta_1 \text{Height}_i + \epsilon_i$$

The observation errors, ϵ_i, are assumed to be independent and identically distributed with a normal distribution with mean zero and variance σ^2.

$$\text{Weight}_i \sim \text{normal}(\beta_0 + \beta_1 \text{Height}_i, \sigma^2)$$

The likelihood function for each of the Weight, which is specified in the MODEL statement, is as follows:

$$p(\text{Weight}|\beta_0, \beta_1, \sigma^2, \text{Height}_i) = \phi(\beta_0 + \beta_1 \text{Height}_i, \sigma^2)$$

where $p(\cdot|\cdot)$ denotes a conditional probability density and ϕ is the normal density. There are three parameters in the likelihood: β_0, β_1, and σ^2. You use the PARMS statement to indicate that these are the parameters in the model.

Suppose that you want to use the following three prior distributions on each of the parameters:

$$
\begin{aligned}
\pi(\beta_0) &= \phi(0, \text{var} = 1e6) \\
\pi(\beta_1) &= \phi(0, \text{var} = 1e6) \\
\pi(\sigma^2) &= f_{i\Gamma}(\text{shape} = 3/10, \text{scale} = 10/3)
\end{aligned}
$$

where $\pi(\cdot)$ indicates a prior distribution and $f_{i\Gamma}$ is the density function for the inverse-gamma distribution. The normal priors on β_0 and β_1 have large variances, expressing your lack of knowledge about the regression coefficients. The priors correspond to an equal-tail 95% credible intervals of approximately $(-2000, 2000)$ for β_0 and β_1. Priors of this type are often called *vague* or *diffuse* priors. See the section "Prior Distributions" on page 144 for more information. Typically diffuse prior distributions have little influence on the posterior distribution and are appropriate when stronger prior information about the parameters is not available.

A frequently used diffuse prior for the variance parameter σ^2 is the *inverse-gamma* distribution. With a shape parameter of $3/10$ and a scale parameter of $10/3$, this prior corresponds to an equal-tail 95% credible interval of $(1.7, 1e6)$, with the mode at 2.5641 for σ^2. Alternatively, you can use any other positive prior, meaning that the density support is positive on this variance component. For example, you can use the gamma prior.

According to Bayes' theorem, the likelihood function and prior distributions determine the posterior (joint) distribution of β_0, β_1, and σ^2 as follows:

$$\pi(\beta_0, \beta_1, \sigma^2 | \text{Weight}, \text{Height}) \quad \propto \quad \pi(\beta_0)\pi(\beta_1)\pi(\sigma^2)p(\text{Weight}|\beta_0, \beta_1, \sigma^2, \text{Height})$$

You do not need to know the form of the posterior distribution when you use PROC MCMC. PROC MCMC automatically obtains samples from the desired posterior distribution, which is determined by the prior and likelihood you supply.

The following statements fit this linear regression model with diffuse prior information:

```
ods graphics on;
proc mcmc data=class outpost=classout nmc=50000 thin=5 seed=246810;
   parms beta0 0 beta1 0;
   parms sigma2 1;
   prior beta0 beta1 ~ normal(mean = 0, var = 1e6);
   prior sigma2 ~ igamma(shape = 3/10, scale = 10/3);
   mu = beta0 + beta1*height;
   model weight ~ n(mu, var = sigma2);
run;
ods graphics off;
```

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 MCMC statement invokes the procedure and specifies the input data set class. The output data set classout contains the posterior samples for all of the model parameters. 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 5 samples is kept. Thinning is often used to reduce the correlations among posterior sample draws. In this example, 10,000 simulated values are saved in the classout data set. The SEED= option specifies a seed for the random number generator, which guarantees the reproducibility of the random stream. For more information about Markov chain sample size, burn-in, and thinning, see the section "Burn-in, Thinning, and Markov Chain Samples" on page 155.

The PARMS statements identify the three parameters in the model: beta0, beta1, and sigma2. Each statement also forms a block of parameters, where the parameters are updated simultaneously in each iteration. In this example, beta0 and beta1 are sampled jointly, conditional on sigma2; and sigma2 is sampled conditional on fixed values of beta0 and beta1. In simple regression models such as this, you expect the parameters beta0 and beta1 to have high posterior correlations, and placing them both in the same block improves the mixing of the chain—that is, the efficiency that the posterior parameter space is explored by the Markov chain. For more information, see the section "Blocking of Parameters" on page 3522. The PARMS statements also assign initial values to the parameters (see the section "Initial Values of the Markov Chains" on page 3527). The regression parameters are given 0 as their initial values, and the scale parameter sigma2 starts at value 1. If you do not provide initial values, the procedure chooses starting values for every parameter.

The PRIOR statements specify prior distributions for the parameters. The parameters beta0 and beta1 both share the same prior—a normal prior with mean 0 and variance $1e6$. The parameter sigma2 has an inverse-gamma distribution with a shape parameter of 3/10 and a scale parameter of 10/3. For a list of standard distributions that PROC MCMC supports, see the section "Standard Distributions" on page 3529.

The mu assignment statement calculates the expected value of Weight as a linear function of Height. The MODEL statement uses the shorthand notation, **n**, for the normal distribution to indicate that the response variable, Weight, is normally distributed with parameters mu and sigma2. The functional argument MEAN= in the normal distribution is optional, but you have to indicate whether sigma2 is a variance (VAR=), a standard deviation (SD=), or a precision (PRECISION=) parameter. See Table 52.2 in the section "MODEL Statement" on page 3511 for distribution specifications.

The distribution parameters can contain expressions. For example, you can write the MODEL statement as follows:

```
model weight ~ n(beta0 + beta1*height, var = sigma2);
```

Before you do any posterior inference, it is essential that you examine the convergence of the Markov chain (see the section "Assessing Markov Chain Convergence" on page 156). You cannot make valid inferences if the Markov chain has not converged. A very effective convergence diagnostic tool is the trace plot. Although PROC MCMC produces graphs at the end of the procedure output (see Figure 52.6), you should visually examine the convergence graph first.

The first table that PROC MCMC produces is the "Number of Observations" table, as shown in Figure 52.1. This table lists the number of observations read from the DATA= data set and the number of non-missing observations used in the analysis.

Figure 52.1 Observation Information

```
                    Simple Linear Regression

                       The MCMC Procedure

            Number of Observations Read        19
            Number of Observations Used        19
```

The "Parameters" table, shown in Figure 52.2, lists the names of the parameters, the blocking information (see the section "Blocking of Parameters" on page 3522), the sampling method used, the starting values (the section "Initial Values of the Markov Chains" on page 3527), and the prior distributions. You should to check this table to ensure that you have specified the parameters correctly, especially for complicated models.

Figure 52.2 Parameter Information

```
                              Parameters

                   Sampling         Initial
 Block   Parameter  Method            Value   Prior Distribution

    1    beta0      N-Metropolis          0   normal(mean = 0, var = 1e6)
    1    beta1      N-Metropolis          0   normal(mean = 0, var = 1e6)
    2    sigma2     N-Metropolis     1.0000   igamma(shape = 3/10, scale =
                                              10/3)
```

The "Tuning History" table, shown in Figure 52.3, shows how the tuning stage progresses for the multivariate random walk Metropolis algorithm used by PROC MCMC to generate samples from the posterior distribution. An important aspect of the algorithm is the calibration of the proposal distribution. The tuning of the Markov chain is broken into a number of phases. In each phase, PROC MCMC generates trial samples and automatically modifies the proposal distribution as a result of the acceptance rate (see the section "Tuning the Proposal Distribution" on page 3524). In this example, PROC MCMC found an acceptable proposal distribution after 7 phases, and this distribution is used in both the burn-in and sampling stages of the simulation.

The "Burn-In History" table shows the burn-in phase, and the "Sampling History" table shows the main phase sampling.

Figure 52.3 Tuning, Burn-In and Sampling History

```
                         Tuning History

                                          Acceptance
            Phase       Block      Scale        Rate

              1           1       2.3800      0.0420
                          2       2.3800      0.8860
              2           1       1.0938      0.2180
                          2      15.5148      0.3720
              3           1       0.8299      0.4860
                          2      15.5148      0.1260
              4           1       1.1132      0.4840
                          2       9.4767      0.0880
              5           1       1.4866      0.5420
                          2       5.1914      0.2000
              6           1       2.2784      0.4600
                          2       3.7859      0.3900
              7           1       2.8820      0.3360
                          2       3.7859      0.4020
```

Figure 52.3 *continued*

```
                         Burn-In History

                                     Acceptance
                  Block       Scale        Rate

                    1        2.8820      0.3400
                    2        3.7859      0.4150


                        Sampling History

                                     Acceptance
                  Block       Scale        Rate

                    1        2.8820      0.3284
                    2        3.7859      0.4008
```

For each posterior distribution, PROC MCMC also reports summary statistics (posterior means, standard deviations, and quantiles) and interval statistics (95% equal-tail and highest posterior density credible intervals), as shown in Figure 52.4. For more information about posterior statistics, see the section "Summary Statistics" on page 170.

Figure 52.4 MCMC Summary and Interval Statistics

```
                        Simple Linear Regression

                          The MCMC Procedure

                          Posterior Summaries

                              Standard            Percentiles
   Parameter        N      Mean   Deviation      25%       50%       75%

   beta0         10000   -142.6    33.9390    -164.5    -142.4    -120.5
   beta1         10000    3.8917    0.5427     3.5406    3.8906    4.2402
   sigma2        10000    136.8    51.7417     101.8     126.0     159.9


                          Posterior Intervals

     Parameter    Alpha    Equal-Tail Interval       HPD Interval

      beta0       0.050    -209.3    -76.1692     -209.7    -77.1624
      beta1       0.050     2.8317    4.9610       2.8280    4.9468
      sigma2      0.050    69.2208    265.5       58.2627    233.8
```

By default, PROC MCMC also computes a number of convergence diagnostics to help you determine whether the chain has converged. These are the Monte Carlo standard errors, the autocorrelations at selected lags, the Geweke diagnostics, and the effective sample sizes. These statistics are shown in Figure 52.5. For details and interpretations of these diagnostics, see the section "Assessing Markov Chain Convergence" on page 156.

The "Monte Carlo Standard Errors" table indicates that the standard errors of the mean estimates for each of the parameters 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, around 0.01. This means that only a fraction of the posterior variability is due to the simulation. The "Autocorrelations of the Posterior Samples" table shows that the autocorrelations among posterior samples reduce quickly and become almost nonexistent after lag 5. The "Geweke Diagnostics" table indicates that no parameter failed the test, and the "Effective Sample Sizes" table reports the number of effective sample sizes of the Markov chain.

Figure 52.5 MCMC Convergence Diagnostics

```
                      Simple Linear Regression

                         The MCMC Procedure

                    Monte Carlo Standard Errors

                                  Standard
              Parameter      MCSE    Deviation      MCSE/SD

                beta0       0.4576    33.9390        0.0135
                beta1      0.00731     0.5427        0.0135
                sigma2      0.7151    51.7417        0.0138


                      Posterior Autocorrelations

         Parameter    Lag 1      Lag 5      Lag 10     Lag 50

           beta0      0.2986    -0.0008    0.0162     0.0193
           beta1      0.2971     0.0000    0.0135     0.0161
           sigma2     0.2966     0.0062    0.0008    -0.0068


                         Geweke Diagnostics

               Parameter           z     Pr > |z|

                 beta0          0.1105    0.9120
                 beta1         -0.1701    0.8649
                 sigma2        -0.2175    0.8278


                      Effective Sample Sizes

                                Correlation
              Parameter    ESS         Time     Efficiency

                beta0     5501.1      1.8178       0.5501
                beta1     5514.8      1.8133       0.5515
                sigma2    5235.4      1.9101       0.5235
```

PROC MCMC produces a number of graphs, shown in Figure 52.6, which also aid convergence diagnostic checks. With the trace plots, there are 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 should 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.

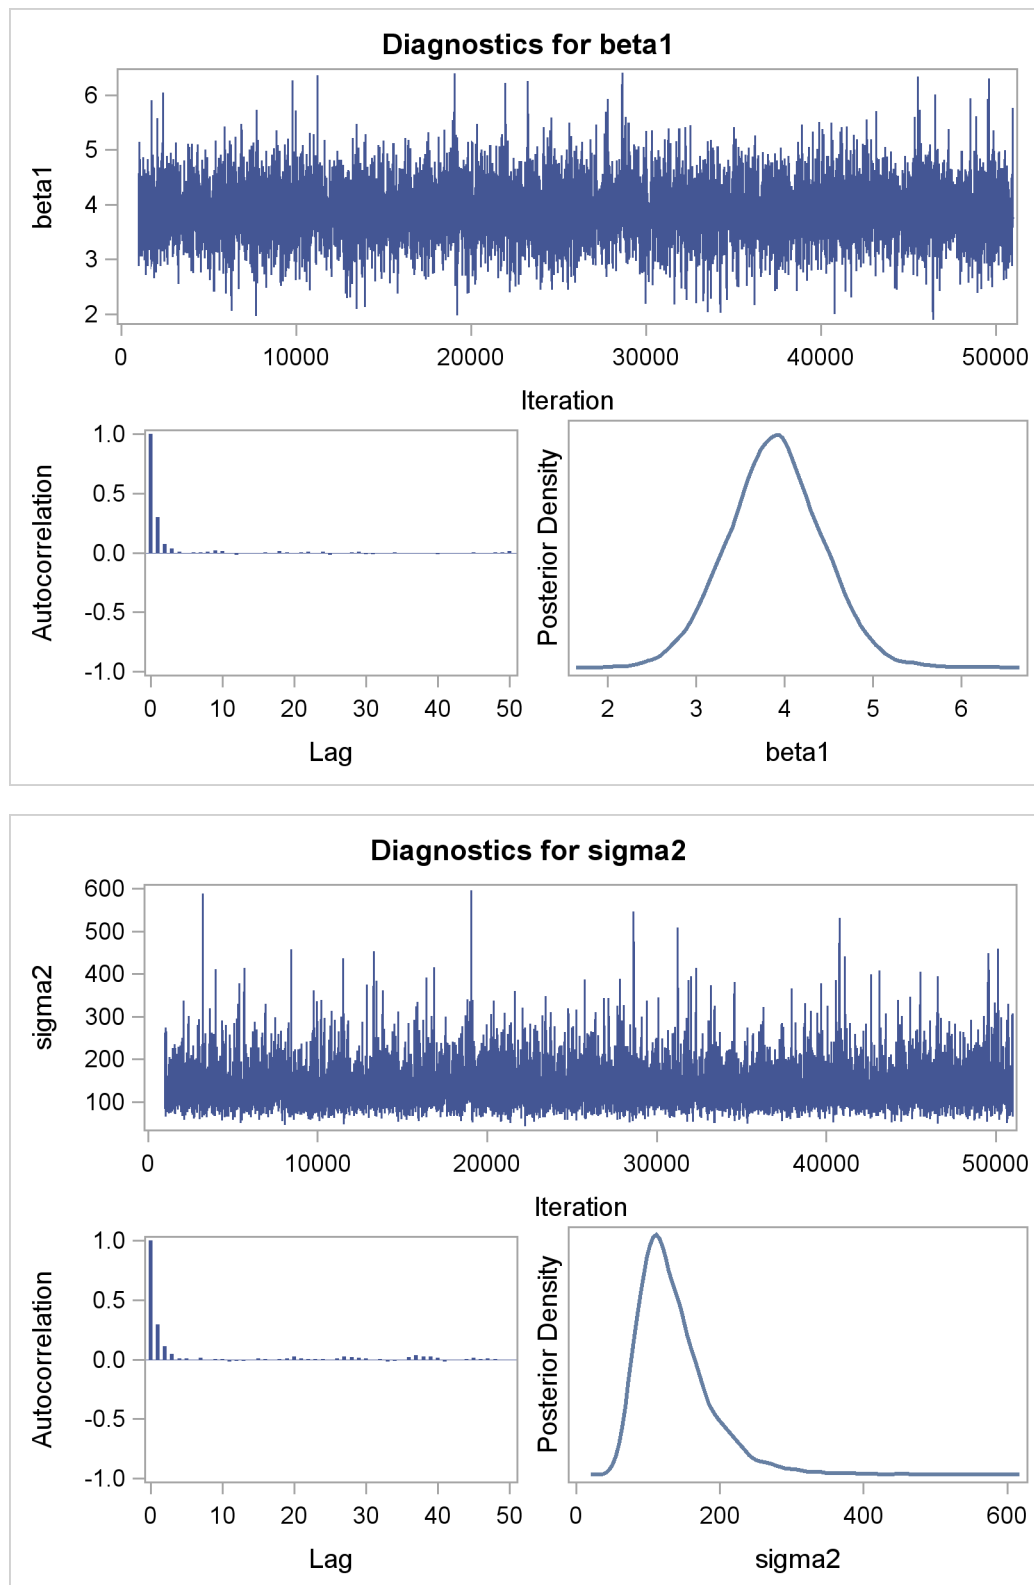**Figure 52.6** Diagnostic Plots for β_0, β_1 and σ^2

Figure 52.6 *continued*

In regression models such as this, you expect the posterior estimates to be very similar to the maximum likelihood estimators with noninformative priors on the parameters, The REG procedure produces the following fitted model (code not shown):

$$\text{Weight} = -143.0 + 3.9 \times \text{Height}$$

These are very similar to the means show in Figure 52.4. With PROC MCMC, you can carry out informative analysis that uses specifications to indicate prior knowledge on the parameters. Informative analysis is likely to produce different posterior estimates, which are the result of information from both the likelihood and the prior distributions. Incorporating additional information in the analysis is one major difference between the classical and Bayesian approaches to statistical inference.

The Behrens-Fisher Problem

One of the famous examples in the history of statistics is the Behrens-Fisher problem (Fisher 1935). Consider the situation where there are two independent samples from two different normal distributions:

$$y_{11}, y_{12}, \cdots, y_{1n_1} \sim \text{normal}(\mu_1, \sigma_1^2)$$
$$y_{21}, y_{22}, \cdots, y_{2n_2} \sim \text{normal}(\mu_2, \sigma_2^2)$$

Note that $n_1 \neq n_2$. When you cannot assume that the variances are equal, testing the hypothesis $H_0 : \mu_1 = \mu2$ is a difficult problem in the classical statistics framework, because the distribution under H_0 is not known. Within the Bayesian framework, this problem is straightforward because you can estimate the posterior distribution of $\mu_1 - \mu_2$ while taking into account the uncertainties in all of parameters by treating them as random variables.

Suppose that you have the following set of data:

```
title 'The Behrens-Fisher Problem';

data behrens;
   input y ind @@;
   datalines;
121 1   94 1 119 1 122 1 142 1 168 1 116 1
172 1 155 1 107 1 180 1 119 1 157 1 101 1
145 1 148 1 120 1 147 1 125 1 126 2 125 2
130 2 130 2 122 2 118 2 118 2 111 2 123 2
126 2 127 2 111 2 112 2 121 2
;
```

The response variable is y, and the ind variable is the group indicator, which takes two values: 1 and 2. There are 19 observations that belong to group 1 and 14 that belong to group 2.

The likelihood functions for the two samples are as follows:

$$
\begin{aligned}
p(y_{1i}|\mu_1,\sigma_1^2) &= \phi(y_{1i};\mu_1,\sigma_1^2) \quad \text{for} \quad i=1,\cdots,19 \\
p(y_{2j}|\mu_2,\sigma_2^2) &= \phi(y_{2j};\mu_2,\sigma_2^2) \quad \text{for} \quad j=1,\cdots,14
\end{aligned}
$$

Berger (1985) showed that a uniform prior on the support of the location parameter is a noninformative prior. The distribution is invariant under location transformations—that is, $\theta = \mu + c$. You can use this prior for the mean parameters in the model:

$$
\begin{aligned}
\pi(\mu_1) &\propto 1 \\
\pi(\mu_2) &\propto 1
\end{aligned}
$$

In addition, Berger (1985) showed that a prior of the form $1/\sigma^2$ is noninformative for the scale parameter, and it is invariant under scale transformations (that is $\tau = c\sigma^2$). You can use this prior for the variance parameters in the model:

$$
\begin{aligned}
\pi(\sigma_1^2) &\propto 1/\sigma_1^2 \\
\pi(\sigma_2^2) &\propto 1/\sigma_2^2
\end{aligned}
$$

The log densities of the prior distributions on σ_1^2 and σ_2^2 are:

$$
\begin{aligned}
\log(\pi(\sigma_1^2)) &= -\log(\sigma_1^2) \\
\log(\pi(\sigma_2^2)) &= -\log(\sigma_2^2)
\end{aligned}
$$

The following statements generate posterior samples of $\mu_1, \mu_2, \sigma_1^2, \sigma_2^2$, and the difference in the means: $\mu_1 - \mu_2$:

```
proc mcmc data=behrens outpost=postout seed=123
          nmc=40000 thin=10 monitor=(_parms_ mudif)
          statistics(alpha=0.01)=(summary interval);
   ods select PostSummaries PostIntervals;
   parm mu1 0 mu2 0;
   parm sig21 1;
   parm sig22 1;
   prior mu: ~ general(0);
   prior sig21 ~ general(-log(sig21));
   prior sig22 ~ general(-log(sig22));
   mudif = mu1 - mu2;
   if ind = 1 then
      llike = lpdfnorm(y, mu1, sqrt(sig21));
   else
      llike = lpdfnorm(y, mu2, sqrt(sig22));
   model general(llike);
run;
```

The PROC MCMC statement specifies an input data set (behrens), an output data set containing the posterior samples (postout), a random number seed, the simulation size, and the thinning rate. The MONITOR= option specifies a list of symbols, which can be either parameters or functions of

the parameters in the model, for which inference is to be done. The symbol _parms_ is a shorthand for all model parameters—in this case, mu1, mu2, sig21, and sig22. The symbol mudif is defined in the program as the difference between μ_1 and μ_2.

The **ods select** statement displays the summary statistics and interval statistics tables while excluding all other output. For a complete list of ODS tables that PROC MCMC can produce, see the sections "Displayed Output" on page 3565 and "ODS Table Names" on page 3569.

The STATISTICS= option calculates summary and interval statistics. The global suboption ALPHA=0.01 specifies 99% equal-tail and highest posterior density (HPD) credible intervals for all parameters.

The PARMS statements assign the parameters mu1 and mu2 to the same block, and sig21 and sig22 each to their own separate blocks. There are a total of three blocks. The PARMS statements also assign an initial value to each parameter.

The PRIOR statements specify prior distributions for the parameters. Because the priors are all nonstandard (uniform on the real axis for μ_1 and μ_2 and $1/\sigma^2$ for σ_1^2 and σ_2^2), you must use the GENERAL function here. The argument in the GENERAL function is an expression for the log of the distribution, up to an additive constant. This distribution can have any functional form, as long as it is programmable using SAS functions and expressions. Note that the function specifies a distribution on the log scale, not the original scale. The log of the prior on mu1 and mu2 is 0, and the log of the priors on sig21 and sig22 are -log(sig21) and -log(sig22) respectively. See the section "Specifying a New Distribution" on page 3540 for more information about how to specify an arbitrary distribution.

The mudif assignment statement calculates the difference between mu1 and mu2. The IF-ELSE statements enable different y's to have different log-likelihood functions, depending on their group indicator ind. The function LPDFNORM is a PROC MCMC function that calculates the log density of a normal distribution. See the section "Using Density Functions in the Programming Statements" on page 3540 for more details. The MODEL statement specifies that llike is the log likelihood for each observation in the model.

Figure 52.7 displays the posterior summary and interval statistics.

Figure 52.7 Posterior Summary and Interval Statistics

```
                    The Behrens-Fisher Problem

                       The MCMC Procedure

                       Posterior Summaries

                                Standard              Percentiles
   Parameter        N      Mean   Deviation        25%       50%       75%

   mu1           4000     134.8      6.0065      130.9     134.7     138.7
   mu2           4000     121.4      1.9150      120.2     121.4     122.7
   sig21         4000     683.2       259.9      507.8     630.1     792.3
   sig22         4000   51.3975     24.2881    35.0212   45.7449   61.2582
   mudif         4000   13.3596      6.3335     9.1732   13.4078   17.6332
```

Figure 52.7 *continued*

```
                      Posterior Intervals

     Parameter    Alpha      Equal-Tail Interval          HPD Interval

     mu1          0.010         118.7       150.6        119.3       151.0
     mu2          0.010         115.9       126.6        116.2       126.7
     sig21        0.010         292.0      1821.1        272.8      1643.7
     sig22        0.010       18.5883       158.8      16.3730       140.5
     mudif        0.010       -3.2537     29.9987      -3.1915     30.0558
```

The mean difference has a posterior mean value of 13.36, and the lower endpoints of the 99% credible intervals are negative. This suggests that the mean difference is positive with a high probability. However, if you want to estimate the probability that $\mu_1 - \mu_2 > 0$, you can do so as follows.

The following statements produce Figure 52.8:

```
proc format;
    value diffmt low-0 = 'mu1 - mu2 <= 0' 0<-high = 'mu1 - mu2 > 0';
run;

proc freq data = postout;
    tables mudif /nocum;
    format mudif diffmt.;
run;
```

The sample estimate of the posterior probability that $\mu_1 - \mu_2 > 0$ is 0.98. This example illustrates an advantage of Bayesian analysis. You are not limited to making inferences based on model parameters only. You can accurately quantify uncertainties with respect to any function of the parameters, and this allows for flexibility and easy interpretations in answering many scientific questions.

Figure 52.8 Estimated Probability of $\mu_1 - \mu_2 > 0$.

```
                    The Behrens-Fisher Problem

                       The FREQ Procedure

                 mudif      Frequency      Percent
          ---------------------------------------
          mu1 - mu2 <= 0          77         1.93
          mu1 - mu2 > 0         3923        98.08
```

Mixed-Effects Model

This example illustrates how you can fit a mixed-effects model in PROC MCMC. PROC MCMC offers you the ability to model beyond the normal likelihood (see "Example 52.5: Random-Effects Models" on page 3608), and you can model as many levels of random effects as are needed with this procedure.

Consider a scenario in which data are collected in groups and you wish to model group-specific effects. You can use a mixed-effects model (sometimes also known as a random-effects model or a variance-components model):

$$y_{ij} = \beta_0 + \beta_1 x_{ij} + \gamma_i + e_{ij}, \quad e_{ij} \sim \text{normal}(0, \sigma^2)$$

where $i = 1, 2, \cdots, I$ is the group index and $j = 1, 2, \cdots, n_i$ indexes the observations in the ith group. In the regression model, the fixed effects β_0 and β_1 are the intercept and the coefficient for variable x_{ij}, respectively. The random effects γ_i is the mean for the ith group, and e_{ij} are the error term.

Consider the following SAS data set:

```
title 'Mixed-Effects Model';

data heights;
   input Family G$ Height @@;
   datalines;
1 F 67    1 F 66    1 F 64    1 M 71    1 M 72    2 F 63
2 F 63    2 F 67    2 M 69    2 M 68    2 M 70    3 F 63
3 M 64    4 F 67    4 F 66    4 M 67    4 M 67    4 M 69
;

data input;
   set heights;
   if g eq 'F' then gender = 1;
   else gender = 0;
   drop g;
run;
```

The response variable Height measures the heights (in inches) of 18 individuals. The individuals are classified according to Family and Gender.

Height is assumed to be normally distributed:

$$y_{ij} \sim \text{normal}(\mu_{ij}, \sigma^2), \quad \mu_{ij} = \beta_0 + \beta_1 x_{ij} + \gamma_i$$

which corresponds to a normal likelihood as follows:

$$p(y_{ij} | \mu_{ij}, \sigma^2) = \phi(\mu_{ij}, \text{var} = \sigma^2)$$

The priors on the parameters $\beta_0, \beta_1, \gamma_i$ are assumed to be normal as well:

$$
\begin{aligned}
\pi(\beta_0) &= \phi(0, \text{var} = 1e5) \\
\pi(\beta_1) &= \phi(0, \text{var} = 1e5) \\
\pi(\gamma_i) &= \phi(0, \text{var} = \sigma_\gamma^2)
\end{aligned}
$$

Priors on the variance terms, σ^2 and σ_γ^2, are inverse-gamma:

$$
\begin{aligned}
\pi(\sigma^2) &= f_{i\Gamma}(\text{shape} = 0.001, \text{scale} = 1000) \\
\pi(\sigma_\gamma^2) &= f_{i\Gamma}(\text{shape} = 0.001, \text{scale} = 1000)
\end{aligned}
$$

where $f_{i\Gamma}$ denotes the density function of an inverse-gamma distribution.

The following statements fit a linear random-effects model to the data and produce the output shown in Figure 52.9 and Figure 52.10:

```
ods graphics on;
proc mcmc data=input outpost=postout thin=10 nmc=50000 seed=7893
          monitor=(b0 b1);
   ods select PostSummaries PostIntervals tadpanel;
   array gamma[4];

   parms b0 0 b1 0 gamma: 0;
   parms s2 1 ;
   parms s2g 1;

   prior b: ~ normal(0, var = 10000);
   prior gamma: ~ normal(0, var = s2g);
   prior s2: ~ igamma(0.001, scale = 1000);

   mu = b0 + b1 * gender + gamma[family];
   model height ~ normal(mu, var = s2);
run;
ods graphics off;
```

The statements are very similar to those shown in the previous two examples. The **ods graphics on** statement requests ODS Graphics. The PROC MCMC statement specifies the input and output data sets, the simulation size, the thinning rate, and a random number seed. The MONITOR= option indicates that the model parameters b0 and b1 are the quantities of interest. The **ods select** statement displays the summary statistics table, the interval statistics table, and the diagnostics plots.

The ARRAY statement defines a one-dimensional array, gamma, with 4 elements. You can refer to the array elements with variable names (gamma1 to gamma4 by default) or with subscripts, such as gamma[2]. To indicate subscripts, you must use either brackets [ ] or braces { }, but not parentheses ( ). Note that this is different from the way subscripts are indicated in the DATA step. See the section "ARRAY Statement" on page 3506 for more information.

The PRIOR statements specify priors for all the parameters. The notation b: is a shorthand for all symbols that start with the letter 'b'. In this example, it includes b0 and b1. Similarly, gamma: stands for all four gamma parameters, and s2: stands for both s2 and s2g. This shorthand notation can save you some typing, and it keeps your statements tidy.

The mu assignment statement calculates the expected value of height in the random-effects model. The symbol family is a data set variable that indexes family. Here gamma[family] is the random effect for the value of family.

Finally, the MODEL statement specifies the likelihood function for height.

The posterior summary and interval statistics for b0 and b1 are shown in Figure 52.9.

Figure 52.9 Posterior Summary and Interval Statistics

```
                           Mixed-Effects Model

                            The MCMC Procedure

                            Posterior Summaries

                                Standard                 Percentiles
    Parameter           N      Mean    Deviation        25%        50%        75%

    b0               5000   66.2685      19.1176    56.0024    66.7260    77.2356
    b1               5000   -3.3492       6.3886    -7.4268    -3.2799     0.6078

                            Posterior Intervals

       Parameter     Alpha      Equal-Tail Interval          HPD Interval

       b0            0.050       26.2226       103.3      27.1749      103.6
       b1            0.050      -16.2018      9.6267     -17.0757      8.5265
```

Trace plots, autocorrelation plots, and posterior density plots for b1 and logpost are shown in Figure 52.10. The mixing of b1 looks good. The convergence plots for the other parameters also look reasonable, and are not shown here.

**Figure 52.10** Plots for b_1 and Log of the Posterior Density

Figure 52.10 *continued*

From the interval statistics table, you see that both of the equal-tail and HPD intervals for β_0 are positive, strongly indicating the positive effect of the parameter. On the other hand, both intervals for β_1 cover the value zero, indicating that gender does not have a strong impact on predicting height in this model.

Syntax: MCMC Procedure

The following statements can be used with PROC MCMC:

> **PROC MCMC** *options* **;**
>   **ARRAY** *array specification* **;**
>   **BEGINCNST/ENDCNST** *constant segment specification* **;**
>   **BEGINHYPER/ENDHYPER** *hyperprior segment specification* **;**
>   **BEGINPRIOR/ENDPRIOR** *prior segment specification* **;**
>   **BY** *variables* **;**
>   **MODEL** *statistical model specification* **;**
>   **PARMS** *parameters and starting values* **;**
>   **PRIOR/HYPERPRIOR** *prior or hyperprior specification* **;**
>   **Program statements** **;**
>   **UDS** *user defined sampler specification* **;**

The PARMS statements declare parameters in the model and assign optional starting values for the Markov chain. The PRIOR/HYPERPRIOR statements specify the prior distributions of the parameters. The MODEL statements specify the log-likelihood functions for the response variables. These statements form the basis of every Bayesian model.

In addition, you can use the ARRAY statement to define constant or parameter arrays, the BEGINCNST/ENDCNST and similar statements to save unnecessary evaluation and reduce simulation time, the program statements to specify more complicated models that you wish to fit, and finally the UDS statements to define your own Gibbs samplers to sample any parameters in the model.

The following sections provide a description of each of these statements.

PROC MCMC Statement

PROC MCMC *options* **;**

This statement invokes PROC MCMC.

A number of options are available in the PROC MCMC statement; the following table categorizes them according to function.

Table 52.1   PROC MCMC Statement Options

Option	Description
Basic options	
DATA=	names the input data set
OUTPOST=	names the output data set for posterior samples of parameters
Debugging output	
LIST	displays model program and variables
LISTCODE	displays compiled model program
TRACE	displays detailed model execution messages
Frequently used MCMC options	
MAXTUNE=	specifies the maximum number of tuning loops
MINTUNE=	specifies the minimum number of tuning loops
NBI=	specifies the number of burn-in iterations
NMC=	specifies the number of MCMC iterations, excluding the burn-in iterations
NTU=	specifies the number of tuning iterations
PROPCOV=	controls options for constructing the initial proposal covariance matrix
SEED=	specifies the random seed for simulation
THIN=	specifies the thinning rate
Less frequently used MCMC options	
ACCEPTTOL=	specifies the acceptance rate tolerance
DISCRETE=	controls sampling discrete parameters
INIT=	controls generating initial values

Table 52.1 (continued)

Option	Description
PROPDIST=	specifies the proposal distribution
SCALE=	specifies the initial scale applied to the proposal distribution
TARGACCEPT=	specifies the target acceptance rate for random walk sampler
TARGACCEPTI=	specifies the target acceptance rate for independence sampler
TUNEWT=	specifies the weight used in covariance updating
Summary, diagnostics, and plotting options	
AUTOCORLAG=	specifies the number of autocorrelation lags used to compute effective sample sizes and Monte Carlo errors
DIAGNOSTICS=	controls the convergence diagnostics
DIC	computes deviance information criterion (DIC)
MONITOR=	monitors a list of parameters of interest
PLOTS=	controls plotting
STATISTICS=	controls posterior statistics
Other Options	
INF=	specifies the machine numerical limit for infinity
JOINTMODEL	specifies joint log-likelihood function
MISSING=	indicates how missing values are handled.
SIMREPORT=	controls the frequency of report for expected run time
SINGDEN=	specifies the singularity tolerance

These options are described in alphabetical order.

ACCEPTTOL=_n_

specifies a tolerance for acceptance probabilities. By default, ACCEPTTOL=0.075.

AUTOCORLAG=_n_

ACLAG=_n_

specifies the maximum number of autocorrelation lags used in computing the effective sample size; see the section "Effective Sample Size" on page 169 for more details. The value is used in the calculation of the Monte Carlo standard error; see the section "Standard Error of the Mean Estimate" on page 170. By default, AUTOCORLAG=MIN(500, MCsample/4), where MCsample is the Markov chain sample size kept after thinning—that is, MCsample $= \left\lceil \frac{\text{NMC}}{\text{NTHIN}} \right\rceil$. If AUTOCORLAG= is set too low, you might observe significant lags, and the effective sample size cannot be calculated accurately. A WARNING message appears, and you can either increase AUTOCORLAG= or NMC=, accordingly.

DISCRETE=_keyword_

specifies the proposal distribution used in sampling discrete parameters. The default is DISCRETE=BINNING.

The _keyword_ values are as follows:

BINNING

uses continuous proposal distributions for all discrete parameter blocks. The proposed

sample is then discretized (binned) before further calculations. This sampling method approximates the correlation structure among the discrete parameters in the block and could improve mixing in some cases.

GEO

uses independent symmetric geometric proposal distributions for all discrete parameter blocks. This proposal does not take parameter correlations into account. However, it can work better than the BINNING option in cases where the range of the parameters is relatively small and a normal approximation can perform poorly.

DIAGNOSTICS=NONE | *(keyword-list)*
DIAG=NONE | *(keyword-list)*

specifies options for MCMC convergence diagnostics. By default, PROC MCMC computes the Geweke test, sample autocorrelations, effective sample sizes, and Monte Carlo errors. The Raftery-Lewis and Heidelberger-Welch tests are also available. See the section "Assessing Markov Chain Convergence" on page 156 for more details on convergence diagnostics. You can request all of the diagnostic tests by specifying DIAGNOSTICS=ALL. You can suppress all the tests by specifying DIAGNOSTICS=NONE.

The following options are available.

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 AUTO-CORR(LAGS=(1 5 35))) computes all tests with default settings and autocorrelations at lags 1, 5, and 35.

AUTOCORR *< (autocorr-options) >*

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

LAGS | AC=*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. See the section "Effective Sample Size" on page 169 for more details.

GEWEKE *< (Geweke-options) >*

computes the Geweke spectral density diagnostics; this is a two-sample t-test between the first f_1 portion and the last f_2 portion of the chain. See the section "Geweke Diagnostics" on page 163 for more details. The default is FRAC1=0.1 and FRAC2=0.5, but you can choose other fractions by using the following *Geweke-options*:

FRAC1 | F1=*value*

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

FRAC2 | F2=*value*

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

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

>   computes the Heidelberger and Welch diagnostic (which consists of a stationarity test and a halfwidth 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 halfwidth test is then carried out. See the section "Heidelberger and Welch Diagnostics" on page 165 for more details.
>
>   These diagnostics are not performed by default. You can specify the DIAGNOS- TICS=HEIDELBERGER option to request these diagnostics, and you can also specify suboptions, such as DIAGNOSTICS=HEIDELBERGER(EPS=0.05), as follows:

SALPHA=*value*

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

HALPHA=*value*

>   specifies the α level ($0 < \alpha < 1$) for the halfwidth test. By default, HAL-PHA=0.05.

EPS=*value*

>   specifies a small positive number ϵ such that if the halfwidth is less than ϵ times the sample mean of the retaining iterates, the halfwidth test is passed. By default, EPS=0.1.

MCSE
MCERROR

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

NONE

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

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

>   computes the Raftery and Lewis diagnostics, which evaluate the accuracy of the esti- mated 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, $\Pr(Q - R \leq \hat{P}_Q \leq Q + R) = S$. See the section "Raftery and Lewis Diagnostics" on page 166 for more details. The *Raftery-options* enable you to specify Q, R, S, and a precision level ϵ for a stationary test.
>
>   These diagnostics are not performed by default. You can specify the DIAGNOS- TICS=RAFERTY option to request these diagnostics, and you can also specify sub- options, such as DIAGNOSTICS=RAFERTY(QUANTILE=0.05), as follows:

QUANTILE | Q=*value*

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

ACCURACY | 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.

PROB | S=*value*

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

EPS=*value*

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

DIC

> computes the Deviance Information Criterion (DIC). DIC is calculated using the posterior mean estimates of the parameters. See the section "Deviance Information Criterion (DIC)" on page 172 for more details.

DATA=*SAS-data-set*

> specifies the input data set. Observations in this data set are used to compute the log-likelihood function that you specify with PROC MCMC statements.

INF=*value*

> specifies the numerical definition of infinity in the procedure. The default is INF= 1E15. For example, PROC MCMC considers 1E16 to be outside of the support of the normal distribution and assigns a missing value to the log density evaluation. You can select a larger value with the INF= option. The minimum value allowed is 1E10.

INIT=*(keyword-list)*

> specifies options for generating the initial values for the parameters. These options apply only to prior distributions that are recognized by PROC MCMC. See the section "Standard Distributions" on page 3529 for a list of these distributions. If either of the functions GENERAL or DGENERAL is used, you must supply explicit initial values for the parameters. By default, INIT=MODE. The following keywords are used:

MODE

> > uses the mode of the prior density as the initial value of the parameter, if you did 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 the mean is outside of the support or on the boundary, which can happen if the prior distribution is truncated, a random number drawn from the prior is used as the initial value.

PINIT

tabulates parameter values after the tuning phase. This option also tabulates the tuned proposal parameters used by the Metropolis algorithm. These proposal parameters include covariance matrices for continuous parameters and probability vectors for discrete parameters for each block. By default, PROC MCMC does not display the initial values or the tuned proposal parameters after the tuning phase.

RANDOM

generates a random number from the prior density and uses it as the initial value of the parameter, if you did not provide one.

REINIT

resets the parameters, after the tuning phase, with the initial values that you provided explicitly or that were assigned by the procedure. By default, PROC MCMC does not reset the parameters because the tuning phase usually moves the Markov chains to a more favorable place in the posterior distribution.

LIST

displays the model program and variable lists. The LIST option is a debugging feature and is not normally needed.

LISTCODE

displays the compiled program code. The LISTCODE option is a debugging feature and is not normally needed.

JOINTMODEL

JOINTLLIKE

specifies how the likelihood function is calculated. By default, PROC MCMC assumes that the observations in the data set are independent so that the joint log-likelihood function is the sum of the individual log-likelihood functions for the observations, where the individual log-likelihood function is specified in the MODEL statement. When your data are not independent, you can specify the JOINTMODEL option to modify the way that PROC MCMC computes the joint log-likelihood function. In this situation, PROC MCMC no longer steps through the input data set to sum the individual log likelihood.

To use this option correctly, you need to do the following two things:

- create ARRAY symbols to store all data set variables that are used in the program. This can be accomplished with the BEGINCNST and ENDCNST statements.
- program the joint log-likelihood function by using these ARRAY symbols only. The MODEL statement specifies the joint log-likelihood function for the entire data set. Typically, you use the function GENERAL in the MODEL statement.

See the sections "BEGINCNST/ENDCNST Statement" on page 3507 and "Modeling Dependent Data" on page 3549 for details.

MAXTUNE=_n_

specifies an upper limit for the number of proposal tuning loops. By default, MAXTUNE=24. See the section "Covariance Tuning" on page 3525 for more details.

MINTUNE=*n*

> specifies a lower limit for the number of proposal tuning loops. By default, MINTUNE=2.
> See the section "Covariance Tuning" on page 3525 for more details.

MISSING=*keyword*

MISS=*keyword*

> specifies how missing values are handled (see the section "Handling of Missing Data" on
> page 3559 for more details). The default is MISSING=COMPLETECASE.

> ### ALLCASE | AC
>
> > gives you the option to model the missing values in an all-case analysis. You can use
> > any techniques that you see fit, for example, fully Bayesian or multiple imputation.
>
> ### COMPLETECASE | CC
>
> > assumes a complete case analysis, so all observations with missing variable values are
> > discarded prior to the simulation.

MONITOR= (*symbol-list***)**

> specifies a list of program symbols for which a posterior analysis is to be done. The symbols
> can be any of the following: model parameters (symbols in the PARMS statement), secondary
> parameters (assigned using the operator "="), the log of the posterior density (LOGPOST),
> the log of the prior density (LOGPRIOR), the log of the hyperprior density (LOGHYPER),
> if the HYPER statement is used, or the log of the likelihood function (LOGLIKE). You can
> use the keyword _PARMS_ as a shorthand for all of the model parameters. PROC MCMC
> performs only posterior analyses, such as plotting, diagnostics and summaries, on the symbols
> selected with the MONITOR= option. By default MONITOR=_PARMS_.
>
> Posterior samples of any secondary parameters listed in the MONITOR= option are saved in
> the OUTPOST= data set. Posterior samples of any model parameters are always saved to the
> OUTPOST= data set, regardless whether they are listed in the MONITOR= option.
>
> You can also choose to monitor an entire array by specifying the name of the array. For this
> to work, at least one of the array elements must be referenced by its variable name in either
> the programming statements or the PARMS statements. Otherwise, PROC MCMC treats the
> array as a temporary array and does not keep track of any elements, meaning that you cannot
> monitor an array by specifying the name only.

NBI=*n*

> specifies the number of burn-in iterations to perform before beginning to save parameter esti-
> mate chains. By default, NBI=1000. See the section "Burn-in, Thinning, and Markov Chain
> Samples" on page 155 for more details.

NMC=*n*

> specifies the number of iterations in the main simulation loop. This is the MCMC sample size
> if THIN=1. By default, NMC=1000.

NTU=*n*

> specifies the number of iterations to use in each proposal tuning phase. By default, NTU=500.

OUTPOST=_SAS-data-set_

> specifies an output data set that contains the posterior samples of all model parameters, the iteration numbers (variable name ITERATION), the log of the posterior density (LOGPOST), the log of the prior density (LOGPRIOR), the log of the hyperprior density (LOGHYPER), if the HYPER statement is used, and the log likelihood (LOGLIKE). Any secondary parameters (assigned using the operator "=") listed in the MONITOR= option are saved to this data set. By default, no OUTPOST= data set is created.

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

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

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

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

> You must enable ODS Graphics before requesting plots—for example, like this:

```
ods graphics on;

proc mcmc;
   ...;
run;

ods graphics off;
```

> If you have enabled ODS Graphics but do not specify the PLOTS= option, then PROC MCMC 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).
>
> The _global-plot-options_ include the following:

> **FRINGE**
>
> > adds a fringe plot to the horizontal axis of the density plot.

> **GROUPBY|GROUP=PARAMETER | TYPE**
>
> > specifies how the plots are grouped when there is more than one type of plot. GROUPBY=PARAMETER is the default. The choices are as follows:
> >
> > **TYPE**
> >
> > > specifies that the plots are grouped by type.
> >
> > **PARAMETER**
> >
> > > specifies that the plots are grouped by parameter.

LAGS=n

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

SMOOTH

> smoothes the trace plot with a fitted penalized B-spline curve (Eilers and Marx 1996).

UNPACKPANEL

UNPACK

> specifies that all paneled plots are to be unpacked, so that each plot in a panel is displayed separately.

The *plot-requests* are as follows:

ALL

> requests all types of plots. PLOTS=ALL is equivalent to specifying PLOTS=(TRACE AUTOCORR 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 with four parameters, X1–X4. Displays for various specifications are depicted as follows.

- 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)

PROPCOV=*value*

> specifies the method used in constructing the initial covariance matrix for the Metropolis-Hastings algorithm. The QUANEW and NMSIMP methods find numerically approximated covariance matrices at the optimum of the posterior density function with respect to all continuous parameters. The optimization does not apply to discrete parameters. The tuning phase starts at the optimized values; in some problems, this can greatly increase convergence performance. If the approximated covariance matrix is not positive definite, then an identity matrix is used instead. Valid values are as follows:

> **IND**
>
>> uses the identity covariance matrix. This is the default. See the section "Tuning the Proposal Distribution" on page 3524.

> **QUANEW<(***optimize-options***)>**
>
>> performs a quasi-Newton optimization. See following for a description of the *optimize-options*.

> **NMSIMP | SIMPLEX<(***optimize-options***)>**
>
>> performs a Nelder-Mead simplex optimization. See following for a description of the *optimize-options*.

> The *optimize-options* are as follows:

> **ITPRINT**
>
>> prints optimization iteration steps and results.

PROPDIST=*value*

> specifies a proposal distribution for the Metropolis algorithm. See the section "Metropolis and Metropolis-Hastings Algorithms" on page 152. You can also use PARMS statement option (see the section "PARMS Statement" on page 3514) to change the proposal distribution for a particular block of parameters. Valid values are as follows:

> **NORMAL**
> **N**
>
>> specifies a normal distribution as the proposal distribution. This is the default.

> **T<(***df***)>**
>
>> specifies a *t*-distribution with the degrees of freedom *df*. By default, *df*=3. If *df* > 100, the normal distribution is used since the two distributions are almost identical.

SCALE=*value*

> controls the initial multiplicative scale to the covariance matrix of the proposal distribution. By default, SCALE=2.38. See the section "Scale Tuning" on page 3525 for more details.

SEED=*n*

> specifies the random number seed. By default, SEED=0, and PROC MCMC gets a random number seed from the clock.

SIMREPORT=*n*

> controls the number of times that PROC MCMC reports the expected run time of the simulation. This can be useful for monitoring the progress of CPU-intensive programs. For example, with SIMREPORT=2, PROC MCMC reports the simulation progress twice. By default, SIMREPORT=0, and there is no reporting. The expected run times are displayed in the log file.

SINGDEN=*value*

> defines the singularity criterion in the procedure. By default, SINGDEN=1E-11. The *value* indicates the exclusion of an endpoint in an interval. The mathematical notation "(0" is equivalent to "[*value*" in PROC MCMC—that is, $x < 0$ is treated as $x \leq value$ in the procedure. The maximum SINGDEN allowed is $1E - 6$.

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

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

> specifies options for posterior statistics. By default, PROC MCMC 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. See the section "Summary Statistics" on page 170 for more details. You can request all of the posterior statistics by specifying STATS=ALL. You can suppress all the calculations by specifying STATS=NONE.
>
> The *global-stats-options* includes the following:

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

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

> The *stats-requests* include the following:

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

> **CORR**
>
> > computes the posterior correlation matrix.

> **COV**
>
> > computes the posterior covariance matrix.

> **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.

INTERVAL

INT

computes the $100(1-\alpha)\%$ equal-tail and HPD credible intervals for each variable. See the sections "Equal-Tail Credible Interval" on page 171 and "Highest Posterior Density (HPD) Interval" on page 171 for details. By default, ALPHA=0.05, but you can use the global ALPHA= option to request other intervals of any probabilities.

NONE

suppresses all of the statistics.

TARGACCEPT=value

specifies the target acceptance rate for the random walk based Metropolis algorithm. See the section "Metropolis and Metropolis-Hastings Algorithms" on page 152. The numeric *value* must be between 0.01 and 0.99. By default, TARGACCEPT=0.45 for models with 1 parameter; TARGACCEPT=0.35 for models with 2, 3, or 4 parameters; and TARGACCEPT=0.234 for models with more than 4 parameters (Roberts, Gelman, and Gilks 1997; Roberts and Rosenthal 2001).

TARGACCEPTI=value

specifies the target acceptance rate for the independence sampler algorithm. The independence sampler is used for blocks of binary parameters. See the section "Independence Sampler" on page 153 for more details. The numeric *value* must be between 0 and 1. By default, TARGACCEPTI=0.6.

THIN=n

NTHIN=n

controls the thinning rate of the simulation. PROC MCMC keeps every *n*th simulation sample and discards the rest. All of the posterior statistics and diagnostics are calculated using the thinned samples. By default, THIN=1. See the section "Burn-in, Thinning, and Markov Chain Samples" on page 155 for more details.

TRACE

displays the result of each operation in each statement in the model program as it is executed. This debugging option is very rarely needed, and it produces voluminous output. If you use this option, also use small NMC=, NBI=, MAXTUNE=, and NTU= numbers.

TUNEWT=value

specifies the multiplicative weight used in updating the covariance matrix of the proposal distribution. The numeric *value* must be between 0 and 1. By default, TUNEWT=0.75. See the section "Covariance Tuning" on page 3525 for more details.

ARRAY Statement

ARRAY *arrayname <{ dimensions }> <$> <variables and constants>* **;**

The ARRAY statement is used to associate a name (of no more than eight characters) with a list of variables and constants. The ARRAY statement is similar to, but not the same as, the ARRAY statement in the DATA step, and it is the same as the ARRAY statements in the NLIN, NLP, NLMIXED, and MODEL procedures. The array name is used with subscripts in the program to refer to the array elements, as illustrated in the following statements:

```
array r[8] r1–r8;

do i = 1 to 8;
   r[i] = 0;
end;
```

The ARRAY statement does not support all the features of the ARRAY statement in the DATA step. Implicit indexing of variables cannot be used; all array references must have explicit subscript expressions. Only exact array dimensions are allowed; lower-bound specifications are not supported. A maximum of six dimensions is allowed.

Both variables and constants can be array elements. Constant array elements cannot have values assigned to them while variables can. Both the dimension specification and the list of elements are optional, but at least one must be specified. When the list of elements is not specified or fewer elements than the size of the array are listed, array variables are created by appending element numbers to the array name to complete the element list. You can index array elements by enclosing a subscript in braces ({}) or brackets ([]), but not in parentheses (()). The parentheses are reserved for function calls only.

For example, the following statement names an array day:

```
array day[365];
```

By default, the variables names are day1 to day365. However, since **day** is a SAS function, any subscript that uses parentheses gives you the wrong results. The expression **day(4)** returns the value 5 and does not reference the array element day4.

BEGINCNST/ENDCNST Statement

BEGINCNST ;

ENDCNST ;

The BEGINCNST statement indicates that PROC MCMC processes the programming statements that follow it once for every observation in the input data set before the first iteration. A subsequent ENDCNST statement indicates a return to normal processing. These statements are used, for example, to define constants and read data set variables into arrays. Reading data set variables into arrays is particularly useful when you use the JOINTMODEL option (see the section "Modeling Dependent Data" on page 3549) and the UDS statement (see the section "UDS Statement" on page 3517). You can also use the BEGINCNST and ENDCNST statements to assign initial values to the parameters (see the section "Assignments of Parameters" on page 3527).

Like the DATA step, PROC MCMC processes every programming statement for each observation in the input data set. Statements that calculate constants do not need to be evaluated multiple times throughout the Markov simulation. The BEGINCNST and ENDCNST statements can reduce redundant processing. For example, you can set up an array of constants in this way:

```
array cnst[17];
begincnst;
   do i = 1 to 17;
      cnst[i] = i * i;
   end;
endcnst;
```

Note that the DO-loop is evaluated once per observation in the data set before the first Markov chain iteration. If you have many observations in the data set, this DO-loop is evaluated many times over. A more efficient way is to create an observation index variable in the input data set. Call it ind for example. And suppose that you have a total number of five observations. You can use the IF statement to further reduce redundant processing:

```
array cnst[17];
begincnst;
   if ind eq 5 then do;
      do i = 1 to 17;
         cnst[i] = i * i;
      end;
   end;
endcnst;
```

The DO-loop is executed only at the last observation of the input data set.

You can also use the BEGINCNST and ENDCNST statements to store data set variables in arrays, which is useful when you model data that are not independent. See the section "Modeling Dependent Data" on page 3549 for more details. To save a SAS data set variable in an array, you need an array of the size of the variable and an observation index variable. The following statement is a simple example that creates an artificial data set with 4 observations:

```
data a;
   input y1 @@;
   ind = _n_;
   datalines;
   2 2.7 8.7 -3.4
;
```

In the data set a, y1 is the variable that is saved in an array in PROC MCMC, and ind is the observation index. In PROC MCMC, you can use the BEGINCNST and ENDCNST statements in the following way:

```
proc mcmc data=a;
   array m[4];
   begincnst;
      m[ind] = y1;
   endcnst;
   ...
```

The ARRAY statement in PROC MCMC names an array m of size 4, the same length as the variable y1. When PROC MCMC steps through the data set, it copies y1 to elements of m, as indexed by ind. After the BEGINCNST and ENDCNST statements are processed, the array m now stores all values of the variable y1.

Note that you want to avoid naming the storage array y if there are variables named y1 in the input data set. When no variable list is specified, the ARRAY statement creates variable names by appending index numbers to the array name. For example,

```
array y[4];
```

assigns y1, y2, y3 and y4 to the elements of y. Statements like

```
begincnst;
    y[ind] = y1;
endcnst;
```

produce incorrect results. Regardless of the value of ind, y[1] always takes the current value of y1. To see this, consider the following program:

```
title 'Avoid Conflicting Data Set and ARRAY Variable Names.';

data a;
    input y1 @@;
    ind = _n_;
    datalines;
    2 2.7 8.7 -3.4
;

proc mcmc data=a;
    ods exclude all;
    array y[4];
    begincnst;
        y[ind] = y1;
        if ind eq 4 then
            put y=;
    endcnst;
    parm b;
    prior b ~ n(0, sd=1);
    model general(0);
run;
ods exclude none;
```

You can ignore the PARMS, PRIOR, and MODEL statements. The assignment statement to y[ind] in the BEGINCNST and ENDCNST statements saves the variable y1 to array y. At the last observation of the data set, PROC MCMC displays the values of y in the log file. The output is as follows:

```
y[1]=-3.4 y[2]=2.7 y[3]=8.7 y[4]=-3.4
```

The first element has an incorrect value of -3.4 instead of 2. As the procedure steps through the data set, y1 points to different observations, but whichever observation it points to is also assigned to y[1]. At the end, y[1] takes on the value of -3.4 since it is the last observation in the data set.

BEGINPRIOR/ENDPRIOR BEGINHYPER/ENDHYPER Statements

BEGINPRIOR ;

ENDPRIOR ;

BEGINHYPER ;

ENDHYPER ;

The BEGINPRIOR statement indicates that PROC MCMC processes the next statements twice at every Markov chain simulation—that is, twice per pass through the data set. The programming statements are called at the first and the last observation of the data set. A subsequent ENDPRIOR statement indicates a return to normal processing. The BEGINPRIOR and ENDPRIOR statements are best used to reduce unnecessary observation-level computations. Any computations that are identical to every observation, such as transformation of parameters, should be enclosed in these statements.

At the first observation, PROC MCMC executes the programming statements. This allows for a quick update of all the symbols enclosed in the statements. At the last observation, PROC MCMC executes the programming statements again and adds the log of the prior density to the log of the posterior density.

The BEGINHYPER and ENDHYPER statements work in the same way as the BEGINPRIOR and ENDPRIOR statements. The BEGINHYPER and ENDHYPER statements are aliases for the BEGINPRIOR and ENDPRIOR statements, and they are treated the same.

BY Statement

BY *variables* ;

You can specify a BY statement with PROC MCMC to obtain separate analyses on observations in groups 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 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 BY statement option NOTSORTED or DESCENDING in the BY statement for PROC MCMC. 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.

For more information about the BY statement, see *SAS Language Reference: Concepts*. For more information about the DATASETS procedure, see the *Base SAS Procedures Guide*.

MODEL Statement

> **MODEL** *dependent-variable-list* $\sim$ *distribution* ;

The MODEL statement is used to specify the conditional distribution of the data given the parameters (the likelihood function). You must specify a single dependent variable or a list of dependent variables, a tilde ($\sim$), and then a distribution with its arguments. The dependent variables can be variables from the input data set or functions of the symbols in the program. The dependent variables must be specified unless the functions GENERAL or DGENERAL are used (see the section "Specifying a New Distribution" on page 3540 for more details). Multiple MODEL statements are allowed for defining models with multiple independent components. The log likelihood value is the sum of the log likelihood values from each MODEL statement.

PROC MCMC is a programming language that is similar to the DATA step, and the order of statement evaluation is important. For example, the MODEL statement must come after any SAS programming statements that define or modify arguments used in the construction of the log likelihood. In PROC MCMC, a symbol is allowed to be defined multiple times and used at different places. Using an expression out of order produces erroneous results that can also be hard to detect.

Standard distributions that the MODEL statement supports are listed in the Table 52.2 (see the section "Standard Distributions" on page 3529 for density specification). These distributions can also be used in the PRIOR and HYPERPRIOR statements. PROC MCMC allows some distributions to be parameterized in multiple ways. For example, you can specify a normal distribution with variance (VAR=), standard deviation (SD=), or precision (PRECISION=) parameter. For distributions that have different parameterizations, you must specify an option to clearly name the ambiguous parameter. In the normal distribution, for example, you must indicate whether the second argument is a variance, a standard deviation, or a precision.

All distributions, with the exception of binary and uniform, can have the optional arguments of LOWER= and UPPER=, which specify a truncated density. See the section "Truncation and Censoring" on page 3542 for more details.

Table 52.2 Valid Distributions

Distribution Name	Definition
beta($<$**a=**$>\alpha$, $<$**b=**$>\beta$)	beta distribution with shape parameters α and β
binary($<$**prob|p=**$> p$)	binary (Bernoulli) distribution with probability of success p. You can use the alias **bern** for this distribution.
binomial ($<$**n=**$> n$, $<$**prob|p=**$> p$)	binomial distribution with count n and probability of success p
cauchy ($<$**location|loc|l=**$>\theta$, $<$**scale|s=**$>\lambda$)	Cauchy distribution with location θ and scale λ
chisq($<$**df=**$> \nu$)	χ^2 distribution with ν degrees of freedom

Table 52.2 (continued)

Distribution Name	Definition
dgeneral(*ll*)	general log-likelihood function that you construct using SAS programming statements for single or multiple *discrete* variables. Also see the function **general**. The name **dlogden** is an alias for this function.
expchisq(<**df=**> ν)	log transformation of a χ^2 distribution with ν degrees of freedom: $\theta \sim$ **chisq**(ν) $\Leftrightarrow$ log(θ) $\sim$ **expchisq**(ν). You can use the alias **echisq** for this distribution.
expexpon(**scale|s=** λ) **expexpon**(**iscale|is=** λ)	log transformation of an exponential distribution with scale or inverse-scale parameter λ: $\theta \sim$ **expon**(λ) $\Leftrightarrow$ log(θ) $\sim$ **expexpon**(λ). You can use the alias **eexpon** for this distribution.
expGamma(<**shape|sp=**> a, **scale|s=** λ) **expGamma**(<**shape|sp=**> a, **iscale|is=** λ)	log transformation of a gamma distribution with shape a and scale or inverse-scale λ: $\theta \sim$ **gamma**(a, λ) $\Leftrightarrow$ log(θ) $\sim$ **expgamma**(a, λ). You can use the alias **egamma** for this distribution.
expichisq(<**df=**> ν)	log transformation of an inverse χ^2 distribution with ν degrees of freedom: $\theta \sim$ **ichisq**(ν) $\Leftrightarrow$ log(θ) $\sim$ **expichisq**(ν). You can use the alias **eichisq** for this distribution.
expiGamma(<**shape|sp=**> a, **scale|s=** λ) **expiGamma**(<**shape|sp=**> a, **iscale|is=** λ)	log transformation of an inverse-gamma distribution with shape a and scale or inverse-scale λ: $\theta \sim$ **igamma**(a, λ) $\Leftrightarrow$ log(θ) $\sim$ **expigamma**(a, λ). You can use the alias **eigamma** for this distribution.
expsichisq(<**df=**> ν, <**scale|s=**> s)	log transformation of a scaled inverse χ^2 distribution with ν degrees of freedom and scale parameter s: $\theta \sim$ **sichisq**(ν) $\Leftrightarrow$ log(θ) $\sim$ **expsichisq**(ν). You can use the alias **esichisq** for this distribution.
expon(**scale|s=** λ) **expon**(**iscale|is=** λ)	exponential distribution with scale or inverse-scale parameter λ
gamma(<**shape|sp=**> a, **scale|s=** λ) **gamma**(<**shape|sp=**> a, **iscale|is=** λ)	gamma distribution with shape a and scale or inverse-scale λ
geo(<**prob|p=**> p)	geometric distribution with probability p

Table 52.2 (continued)

Distribution Name	Definition
general(*ll*)	general log likelihood function that you construct using SAS programming statements for a single or multiple continuous variables. The argument *ll* is an expression for the log of the distribution. If there are multiple variables specified before the tilde in a MODEL, PRIOR, or HYPERPRIOR statement, *ll* is interpreted as the log of the joint distribution for these variables. Note that in the MODEL statement, the response variable specified before the tilde is just a place holder and is of no consequence; the variable must have appeared in the construction of *ll* in the programming statements. **general**(*constant*) is equivalent to a uniform distribution on the real line. You can use the alias **logden** for this distribution.
ichisq(< **df=** >ν)	inverse χ^2 distribution with ν degrees of freedom
igamma(< **shape|sp=** >a, **scale|s=** λ) **igamma**(< **shape|sp=** >a, **iscale|is=** λ)	inverse-gamma distribution with shape a and scale or inverse-scale λ
laplace(< **location|loc|l=** >θ, **scale|s=** λ) **laplace**(< **location|loc|l=** >θ, **iscale|is=** λ)	Laplace distribution with location θ and scale or inverse-scale λ. This is also known as the *double exponential* distribution. You can use the alias **dexpon** for this distribution.
logistic(< **location|loc|l=** >a, < **scale|s=** >b)	logistic distribution with location a and scale b
lognormal(< **mean|m=** >μ, **sd=** λ) **lognormal**(< **mean|m=** >μ, **var|v=** λ) **lognormal**(< **mean|m=** >μ, **prec=** λ)	log-normal distribution with mean μ and standard deviation or variance or precision λ. You can use the aliases **lognormal** or **lnorm** for this distribution.
negbin(< **n=** >n, < **prob|p=** >p)	negative binomial distribution with count n and probability of success p. You can use the alias **nb** for this distribution.
normal(< **mean|m=** >μ, **sd=** λ) **normal**(< **mean|m=** >μ, **var|v=** λ) **normal**(< **mean|m=** >μ, **prec=** λ)	normal (Gaussian) distribution with mean μ and standard deviation or variance or precision λ. You can use the aliases **gaussian**, **norm**, or **n** for this distribution.
pareto(< **shape|sp=** >a, < **scale|s=** >b)	Pareto distribution with shape a and scale b
poisson(< **mean|m=** >λ)	Poisson distribution with mean λ

Table 52.2 (continued)

Distribution Name	Definition
sichisq(<**df=**> ν, <**scale|s=**> s)	scaled inverse χ^2 distribution with ν degrees of freedom and scale parameter s
t(<**mean|m=**> μ, **sd=** λ, <**df=**> ν) **t**(<**mean|m=**> μ, **var|v=** λ, <**df=**> ν) **t**(<**mean|m=**> μ, **prec=** λ, <**df=**> ν)	t distribution with mean μ, standard deviation or variance or precision λ, and ν degrees of freedom
uniform(<**left|l=**> a, <**right|r=**> b)	uniform distribution with range a and b. You can use the alias **unif** for this distribution.
wald(<**mean|m=**> μ, <**iscale|is=**> λ)	Wald distribution with mean parameter μ and inverse scale parameter λ. This is also known as the *Inverse Gaussian* distribution. You can use the alias **igaussian** for this distribution.
weibull(μ, c, σ)	Weibull distribution with location (threshold) parameter μ, shape parameter c, and scale parameter σ.

PARMS Statement

> **PARMS** *name | (name-list) < => number < name | (name-list) <=> number . . . >< / NORMAL | T < (df)> | UDS > ;*

The PARMS statement lists the names of the parameters in the model and specifies optional initial values for these parameters. Multiple PARMS statements are allowed. Each PARMS statement defines a block of parameters, and the blocked Metropolis algorithm updates the parameters in each block simultaneously. See the section "Blocking of Parameters" on page 3522 for more details. PROC MCMC generates missing initial values from the prior distributions whenever needed, as long as they are the standard distributions and not the functions GENERAL or DGENERAL.

Every parameter in the PARMS statement must have a corresponding prior distribution in the PRIOR statement. The program exits if the one-to-one requirement is not satisfied.

The optional arguments give you control over different samplers explicitly for that block of parameters. The normal proposal distribution in the random walk Metropolis is the default. You can also choose a *t*-distribution with *df* degrees of freedom. If *df* > 100, the normal distribution is used instead.

The user defined sampler (UDS, see the section "UDS Statement" on page 3517) option allows you to implement a new sampler for any of the parameters in the block. PROC MCMC does not use the Metropolis sampler on these parameters and incorporates your sampler to draw posterior samples. This can sometimes greatly improve the convergence and mixing of the Markov chain. This functionality is for advanced users, and you should proceed with caution.

PRIOR/HYPERPRIOR Statement

> **PRIOR** *parameter-list ~ distribution* **;**
>
> **HYPERPRIOR** *parameter-list ~ distribution* **;**
>
> **HYPER** *parameter-list ~ distribution* **;**

The PRIOR statement is used to specify the prior distribution of the model parameters. You must specify a single parameter or a list of parameters, a tilde ($\sim$), and then a distribution with its parameters. Multiple PRIOR statements are allowed for defining models with multiple independent prior components. The log of the prior is the sum of the log prior values from each of the PRIOR statements. See the section "MODEL Statement" on page 3511 for the names of the standard distributions and the section "Standard Distributions" on page 3529 for density specification.

The PRIOR statements are processed twice at every Markov chain simulation—that is, twice per pass through the data set. The statements are called at the first and the last observation of the data set. This is the same as how the BEGINPRIOR and ENDPRIOR statements are processed.

The HYPERPRIOR statement is internally treated the same as the PRIOR statement. It provides a notational convenience in case you wish to fit a multilevel hierarchical model. It is used to specify the hyperprior distribution of the prior distribution parameters. The log of the hyperprior is the sum of the log hyperprior values from each of the HYPERPRIOR statements.

If you want to specify a multilevel hierarchical model, you can use either a PRIOR or a HYPERPRIOR statement as if it were a hyper-HYPERPRIOR statement. Your model can have as many hierarchical levels as desired.

Programming Statements

This section lists the programming statements available in PROC MCMC to compute the priors and log-likelihood functions. This section also documents the differences between programming statements in PROC MCMC and programming statements in the DATA step. The syntax of programming statements used in PROC MCMC is identical to that used in the NLMIXED procedure (see Chapter 61, "The NLMIXED Procedure") and the MODEL procedure (see Chapter 18, "The MODEL Procedure" (*SAS/ETS User's Guide*),). Most of the programming statements that can be used in the DATA step can also be used in PROC MCMC. Refer to *SAS Language Reference: Dictionary* for a description of SAS programming statements.

There are also a number of unique functions in PROC MCMC that calculate the log density of various distributions in the procedure. You can find them at the section "Using Density Functions in the Programming Statements" on page 3540.

For the list of matrix-based functions that is supported in PROC MCMC, see the section "Matrix Functions in PROC MCMC" on page 3529.

The following are valid statements:

ABORT;
CALL *name [(expression [, expression . . .])]*;
DELETE;
DO *[variable = expression*
   *[***TO** *expression] [***BY** *expression]*
   *[, expression [***TO** *expression] [***BY** *expression] . . .]*
   ]
   *[***WHILE** *expression] [***UNTIL** *expression]*;
END;
GOTO *statement_label*;
IF *expression*;
IF *expression* **THEN** *program_statement*;
   ELSE *program_statement*;
variable = expression;
variable + expression;
LINK *statement_label*;
PUT *[variable] [=] [...]*;
RETURN;
SELECT[(*expression* **)]**;
STOP;
SUBSTR(*variable, index, length* **)=** *expression*;
WHEN *(expression) program_statement*;
      OTHERWISE *program_statement*;

For the most part, the SAS programming statements work the same as they do in the DATA step, as documented in *SAS Language Reference: Concepts*. However, there are several differences:

- The ABORT statement does not allow any arguments.

- The DO statement does not allow a character index variable. Thus

  ```
  do i = 1,2,3;
  ```

  is supported; however, the following statement is not supported:

  ```
  do i = 'A','B','C';
  ```

- The PUT statement, used mostly for program debugging in PROC MCMC (see the section "Handling Error Messages" on page 3562), supports only some of the features of the DATA step PUT statement, and it has some features that are not available with the DATA step PUT statement:

  - The PROC MCMC PUT statement does not support line pointers, factored lists, iteration factors, overprinting, _INFILE_, _OBS_, the colon (:) format modifier, or "$".
  - The PROC MCMC PUT statement does support expressions, but the expression must be enclosed in parentheses. For example, the following statement displays the square root of x:

```
        put (sqrt(x));
```

- The WHEN and OTHERWISE statements enable you to specify more than one target statement. That is, DO/END groups are not necessary for multiple statement WHENs. For example, the following syntax is valid:

```
select;
    when (exp1) stmt1;
                stmt2;
    when (exp2) stmt3;
                stmt4;
end;
```

You should avoid defining variables that begin with an underscore (_). They might conflict with internal variables created by PROC MCMC. The MODEL statement must come after any SAS programming statements that define or modify terms used in the construction of the log likelihood.

UDS Statement

> **UDS** *subroutine-name (subroutine-argument-list)* ;

UDS stands for user defined sampler. The UDS statement allows you to use a separate algorithm, other than the default random walk Metropolis, to update parameters in the model. The purpose of the UDS statement is to give you a greater amount of flexibility and better control over the updating schemes of the Markov chain. Multiple UDS statements are allowed.

For the UDS statement to work properly, you have to do the following:

- write a subroutine by using PROC FCMP (see the FCMP Procedure in the *Base SAS Procedures Guide*) and save it to a SAS catalog (see the example in this section). The subroutine must update some parameters in the model. These are the UDS parameters. The subroutine is called the UDS subroutine.

- declare any UDS parameters in the PARMS statement with a sampling option, as in </ UDS > (see the section "PARMS Statement" on page 3514).

- specify the prior distributions for all UDS parameters, using the PRIOR statements.

NOTE: All UDS parameters must appear in three places: the UDS statement, the PARMS statement, and the PRIOR statement. Otherwise, PROC MCMC exits.

To obtain a valid Markov chain, a UDS subroutine must update a parameter from its full posterior conditional distribution and not the posterior marginal distribution. The posterior conditional is something that you need to provide. This conditional is implicitly based on a prior distribution. PROC MCMC has no means to verify that the implied prior in the UDS subroutine is the same as the prior that you specified in the PRIOR statement. You need to make sure that the two distributions agree; otherwise, you will get misleading results.

The priors in the PRIOR statements do not directly affect the sampling of the UDS parameters. They could affect the sampling of the other parameters in the model, which, in turn, changes the behavior of the Markov chain. You can see this by noting cases where the hyperparameters of the UDS parameters are model parameters; the priors should be part of the posterior conditional distributions of these hyperparameters, and they cannot be omitted.

Some additional information is listed to help you better understand the UDS statement:

- Most features of the SAS programming language can be used in subroutines processed by PROC FCMP (see the FCMP Procedure in the *Base SAS Procedures Guide*).

- The UDS statement does not support FCMP functions—a FCMP function returns a value, while a subroutine does not. A subroutine updates some of its subroutine arguments. These arguments are called OUTARGS arguments.

- The UDS parameters cannot be in the same block as other parameters. The optional argument < / UDS > in the PARMS statement prevents parameters that use the default Metropolis from being mixed with those that are updated by the UDS subroutines.

- You can put all the UDS parameters in the same PARMS statement or have a separate UDS statement for each of them.

- The same subroutine can be used in multiple UDS statements. This feature comes in handy if you have a generic sampler that can be applied to different parameters.

- PROC MCMC updates the UDS parameters by calling the UDS subroutines directly. At every iteration, PROC MCMC first samples parameters that use the Metropolis algorithm, then the UDS parameters. Sampling of the UDS parameters proceeds in the order in which the UDS statements are listed.

- A UDS subroutine accepts any symbols in the program as well as any input data set variables as its arguments.

- Only the OUTARGS arguments in a UDS subroutine are updated in PROC MCMC. You can modify other arguments in the subroutine, but the changes are not global in the procedure.

- If a UDS subroutine has an argument that is a SAS data set variable, PROC MCMC steps through the data set while updating the UDS parameters. The subroutine is called once per observation in the data set for every iteration.

- If a UDS subroutine does not have any arguments that are data set variables, PROC MCMC does not access the data set while executing the subroutine. The subroutine is called once per iteration.

- To reduce the overhead in calling the UDS subroutine and accessing the data set repeatedly, you might consider reading all the input data set variables into arrays and using the arrays as the subroutine arguments. See the section "BEGINCNST/ENDCNST Statement" on page 3507 about how to use the BEGINCNST and ENDCNST statements to store data set variables.

An Example that Uses the UDS Statement

Suppose that you are interested in modeling normal data with conjugate prior distributions. The data are as follows:

```
title 'An Example that uses the UDS Statement';

data a;
   input y @@;
   i = _n_;
   datalines;
-0.651   17.435  -5.943   -2.543 -10.444
-5.754   -5.002  -2.545   -1.743   0.998
;
```

The likelihood for each observation is as follows:

$$f(y_i | \mu, \sigma) = \phi(\mu, \text{var} = \sigma^2)$$

The prior distributions on μ and σ^2 are as follows:

$$
\begin{aligned}
\pi(\mu | \mu_0, \tau_0^2) &= \phi(\mu_0, \text{var} = \tau_0^2) \\
\pi(\sigma^2 | v_0, \sigma_0^2) &= f_{si\chi^2}(\text{shape} = v_0, \text{scale} = \sigma_0^2)
\end{aligned}
$$

where $f_{si\chi^2}$ is the density function for a scaled inverse chi-square distribution. To sample μ and σ^2 without using any UDS statements, you can use the following program:

```
proc mcmc data=a seed=17;
   parm mu;
   parm s2;
   begincnst;
      mu0 = 0;   t0 = 20;
      nu0 = 10; s0 = 10;
   endcnst;

   prior mu ~ normal(mu0, var=t0);
   prior s2 ~ sichisq(nu0, s0);
   model y ~ normal(mu, var = s2);
run;
```

This is a case where the full posterior conditional distribution of μ given σ^2 and y has a closed form. It is also a normal distribution:

$$p(\mu | \sigma^2, y) = \phi \left(\frac{\frac{\mu_0}{\tau_0^2} + \frac{n\bar{y}}{\sigma^2}}{\frac{1}{\tau_0^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\tau_0} + \frac{n}{\sigma^2}} \right)$$

You can define a subroutine, **muupdater**, which generates a random normal sample from the posterior conditional distribution described previously.

```
proc fcmp outlib=sasuser.funcs.uds;
   subroutine muupdater(mu, s2, mu0, t0, n, sumy);
   outargs mu;
   sigma2 = 1 / (1/t0 + n/s2);
   mean = (mu0/t0 + sumy/s2) * sigma2;
   mu = rand("normal", mean, sqrt(sigma2));
   endsub;
run;
```

The subroutine is saved in the OUTLIB= library. The declaration of any subroutine begins with a SUBROUTINE statement and ends with an ENDSUB statement. The OUTARGS statement in the subroutine indicates that mu is updated. Others, such as sigma2, mu0, and so on, are arguments that are needed in the full conditional distribution. Here the rand and sqrt are two of the many SAS functions that you can use.

You specify a CMPLIB option to let SAS search each of the catalogs that are specified in the option for a package that contains **muupdater**.

```
options cmplib=sasuser.funcs;
```

To use the subroutine in the UDS statement, you can use the following statements:

```
proc mcmc data=a seed=17;
   UDS muupdater(mu, s2, mu0, t0, n, sumy);
   parm mu /uds;
   parm s2;
   begincnst;
      mu0 = 0;   t0 = 20;
      nu0 = 10; s0 = 10;
      n = 10;
      if i eq 1 then sumy = 0;
      sumy = sumy + y;
      call streaminit(1);
   endcnst;

   prior mu ~ normal(mu0, var=t0);
   prior s2 ~ sichisq(nu0, s0);
   model y ~ normal(mu, var = s2);
run;
```

These statements are very similar to the previous program. The differences are the UDS statement, the < / UDS > option in the PARMS statement, and a few lines that computes the values of sumy and n.

The symbol sumy is the sum of y. The value is obtained by taking advantage of the BEGINCNST and ENDCNST statements. See the example in the section "BEGINCNST/ENDCNST Statement" on page 3507. The symbol n is the sample size in the data set.

The CALL STREAMINIT routine ensures that the RAND function in **muupdater** creates a reproducible stream of random numbers. The SEED= option specifies a seed for the random number generator in PROC MCMC, which does not control the random number generator in the RAND function in the subroutine. You need to set both to reproduce the same stream of Markov chain

samples.

The two programs produce different but similar numbers (results not shown) for the posterior distributions of μ and σ^2.

For a more realistic example that uses the UDS statement, see "Example 52.11: Implement a New Sampling Algorithm" on page 3666.

Details: MCMC Procedure

How PROC MCMC Works

PROC MCMC uses a random walk Metropolis algorithm to obtain posterior samples. For details on the Metropolis algorithm, see the section "Metropolis and Metropolis-Hastings Algorithms" on page 152. For the actual implementation details of the Metropolis algorithm in PROC MCMC, such as the blocking of the parameters and tuning of the covariance matrices, see the section "Tuning the Proposal Distribution" on page 3524. By default, PROC MCMC assumes that all observations in the data set are independent, and the logarithm of the posterior density is calculated as follows:

$$\log(p(\theta|\mathbf{y})) = \log(\pi(\theta)) + \sum_{i=1}^{n} \log(f(y_i|\theta))$$

where θ is a parameter or a vector of parameters. The term $\log(\pi(\theta))$ is the sum of the log of the prior densities specified in the PRIOR and HYPERPRIOR statements. The term $\log(f(y_i|\theta))$ is the log likelihood specified in the MODEL statement. The MODEL statement specifies the log likelihood for a single observation in the data set.

The statements in PROC MCMC are in many ways like DATA step statements; PROC MCMC evaluates every statement in order for each observation. The procedure cumulatively adds the log likelihood for each observation. Statements between the BEGINPRIOR and ENDPRIOR statements are evaluated only at the first and the last observations. At the last observation, the log of the prior and hyperprior distributions is added to the sum of the log likelihood to obtain the log of the posterior distribution.

With multiple PARMS statements (multiple blocks of parameters), PROC MCMC updates each block of parameters while holding the others constants. The procedure still steps through all of the programming statements to calculate the log of the posterior distribution, given the current or the proposed values of the updating block of parameters. In other words, the procedure does not calculate the conditional distribution explicitly for each block of parameters, and it uses the full joint distribution in the Metropolis step for every block update. If you wish to model dependent data—that is, $\log(f(\mathbf{y}|\theta)) \neq \sum_i \log(f(y_i|\theta))$—you can use the PROC option JOINTMODEL. See the section "Modeling Dependent Data" on page 3549 for more details.

Blocking of Parameters

In a multivariate parameter model, if all k parameters are proposed with one joint distribution $q(\cdot|\cdot)$, acceptance or rejection would occur for all of them. This can be rather inefficient, especially when parameters have vastly different scales. A way to avoid this difficulty is to allocate the k parameters into d blocks and update them separately. The PARMS statement is used to specify model parameters. It also puts parameters in separate blocks, and each block of parameters is updated sequentially in the procedure.

Suppose that you wish to sample from a multivariate distribution with probability density function $p(\theta|\mathbf{y})$ where $\theta = \{\theta_1, \theta_2, \ldots, \theta_k\}$. Now suppose that these k parameters are separated into d blocks—for example, $p(\theta|\mathbf{x}) = f_d(z)$ where $z = \{z_1, z_2, \ldots, z_d\}$, where each z_j contains a nonempty subset of the $\{\theta_i\}$, and where each θ_i is contained in one and only one z_j. In the MCMC context, the z's are blocks of parameters. In the blocked algorithm, a proposal is composed of several parts. Instead of proposing a simultaneous move for all the θ's, a proposal is made for the θ_i's in z_1 only, then for the θ_i's in z_2, and so on for d subproposals. Any accepted proposal can involve any number of the blocks moving. Not necessarily all of the parameters move at once as in the all-at-once Metropolis algorithm.

Formally, the blocked Metropolis algorithm is as follows. Let w_j be the collection of θ_i that are in block z_j and let $q_j(\cdot|w_j)$ be a symmetric multivariate distribution centered at the current values of w_j.

1. Let $t = 0$. Choose points for all w_j^t. This can be an arbitrary point as long as $p(w_j^t|\mathbf{y}) > 0$.

2. For $j = 1, \cdots, d$:

    a) Generate a new sample, $w_{j,new}$, using the proposal distribution $q_j(\cdot|w_j^t)$.

    b) Calculate the following quantity:

    $$r = \min\left\{\frac{p(w_{j,new}|w_1^t, \cdots, w_{j-1}^t, w_{j+1}^{t-1}, \cdots, w_d^t, \mathbf{y})}{p(w_j^t|w_1^t, \cdots, w_{j-1}^t, w_{j-1}^{t+1}, \cdots, w_d^t, \mathbf{y})}, 1\right\}.$$

    c) Sample u from the uniform distribution $U(0, 1)$.

    d) Set $w_j^{t+1} = w_{j,new}$ if $r < a$; $w_j^{t+1} = w_j^t$ otherwise.

3. Set $t = t + 1$. If $t < T$, the number of desired samples, go back to Step 2; otherwise, stop.

With PROC MCMC, you can sample all parameters simultaneously by putting them all in a single PARMS statement, you can sample parameters individually by putting each parameter in its own PARMS statement, or you can sample certain subsets of parameters together by grouping each subset in its own PARMS statements. For example, if the model you are interested in has five parameters, alpha, beta, gamma, phi, sigma, the all-at-once strategy is as follows:

```
parms alpha beta gamma phi sigma;
```

The one-at-a-time strategy is as follows:

```
parms alpha;
parms beta;
parms gamma;
parms phi;
parms sigma;
```

A two-block strategy could be as follows:

```
parms alpha beta gamma;
parms phi sigma;
```

One of the greatest challenges in MCMC sampling is achieving good mixing of the chains—the chains should quickly traverse the support of the stationary distribution. A number of factors determine the behavior of a Metropolis sampler; blocking is one of them, so you want to be extra careful when it comes to choosing a good design. Generally speaking, forming blocks of parameters has its advantages, but it is not true that the larger the block the faster the convergence.

When simultaneously sampling a large number of parameters, the algorithm might find it difficult to achieve good mixing. As the number of parameters gets large, it is much more likely to have (proposal) samples that fall well into the tails of the target distribution, producing too small a test ratio. As a result, few proposed values are accepted and convergence is slow. On the other hand, when sampling each parameter individually, the chain might mix far too slowly because the conditional distributions (of θ_i given all other θ's) might be very "narrow." Hence, it takes a long time for the chain to explore fully that dimension alone. There are no theoretical results that can help determine an optimal "blocking" for an arbitrary parametric model. A rule followed in practice is to form small groups of correlated parameters that belong to the same context in the formulation of the model. The best mixing is usually obtained with a blocking strategy somewhere between the all-at-once and one-at-a-time strategies.

Samplers

This section describes the sampling methods used in PROC MCMC. Each block of parameters is classified by the nature of the prior distributions. "Continuous" means all priors of the parameters in the same block are continuous distribution. "Discrete" means all priors are discrete. "Mixed" means that some parameters are continuous and others are discrete. Parameters that have binary priors are treated differently, as indicated in the table. MVN stands for the multivariate normal distribution, and MVT is short for the multivariate *t*-distribution.

Blocks	Default Method	Alternative Method
continuous	MVN	MVT
discrete (other than binary)	binned MVN	binned MVT or symmetric geometric
mixed	MVN	MVT
binary (single dimensional)	inverse CDF	
binary (multi-dimensional)	independence sampler	

For a block of continuous parameters, PROC MCMC uses a multivariate normal distribution as the default proposal distribution. In the tuning phase, the procedure finds an optimal scale c and a tuning covariance matrix Σ.

For a discrete block of parameters, PROC MCMC uses a discretized multivariate normal distribution as the default proposal distribution. The scale c and covariance matrix Σ are tuned. Alternatively, you can use an independent symmetric geometric proposal distribution. The density has form $\frac{p(1-p)^{|\theta|}}{2(1-p)}$ and has variance $\frac{(2-p)(1-p)}{p^2}$. In the tuning phase, the procedure finds an optimal proposal probability p for every parameter in the block.

You can change the proposal distribution, from the normal to a t-distribution. You can either use the PROC option PROPDIST=T(df) or PARMS statement option </ T(df)> to make the change. The t-distributions have thicker tails, and they can propose to the tail areas more efficiently than the normal distribution. It can help with the mixing of the Markov chain if some of the parameters have a skewed tails. See "Example 52.4: Nonlinear Poisson Regression Models" on page 3599. The independence sampler (see the section "Independence Sampler" on page 153) is used for a block of binary parameters. The inverse CDF method is used for a block that consists of a single binary parameter.

Tuning the Proposal Distribution

One key factor in achieving high efficiency of a Metropolis-based Markov chain is finding a good proposal distribution for each block of parameters. This process is referred to as tuning. The tuning phase consists of a number of loops. The minimum number of loops is controlled by the option MINTUNE=, with a default value of 2. The option MAXTUNE= controls the maximum number of tuning loops, with a default value of 24. Each loop lasts for NTU= iterations, where by default NTU= 500. At the end of every loop, PROC MCMC examines the acceptance probability for each block. The acceptance probability is the percentage of NTU= proposals that have been accepted. If the probability falls within the acceptance tolerance range (see the section "Scale Tuning" on page 3525), the current configuration of c/Σ or p is kept. Otherwise, these parameters are modified before the next tuning loop.

Continuous Distribution: Normal or *t*-Distribution

A good proposal distribution should resemble the actual posterior distribution of the parameters. Large sample theory states that the posterior distribution of the parameters approaches a multivariate normal distribution (see Gelman et al. 2004, Appendix B, and Schervish 1995, Section 7.4). That is why a normal proposal distribution often works well in practice. The default proposal distribution in PROC MCMC is the normal distribution: $q_j(\theta_{\text{new}}|\theta^t) = \text{MVN}(\theta_{\text{new}}|\theta^t, c^2\Sigma)$. As an alternative, you can choose a multivariate t-distribution as the proposal distribution. It is a good distribution to use if you think that the posterior distribution has thick tails and a t-distribution can improve the mixing of the Markov chain. See "Example 52.4: Nonlinear Poisson Regression Models" on page 3599.

Scale Tuning

The acceptance rate is closely related to the sampling efficiency of a Metropolis chain. For a random walk Metropolis, high acceptance rate means that most new samples occur right around the current data point. Their frequent acceptance means that the Markov chain is moving rather slowly and not exploring the parameter space fully. On the other hand, a low acceptance rate means that the proposed samples are often rejected; hence the chain is not moving much. An efficient Metropolis sampler has an acceptance rate that is neither too high nor too low. The scale c in the proposal distribution $q(\cdot|\cdot)$ effectively controls this acceptance probability. Roberts, Gelman, and Gilks (1997) showed that if both the target and proposal densities are normal, the optimal acceptance probability for the Markov chain should be around 0.45 in a single dimensional problem, and asymptotically approaches 0.234 in higher dimensions. The corresponding optimal scale is 2.38, which is the initial scale set for each block.

Due to the nature of stochastic simulations, it is impossible to fine-tune a set of variables such that the Metropolis chain has the exact desired acceptance rate. In addition, Roberts and Rosenthal (2001) empirically demonstrated that an acceptance rate between 0.15 and 0.5 is at least 80% efficient, so there is really no need to fine-tune the algorithms to reach acceptance probability that is within small tolerance of the optimal values. PROC MCMC works with a probability range, determined by the PROC options TARGACCEPT $\pm$ ACCEPTTOL. The default value of TARGACCEPT is a function of the number of parameters in the model, as outlined in Roberts, Gelman, and Gilks (1997). The default value of ACCEPTTOL is 0.075. If the observed acceptance rate in a given tuning loop is less than the lower bound of the range, the scale is reduced; if the observed acceptance rate is greater than the upper bound of the range, the scale is increased. During the tuning phase, a scale parameter in the normal distribution is adjusted as a function of the observed acceptance rate and the target acceptance rate. The following updating scheme is used in PROC MCMC [1]:

$$c_{new} = \frac{c_{cur} \cdot \Phi^{-1}(p_{opt}/2)}{\Phi^{-1}(p_{cur}/2)}$$

where c_{cur} is the current scale, p_{cur} is the current acceptance rate, p_{opt} is the optimal acceptance probability.

Covariance Tuning

To tune a covariance matrix, PROC MCMC takes a weighted average of the old proposal covariance matrix and the recent observed covariance matrix, based on NTU samples in the current loop. The TUNEWT=w option determines how much weight is put on the recently observed covariance matrix. The formula used to update the covariance matrix is as follows:

$$COV_{new} = w\,COV_{cur} + (1-w)COV_{old}$$

There are two ways to initialize the covariance matrix:

---

[1] Roberts, Gelman, and Gilks (1997) and Roberts and Rosenthal (2001) demonstrate that the relationship between acceptance probability and scale in a random walk Metropolis is $p = 2\Phi\left(-\sqrt{I}c/2\right)$, where c is the scale, p is the acceptance rate, Φ is the CDF of a standard normal, and $I \equiv E_f[(f'(x)/f(x))^2]$, $f(x)$ is the density function of samples. This relationship determines the updating scheme, with I being replaced by the identity matrix to simplify calculation.

- The default is an identity matrix multiplied by the initial scale of 2.38 (controlled by the PROC option SCALE=) and divided by the square root of the number of estimated parameters in the model. It can take a number of tuning phases before the proposal distribution is tuned to its optimal stage, since the Markov chain needs to spend time learning about the posterior covariance structure. If the posterior variances of your parameters vary by more than a few orders of magnitude, if the variances of your parameters are much different from 1, or if the posterior correlations are high, then the proposal tuning algorithm might have difficulty with forming an acceptable proposal distribution.

- Alternatively, you can use a numerical optimization routine, such as the quasi-Newton method, to find a starting covariance matrix. The optimization is performed on the joint posterior distribution, and the covariance matrix is a quadratic approximation at the posterior mode. In some cases this is a better and more efficient way of initializing the covariance matrix. However, there are cases, such as when the number of parameters is large, where the optimization could fail to find a matrix that is positive definite. In that case, the tuning covariance matrix is reset to the identity matrix.

A side product of the optimization routine is that it also finds the *maximum a posteriori* (MAP) estimates with respect to the posterior distribution. The MAP estimates are used as the initial values of the Markov chain.

If any of the parameters are discrete, then the optimization is performed conditional on these discrete parameters at their respective fixed initial values. On the other hand, if all parameters are continuous, you can in some cases skip the tuning phase (by setting MAXTUNE=0) or the burn-in phase (by setting NBI=0).

Discrete Distribution: Symmetric Geometric

By default, PROC MCMC uses the normal density as the proposal distribution in all Metropolis random walks. For parameters that have discrete prior distributions, PROC MCMC discretizes proposed samples. You can choose an alternative symmetric geometric proposal distribution by specifying the option DISCRETE=GEO.

The density of the symmetric geometric proposal distribution is as follows:

$$\frac{p_g(1-p_g)^{|\theta|}}{2(1-p_g)}$$

where the symmetry centers at θ. The distribution has a variance of

$$\sigma^2 = \frac{(2-p_g)(1-p_g)}{p_g^2}$$

Tuning for the proposal p_g uses the following formula:

$$\frac{\sigma_{\text{new}}}{\sigma_{\text{cur}}} = \frac{\Phi^{-1}(p_{\text{opt}}/2)}{\Phi^{-1}(p_{\text{cur}}/2)}$$

where σ_{new} is the standard deviation of the new proposal geometric distribution, σ_{cur} is the standard deviation of the current proposal distribution, p_{opt} is the target acceptance probability, and p_{cur} is the current acceptance probability for the discrete parameter block.

The updated p_g is the solution to the following equation that is between 0 and 1 :

$$\sqrt{\frac{(2 - p_g)(1 - p_g)}{p_g^2}} = \frac{\sigma_{\text{cur}} \cdot \Phi^{-1}(p_{\text{opt}}/2)}{\Phi^{-1}(p_{\text{cur}}/2)}$$

Binary Distribution: Independence Sampler

Blocks consisting of a single parameter with a binary prior do not require any tuning; the inverse-CDF method applies. Blocks that consist of multiple parameters with binary prior are sampled by using an independence sampler with binary proposal distributions. See the section "Independence Sampler" on page 153. During the tuning phase, the success probability p of the proposal distribution is taken to be the probability of acceptance in the current loop. Ideally, an independence sampler works best if the acceptance rate is 100%, but that is rarely achieved. The algorithm stops when the probability of success exceeds the TARGACCEPTI=*value*, which has a default value of 0.6.

Initial Values of the Markov Chains

You can assign initial values to any parameters. To assign initial values, you can either use the PARMS statements or use programming statements within the BEGINCNST and ENDCNST statements. For the latter approach, see the section "BEGINCNST/ENDCNST Statement" on page 3507.

When parameters have missing initial values, PROC MCMC tries to generate them from the respective prior distributions, as long as the distributions are listed in the section "Standard Distributions" on page 3529. PROC MCMC either uses the mode from the prior distribution or draws a random number from it. For distributions that do not have modes, such as the uniform distribution, PROC MCMC uses the mean instead. In general, PROC MCMC avoids using starting values that are close to the boundary of support of the prior distribution. For example, the exponential prior has a mode at 0, and PROC MCMC starts an initial value at the mean. This avoids some potential numerical problems. If you use the GENERAL or DGENERAL functions in the PRIOR statements, you must provide initial values for those parameters.

If you use the optimization option PROPCOV, PROC MCMC starts the tuning at the optimized values. The procedure overwrites the initial values that you provided unless you use the option INIT=REINIT.

Assignments of Parameters

In PROC MCMC, the values of any model parameters cannot be altered outside of the BEGINCNST and ENDCNST statements. The following statements produce an error:

```
parms alpha;
alpha = 27;
```

All programming statements within the BEGINCNST and ENDCNST statements are executed be-
fore the first Markov chain iteration, so they can be used to give parameters initial values. The
following syntax is allowed:

```
parms alpha;
begincnst;
   alpha = 27;
endcnst;
```

The value 27 is alpha's initial value. This overrides any value that you give to alpha in the PARMS
statement. For example, with the following statements,

```
parms alpha 23;
begincnst;
   alpha = 27;
endcnst;
```

the Markov chain starts at alpha=27, not 23.

This feature gives you a way to systematically assign initial values. Suppose that z is an array
parameter of the same length as the number of observations in the input data set. You want to
start the Markov chain with each z_i having a different value depending on the data set variable y.
Suppose that you want set $z_i = |y|$ for the first half of the observations and $z_i = 2.3$ for the rest.
This is what you can do:

```
/* a rather artificial input data set. */
data inputdata;
   do ind = 1 to 10;
      y = rand('normal');
      output;
   end;
run;

proc mcmc data=inputdata;
   array z[10];
   begincnst;
      if ind <= 5 then z[ind] = abs(y);
      else z[ind] = 2.3;
   endcnst;
   parms z:;
   prior z: ~ normal(0, sd=1);
   model general(0);
run;
```

Elements of z are modified as PROC MCMC executes the programming statements within the
BEGINCNST and ENDCNST. This feature is use most when you construct a prior on z using the
GENERAL function and find that the PARMS statements are too cumbersome for assigning starting
values. Other symbols, including the input data set variables, can be altered in the program.

Matrix Functions in PROC MCMC

Table 52.4 displays the matrix functions that PROC MCMC supports. For detailed descriptions of these functions, see the FCMP Procedure in the *Base SAS Procedures Guide*.

Table 52.4   Matrix Functions in PROC MCMC

Function Call	Description
`call addmatrix(m1, m2, outm);`	performs element-wise addition of m1 and m2 and saves the result to outm
`call chol(m, outm);`	calculates the Cholesky decomposition of m and saves the result to outm
`call det(m,d);`	calculates the determinant of m and saves the result to d
`call elemmult(m1, m2, outm);`	performs element-wise multiplication of m1 and m2 and saves the result to outm
`call identity(m);`	converts m to an identity matrix of the same order
`call inv(m,im);`	calculates the inverse of m and saves the result to im
`call mult(m1,m2,outm);`	performs matrix multiplication of m1 and m2 and saves the result to outm
`call subtractmatrix(m1,m2,outm);`	performs element-wise subtraction of m1 and m2 and saves the result to outm
`call transpose(m,tm);`	transposes m and saves it to tm
`call zeromatrix(m);`	replaces all the element values of m with 0 m

Standard Distributions

Table 52.5 to Table 52.32 show all densities that PROC MCMC recognizes. These densities can be used in the MODEL, PRIOR, and HYPERPRIOR statements. See the section "Using Density Functions in the Programming Statements" on page 3540 for information about how to use distributions in the programming statements. To specify an arbitrary distribution, you can use the functions GENERAL and DGENERAL. See the section "Specifying a New Distribution" on page 3540 for more details. See the section "Truncation and Censoring" on page 3542 for tips on how to work with truncated distributions and censoring data.

Table 52.5 Beta Distribution

PROC specification	**beta**(a, b)
density	$\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}\theta^{a-1}(1-\theta)^{b-1}$
parameter restriction	$a > 0, b > 0$
range	$\begin{cases} [0,1] & \text{when } a = 1, b = 1 \\ [0,1) & \text{when } a = 1, b \neq 1 \\ (0,1] & \text{when } a \neq 1, b = 1 \\ (0,1) & \text{otherwise} \end{cases}$
mean	$\frac{a}{a+b}$
variance	$\frac{ab}{(a+b)^2(a+b+1)}$
mode	$\begin{cases} \frac{a-1}{a+b-2} & a > 1, b > 1 \\ 0 \text{ and } 1 & a < 1, b < 1 \\ 0 & \begin{cases} a < 1, b \geq 1 \\ a = 1, b > 1 \end{cases} \\ 1 & \begin{cases} a \geq 1, b < 1 \\ a > 1, b = 1 \end{cases} \\ \text{does not exist uniquely} & a = b = 1 \end{cases}$
random number	if $\min(a,b) > 1$, see (Cheng 1978); if $\max(a,b) < 1$, see (Atkinson and Whittaker 1976) and (Atkinson 1979); if $\min(a,b) < 1$ and $\max(a,b) > 1$, see (Cheng 1978); if $a = 1$ or $b = 1$, inversion method; if $a = b = 1$, uniform random variable

Table 52.6 Binary Distribution

PROC specification	**binary**(p)
density	$p^{\theta}(1-p)^{1-\theta}$
parameter restriction	$0 \leq p \leq 1$
range	$\begin{cases} \{0\} & \text{when } p = 0 \\ \{1\} & \text{when } p = 1 \\ \{0,1\} & \text{otherwise} \end{cases}$
mean	round(p)
variance	$p(1-p)$
mode	$\begin{cases} \{1\} & \text{when } p = 1 \\ \{0\} & \text{otherwise} \end{cases}$
random number	generate $u \sim \text{uniform}(0, 1)$. If $u \leq p$, $\theta = 1$; else, $\theta = 0$

Table 52.7 **Binomial Distribution**

PROC specification	**binomial**(n, p)
density	$\dbinom{n}{\theta} p^\theta (1-p)^{n-\theta}$
parameter restriction	$n = 0, 1, 2, \cdots 0 \le p \le 1$
range	$\theta \in \{0, \cdots, n\}$
mean	$\lfloor np \rfloor$
variance	$np(1-p)$
mode	$\lfloor (n+1)p \rfloor$

Table 52.8 **Cauchy Distribution**

PROC specification	**cauchy**(a, b)
density	$\frac{1}{\pi} \left(\frac{b}{b^2 + (\theta - a)^2} \right)$
parameter restriction	$b > 0$
range	$\theta \in (-\infty, \infty)$
mean	does not exist
variance	does not exist
mode	a
random number	generate $u_1, u_2 \sim$ uniform$(0, 1)$, let $v = 2u_2 - 1$. Repeat the procedure until $u_1^2 + v^2 < 1$. $y = v/u_1$ is a draw from the standard Cauchy, and $\theta = a + by$ (Ripley 1987)

Table 52.9 χ^2 **Distribution**

PROC specification	**chisq**(ν)
density	$\frac{1}{\Gamma(\nu/2) 2^{\nu/2}} \theta^{(\nu/2)-1} e^{-\theta/2}$
parameter restriction	$\nu > 0$
range	$\theta \in [0, \infty)$ if $\nu = 2$; $(0, \infty)$ otherwise
mean	ν
variance	2ν
mode	$\nu - 2$ if $\nu \ge 2$; does not exist otherwise
random number	χ^2 is a special case of the gamma distribution: $\theta \sim$ gamma$(\nu/2, \text{scale}=2)$ is a draw from the χ^2 distribution

Table 52.10 **Exponential χ^2 Distribution**

PROC specification	**expchisq**(ν)
density	$\frac{1}{\Gamma(\nu/2)2^{\nu/2}} \exp(\theta)^{\nu/2} \exp(-\exp(\theta)/2)$
parameter restriction	$\nu > 0$
range	$\theta \in (-\infty, \infty)$
mode	$\log(\nu)$
random number	generate $x_1 \sim \chi^2(\nu)$, and $\theta = \log(x_1)$ is a draw from the exponential χ^2 distribution
relationship to the χ^2 distribution	$\theta \sim \chi^2(\nu) \Leftrightarrow \log(\theta) \sim \exp \chi^2(\nu)$

Table 52.11 **Exponential Exponential Distribution**

PROC specification	**expexpon**(scale = b)	**expexpon**(iscale = β)
density	$\frac{1}{b} \exp(\theta) \exp(-\exp(\theta)/b)$	$\beta \exp(\theta) \exp(-\exp(\theta) \cdot \beta)$
parameter restriction	$b > 0$	$\beta > 0$
range	$\theta \in (-\infty, \infty)$	same
mode	$\log(b)$	$\log(1/\beta)$
random number	generate $x_1 \sim$ expon(scale=b), and $\theta = \log(x_1)$ is a draw from the exponential exponential distribution. Note that an exponential exponential distribution is not the same as the double exponential distribution.	
relationship to the Expon distribution	$\theta \sim$ expon(b) $\Leftrightarrow \log(\theta) \sim$ expExpon(b)	

Table 52.12 **Exponential Gamma Distribution**

PROC specification	**expgamma**(a, scale = b)	**expgamma**(a, iscale = β)
density	$\frac{1}{b^a \Gamma(a)} e^{a\theta} \exp(-e^\theta/b)$	$\frac{\beta^a}{\Gamma(a)} e^{a\theta} \exp(-e^\theta \cdot \beta)$
parameter restriction	$a > 0, b > 0$	$a > 0, \beta > 0$
range	$\theta \in (-\infty, \infty)$	same
mode	$\log(ab)$	$\log(a/\beta)$
random number	generate $x_1 \sim$ gamma(a, scale $= b$), and $\theta = \log(x_1)$ is a draw from the exponential gamma distribution	
relationship to the Γ distribution	$\theta \sim$ gamma(a, b) $\Leftrightarrow \log(\theta) \sim$ expGamma(a, b)	

Table 52.13 **Exponential Inverse χ^2 Distribution**

PROC specification	**expichisq**(ν)
density	$\frac{1}{\Gamma(\frac{\nu}{2})2^{\nu/2}} \exp(-\nu\theta/2) \exp(-1/(2\exp(\theta)))$
parameter restriction	$\nu > 0$
range	$\theta \in (-\infty, \infty)$
mode	$-\log(\nu)$
random number	generate $x_1 \sim i\chi^2(\nu)$, and $\theta = \log(x_1)$ is a draw from the exponential inverse χ^2 distribution
relationship to the $i\chi^2$ distribution	$\theta \sim i\chi^2(\nu) \Leftrightarrow \log(\theta) \sim \exp i\chi^2(\nu)$

Table 52.14 **Exponential Inverse-Gamma Distribution**

PROC specification	**expigamma**(a, scale = b)	**expigamma**(a, iscale = β)
density	$\frac{b^a}{\Gamma(a)} \exp(-\alpha\theta) \exp(-b/\exp(\theta))$	$\frac{1}{\beta^\alpha \Gamma(a)} \exp(-\alpha\theta) \exp(-\frac{1}{\beta\exp(\theta)})$
parameter restriction	$a > 0, b > 0$	$a > 0, \beta > 0$
range	$\theta \in (-\infty, \infty)$	same
mode	$-\log(a/b)$	$-\log(a\beta)$
random number	generate $x_1 \sim$ igamma(a, scale $= b$), and $\theta = \log(x_1)$ is a draw from the exponential inverse-gamma distribution	
relationship to the $i\Gamma$ distribution	$\theta \sim$ igamma(a, b) $\Leftrightarrow \log(\theta) \sim$ eigamma(a, b)	

Table 52.15 **Exponential Scaled Inverse chi^2 Distribution**

PROC specification	**expsichisq**(ν, s)
density	$\frac{(\frac{\nu}{2})^{\nu/2}}{\Gamma(\frac{\nu}{2})} s^\nu \exp(-\nu\theta/2) \exp(-\nu s^2/(2\exp(\theta)))$
parameter restriction	$\nu > 0, s > 0$
range	$\theta \in (-\infty, \infty)$
mode	$\log(s^2)$
random number	generate $x_1 \sim si\chi^2(\nu)$, and $\theta = \log(x_1)$ is a draw from the exponential scaled inverse χ^2 distribution
relationship to the $si\chi^2$ distribution	$\theta \sim si\chi^2(\nu) \Leftrightarrow \log(\theta) \sim \exp si\chi^2(\nu)$

Table 52.16 **Exponential Distribution**

PROC specification	**expon**(scale = b)	**expon**(iscale = β)
density	$\frac{1}{b}e^{-\theta/b}$	$\beta e^{-\beta\theta}$
parameter restriction	$b > 0$	$\beta > 0$
range	$\theta \in [0, \infty)$	same
mean	b	$1/\beta$
variance	b^2	$1/\beta^2$
mode	0	0
random number	the exponential distribution is a special case of the gamma distribution: $\theta \sim \text{gamma}(1, \text{scale} = b)$ is a draw from the exponential distribution	

Table 52.17 **Gamma Distribution**

PROC specification	**gamma**(a, scale = b)	**gamma**(a, iscale = β)
density	$\frac{1}{b^a \Gamma(a)}\theta^{a-1}e^{-\theta/b}$	$\frac{\beta^a}{\Gamma(a)}\theta^{a-1}e^{-\beta\theta}$
parameter restriction	$a > 0, b > 0$	$a > 0, \beta > 0$
range	$\theta \in [0, \infty)$ if $a = 1$; $(0, \infty)$ otherwise	same
mean	ab	a/β
variance	ab^2	a/β^2
mode	$(a - 1)b$ if $a \geq 1$	$(a - 1)/\beta$ if $a \geq 1$
random number	see (McGrath and Irving 1973)	

Table 52.18 **Geometric Distribution**

PROC specification	**geo**(p)
density [2]	$p(1 - p)^\theta$
parameter restriction	$0 < p \leq 1$
range	$\theta \in \begin{cases} \{0, 1, 2, \ldots\} & 0 < p < 1 \\ \{0\} & p = 1 \end{cases}$
mean	round($\frac{1-p}{p}$)
variance	$\frac{1-p}{p^2}$
mode	0
random number	based on samples obtained from a Bernoulli distribution with probability p until the first success

---

[2]The random variable θ is the total number of failures in an experiment *before* the first success. This density function is not to be confused with another popular formulation, $p(1 - p)^{\theta-1}$, which counts the total number of trials *until* the first success.

Table 52.19 Inverse χ^2 **Distribution**

PROC specification	**ichisq**(ν)
density	$\frac{1}{\Gamma(\nu/2)2^{\nu/2}}\theta^{-(\nu/2+1)}e^{-1/(2\theta)}$
parameter restriction	$\nu > 0$
range	$\theta \in (0, \infty)$
mean	$\frac{1}{\nu-2}$ if $\nu > 2$
variance	$\frac{2}{(\nu-2)^2(\nu-4)}$ if $\nu > 4$
mode	$\frac{1}{\nu+2}$
random number	inverse χ^2 is a special case of the inverse-gamma distribution: $\theta \sim$ igamma($\nu/2$, scale $= 2$) is a draw from the inverse χ^2 distribution

Table 52.20 **Inverse-Gamma Distribution**

PROC specification	**igamma**(a, scale $= b$)	**igamma**(a, iscale $= \beta$)
density	$\frac{b^a}{\Gamma(a)}\theta^{-(a+1)}e^{-b/\theta}$	$\frac{1}{\beta^a\Gamma(a)}\theta^{-(a+1)}e^{-1/\beta\theta}$
parameter restriction	$a > 0, b > 0$	$a > 0, \beta > 0$
range	$\theta \in (0, \infty)$	same
mean	$\frac{b}{a-1}$ if $a > 1$	$\frac{1}{\beta(a-1)}$ if $a > 1$
variance	$\frac{b^2}{(a-1)^2(a-2)}$	$\frac{1}{\beta^2(a-1)^2(a-2)}$
mode	$\frac{b}{a+1}$	$\frac{1}{\beta(a+1)}$
random number	generate $x_1 \sim$ gamma(a, scale $= b$), and $\theta = 1/x_1$ is a draw from the igamma(a, iscale $= b$) distribution	
relationship to the gamma distribution	$\theta \sim$ gamma(a, scale $= b$) $\Leftrightarrow$ $1/\theta \sim$ igamma(a, iscale $= b$)	

Table 52.21 **Laplace (Double Exponential) Distribution**

PROC specification	**laplace**(a, scale $= b$)	**laplace**(a, iscale $= \beta$)
density	$\frac{1}{2b}e^{-|\theta-a|/b}$	$\frac{\beta}{2}e^{-\beta|\theta-a|}$
parameter restriction	$b > 0$	$\beta > 0$
range	$\theta \in (-\infty, \infty)$	same
mean	a	a
variance	$2b^2$	$2/\beta^2$
mode	a	a
random number	inverse CDF. $F(\theta) = \begin{cases} \frac{1}{2}\exp\left(-\frac{a-\theta}{b}\right) & \theta < a \\ 1 - \frac{1}{2}\exp\left(-\frac{\theta-a}{b}\right) & \theta \geq a \end{cases}$. Generate $u_1, u_2 \sim$ uniform(0, 1). If $u_1 < 0.5, \theta = a + b\log(u_2)$; else $\theta = a - b\log(u_2)$. θ is a draw from the Laplace distribution	

Table 52.22 **Logistic Distribution**

PROC specification	**logistic**(a, b)
density	$\dfrac{\exp\left(-\frac{\theta-a}{b}\right)}{b\left(1+\exp\left(-\frac{\theta-a}{b}\right)\right)^2}$
parameter restriction	$b > 0$
range	$\theta \in (-\infty, \infty)$
mean	a
variance	$\dfrac{\pi^2 b^2}{3}$
mode	a
random number	inverse CDF method with $F(\theta) = \left(1 + \exp(-\frac{\theta-a}{b})\right)^{-1}$. Generate $u \sim \text{uniform}(0, 1)$, and $\theta = a - b \log(1/u - 1)$ is a draw from the logistic distribution

Table 52.23 **LogNormal Distribution**

	lognormal$(\mu, \text{sd} = s)$	lognormal$(\mu, \text{var} = v)$	lognormal$(\mu, \text{prec} = \tau)$
PROC specification	**lognormal**$(\mu, \text{sd} = s)$	**lognormal**$(\mu, \text{var} = v)$	**lognormal**$(\mu, \text{prec} = \tau)$
density	$\dfrac{1}{\theta s \sqrt{2\pi}} \exp\left(-\dfrac{(\log\theta - \mu)^2}{2s^2}\right)$	$\dfrac{1}{\theta \sqrt{2\pi v}} \exp\left(-\dfrac{(\log\theta - \mu)^2}{2v}\right)$	$\dfrac{1}{\theta}\sqrt{\dfrac{\tau}{2\pi}} \exp\left(-\dfrac{\tau(\log\theta - \mu)^2}{2}\right)$
parameter restriction	$s > 0$	$v > 0$	$\tau > 0$
range	$\theta \in (0, \infty)$	same	same
mean	$\exp(\mu + s^2/2)$	$\exp(\mu + v/2)$	$\exp(\mu + 1/(2\tau))$
variance	$\exp\left(2(\mu + s^2)\right)$ $- \exp\left(2\mu + s^2\right)$	$\exp\left(2(\mu + v)\right)$ $- \exp\left(2\mu + v\right)$	$\exp\left(2(\mu + 1/\tau)\right)$ $- \exp\left(2\mu + 1/\tau\right)$
mode	$\exp(\mu - s^2)$	$\exp(\mu - v)$	$\exp(\mu - 1/\tau)$
random number	generate $x_1 \sim \text{normal}(0, 1)$, and $\theta = \exp(\mu + sx_1)$ is a draw from the lognormal distribution		

Table 52.24 **Negative Binomial Distribution**

PROC specification	**negbin**(n, p)
density	$\begin{pmatrix} \theta + n - 1 \\ \theta \end{pmatrix} p^n (1 - p)^\theta$
parameter restriction	$n = 1, 2, \cdots 0 < p \le 1$
range	$\theta \in \begin{cases} \{0, 1, 2, \ldots\} & 0 < p < 1 \\ \{0\} & p = 1 \end{cases}$

mean	round$\left(\frac{n(1-p)}{p}\right)$
variance	$\frac{n(1-p)}{p^2}$
mode	$\begin{cases} 0 & n = 1 \\ \text{round}\left(\frac{(n-1)(1-p)}{p}\right) & n > 1 \end{cases}$
random number	generate $x_1 \sim \text{gamma}(n, 1)$, and $\theta \sim \text{Poisson}(x_1 \cdot (1-p)/p)$ (Fishman 1996).

Table 52.25   **Normal Distribution**

	normal(μ, sd = s)	**normal**(μ, var = v)	**normal**(μ, prec = τ)
PROC specification			
density	$\frac{1}{s\sqrt{2\pi}} \exp\left(-\frac{(\theta-\mu)^2}{2s^2}\right)$	$\frac{1}{\sqrt{2\pi v}} \exp\left(-\frac{(\theta-\mu)^2}{2v}\right)$	$\sqrt{\frac{\tau}{2\pi}} \exp\left(-\frac{\tau(\theta-\mu)^2}{2}\right)$
parameter restriction	$s > 0$	$v > 0$	$\tau > 0$
range	$\theta \in (-\infty, \infty)$	same	same
mean	μ	same	same
variance	s^2	v	$1/\tau$
mode	μ	same	same

Table 52.26   **Pareto Distribution**

PROC specification	**pareto**(a, b)
density	$\frac{a}{b}\left(\frac{b}{\theta}\right)^{a+1}$
parameter restriction	$a > 0, b > 0$
range	$\theta \in [b, \infty)$
mean	$\frac{ab}{a-1}$ if $a > 1$
variance	$\frac{b^2 a}{(a-1)^2(a-2)}$ if $a > 2$
mode	b
random number	inverse CDF method with $F(\theta) = 1 - (b/\theta)^a$. Generate $u \sim$ uniform(0, 1), and $\theta = \frac{b}{u^{1/a}}$ is a draw from the Pareto distribution.
useful transformation	$x = 1/\theta$ is Beta(a, 1)$\mathbf{I}\{x < 1/b\}$.

Table 52.27 **Poisson Distribution**

PROC specification	**poisson**(λ)
density	$\frac{\lambda^{\theta}}{\theta!}\exp(-\lambda)$
parameter restriction	$\lambda \geq 0$
range	$\theta \in \begin{cases} \{0, 1, \ldots\} & \text{if } \lambda > 0 \\ \{0\} & \text{if } \lambda = 0 \end{cases}$
mean	λ
variance	λ, if $\lambda > 0$
mode	$\text{round}(\lambda)$

Table 52.28 **Scaled Inverse χ^2 Distribution**

PROC specification	**sichisq**(ν, s^2)
density	$\frac{(\nu/2)^{\nu/2}}{\Gamma(\nu/2)}s^{\nu}\theta^{-(\nu/2+1)}e^{-\nu s^2/(2\theta)}$
parameter restriction	$\nu > 0, s > 0$
range	$\theta \in (0, \infty)$
mean	$\frac{\nu}{\nu-2}s^2$ if $\nu > 2$
variance	$\frac{2\nu^2}{(\nu-2)^2(\nu-4)}s^4$ if $\nu > 4$
mode	$\frac{\nu}{\nu+2}s^2$
random number	scaled inverse χ^2 is a special case of the inverse-gamma distribution: $\theta \sim \text{igamma}(\nu/2, \text{scale} = (\nu s^2)/2)$ is a draw from the scaled inverse χ^2 distribution.

Table 52.29 **T Distribution**

PROC specification	**t**(μ, sd = s, ν)	**t**(μ, var = v, ν)	**t**(μ, prec = τ, ν)
density	$\frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})s\sqrt{\nu\pi}}(1 + \frac{(\theta-\mu)^2}{\nu s^2})^{-\frac{\nu+1}{2}}$	$\frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\nu\pi v}}(1 + \frac{(\theta-\mu)^2}{\nu v})^{-\frac{\nu+1}{2}}$	$\frac{\Gamma(\frac{\nu+1}{2})\sqrt{\tau}}{\Gamma(\frac{\nu}{2})\sqrt{\nu\pi}}(1 + \frac{\tau(\theta-\mu)^2}{\nu})^{-\frac{\nu+1}{2}}$
parm restriction	$s > 0, \nu > 0$	$v > 0, \nu > 0$	$\tau > 0, \nu > 0$
range	$\theta \in (-\infty, \infty)$	same	same
mean	μ if $\nu > 1$	same	same
variance	$\frac{\nu}{\nu-2}s^2$ if $\nu > 2$	$\frac{\nu}{\nu-2}v$ if $\nu > 2$	$\frac{\nu}{\nu-2}\frac{1}{\tau}$ if $\nu > 2$
mode	μ	same	same
random number	$x_1 \sim \text{normal}(0, 1)$, $x_2 \sim \chi^2(d)$, and $\theta = m + \sigma x_1\sqrt{d/x_2}$ is a draw from the t-distribution.		

Table 52.30 **Uniform Distribution**

PROC specification	**uniform**(a, b)
density	$\begin{cases} \frac{1}{a-b} & \text{if } a > b \\ \frac{1}{b-a} & \text{if } b > a \\ 1 & \text{if } a = b \end{cases}$
parameter restriction	none
range	$\theta \in [a, b]$
mean	$\frac{a+b}{2}$
variance	$\frac{|b-a|^2}{12}$
mode	does not exist
random number	Mersenne Twister (Matsumoto and Kurita 1992, 1994; Matsumoto and Nishimura 1998)

Table 52.31 **Wald Distribution**

PROC specification	**wald**(μ, λ)
density	$\sqrt{\frac{\lambda}{2\pi\theta^3}} \exp\left(\frac{-\lambda(\theta-\mu)^2}{2\mu^2\theta}\right)$
parameter restriction	$\mu > 0, \lambda > 0$
range	$\theta \in (0, \infty)$
mean	μ
variance	μ^3/λ
mode	$\mu\left[\left(1 + \frac{9\mu^2}{4\lambda^2}\right)^{1/2} - \frac{3\mu}{2\lambda}\right]$
random number	generate $v_0 \sim \chi^2_{(1)}$. Let $x_1 = \mu + \frac{\mu^2 v_0}{2\lambda} - \frac{\mu}{2\lambda}\sqrt{4\mu\lambda v_0 + \mu^2 v_0^2}$ and $x_2 = \mu^2/x_1$. Perform a Bernoulli trial, $w \sim$ Bernoulli$(\frac{\mu}{\mu+x_1})$. If $w = 1$, choose $\theta = x_1$; otherwise, choose $\theta = x_2$ (Michael, Schucany, and Haas 1976).

Table 52.32 **Weibull Distribution**

PROC specification	**weibull**(μ, c, σ)
density	$\exp\left(-\left(\frac{\theta-\mu}{\sigma}\right)^c\right) \frac{c}{\sigma}\left(\frac{\theta-\mu}{\sigma}\right)^{c-1}$
parameter restriction	$c > 0, \sigma > 0$
range	$\theta \in [\mu, \infty)$ if $c = 1$; (μ, ∞) otherwise
mean	$\mu + \sigma\Gamma(1 + 1/c)$
variance	$\sigma^2[\Gamma(1 + 2/c) - \Gamma^2(1 + 1/c)]$
mode	$\mu + \sigma(1 - 1/c)^{1/c}$ if $c > 1$
random number	inverse CDF method with $F(\theta) = 1 - \exp\left(-\left(-\frac{\theta-\mu}{\sigma}\right)^c\right)$. Generate $u \sim$ uniform$(0, 1)$, and $\theta = \mu + \sigma \cdot (-\ln u)^{1/c}$ is a draw from the Weibull distribution.

Specifying a New Distribution

To work with a new density that is not listed in the section "Standard Distributions" on page 3529, you can use the GENERAL and DGENERAL functions. The letter "D" stands for discrete. The new distributions have to be specified on the logarithm scale.

Suppose that you want to use the inverse-beta distribution:

$$p(\alpha|a, b) = \frac{\Gamma(a + b)}{\Gamma(a) + \Gamma(b)} \cdot \alpha^{(a-1)} \cdot (1 + \alpha)^{-(a+b)}$$

The following statements in PROC MCMC define the density on its log scale:

```
a = 3; b = 5;
const = lgamma(a + b) - lgamma(a) - lgamma(b);
lp = const + (a - 1) * log(alpha) - (a + b) * log(1 + alpha);
prior alpha ~ general(lp);
```

The symbol lp is the expression for the log of an inverse-beta (a = 3, b = 5). The function `general(lp)` assigns that distribution to alpha. Note that the constant term, const, can be omitted as the Markov simulation requires only the log of the density kernel.

When you use the GENERAL function in the MODEL statement, you do not need to specify the dependent variable on the left of the tilde (~). The log-likelihood function takes the dependent variable into account; hence there is no need to explicitly state the dependent variable in the MODEL statement. However, in the PRIOR statements, you need to explicitly state the parameter names and a tilde with the GENERAL and DGENERAL functions.

You can specify any distribution function by using the GENERAL and DGENERAL functions as long as they are programmable with SAS statements. When the function is used in the PRIOR statements, you must supply initial values. This can be done in either the PARMS statement ("PARMS Statement" on page 3514) or within the BEGINCNST and ENDCNST statements ("BEGINCNST/ENDCNST Statement" on page 3507).

It is important to remember that PROC MCMC does not verify that the GENERAL function you specify is a valid distribution—that is, an integrable density. You must use the function with caution.

Using Density Functions in the Programming Statements

Density Functions in PROC MCMC

PROC MCMC also has a number of internally defined log-density functions. The functions have the basic form of lpdf*dist*(*x*, *parm-list*, < *lower* >, < *upper* >), where *dist* is the name of the distribution (see Table 52.33). The argument *x* is the random variable, *parm-list* is the list of parameters, and *lower* and *upper* are boundary arguments. The *lower* and *upper* arguments are optional but positional. With the exception of the Bernoulli and uniform distribution, you can specify limits on all distributions.

To set a lower bound on the normal density:

```
lpdfnorm(x, 0, 1, -2);
```

To set just an upper bound, specify a missing value for the lower bound argument:

```
lpdfnorm(x, 0, 1, ., 2);
```

Leaving both limits out gives you the unbounded density, and you can also specify both bounds:

```
lpdfnorm(x, 0, 1);
lpdfnorm(x, 0, 1, -3, 4);
```

See the following table for a list of distributions and their corresponding lpdf functions.

Table 52.33 Logarithm of Density Functions in PROC MCMC

Distribution Name	Function Call
beta	`lpdfbeta(x, a, b, <lower>, <upper>);`
binary	`lpdfbern(x, p);`
binomial	`lpdfbin(x, n, p, <lower>, <upper>);`
Cauchy	`lpdfcau(x, loc, scale, <lower>, <upper>);`
χ^2	`lpdfchisq(x, df, <lower>, <upper>);`
exponential χ^2	`lpdfechisq(x, df, <lower>, <upper>);`
exponential gamma	`lpdfegamma(x, sp, scale, <lower>, <upper>);`
exponential expo-nential	`lpdfeexpon(x, scale, <lower>, <upper>);`
exponential inverse χ^2	`lpdfeichisq(x, df, <lower>, <upper>);`
exponential inverse-gamma	`lpdfeigamma(x, sp, scale, <lower>, <upper>);`
exponential scaled inverse χ^2	`lpdfesichisq(x, df, scale, <lower>, <upper>);`
exponential	`lpdfexpon(x, scale, <lower>, <upper>);`
gamma	`lpdfgamma(x, sp, scale, <lower>, <upper>);`
geometric	`lpdfgeo(x, p, <lower>, <upper>);`
inverse χ^2	`lpdfichisq(x, df, <lower>, <upper>);`
inverse-gamma	`lpdfigamma(x, sp, scale, <lower>, <upper>);`
Laplace	`lpdfdexp(x, loc, scale, <lower>, <upper>);`
logistic	`lpdflogis(x, loc, scale, <lower>, <upper>);`
lognormal	`lpdflnorm(x, loc, sd, <lower>, <upper>);`
negative binomial	`lpdfnegbin(x, n, p, <lower>, <upper>);`
normal	`lpdfnorm(x, mu, sd, <lower>, <upper>);`
Pareto	`lpdfpareto(x, sp, scale, <lower>, <upper>);`
Poisson	`lpdfpoi(x, mean, <lower>, <upper>);`
scaled inverse χ^2	`lpdfsichisq(x, df, scale, <lower>, <upper>);`
T	`lpdft(x, mu, sd, df, <lower>, <upper>);`
uniform	`lpdfunif(x, a, b);`

Table 52.33 (continued)

Distribution Name	Function Call
Wald	`lpdfwald(`*x*`, `*mean*`, `*scale*`, <lower>,` `<upper>);`
Weibull	`lpdfwei(`*x*`, `*loc*`, `*sp*`, `*scale*`, <lower>,` `<upper>);`

Standard Distributions, the logpdf Functions, and the lpdf*dist* Functions

Standard distributions listed in the section "Standard Distributions" on page 3529 are *names* only, and they can only be used in the MODEL, PRIOR, and HYPERPRIOR statements to specify either a prior distribution or a conditional distribution of the data given parameters. They do not return any values, and you cannot use them in the programming statements.

The LOGPDF functions are DATA step functions that compute the logarithm of various probability density (mass) functions. For example, `logpdf("beta", x, 2, 15)` returns the log of a beta density with parameters a = 2 and b = 15, evaluated at x. All the LOGPDF functions are supported in PROC MCMC.

The lpdf*dist* functions are unique to PROC MCMC. They compute the logarithm of various probability density (mass) functions. The functions are the same as the LOGPDF functions when it comes to calculating the log density. For example, `lpdfbeta(x, 2, 15)` returns the same value as `logpdf("beta", x, 2, 15)`. The lpdf*dist* functions cover a greater class of probability density functions, and they take the optional but positional boundary arguments. There are no corresponding `lcdf`*dist* or `lsdf`*dist* functions in PROC MCMC. To work with the cumulative probability function or the survival functions, you need to use the LOGCDF and the LOGSDF DATA step functions.

Truncation and Censoring

Truncated Distributions

To specify a truncated distribution, you can use the LOWER= and/or UPPER= options. Almost all of the standard distributions, including the GENERAL and DGENERAL functions, take these optional truncation arguments. The exceptions are the binary and uniform distributions.

For example, you can specify the following:

```
prior alpha ~ normal(mean = 0, sd = 1, lower = 3, upper = 45);
```

or

```
parms beta;
a = 3; b = 7;
ll = (a + 1) * log(b / beta);
prior beta ~ general(ll, upper = b + 17);
```

The preceding statements state that if beta is less than b+17, the log of the prior density is ll, as calculated by the equation; otherwise, the log of the prior density is missing—the log of zero.

When the same distribution is applied to multiple parameters in a PRIOR statement, the LOWER= and UPPER= truncations apply to all parameters in that statement. For example, the following statements define a Poisson density for theta and gamma:

```
parms theta gamma;
lambda = 7;
l1 = theta * log(lambda) - lgamma(1 + theta);
l2 = gamma * log(lambda) - lgamma(1 + gamma);
ll = l1 + l2;
prior theta gamma ~ dgeneral(ll, lower = 1);
```

The LOWER=1 condition is applied to both theta and gamma, meaning that for the assignment to ll to be meaningful, both theta and gamma have to be greater than 1. If either of the parameters is less than 1, the log of the joint prior density becomes a missing value.

With the exceptions of the normal distribution and the GENERAL and DGENERAL functions, the LOWER= and UPPER= options cannot be parameters or functions of parameters. The reason is that most of the truncated distributions are not normalized. Unnormalized densities do not lead to wrong MCMC answers as long as the bounds are constants. However if the bounds involve model parameters, then the normalizing constant, which is a function of these parameters, has to be taken into account in the posterior. Without specifying the normalizing constant, inferences on these boundary parameters are incorrect.

It is not difficult to construct a truncated distribution with a normalizing constant. Any truncated distribution has the probability distribution:

$$p(\theta|a < \theta < b) = \frac{p(\theta)}{F(a) - F(b)}$$

where $p(\cdot)$ is the density function and $F(\cdot)$ is the cumulative distribution function. In SAS functions, $p(\cdot)$ is probability density function and $F(\cdot)$ is cumulative distribution function. The following example shows how to construct a truncated gamma prior on theta, with SHAPE = 3, SCALE = 2, LOWER = a, and UPPER = b:

```
lp = logpdf('gamma', theta, 3, 2)
        - log(cdf('gamma', a, 3, 2) - cdf('gamma', b, 3, 2));
prior theta ~ general(lp);
```

Note the difference from a naive definition of the density, without taking into account of the normalizing constant:

```
lp = logpdf('gamma', theta, 3, 2);
prior theta ~ general(lp, lower=a, upper=b);
```

If a or b are parameters, you get very different results from the two formulations.

Censoring

There is no built-in mechanism in PROC MCMC that models censoring automatically. You need to construct the density function (using a combination of the LOGPDF, LOGCDF, and LOGSDF

functions and IF-ELSE statements) for the censored data.

Suppose that you partition the data into four categories: uncensored (with observation x), left censored (with observation xl), right censored (with observation xr), and interval censored (with observations xl and xr). The likelihood is the normal with mean mu and standard deviation s. The following statements construct the corresponding log likelihood for the observed data:

```
if uncensored then
   ll = logpdf('normal', x, mu, s);
else if leftcensored then
   ll = logcdf('normal', xl, mu, s);
else if rightcensored then
   ll = logsdf('normal', xr, mu, s);
else /* this is the case of interval censored. */
   ll = log(cdf('normal', xr, mu, s) - cdf('normal', xl, mu, s));
model general(ll);
```

See "Example 52.9: Normal Regression with Interval Censoring" on page 3658.

Multivariate Density Functions

The DATA step has functions that compute the logarithm of the density of some multivariate distributions. You can use them in PROC MCMC. For a complete listing of multivariate functions, see *SAS Language Reference: Dictionary*.

Some commonly used multivariate functions in Bayesian analysis are as follows:

- LOGMPDFNORMAL, the logarithm of the multivariate normal

- LOGMPDFWISHART, the logarithm of the Wishart

- LOGMPDFIWISHART, the logarithm of the inverted-Wishart

- LOGMPDFDIR1, the logarithm of the Dirichlet distribution of Type I

- LOGMPDFDIR2, the logarithm of the Dirichlet distribution of Type II

- LOGMPDFMULTINOM, the logarithm of the multinomial

Other multivariate density functions include: LOGMPDFT (*t*-distribution), LOGMPDFGAMMA (gamma distribution), LOGMPDFBETA1 (beta of type I), and LOGMPDFBETA2 (beta of type II).

Density Function Definition

LOGMPDFNORMAL

Let x be an n-dimensional random vector with mean vector μ and covariance matrix Σ. The density is

$$pdf(x; \mu, \Sigma) = \frac{\exp(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu))}{\sqrt{(2\pi)^n |\Sigma|}}$$

where $|\mathbf{\Sigma}|$ is the determinant of the covariance matrix $\mathbf{\Sigma}$.

The function has syntax:

$$y = \text{LOGPDFMNORM}(x\_list, \mu\_list, cov\_name);$$

WARNING: you must set up the cov_name covariance matrix before using the LOGPDFMNORM function and free the memory after PROC MCMC exits. See the section "Set Up the Covariance Matrices and Free Memory" on page 3546.

LOGMPDFWISHART and LOGMPDFIWISHART

The density function from the Wishart distribution is:

$$pdf(x; \mu, \mathbf{\Sigma}) = \frac{1}{C_n(\mu)} |\mathbf{\Sigma}|^{-\frac{\mu}{2}} |x|^{\frac{\mu-n-1}{2}} \exp\left(-\frac{1}{2}tr(\mathbf{\Sigma}^{-1}x)\right)$$

with $\mu > n$, and the trace of a square matrix A is given by:

$$tr(A) = \sum_i a_{ii} \quad C_n(\mu) = 2^{\frac{\mu n}{2}} \Gamma_n\left(\frac{\mu}{2}\right) \quad \Gamma_n(z) = \pi^{\frac{n(n-1)}{4}} \prod_{i=1}^{n} \Gamma\left(z - \frac{i-1}{2}\right)$$

The density function from the inverse-Wishart distribution is:

$$pdf(x; \mu, \mathbf{\Sigma}) = \frac{1}{D_n(\mu)} |\mathbf{\Sigma}|^{\frac{\mu-n-1}{2}} |x|^{-\frac{\mu}{2}} \exp\left(-\frac{1}{2}tr(\mathbf{\Sigma}x^{-1})\right)$$

for $\mu > 2n$, and

$$D_n(\mu) = 2^{\frac{(\mu-n-1)n}{2}} \Gamma_n\left(\frac{\mu-n-1}{2}\right)$$

If $V \sim IW_n(\mu, \mathbf{\Sigma})$ then $V^{-1} \sim W_n(\mu - n - 1, \mathbf{\Sigma}^{-1})$

The functions have syntax:

$$y = \text{LOGMPDFWISHART}('name'_V, \mu, 'name'_\mathbf{\Sigma});$$

and for the inverted Wishart:

$$y = \text{LOGMPDFIWISHART}('name'_V, \mu, 'name'_\mathbf{\Sigma});$$

The three arguments are the multivariate matrix $'name'_V$, the degrees of freedom μ, and the covariance matrix $'name'_\mathbf{\Sigma}$ k

WARNING: you must set up the cov_name covariance matrix before using these functions and free the memory after PROC MCMC exits. See the section "Set Up the Covariance Matrices and Free Memory" on page 3546.

LOGMPDFDIR1 and LOGMPDFDIR2

The random variables $u_1...u_k$, with $u_i > 0$ and $\sum_{i=1}^{k} u_i < 1$, are said to have a Dirichlet Type I distribution with parameters $a_1...a_{k+1}$ if their joint pdf is given by:

$$pdf_1(u_1, u_2, ..., u_k, a_1, a_2, ..., a_{k+1}) = \frac{\Gamma(\sum_{i=1}^{k+1} a_i)}{\prod_{i=1}^{r+1} \Gamma(a_i)} \left(\prod_{i=1}^{k} u_i^{a_i-1}\right) \left(1 - \sum_{i=1}^{k} u_i\right)^{a_{k+1}-1}$$

The variables are said to have a Dirichlet type II distribution with parameters $a_1...a_{k+1}$ if their joint pdf is given by the following:

$$pdf_2(u_1, u_2, ..., u_k, a_1, a_2, ..., a_{k+1}) = \frac{\Gamma(\sum_{i=1}^{k+1} a_i)}{\prod_{i=1}^{r+1} \Gamma(a_i)} \left(\prod_{i=1}^{k} u_i^{a_i-1}\right) \left(1 + \sum_{i=1}^{k} u_i\right)^{-\sum_{i=1}^{k+1} a_i}$$

The functions have syntax:

$$y = \text{LOGMPDFDIR1}(u\_list, a\_list);$$

and

$$y = \text{LOGMPDFDIR2}(u\_list, a\_list);$$

LOGMPDFMULTINOM

Let $n_1...n_k$ be random variables that denote the number of occurring of the events $E_1, ...E_k$ respectively occurring with probabilities $p_1...p_k$. Let $\sum_{i=1}^{k} p_i = 1$ and let $n = \sum_{i=1}^{k} n_i$. Then the joint distribution of $n_1, ...n_k$ is the following:

$$pdf(n_1, n_2, ...n_k, p_1, p_2, ..., p_k) = n! \prod_{i=1}^{k} \left(\frac{p_i^{n_i}}{n_i!}\right)$$

The function has syntax:

$$y = \text{LOGMPDFMULTINOM}(n\_list, p\_list);$$

Set Up the Covariance Matrices and Free Memory

For distributions that require symmetric positive definite matrices, such as the LOGMPDFNORMAL, LOGPDFMWISHART and LOGMPDFIWISHART functions, you need to set up these matrices by using the following functions:

- Use LOGMPDFSETSQ to set up a symmetric positive definite matrix from all its elements:

$$rc = \text{LOGMPDFSETSQ}(name, num1, num2, ...);$$

  rc is set to 0 when the numeric arguments describe a symmetric positive definite matrix, otherwise it is set to a nonzero value.

- Use LOGMPDFSET to set up a symmetric positive definite matrix from its lower triangular elements:

$$rc = \text{LOGMPDFSET}(name, num1, num2, ...);$$

When the numeric arguments describe a symmetric positive definite matrix, the returned value rc is set to 0. Otherwise, a nonzero value for rc is returned.

- Use LOGMPFFREE to free the workspace previously allocated with either LOGMPDFSET or LOGMPDFSETSQ:

$$rc = \text{LOGMPDFFREE}(< ... < 'name' >, 'name2' > ...);$$

When called without arguments, the LOGMPDFFREE frees all the symbols previously allocated by LOGMPDFSETSQ or LOGMPDFSET. Each freed symbol is reported back in the SAS log.

The parameters used in these functions are defined as follows:

name   is a string containing the name of the work space that stores the matrix by the numeric parameters $num1, ...$

num1, ...   are numeric arguments that represent the elements of a symmetric positive definite matrix.

You would set up this matrix under the DATA step by using the following syntax:

$$rc = \text{LOGMPDFSETSQ}(name, \sigma_{11}, \sigma_{12}, \sigma_{21}, \sigma_{22});$$

or the syntax:

$$rc = \text{LOGMPDFSET}(name, \sigma_{11}, \sigma_{21}, \sigma_{22});$$

If the matrix is positive definite, the returned value rc is zero.

Some Useful SAS Functions

Table 52.34   Some Useful SAS Functions

SAS Function	Definition
`abs(x)`	$|x|$
`airy(x)`	returns the value of the AIRY function.
`beta(x1, x2)`	$\int_0^1 z^{x1-1}(1-z)^{x2-1}dz$
`call logistic(x)`	$\ln(x/(1-x))$
`call softmax(x1,...,xn)`	each element is replaced by $\exp(x_j)/\sum \exp(x_j)$
`call stdize(x1,...,xn)`	standardize values

Table 52.34 (continued)

SAS Function	Definition
`cdf`	cumulative distribution function
`cdf('normal', x, 0, 1)`	standard normal cumulative distribution function
`comb(x1, x2)`	$\frac{x1!}{x2!(x1-x2)!}$
`constant('.')`	calculate commonly used constants
`cos(x)`	cosine(x)
`css(x1, ..., xn)`	$\sum_i (x_i - \bar{x})^2$
`cv(x1, ..., xn)`	std(x) / mean(x) * 100
`dairy(x)`	derivative of the AIRY function
`dimN(m)`	returns the numbers of elements in the Nth dim of array m
`(x1 eq x2)`	returns 1 if x1 = x2; 0 otherwise
`x1**x2`	$x1^{x2}$
`geomean(x1, ..., xn)`	$\exp\left(\frac{\log(x_1)+\cdots+\log(x_n)}{n}\right)$
`difN(x)`	returns differences between the argument and its Nth lag
`digamma(x1)`	$\frac{\Gamma'(x1)}{\Gamma(x1)}$
`erf(x)`	$\frac{2}{\sqrt{\pi}} \int_0^x exp(-z^2)dz$
`erfc(x)`	1 - erf(x)
`fact(x)`	$x!$
`floor(x)`	greatest integer $\leq x$
`gamma(x)`	$\int_0^\infty z^{x-1} \exp(-1)dz$
`harmean(x1, ..., xn)`	$\frac{n}{1/x_1+\cdots 1/x_n}$
`ibessel(nu, x, kode)`	modified Bessel function of order nu evaluated at x
`jbessel(nu, x)`	Bessel function of order nu evaluated at x
`lagN(x)`	returns values from a queue
`largest(k, x1, ..., xn)`	the k^{th} largest element
`lgamma(x)`	$\ln(\Gamma(x))$
`lgamma(x+1)`	$\ln(x!)$
`log(x), logN(x)`	$\ln(x)$
`logbeta(x1, x2)`	lgamma(x_1) + lgamma(x_2) - lgamma($x_1 + x_2$)
`logcdf`	log of a left cumulative distribution function
`logpdf`	log of a probability density (mass) function
`logsdf`	log of a survival function
`max(x1, x2)`	returns x_1 if $x_1 > x_2$; x_2 otherwise
`mean(of x1-xn)`	$\sum_i x_i/n$
`median(of x1-xn)`	returns the median of nonmissing values
`min(x1, x2)`	returns x_1 if $x_1 < x_2$; x_2 otherwise
`missing(x)`	returns 1 if x is missing; 0 otherwise
`mod(x1, x2)`	returns the remainder from x_1/x_2
`n(x1, ..., xn)`	returns number of nonmissing values
`nmiss(of y1-yn)`	number of missing values
`quantile`	computes the quantile from a specific distribution
`pdf`	probability density (mass) functions
`perm(n, r)`	$\frac{n!}{(n-r)!}$
`put`	returns a value that uses a specified format

Table 52.34 (continued)

SAS Function	Definition
`round(x)`	rounds x
`rms(of x1-xn)`	$\sqrt{\frac{x_1^2 + \cdots x_n^2}{n}}$
`sdf`	survival function
`sign(x)`	returns -1 if $x < 0$; 0 if $x = 0$; 1 if $x > 0$
`sin(x)`	sine(x)
`smallest(s, x1, ..., en )`	the s^{th} smallest component of $x_1, \cdots, x_n$
`sortn(of x1-xn)`	sorts the values of the variables
`sqrt(x)`	$\sqrt{x}$
`std(x1, ..., xn)`	standard deviation of $x_1, \cdots, x_n$ (n-1 in denominator)
`sum(of x:)`	$\sum_i x_i$
`trigamma(x)`	derivative of the DIGAMMA(x) function
`uss(of x1-xn)`	uncorrected sum of squares

Here are examples of some commonly used transformations:

- logit

  ```
  mu = beta0 + beta1 * z1;
  call logistic(mu);
  ```

- log

  ```
  w = beta0 + beta1 * z1;
  mu = exp(w);
  ```

- probit

  ```
  w = beta0 + beta1 * z1;
  mu = cdf('normal', w, 0, 1);
  ```

- cloglog

  ```
  w = beta0 + beta1 * z1;
  mu = 1 - exp(-exp(w));
  ```

Modeling Dependent Data

PROC MCMC assumes that the input observations are independent and that the joint log likelihood is the sum of individual log-likelihood functions. You specify the log likelihood of one observation in the MODEL statement. PROC MCMC evaluates that function for each observation in the data set and cumulatively sums them up. If observations are not independent of each other, this summation produces the incorrect log likelihood.

There are two ways to model dependent data. You can either use the DATA step function LAG or use the PROC option JOINTMODEL. The LAG function returns values of a variable from a queue. As

PROC MCMC steps through the data set, the LAG function queues each data set variable, and you have access to the current value as well as all previous values of any variable. If the log likelihood for observation x_i depends only on observations 1 to i in the data set, you can use this SAS function to construct the log-likelihood function for each observation. Note that the LAG function allows you to access observations from different rows, but the log-likelihood function in the MODEL statement has to be generic enough that it applies to all observations. See "Example 52.8: Cox Models" on page 3642 for how to use this LAG function.

A second option is to create arrays, store all relevant variables in the arrays, and construct the joint log likelihood for the entire data set instead of for each observation. Here is a simple example that uses a normal independent likelihood:

```
/* allocate the sample size. */
%let n = 100;
data exi;
   call streaminit(17);
   do ind = 1 to &n;
      y = rand("normal", 2.3, 1);
      output;
   end;
run;
```

The log-likelihood function for each observation is as follows:

$$\log(f(y_i|\mu,\sigma)) = \log(\phi(y_i;\mu,\text{var} = \sigma^2))$$

The joint log-likelihood function is as follows:

$$\log(f(\mathbf{y}|\mu,\sigma)) = \sum_i \log(\phi(y_i;\mu,\text{var} = \sigma^2))$$

The following statements fit a simple normal model with an unknown mean (mu) in PROC MCMC, with the variance in the likelihood assumed known. The MODEL statement indicates a normal likelihood for each observation y.

```
proc mcmc data=exi seed=7 outpost=p1;
   parm mu;
   prior mu ~ normal(0, sd=10);
   model y ~ normal(mu, sd=1);
run;
```

The following statements show how you can specify the log-likelihood function for the entire data set:

```
proc mcmc data=exi seed=7 outpost=p2 jointmodel;
   array data[&n];
   begincnst;
      data[ind] = y;
   endcnst;

   parm mu;
   prior mu ~ normal(0, sd=10);
   ll = 0;
```

```
      do i = 1 to &n;
         ll = ll + lpdfnorm(data[i], mu, 1);
      end;
      model general(ll);
   run;
```

The option JOINTMODEL indicates that the function used in the MODEL statement calculates the log likelihood for the entire data set, and not just for one observation. Given this option, PROC MCMC interprets the MODEL statement as the log-likelihood function for the data set, and the procedure no longer steps through the input data set to calculate the joint log-likelihood function. Hence, every term in the log-likelihood function must be presented in the program independently of any data set variables.

You need to allocate arrays and store the full data set in the arrays. The ARRAY statement allocates an array (data) of the same size as the input data set (&n). The single statement in the BEGINCNST and ENDCNST statements fills in array data with y. See the section "BEGINCNST/ENDCNST Statement" on page 3507 for an explanation on this line of statement. PROC MCMC executes this line of code prior to the simulation by using the entire data set. As the variable ind progresses, data[ind] takes on the corresponding value of y.

In the programming statements for the model, you construct ll by using a do-loop. The expression ll in the GENERAL function is now the joint log-likelihood function for all of y.

You can run the following statements to see that two PROC MCMC runs produce identical result.

```
   proc compare data=p1 compare=p2;
      var mu;
   run;
```

For a more realistic example that models dependent data, see "Example 52.8: Cox Models" on page 3642.

It is important to remember that, with the JOINTMODEL specified, PROC MCMC does not access the data set in the simulation. The procedure only steps through the data set during the setup stage, where all observations are read and saved in the allocated arrays in the BEGINCNST and ENDCNST statements. After the simulation begins, PROC MCMC uses the arrays to do its posterior density calculations. Any data set variable symbol outside of the BEGINCNST and ENDCNST statements no longer has meaning and is not allowed. The following program produces an error:

```
   proc mcmc data=exi seed=7 outpost=p2 jointmodel;
      array data[&n];
      begincnst;
         data[ind] = y;
      endcnst;

      parm mu;
      prior mu ~ normal(0, sd=10);
      ll = 0;
      do i = 1 to &n;
         ll = ll + lpdfnorm(data[i], mu, 1);
      end;
      v = mu * y;       /* an extra line, where the variable y appears */
      model general(ll);
   run;
```

This program is identical to the previous PROC MCMC run, except for the extra line that contains the data set variable y. PROC MCMC procedure exits with the following error message:

```
ERROR: When the option JOINTMODEL is specified, PROC MCMC calculates the
       log-likelihood function without accessing the data set. Data set
       variables, such as y in your program, cannot appear outside of the
       BEGINCNST/ENDCNST statements. See documentation for instructions and
       examples.
```

Regenerating Diagnostics Plots

By default, PROC MCMC generates three plots: the trace plot, the autocorrelation plot and the kernel density plot. Unless you had requested the display of ODS Graphics (`ods graphics on`) before calling the procedure, it is hard to generate the same graph afterwards. Directly using the template (Stat.MCMC.Graphics.TraceAutocorrDensity) is not feasible. To regenerate the same graph with a Markov chain, you need to define a template and use PROC SGRENDER (see the SGRENDER Procedure in the *SAS/GRAPH: Statistical Graphics Procedures Guide*) to create the graph. The following PROC TEMPLATE (see Chapter 21, "Statistical Graphics Using ODS") statements define a new graph template `mygraphs.mcmc`:

```
proc template;
   define statgraph mygraphs.mcmc;
   dynamic _sim _parm;
   BeginGraph;
      layout gridded /rows=2 columns=1 rowgutter=5;
         seriesplot x=_sim y=_parm;
         layout gridded /rows=1 columns=2 columngutter=15;
            layout overlay /
               yaxisopts=(linearopts=(viewmin=-1 viewmax=1
                                      tickvaluelist=(-1 -0.5 0 0.5 1))
                          label="Autocorrelation")
               xaxisopts=(linearopts=(integer=true)
                          label="Lag" offsetmin=.015);
               needleplot x=eval(lags(_parm,Max=50))
                          y=eval(acf(_parm, NLags=50));
            endlayout;

            layout overlay / xaxisopts=(label=_parm)
                             yaxisopts=(label="Density");
               densityplot _parm /kernel();
            endlayout;
         endlayout;
      endlayout;
   EndGraph;
   end;
```

The `define statgraph` statement tells PROC TEMPLATE that you are defining a new graph template (instead of a table or style template). The template is named `mygraphs.mcmc`. There are two dynamic variables: _sim and _parm. The variable _sim is the iteration number and the

variable _parm is the variable in the data set that stores the posterior sample. All STATGRAPH template definitions must start with a BEGINGRAPH statement and conclude with a ENDGRAPH statement. The first LAYOUT GRIDDED statement assembles the results of nested `statgraph` statements into a grid, with two rows and 1 column. The trace plot (SERIESPLOT) is shown in the first row of the graph. The second LAYOUT GRIDDED statement divides the second row of the graph into two graphs: one an autocorrelation plot (NEEDLEPLOT) and the other a kernel density plot (DENSITYPLOT). For details of other controls, such as the labels, line types, see Chapter 21, "Statistical Graphics Using ODS."

A simple regression example, with three parameters, is used here. For an explanation of the regression model and the data involved, see "Simple Linear Regression" on page 3478. The following statements generate a SAS data set and fit a regression model:

```
title 'Simple Linear Regression';

data Class;
   input Name $ Height Weight @@;
   datalines;
Alfred   69.0 112.5    Alice   56.5  84.0    Barbara 65.3  98.0
Carol    62.8 102.5    Henry   63.5 102.5    James   57.3  83.0
Jane     59.8  84.5    Janet   62.5 112.5    Jeffrey 62.5  84.0
John     59.0  99.5    Joyce   51.3  50.5    Judy    64.3  90.0
Louise   56.3  77.0    Mary    66.5 112.0    Philip  72.0 150.0
Robert   64.8 128.0    Ronald 67.0 133.0     Thomas  57.5  85.0
William 66.5 112.0
;

proc mcmc data=class nmc=50000 thin=5 outpost=classout seed=246810;
   ods select none;
   parms beta0 0 beta1 0;
   parms sigma2 1;
   prior beta0 beta1 ~ normal(0, var = 1e6);
   prior sigma2 ~ igamma(3/10, scale = 10/3);
   mu = beta0 + beta1*height;
   model weight ~ normal(mu, var = sigma2);
run;
ods select all;
```

The output data set classout contains iteration number (Iteration) and posterior draws for beta0, beta1, and sigma2. It also stores the log of the prior density (LogPrior), log of the likelihood (LogLike), and the log of the posterior density (LogPost). If you want to examine the LogPost variable, you can use the following statements to generate the graphs:

```
proc sgrender data=classout template=mygraphs.mcmc;
   dynamic _sim='iteration' _parm='logpost';
run;
```

The SGRENDER procedure takes the classout data set and applies the template MYGRAPHS.MCMC that was defined previously. The DYNAMIC statement needs two arguments, iteration and logpost. The resulting graph is shown in Output 52.11.

Figure 52.11 Regenerate Diagnostics Plots for Log of the Posterior Density
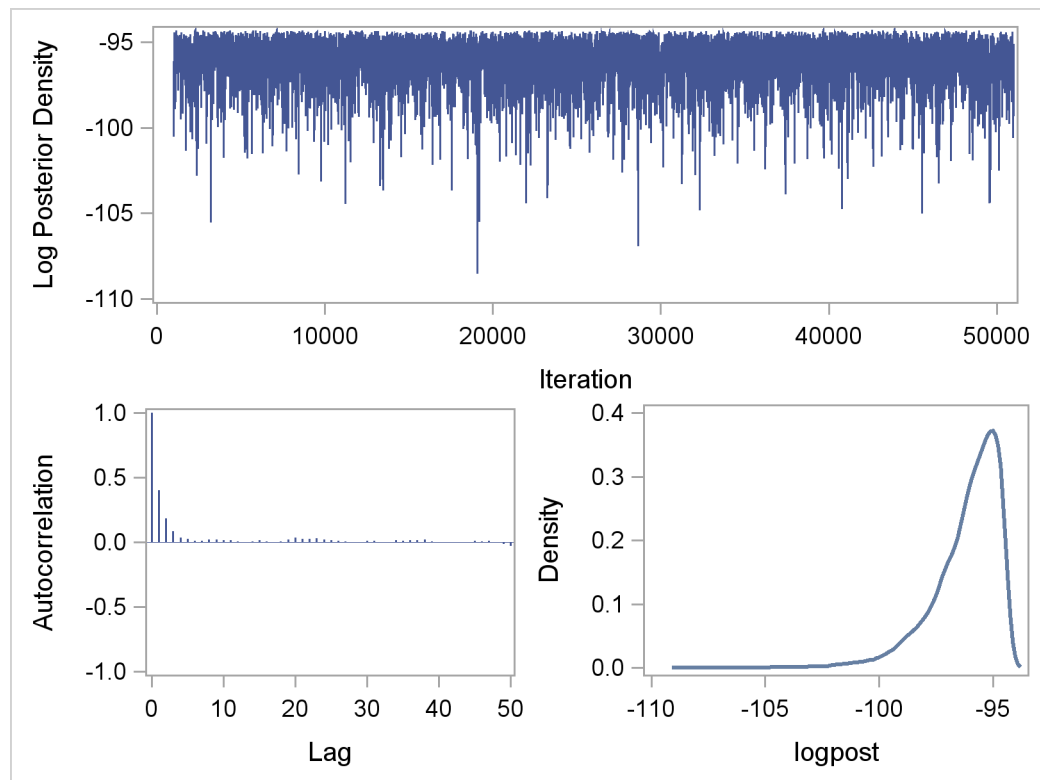
Posterior Predictive Distribution

The posterior predictive distribution is the distribution of unobserved observations (prediction) conditional on the observed data. Let $\mathbf{y}$ be the observed data, θ be the parameter, and $\mathbf{y}_{\text{pred}}$ be the unobserved data; the posterior predictive distribution is defined to be the following:

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

Given the assumption that the observed and unobserved data are conditional independent given θ, the posterior predictive distribution can be further simplified as the following:

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

The posterior predictive distribution is an integral of the likelihood function $p(\mathbf{y}_{\text{pred}}|\theta)$ with respect to the posterior distribution $p(\theta|\mathbf{y})$. You can use PROC MCMC to generate samples from a posterior predictive distribution based on draws from the posterior distribution of θ.

Note that the posterior predictive distribution is not the same as the prior predictive distribution. The prior predictive distribution is $p(\mathbf{y})$, which is also known as the marginal distribution of the

data. The prior predictive distribution is an integral of the likelihood function with respect to the prior distribution:

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

and the distribution is not conditional on observed data.

You can use the posterior predictive distribution to check whether the model is consistent with data. For more information about using predictive distribution as a model checking tool, see Gelman et al. 2004, Chapter 6 and the bibliography in that chapter. The idea is to generate replicate data from $p(\mathbf{y}_{\text{pred}}|\mathbf{y})$—call them $\mathbf{y}_{\text{pred}}^i$, for $i = 1, \cdots, M$, where M is the total number of replicates— compare them to the observed data, and see if there are any large and systematic differences. Large discrepancies suggest possible model misfit. One way to compare the replicate data to the observed data is to first summarize the data to some test quantities, such as the mean, standard deviation, order statistics, and so on. Then compute the tail-area probabilities of the test statistics (based on the observed data) with respect to the estimated posterior predictive distribution using the M replicate $\mathbf{y}_{\text{pred}}$ samples.

Let $T(\cdot)$ denote the function of the test quantity, $T(\mathbf{y})$ the test quantity using the observed data, and $T(\mathbf{y}_{\text{pred}}^i)$ the test quantity using the ith replicate data from the posterior predictive distribution. You calculate the tail-area probability using the following formula:

$$\Pr(T(\mathbf{y}_{\text{pred}}) > T(\mathbf{y})|\theta)$$

The following example shows how you can estimate this probability using PROC MCMC.

An Example for Posterior Predictive Distribution

This example uses a normal mixed model to analyze the effects of coaching programs for the scholastic aptitude test (SAT) in eight high schools. For the original analysis of the data, see Rubin (1981). The presentation here follows the analysis and posterior predictive check presented in Gelman et al. (2004). The data are as follows:

```
title 'An Example for Posterior Predictive Distribution';

data SAT;
   input effect sd @@;
   ind=_n_;
   datalines;
28.39 14.9  7.94 10.2 -2.75 16.3
 6.82 11.0 -0.64  9.4  0.63 11.4
18.01 10.4 12.16 17.6
;
```

The variable effect is the reported test score difference between coached and uncoached students in eight schools. The variable sd is the corresponding estimated standard error for that school. In a normal mixed effect model, the variable effect is assumed to be normally distributed:

$$\text{effect}_i \sim \text{normal}(\mu_i, \text{sd}^2) \quad \text{for } i = 1, \cdots, 8$$

The parameter μ_i has a normal prior with hyperparameters (m, v):

$$\mu_i \sim \text{normal}(m, \text{var} = v)$$

The hyperprior distribution on m is a uniform prior on the real axis, and the hyperprior distribution on v is a uniform prior from 0 to infinity.

The following statements fit a normal mixed model, general draws from the posterior predictive distribution, and calculate relevant test quantities.

```
ods listing close;
proc mcmc data=SAT outpost=pred nmc=50000 thin=10 seed=17
   monitor=(yrep mean sd max min);

   array theta[8];
   array yrep[8];
   begincnst;
   call streaminit(1);
   endcnst;

   parms theta: 0;
   parms m 0;
   parms v 1;

   hyper m ~ general(1);
   hyper v ~ general(1,lower=0);
   prior theta: ~ normal(m,var=v);
   mu = theta[ind];
   model effect ~ normal(mu,sd=sd);

   /* generate predictive data and calculate test statistics. */
   yrep[ind] = rand('normal', mu, sd);
   if (ind eq 8) then do;
      mean = mean(of yrep1-yrep8);
      sd = std(of yrep1-yrep8);
      max = max(of yrep1-yrep8);
      min = min(of yrep1-yrep8);
      end;
   run;
ods listing;
```

Four test quantities constructed are: the average (mean), the sample standard deviation (sd), the maximum effect (max), and the minimum effect (min). The MONITOR= option selects yrep (replicate samples) and the four test quantities and saves them to the OUTPOST= data set. The CALL STREAMINIT routine ensures that the RAND function, used here to generate posterior predictive samples, creates a reproducible stream of random numbers. The `ods listing close` statement disables listing output because you are primarily interested only in the samples of the monitored quantities. The HYPER, PRIOR, and MODEL statements specify the Bayesian model of interest.

The yrep[ind] assignment statement generates a random normal sample for each predictive observation, indexed by ind, with ind $= 1, \cdots, 8$. Note that this normal distribution is the same as the likelihood function specified in the MODEL statement, with the same mean and standard deviation.

To calculate the test quantities, you want to wait until all $\mathbf{y}^i_{\text{pred}}$ is generated—that is at the last observation of the data set.

The following statements compute the corresponding test statistics, the mean, standard deviation, and the minimum and maximum statistics on the real data and store them in macro variables. You then calculate the tail-area probabilities by counting the number of samples in the data set pred that are greater than the observed test statistics based on the real data.

```
proc means data=SAT noprint;
   var effect;
   output out=stat mean=mean max=max min=min stddev=sd;
   run;

data _null_;
   set stat;
   call symputx('mean',mean);
   call symputx('sd',sd);
   call symputx('min',min);
   call symputx('max',max);
   run;

data _null_;
   set pred end=eof nobs=nobs;
   ctmean + (mean>&mean);
   ctmin + (min>&min);
   ctmax + (max>&max);
   ctsd + (sd>&sd);
   if eof then do;
      pmean = ctmean/nobs; call symputx('pmean',pmean);
      pmin = ctmin/nobs; call symputx('pmin',pmin);
      pmax = ctmax/nobs; call symputx('pmax',pmax);
      psd = ctsd/nobs; call symputx('psd',psd);
      end;
   run;
```

You can plot histograms of each test quantity to visualize the posterior predictive distributions. In addition, you can see where the estimated *p*-values fall on these densities. Figure 52.12 shows the histograms. To put all four histograms on the same panel, you need to use PROC TEMPLATE (see Chapter 21, "Statistical Graphics Using ODS") and define a new graph template. The following statements defines the template **twobytwo**:

```
proc template;
   define statgraph twobytwo;
      begingraph;
         layout lattice / rows=2 columns=2;
            layout overlay / yaxisopts=(display=none)
                             xaxisopts=(label="mean");
               layout gridded / columns=2 border=false
                                autoalign=(topleft topright);
                  entry halign=right "p-value =";
                  entry halign=left eval(strip(put(&pmean, 12.2)));
               endlayout;
```

```
               histogram mean / binaxis=false;
               lineparm x=&mean y=0 slope=. /
                       lineattrs=(color=red thickness=5);
            endlayout;
            layout overlay / yaxisopts=(display=none)
                             xaxisopts=(label="sd");
               layout gridded / columns=2 border=false
                                autoalign=(topleft topright);
                  entry halign=right "p-value =";
                  entry halign=left eval(strip(put(&psd, 12.2)));
               endlayout;
               histogram sd / binaxis=false;
               lineparm x=&sd y=0 slope=. /
                       lineattrs=(color=red thickness=5);
            endlayout;
            layout overlay / yaxisopts=(display=none)
                             xaxisopts=(label="max");
               layout gridded / columns=2 border=false
                                autoalign=(topleft topright);
                  entry halign=right "p-value =";
                  entry halign=left eval(strip(put(&pmax, 12.2)));
               endlayout;
               histogram max / binaxis=false;
               lineparm x=&max y=0 slope=. /
                       lineattrs=(color=red thickness=5);
            endlayout;
            layout overlay / yaxisopts=(display=none)
                             xaxisopts=(label="min");
               layout gridded / columns=2 border=false
                                autoalign=(topleft topright);
                  entry halign=right "p-value =";
                  entry halign=left eval(strip(put(&pmin, 12.2)));
               endlayout;
               histogram min / binaxis=false;
               lineparm x=&min y=0 slope=. /
                       lineattrs=(color=red thickness=5);
            endlayout;
         endlayout;
      endgraph;
   end;
run;
```

You call PROC SGRENDER (see the SGRENDER procedure in the *SAS/GRAPH: Statistical Graphics Procedures Guide*) to create the graph, which is shown in Figure 52.12. There are no extreme *p*-values observed; this supports the notion that the predicted results are similar to the actual observations and that the model fits the data.

```
proc sgrender data=pred template=twobytwo;
run;
```

Figure 52.12 Posterior Predictive Distribution Check for the SAT example
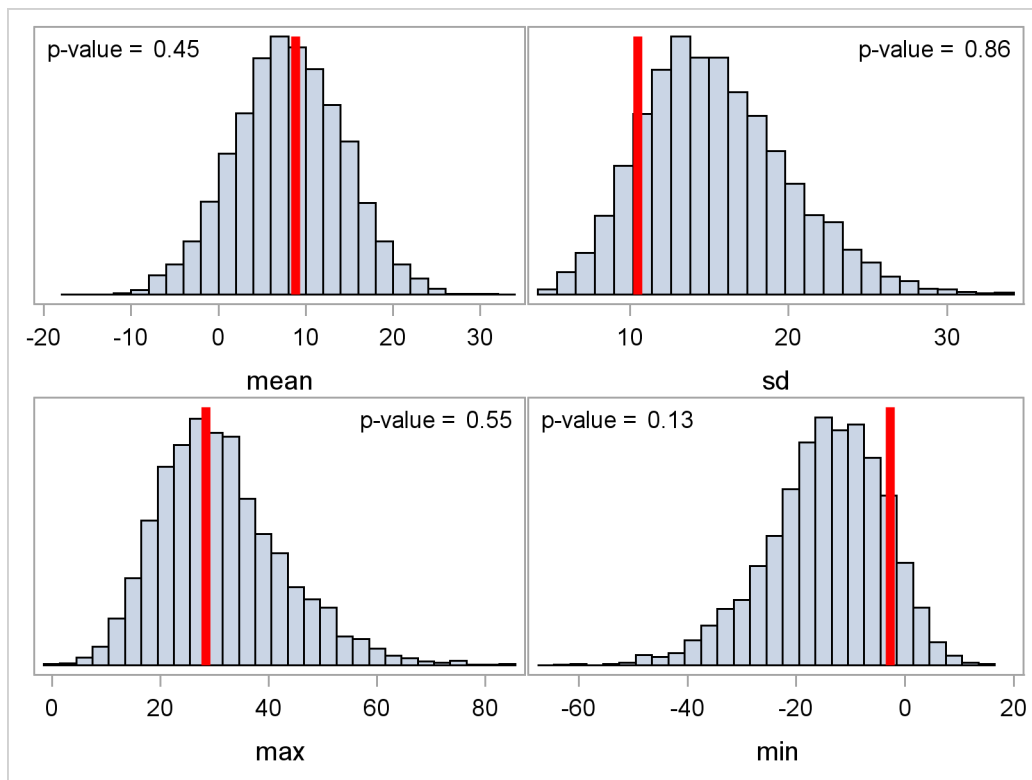
Handling of Missing Data

By default, PROC MCMC discards all observations that have missing values before carrying out the posterior sampling. This corresponds to the option MISSING=CC, where CC stands for complete cases. PROC MCMC does not automatically augment missing data. However, you can choose to model the missing values by using MISSING=AC. Given this option, PROC MCMC does not discard any missing values. It is up to you to specify how the missing values are handled in the program. You can choose to model the missing values as parameters (a fully Bayesian approach) or assign specific values to them (multiple imputation). In general, however, the handling of missing values largely depends on the assumptions you have about the missing mechanism, which is beyond the score of this chapter.

Floating Point Errors and Overflows

When performing a Markov chain Monte Carlo simulation, you must calculate a proposed jump and an objective function (usually a posterior density). These calculations might lead to arithmetic exceptions and overflows. A typical cause of these problems is parameters with widely varying scales. If the posterior variances of your parameters vary by more than a few orders of magni-

tude, the numerical stability of the optimization problem can be severely reduced and can result in computational difficulties. A simple remedy is to rescale all the parameters so that their posterior variances are all approximately equal. Changing the SCALE= option might help if the scale of your parameters is much different than one. Another source of numerical instability is highly correlated parameters. Often a model can be reparameterized to reduce the posterior correlations between parameters.

If parameter rescaling does not help, consider the following actions:

- provide different initial values or try a different seed value

- use boundary constraints to avoid the region where overflows might happen

- change the algorithm (specified in programming statements) that computes the objective function

Problems Evaluating Code for Objective Function

The initial values must define a point for which the programming statements can be evaluated. However, during simulation, the algorithm might iterate to a point where the objective function cannot be evaluated. If you program your own likelihood, priors, and hyperpriors by using SAS statements and the GENERAL function in the MODEL, PRIOR, AND HYPERPRIOR statements, you can specify that an expression cannot be evaluated by setting the value you pass back through the GENERAL function to missing. This tells the PROC MCMC that the proposed set of parameters is invalid, and the proposal will not be accepted. If you use the shorthand notation that the MODEL, PRIOR, AND HYPERPRIOR statements provide, this error checking is done for you automatically.

Long Run Times

PROC MCMC can take a long time to run for problems with complex models, many parameters, or large input data sets. Although the techniques used by PROC MCMC are some of the best available, they are not guaranteed to converge or proceed quickly for all problems. Ill-posed or misspecified models can cause the algorithms to use more extensive calculations designed to achieve convergence, and this can result in longer run times. You should make sure that your model is specified correctly, that your parameters are scaled to the same order of magnitude, and that your data reasonably match the model that you are specifying.

To speed general computations, you should check over your programming statements to minimize the number of unnecessary operations. For example, you can use the proportional kernel in the priors or the likelihood and not add constants in the densities. You can also use the BEGINCNST and ENDCNST to reduce unnecessary computations on constants, and the BEGINPRIOR and ENDPRIOR statements to reduce observation-level calculations.

Reducing the number of blocks (the number of the PARMS statements) can speed up the sampling process. A single-block program is approximately three times faster than a three-block program for the same number of iterations. On the other hand, you do not want to put too many parameters in a single block, because blocks with large size tend not to produce well-mixed Markov chains.

Slow or No Convergence

There are a number of things to consider if the simulator is slow or fails to converge:

- Change the number of Monte Carlo iterations (NMC=), or the number of burn-in iterations (NBI=), or both. Perhaps the chain just needs to run a little longer. Note that after the simulation, you can always use the DATA step or the FIRSTOBS data set option to throw away initial observations where the algorithm has not yet burned in, so it is not always necessary to set NBI= to a large value.

- Increase the number of tuning. The proposal tuning can often work better in large models (models that have more parameters) with larger values of NTU=. The idea of tuning is to find a proposal distribution that is a good approximation to the posterior distribution. Sometimes 500 iterations per tuning phase (the default) is not sufficient to find a good approximating covariance.

- Change the initial values to more feasible starting values. Sometimes the proposal tuning starts badly if the initial values are too far away from the main mass of the posterior density, and it might not be able to recover.

- Use the PROPCOV= option to start the Markov chain at better starting values. With the PROPCOV=QUANEW option, PROC MCMC optimizes the object function and uses the posterior mode as the starting value of the Markov chain. In addition, a quadrature approximation to the posterior mode is used as the proposal covariance matrix. This option works well in many cases and can improve the mixing of the chain and shorten the tuning and burn-in time.

- Change the blocking by using the PARMS statements. Sometimes poor mixing and slow convergence can be attributed to highly correlated parameters being in different parameter blocks.

- Modify the target acceptance rate. A target acceptance rate of about 25% works well for many multi-parameter problems, but if the mixing is slow, a lower target acceptance rate might be better.

- Change the initial scaling or the TUNEWT= option to possibly help the proposal tuning.

- Consider using a different proposal distribution. If from a trace plot you see that a chain traverses to the tail area and sometimes takes quite a few simulations before it comes back, you can consider using a t-proposal distribution. You can do this by either using the PROC option PROPDIST=T or using a PARMS statement option T.

- Transform parameters and sample on a different scale. For example, if a parameter has a gamma distribution, sample on the logarithm scale instead. A parameter a that has a gamma distribution is equivalent to $\log(a)$ that has an egamma distribution, with the same distribution specification. For example, the following two formulations are equivalent:

```
parm a;
prior a ~ gamma(shape = 0.001, iscale = 0.001);
```

and

```
parm la;
prior la ~ egamma(shape = 0.001, iscale = 0.001);
a = exp(la);
```

See "Example 52.4:   Nonlinear Poisson Regression Models" on page 3599 and "Example 52.12: Using a Transformation to Improve Mixing" on page 3677. You can also use the logit transformation on parameters that have uniform$(0, 1)$ priors. This prior is often used on probability parameters. The logit transformation is as follows: $q = \log(\frac{p}{1-p})$. The distribution on q is the Jacobian of the transformation: $\exp(-q)(1+\exp(-q))^{-2}$. Again, the following two formulations are equivalent:

```
parm p;
prior p ~ uniform(0, 1);
```

and

```
parm q;
lp = -q - 2 * log(1 + exp(-q));
prior q ~ general(lp);
p = 1/(1+exp(-q));
```

Precision of Solution

In some applications, PROC MCMC might produce parameter values that are not precise enough. Usually, this means that there were not enough iterations in the simulation. At best, the precision of MCMC estimates increases with the square of the simulation sample size. Autocorrelation in the parameter values deflate the precision of the estimates. For more information about autocorrelations in Markov chains, see the section "Autocorrelations" on page 169.

Handling Error Messages

PROC MCMC does not have a debugger. This section covers a few ways to debug and resolve error messages.

Using the PUT Statement

Adding the **put** statement often helps to find errors in a program. The following program produces an error:

```
data a;
run;

proc mcmc data=a seed=1;
   parms sigma lt w;
```

```
    beginprior;
    prior sigma ~ unif(0.001,100);
    s2 = sigma*sigma;
    prior lt ~ gamma(shape=1, iscale=0.001);
    t = exp(lt);
    c = t/s2;
    d = 1/(s2);
    prior w ~ gamma(shape=c, iscale=d);
    endprior;

    model general(0);
run;
```

```
ERROR: PROC MCMC is unable to generate an initial value for the
        parameter w.  The first parameter in the prior distribution is
        missing.
```

To find out why the shape parameter c is missing, you can add the put statement and examine all the calculations that lead up to the assignment of c:

```
proc mcmc data=a seed=1;
    parms sigma lt w;

    beginprior;
    prior sigma ~ unif(0.001,100);
    s2 = sigma*sigma;
    prior lt ~ gamma(shape=1, iscale=0.001);
    t = exp(lt);
    c = t/s2;
    d = 1/(s2);
    put c= t= s2= lt=; /* display the values of these symbols. */
    prior w ~ gamma(shape=c, iscale=d);
    endprior;

    model general(0);
run;
```

In the log file, you see the following:

```
c=. t=. s2=. lt=.
c=. t=. s2=2500.0500003 lt=1000
c=. t=. s2=2500.0500003 lt=1000
ERROR: PROC MCMC is unable to generate an initial value for the parameter w.
        The first parameter in the prior distribution is missing.
```

You can ignore the first few lines. They are the results of initial set up by PROC MCMC. The last line is important. The variable c is missing because t is the exponential of a very large number, 1000, in lt. The value 1000 is assigned to lt by PROC MCMC because none was given. The gamma prior with shape of 1 and inverse scale of 0.001 has mode 0 (see "Standard Distributions" on page 3529 for more details). PROC MCMC avoids starting the Markov chain at the boundary

of the support of the distribution, and it uses the mean value here instead. The mean of the gamma prior is 1000, hence the problem. You can change how the initial value is generated by using the PROC statement INIT=RANDOM. Do not forget to take out the put statement once you identify the problem. Otherwise, you will see a voluminous output in the log file.

Using the HYPER Statement

You can use the HYPER statement to narrow down possible errors in the prior distribution specification. With multiple PRIOR statements in a program, you might see the following error message if one of the prior distributions is not specified correctly:

```
ERROR: The initial prior parameter specifications must yield log
       of positive prior density values.
```

This message is displayed when PROC MCMC detects an error in the prior distribution calculation but cannot pinpoint the specific parameter at fault. It is frequently, although not necessarily, associated with parameters that have GENERAL or DGENERAL distributions. If you have a complicated model with many PRIOR statements, finding the parameter at fault can be time consuming. One way is to change a subset of the PRIOR statements to HYPER statements. The two statements are treated the same in PROC MCMC and the simulation is not affected, but you get a different message if the hyperprior distributions are calculated incorrectly:

```
ERROR: The initial hyperprior parameter specifications must yield
       log of positive hyperprior density values.
```

This message can help you identify more easily which distributions are producing the error, and you can then use the `put` statement to further investigate.

Computational Resources

It is not possible to estimate how long it will take for a general Markov chain to converge to its stationary distribution. It takes a skilled and thoughtful analysis of the chain to decide if it has converged to the target distribution and if the chain is mixing rapidly enough. It is easier, however, to estimate how long a particular simulation might take. The running time of a program is roughly linear to the following factors: the number of samples in the input data set (nsamples), the number of simulations (nsim), the number of blocks in the program (nblocks), and the speed of your computer. For an analysis that uses a data set of size nsamples, a simulation length of nsim, and a block design of nblocks, PROC MCMC evaluates the log-likelihood function the following number of times, excluding the tuning phase:

$$\text{nsamples} \times \text{nsim} \times \text{nblocks}$$

The faster your computer evaluates a single log-likelihood function, the faster this program runs. Suppose that you have nsamples equal to 200, nsim equal to 55,000, and nblocks equal to 3. PROC

MCMC evaluates the log-likelihood function roughly a total number of 3.3×10^7 times. If your computer can evaluate the log likelihood, for one observation, 10^6 times per second, this program will take approximately a half a minute to run. If you want to increase the number of simulations five-fold, the run time will approximately increase five-fold as well.

Of course, larger problems take longer than shorter ones, and if your model is amenable to frequentist treatment, then one of the other SAS procedures might be more suitable. With "regular" likelihoods and a lot of data, the results of standard frequentist analysis are often asymptotically equivalent to a Bayesian approach. If PROC MCMC requires too much CPU time, then perhaps another tool in SAS/STAT would be suitable.

Displayed Output

This section describes the displayed output from PROC MCMC. For a quick reference of all ODS table names, see the section "ODS Table Names" on page 3569. ODS tables are arranged under four groups, listed in the following sections: "Sampling Related ODS Tables" on page 3565, "Posterior Statistics Related ODS Tables" on page 3566, "Convergence Diagnostics Related ODS Tables" on page 3567, and "Optimization Related ODS Tables" on page 3568.

Sampling Related ODS Tables

Burn-In History

The "Burn-In History" table (ODS table name BurnInHistory) shows the scales and acceptance rates for each parameter block in the burn-in phase. The table is displayed by default.

Number of Observation Table

The "NObs" table (ODS table name NOBS) shows the number of observations that is in the data set and the number of observations that is used in the analysis. By default, observations with missing values are not used (see the section "Handling of Missing Data" on page 3559 for more details). This table is displayed by default.

Parameters

The "Parameters" table (ODS table name Parameters) shows the name of each parameter, the block number of each parameter, the sampling method used for the block, the initial values, and the prior or hyperprior distributions. This table is displayed by default.

Parameters Initial Value Table

The "Parameters Initial" table (ODS table name ParametersInit) shows the value of each parameter after the tuning phase. This table is not displayed by default and can be requested by specifying the option INIT=PINIT.

Posterior Samples

The "Posterior Samples" table (ODS table name PosteriorSample) stores posterior draws of all parameters. It is not printed by PROC MCMC. You can create an ODS output data set of the chain by specifying the following:

```
ODS OUTPUT PosteriorSample = SAS-data-set;
```

Sampling History

The "Sampling History" table (ODS table name SamplingHistory) shows the scales and acceptance rates for each parameter block in the main sampling phase. The table is displayed by default.

Tuning Covariance

The "Tuning Covariance" table (ODS table name TuneCov) shows the proposal covariance matrices for each parameter block after the tuning phase. The table is not displayed by default and can be requested by specifying the option INIT=PINIT. For more details about proposal tuning, see the section "Tuning the Proposal Distribution" on page 3524.

Tuning History

The "Tuning History" table (ODS table name TuningHistory) shows the number of tuning phases used in establishing the proposal distribution. The table also displays the scales and acceptance rates for each parameter block at each of the tuning phases. For more information about the self-adapting proposal tuning algorithm used by PROC MCMC, see the section "Tuning the Proposal Distribution" on page 3524. The table is displayed by default.

Tuning Probability Vector

The "Tuning Probability" table (ODS table name TuneP) shows the proposal probability vector for each discrete parameter block (when the option DISCRETE=GEO is specified and the geometric proposal distribution is used for discrete parameters) after the tuning phase. The table is not displayed by default and can be requested by specifying the option INIT=PINIT. For more information about proposal tuning, see the section "Tuning the Proposal Distribution" on page 3524.

Posterior Statistics Related ODS Tables

PROC MCMC calculates some essential posterior statistics and outputs them to a number of ODS tables that you can request and save individually. For details of the calculations, see the section "Summary Statistics" on page 170.

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 displayed by default.

Correlation Matrix

The "Posterior Correlation Matrix" table (ODS table name Corr) contains the posterior correlation of model parameters. The table is not displayed by default and can be requested by specifying the option STATS=CORR.

Covariance Matrix

The "Posterior Covariance Matrix" table (ODS table name Cov) contains the posterior covariance of model parameters. The table is not displayed by default and can be requested by specifying the option STATISTICS=COV.

Deviance Information Criterion

The "Deviance Information Criterion" table (ODS table name DIC) contains the DIC of the model. The table is not displayed by default and can be requested by specifying the option DIC. For details of the calculations, see the section "Deviance Information Criterion (DIC)" on page 172.

Interval Statistics

The "Posterior Intervals" table (ODS table name PostIntervals) contains two 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 displayed by default.

Convergence Diagnostics Related ODS Tables

PROC MCMC has 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 details in calculation, see the section "Statistical Diagnostic Tests" on page 160.

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 MCMC 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. See the section "Effective Sample Size" on page 169 for more details. The 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. See the section "Standard Error of the Mean Estimate" on page 170 for more details. The table is displayed by default.

Geweke Diagnostics

The "Geweke Diagnostics" table (ODS table name Geweke) lists the result of the Geweke diagnostic test. See the section "Geweke Diagnostics" on page 163 for more details. The table is displayed by default.

Heidelberger-Welch Diagnostics

The "Heidelberger-Welch Diagnostics" table (ODS table name Heidelberger) lists the result of the Heidelberger-Welch diagnostic test. The test is consisted of two parts: a stationary test and a half-width test. See the section "Heidelberger and Welch Diagnostics" on page 165 for more details. The table is not displayed by default and can be requested by specifying DIAGNOSTICS = HEIDEL.

Raftery-Lewis Diagnostics

The "Raftery-Lewis Diagnostics" table (ODS table name Raftery) lists the result of the Raftery-Lewis diagnostic test. See the section "Raftery and Lewis Diagnostics" on page 166 for more details. The table is not displayed by default and can be requested by specifying DIAGNOSTICS = RAFTERY.

Optimization Related ODS Tables

PROC MCMC can perform optimization on the joint posterior distribution. This is requested by the PROPCOV= option. The most commonly used optimization method is the quasi-Newton method: PROPCOV=QUANEW(ITPRINT). The ITPRINT option displays the ODS tables, listed as follows:

Input Options

The "Input Options" table (ODS table name InputOptions) lists optimization options used in the procedure.

Optimization Start

The "Optimization Start" table (ODS table name ProblemDescription) shows the initial state of the optimization.

Iteration History

The "Iteration History" table (ODS table name IterHist) shows iteration history of the optimization.

Optimization Results

The "Optimization Results" table (ODS table name IterStop) shows the results of the optimization, includes information about the number of function calls, and the optimized objective function, which is the joint log posterior density.

Convergence Status

The "Convergence Status" table (ODS table name ConvergenceStatus) shows whether the convergence criterion is satisfied.

Parameters Value After Optimization Table

The "Parameter Values After Optimization" table (ODS table name OptiEstimates) lists the parameter values that maximize the joint log posterior. These are the maximum a posteriori point estimates, and they are used to start the Markov chain.

Covariance Matrix After Optimization Table

The "Proposal Covariance" table (ODS table name OptiCov) lists covariance matrices for each block parameter by using quadrature approximation at the posterior mode. These covariance matrices are used in the proposal distribution.

ODS Table Names

PROC MCMC assigns a name to each table it creates. You can use these names to reference the table when using the Output Delivery System (ODS) to select tables and create output data sets. These names are listed in the following table. For more information about ODS, see Chapter 21, "Statistical Graphics Using ODS."

Table 52.35 ODS Tables Produced in PROC MCMC

ODS Table Name	Description	Statement or Option
AutoCorr	autocorrelation statistics for each parameter	default
PostSummaries	basic statistics for each parameter, including sample size, mean, standard deviation, and percentiles	default
ConvergenceStatus	optimization convergence status	PROPCOV=method(ITPRINT)
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
MCSE	Monte Carlo standard error for each parameter	default
Geweke	Geweke diagnostics for each parameter	default
Heidelberger	Heidelberger-Welch diagnostics for each parameter	DIAGNOSTICS=HEIDEL
InputOptions	optimization input table	PROPCOV=method(ITPRINT)
PostIntervals	equal-tail and HPD intervals for each parameter	default
IterHist	optimization iteration history	PROPCOV=method(ITPRINT)
IterStop	optimization results table	PROPCOV=method(ITPRINT)
NObs	number of observations	default
OptiEstimates	parameter values after either optimization	PROPCOV=method(ITPRINT)
OptiCov	covariance used in proposal distribution after optimization	PROPCOV=method(ITPRINT)
Parameters	summary of the PARMS, BLOCKING, PRIOR, sampling method, and initial value specification	default
ParametersInit	parameter values after the tuning phase	INIT=PINIT
PosteriorSample	posterior samples for each parameter	(for ODS output data set only)
ProblemDescription	optimization table	PROPCOV=method(ITPRINT)
Raftery	Raftery-Lewis diagnostics for each parameter	DIAGNOSTICS=RAFTERY
SamplingHistory	history of burn-in and main phase sampling	default

Table 52.35 (continued)

ODS Table Name	Description	Statement or Option
TuneCov	proposal covariance matrix (for continuous parameters) after the tuning phase	INIT=PINIT
TuneP	proposal probability vector (for discrete parameters) after the tuning phase	INIT=PINIT and DISCRETE=GEO
TuningHistory	history of proposal distribution tuning	default

ODS Graphics

To request graphics with PROC MCMC, you must first enable ODS Graphics by specifying the **ods graphics on** statement. See Chapter 21, "Statistical Graphics Using ODS," for more information. You can reference every graph produced through ODS Graphics with a name. The names of the graphs that PROC MCMC generates are listed in Table 52.36.

Table 52.36 ODS Graphics Produced by PROC MCMC

ODS Graph Name	Plot Description	Statement & 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: MCMC Procedure

Example 52.1: Simulating Samples From a Known Density

This example illustrates how you can obtain random samples from a known function. The target distributions are the normal distribution and a mixture of the normal distributions. You do not need any input data set to generate samples from a known density. You can set the likelihood function to a constant. The posterior distribution becomes identical to the prior distributions that you specify.

Sampling from a Normal Density

With a constant likelihood, there is no need to input a response variable since no data are relevant to a flat likelihood. However, PROC MCMC requires an input data set, so you can use an empty data set as the input data set. The following statements generate 10000 samples from a standard normal distribution:

```
data x;
run;

ods graphics on;
proc mcmc data=x outpost=simout seed=23 nmc=10000 maxtune=0
          nbi=0 statistics=(summary interval) diagnostics=none;
   ods exclude nobs parameters samplinghistory;
   parm alpha 0;
   prior alpha ~ normal(0, sd=1);
   model general(0);
run;
ods graphics off;
```

The `ods graphics on` statement requests ODS Graphics. The PROC MCMC statement specifies the input and output data sets, a random number seed, and the size of the simulation sample. There is no need for tuning (MAXTUNE=0) because the default scale and the proposal variance are optimal for a standard normal target distribution. For the same reason, no burn-in is needed (NBI=0). The STATISTICS= option is used to display only the summary and interval statistics. The `ods exclude` statement excludes the display of the NObs, Parameters and SamplingHistory tables. The summary statistics (Output 52.1.1) are what you would expect from a standard normal distribution.

Output 52.1.1 MCMC Summary and Interval Statistics from a Normal Target Distribution

```
             An Example for Posterior Predictive Distribution

                          The MCMC Procedure

                          Posterior Summaries

                                  Standard              Percentiles
    Parameter          N       Mean   Deviation       25%         50%         75%

    alpha          10000    -0.0392      1.0194    -0.7198     -0.0403      0.6351

                          Posterior Intervals

       Parameter    Alpha     Equal-Tail Interval        HPD Interval

       alpha        0.050       -2.0746      1.9594      -2.2197      1.7869
```

The trace plot (Output 52.1.2) shows good mixing of the Markov chain, and there is no significant autocorrelation in the lag plot.

Output 52.1.2 Diagnostics Plots for α

You can also overlay the estimated kernel density with the true density to get a visual comparison, as displayed in Output 52.1.3.

To create Output 52.1.3, you first use PROC KDE (see Chapter 45, "The KDE Procedure") to obtain a kernel density estimate of the posterior density on alpha, and then you evaluate a grid of alpha values by using PROC KDE output data set sample on a normal density. The following statements evaluate kernel density and compute corresponding normal density.

```
proc kde data=simout;
   ods exclude inputs controls;
   univar alpha /out=sample;
run;

data den;
   set sample;
   alpha = value;
   true  = pdf('normal', alpha, 0, 1);
   keep alpha density true;
run;
```

Finally, you plot the two curves on top of each other by using PROC SGPLOT (see Chapter 21, "Statistical Graphics Using ODS"); the resulting figure is in Output 52.1.3. You can see that the kernel estimate and the true density are very similar to one another. The following statements produce Output 52.1.3:

```
proc sgplot data=den;
   yaxis label="Density";
   series y=density x=alpha / legendlabel = "MCMC Kernel";
   series y=true x=alpha / legendlabel = "True Density";
   discretelegend;
run;
```

Output 52.1.3 Estimated Density versus the True Density

Sampling from a Mixture of Normal Densities

Suppose that you are interested in generating samples from a three-component mixture of normal distributions, with the density specified as follows:

$$p(\alpha) = 0.3 \cdot \phi(-3, \sigma = 2) + 0.4 \cdot \phi(2, \sigma = 1) + 0.3 \cdot \phi(10, \sigma = 4)$$

The following statements generate random samples from this mixture density:

```
data x;
run;

ods graphics on;
proc mcmc data=x outpost=simout seed=1234 nmc=30000;
   ods select TADpanel;
   parm alpha 0.3;
   lp = logpdf('normalmix', alpha, 3, 0.3, 0.4, 0.3, -3, 2, 10, 2, 1, 4);
   prior alpha  ~ general(lp);
   model general(0);
run;
ods graphics off;
```

The `ods select` statement displays the diagnostic plots. All other tables, such as the NObs tables, are excluded. The PROC MCMC statement uses the input data set x, saves output to the simout data set, sets a random number seed, and simulates 30,000 samples.

The lp assignment statement evaluates the log density of alpha at the mixture density, using the SAS function LOGPDF. The number 3 after alpha in the LOGPDF function indicates that the density is a three-component normal mixture. The following three numbers, 0.3, 0.4, and 0.3, are the weights in the mixture; -3, 2, and 10 are the means; 2, 1, and 4 are the standard deviations. The PRIOR statement assigns this log density function to alpha as its prior. Note that the GENERAL function interprets the density on the log scale, and not the original scale. Hence, you must use the LOGPDF function, not the PDF function. Output 52.1.4 displays the results. The kernel density clearly shows three modes.

Output 52.1.4 Plots of Posterior Samples from a Mixture Normal Distribution

Using the following set of statements similar to the previous example, you can overlay the estimated kernel density with the true density. The comparison is shown in Output 52.1.5.

```
proc kde data=simout;
   ods exclude inputs controls;
   univar alpha /out=sample;
run;

data den;
   set sample;
   alpha = value;
   true  = pdf('normalmix', alpha, 3, 0.3, 0.4, 0.3, -3, 2, 10, 2, 1, 4);
   keep alpha density true;
run;

proc sgplot data=den;
   yaxis label="Density";
   series y=density x=alpha / legendlabel = "MCMC Kernel";
   series y=true x=alpha / legendlabel = "True Density";
   discretelegend;
run;
```

Output 52.1.5 Estimated Density versus the True Density

Example 52.2: Box-Cox Transformation

Box-Cox transformations (Box and Cox 1964) are often used to find a power transformation of a dependent variable to ensure the normality assumption in a linear regression model. This example illustrates how you can use PROC MCMC to estimate a Box-Cox transformation for a linear regression model. Two different priors on the transformation parameter λ are considered: a continuous prior and a discrete prior. You can estimate the probability of λ being 0 with a discrete prior but not with a continuous prior. The IF-ELSE statements are demonstrated in the example.

Using a Continuous Prior on λ

The following statements create a SAS data set with measurements of y (the response variable) and x (a single dependent variable):

```
title 'Box-Cox Transformation, with a Continuous Prior on Lambda';
data boxcox;
   input y x @@;
   datalines;
10.0   3.0   72.6   8.3   59.7   8.1   20.1   4.8   90.1   9.8    1.1   0.9
78.2   8.5   87.4   9.0    9.5   3.4    0.1   1.4    0.1   1.1   42.5   5.1
```

```
   ... more lines ...

   2.6   1.8   58.6   7.9   81.2   8.1   37.2   6.9
   ;
```

The Box-Cox transformation of y takes on the form of:

$$y(\lambda) = \begin{cases} \frac{y^{\lambda}-1}{\lambda} & \text{if } \lambda \neq 0; \\ \log(y) & \text{if } \lambda = 0. \end{cases}$$

The transformed response $y(\lambda)$ is assumed to be normally distributed:

$$y_i(\lambda) \sim \text{normal}(\beta_0 + \beta_1 x_i, \sigma^2)$$

The likelihood with respect to the original response y_i is as follows:

$$f(y_i|\lambda, \beta, \sigma^2, x_i) \propto \phi(y_i|\beta_0 + \beta_1 x_i, \sigma^2) \cdot J(\lambda, y_i)$$

where $J(\lambda, y_i)$ is the Jacobian:

$$J(\lambda, y) = \begin{cases} y_i^{\lambda-1} & \text{if } \lambda \neq 0; \\ 1/y_i & \text{if } \lambda = 0. \end{cases}$$

And on the log-scale, the Jacobian becomes:

$$\log(J(\lambda, y)) = \begin{cases} (\lambda - 1) \cdot \log(y_i) & \text{if } \lambda \neq 0; \\ -\log(y_i) & \text{if } \lambda = 0. \end{cases}$$

There are four model parameters: $\lambda, \beta = \{\beta_0, \beta_1\}$, and σ^2. You can considering using a flat prior on β and a gamma prior on σ^2.

To consider only power transformations ($\lambda \neq 0$), you can use a continuous prior (for example, a uniform prior from -2 to 2) on λ. One issue with using a continuous prior is that you cannot estimate the probability of $\lambda = 0$. To do so, you need to consider a discrete prior that places positive probability mass on the point 0. See "Modeling $\lambda = 0$" on page 3582.

The following statements fit a Box-Cox transformation model:

```
ods graphics on;
proc mcmc data=boxcox nmc=50000 thin=10 propcov=quanew seed=12567
          monitor=(lda);
   ods select PostSummaries PostIntervals TADpanel;

   parms beta0 0  beta1 0  lda 1 s2 1;

   beginprior;
   prior beta: ~ general(0);
   prior s2 ~ gamma(shape=3, scale=2);
   prior lda ~ unif(-2,2);
   sd = sqrt(s2);
   endprior;

   ys = (y**lda-1)/lda;
   mu = beta0+beta1*x;
   ll = (lda-1)*log(y)+lpdfnorm(ys, mu, sd);
   model general(ll);
run;
```

The PROPCOV option initializes the Markov chain at the posterior mode and uses the estimated inverse Hessian matrix as the initial proposal covariance matrix. The MONITOR= option selects λ as the variable to report. The **ods select** statement displays the summary statistics table, the interval statistics table, and the diagnostic plots.

The PARMS statement puts all four parameters, β_0, β_1, λ, and σ^2, in a single block and assigns initial values to each of them. Three PRIOR statements specify previously stated prior distributions for these parameters. The assignment to sd transforms a variance to a standard deviation. It is better to place the transformation inside the BEGINPRIOR and ENDPRIOR statements to save computational time.

The assignment to the symbol ys evaluates the Box-Cox transformation of y, where mu is the regression mean and ll is the log likelihood of the transformed variable ys. Note that the log of the Jacobian term is included in the calculation of ll.

Summary statistics and interval statistics for lda are listed in Output 52.2.1.

Output 52.2.1 Box-Cox Transformation

```
          Box-Cox Transformation, with a Continuous Prior on Lambda

                            The MCMC Procedure

                            Posterior Summaries

                                Standard                  Percentiles
   Parameter         N       Mean    Deviation        25%        50%        75%

   lda            5000     0.4702       0.0284     0.4515     0.4703     0.4884
```

Output 52.2.1 *continued*

```
                          Posterior Intervals

        Parameter     Alpha      Equal-Tail Interval          HPD Interval

          lda         0.050        0.4162      0.5269       0.4197      0.5298
```

The posterior mean of λ is 0.47, with a 95% equal-tail interval of $[0.42, 0.53]$ and a similar HPD interval. The prefered power transformation would be 0.5 (rounding λ up to the square root transformation).

Output 52.2.2 shows diagnostics plots for lda. The chain appears to converge, and you can proceed to make inferences. The density plot shows that the posterior density is relatively symmetric around its mean estimate.

Output 52.2.2 Diagnostic Plots for λ

To verify the results, you can use PROC TRANSREG (see Chapter 90, "The TRANSREG Procedure") to find the estimate of λ.

```
proc transreg data=boxcox details pbo;
   ods output boxcox = bc;
   model boxcox(y / convenient lambda=-2 to 2 by 0.01) = identity(x);
   output out=trans;
run;
ods graphics off;
```

Output from PROC TRANSREG is shown in Output 52.2.5 and Output 52.2.4. PROC TRAN-SREG produces a similar point estimate of $\lambda = 0.46$, and the 95% confidence interval is shown in Output 52.2.5.

Output 52.2.3 Box-Cox Transformation Using PROC TRANSREG

Output 52.2.4 Estimates Reported by PROC TRANSREG

```
        Box-Cox Transformation, with a Continuous Prior on Lambda

                        The TRANSREG Procedure

                   Model Statement Specification Details

   Type  DF Variable     Description       Value

   Dep    1 BoxCox(y)    Lambda Used       0.5
                         Lambda            0.46
                         Log Likelihood    -167.0
                         Conv. Lambda      0.5
                         Conv. Lambda LL  -168.3
                         CI Limit          -169.0
                         Alpha             0.05
                         Options           Convenient Lambda Used


   Ind    1 Identity(x)  DF                1
```

The ODS data set bc contains the 95% confidence interval estimates produced by PROC TRAN-SREG. This ODS table is rather large, and you want to see only the relevant portion. The following statements generate the part of the table that is important and display Output 52.2.5:

```
proc print noobs label data=bc(drop=rmse);
   title2 'Confidence Interval';
   where ci ne ' ' or abs(lambda - round(lambda, 0.5)) < 1e-6;
   label convenient = '00'x ci = '00'x;
run;
```

The estimated 90% confidence interval is [0.41, 0.51], which is very close to the reported Bayesian credible intervals. The resemblance of the intervals is probably due to the noninformative prior that you used in this analysis.

Output 52.2.5 Estimated Confidence Interval on λ

```
        Box-Cox Transformation, with a Continuous Prior on Lambda
                         Confidence Interval


                                                    Log
         Dependent       Lambda        R-Square    Likelihood

         BoxCox(y)        -2.00           0.14      -1030.56
         BoxCox(y)        -1.50           0.17       -810.50
         BoxCox(y)        -1.00           0.22       -602.53
         BoxCox(y)        -0.50           0.39       -415.56
         BoxCox(y)         0.00           0.78       -257.92
         BoxCox(y)         0.41           0.95       -168.40    *
         BoxCox(y)         0.42           0.95       -167.86    *
         BoxCox(y)         0.43           0.95       -167.46    *
         BoxCox(y)         0.44           0.95       -167.19    *
         BoxCox(y)         0.45           0.95       -167.05    *
         BoxCox(y)         0.46           0.95       -167.04    <
         BoxCox(y)         0.47           0.95       -167.16    *
         BoxCox(y)         0.48           0.95       -167.41    *
         BoxCox(y)         0.49           0.95       -167.79    *
         BoxCox(y)         0.50     +     0.95       -168.28    *
         BoxCox(y)         0.51           0.95       -168.89    *
         BoxCox(y)         1.00           0.89       -253.09
         BoxCox(y)         1.50           0.79       -345.35
         BoxCox(y)         2.00           0.70       -435.01
```

Modeling $\lambda = 0$

With a continuous prior on λ, you can get only a continuous posterior distribution, and this makes the probability of $\Pr(\lambda = 0|\text{data})$ equal to 0 by definition. To consider $\lambda = 0$ as a viable solution to the Box-Cox transformation, you need to use a discrete prior that places some probability mass on the point 0 and allows for a meaningful posterior estimate of $\Pr(\lambda = 0|\text{data})$.

This example uses a simulation study where the data are generated from an exponential likelihood. The simulation implies that the correct transformation should be the logarithm and λ should be 0.

Consider the following exponential model:

$$y = \exp(x + \epsilon),$$

where $\epsilon \sim \text{normal}(0, 1)$. The transformed data can be fitted with a linear model:

$$\log(y) = x + \epsilon$$

The following statements generate a SAS data set with a gridded x and corresponding y:

```
title 'Box-Cox Transformation, Modeling Lambda = 0';
data boxcox;
   do x = 1 to 8 by 0.025;
      ly = x + normal(7);
      y = exp(ly);
      output;
   end;
run;
```

The log-likelihood function, after taking the Jacobian into consideration, is as follows:

$$\log p(y_i | \lambda, x_i) = \begin{cases} (\lambda - 1)\log(y_i) - \frac{1}{2}\left(\log\sigma^2 + \frac{\left((y_i^\lambda - 1)/\lambda - x_i\right)^2}{\sigma^2}\right) + C_1 & \text{if } \lambda \neq 0; \\ -\log(y_i) - \frac{1}{2}\left(\log\sigma^2 + \frac{(\log(y_i) - x_i)^2}{\sigma^2}\right) + C_2 & \text{if } \lambda = 0. \end{cases}$$

where C_1 and C_2 are two constants.

You can use the function DGENERAL to place a discrete prior on λ. The function is similar to the function GENERAL, except that it indicates a discrete distribution. For example, you can specify a discrete uniform prior from -2 to 2 using

```
prior lda ~ dgeneral(1, lower=-2, upper=2);
```

This places equal probability mass on five points, -2, -1, 0, 1, and 2. This prior might not work well here because the grid is too coarse. To consider smaller values of λ, you can sample a parameter that takes a wider range of integer values and transform it back to the λ space. For example, set alpha as your model parameter and give it a discrete uniform prior from -200 to 200. Then define λ as alpha/100 so λ can take values between -2 and 2 but on a finer grid.

The following statements fit a Box-Cox transformation by using a discrete prior on λ:

```
proc mcmc data=boxcox outpost=simout nmc=50000 thin=10 seed=12567
          monitor=(lda);

   ods select PostSummaries PostIntervals;
   parms s2 1 alpha 10;

   beginprior;
   prior s2 ~ gamma(shape=3, scale=2);
   if alpha=0 then lp = log(2);
      else lp = log(1);
   prior alpha ~ dgeneral(lp, lower=-200, upper=200);
```

```
      lda = alpha * 0.01;
      sd = sqrt(s2);
      endprior;

      if alpha=0 then
         ll = -ly+lpdfnorm(ly, x, sd);
      else do;
         ys = (y**lda - 1)/lda;
         ll = (lda-1)*ly+lpdfnorm(ys, x, sd);
      end;
      model general(ll);
   run;
```

There are two parameters, s2 and alpha, in the model. They are placed in a single PARMS statement so that they are sampled in the same block.

The parameter s2 takes a gamma distribution, and alpha takes a discrete prior. The IF-ELSE statements state that alpha takes twice as much prior density when it is 0 than otherwise. Note that on the original scale, $\Pr(\text{alpha} = 0) = 2 \cdot \Pr(\text{alpha} \neq 0)$. Translating that to the log scale, the densities become $\log(2)$ and $\log(1)$, respectively. The lda assignment statement transforms alpha to the parameter of interest: lda takes values between -2 and 2. You can model lda on a even smaller scale by dividing alpha by a larger constant. However, an increment of 0.01 in the Box-Cox transformation is usually sufficient. The sd assignment statement calculates the square root of the variance term.

The log-likelihood function uses another set of IF-ELSE statements, separating the case of $\lambda = 0$ from the others. The formulas are stated previously. The output of the program is shown in Output 52.2.6.

Output 52.2.6 Box-Cox Transformation

```
                  Box-Cox Transformation, Modeling Lambda = 0

                            The MCMC Procedure

                            Posterior Summaries

                                   Standard              Percentiles
      Parameter        N      Mean  Deviation        25%        50%        75%

      lda           5000  -0.00002    0.00201          0          0          0

                            Posterior Intervals

          Parameter    Alpha     Equal-Tail Interval        HPD Interval

          lda          0.050           0         0            0         0
```

From the summary statistics table, you see that the point estimate for λ is 0 and both of the 95% equal-tail and HPD credible intervals are 0. This strongly suggests that $\lambda = 0$ is the best estimate for this problem. In addition, you can also count the frequency of λ among posterior samples to get a more precise estimate on the posterior probability of λ being 0.

The following statements use PROC FREQ to produce Output 52.2.7 and Output 52.2.8:

```
ods graphics on;
proc freq data=simout;
   ods select onewayfreqs freqplot;
   tables lda /nocum plot=freqplot(scale=percent);
run;
ods graphics off;
```

Output 52.2.7 shows the frequency count table. An estimate of $\Pr(\lambda = 0|\text{data})$ is 96%. The conclusion is that the log transformation should be the appropriate transformation used here, which agrees with the simulation setup. Output 52.2.8 shows the histogram of λ.

Output 52.2.7 Frequency Counts of λ

```
          Box-Cox Transformation, Modeling Lambda = 0

                       The FREQ Procedure

                lda     Frequency       Percent
             ------------------------------------
             -0.0100         106          2.12
                  0         4798         95.96
              0.0100          96          1.92
```

Output 52.2.8 Histogram of λ

Example 52.3: Generalized Linear Models

This example discusses two examples of fitting generalized linear models (GLM) with PROC MCMC. One uses a logistic regression model and one uses a Poisson regression model. The logistic examples use both a diffuse prior and a Jeffreys' prior on the regression coefficients. You can also use the BAYES statement in PROC GENMOD. See Chapter 37, "The GENMOD Procedure."

Logistic Regression Model with a Diffuse Prior

The following statements create a SAS data set with measurements of the number of deaths, y, among n beetles that have been exposed to an environmental contaminant x:

```
title 'Logistic Regression Model with a Diffuse Prior';
data beetles;
   input n y x @@;
   datalines;
6  0  25.7   8  2  35.9   5  2  32.9   7  7  50.4   6  0  28.3
7  2  32.3   5  1  33.2   8  3  40.9   6  0  36.5   6  1  36.5
6  6  49.6   6  3  39.8   6  4  43.6   6  1  34.1   7  1  37.4
8  2  35.2   6  6  51.3   5  3  42.5   7  0  31.3   3  2  40.6
;
```

You can model the data points y_i with a binomial distribution:

$$y_i \mid p_i \sim \text{binomial}(n_i, p_i)$$

where p_i is the success probability and links to the regression covariate x_i through a logit transformation:

$$logit(p_i) = \log\left(\frac{p_i}{1 - p_i}\right) = \alpha + \beta x_i$$

The priors on α and β are both diffuse normal:

$$\pi(\alpha) = \phi(0, \text{var} = 10000)$$
$$\pi(\beta) = \phi(0, \text{var} = 10000)$$

These statements fit a logistic regression with PROC MCMC:

```
ods graphics on;
proc mcmc data=beetles ntu=1000 nmc=20000 nthin=2 propcov=quanew
          diag=(mcse ess) outpost=beetleout seed=246810;
   ods select PostSummaries PostIntervals mcse ess TADpanel;
   parms (alpha  beta) 0;
   prior alpha beta ~ normal(0, var = 10000);
   p = logistic(alpha + beta*x);
   model y ~ binomial(n,p);
run;
```

The key statement in the program is the assignment to p that calculates the probability of death. The SAS function LOGISTIC does the proper transformation. The MODEL statement specifies that the response variable, y, is binomially distributed with parameters n (from the input data set) and p. The summary statistics table, interval statistics table, the Monte Carlos standard error table, and the effective sample sizes table are shown in Output 52.3.1.

Output 52.3.1 MCMC Results

```
              Logistic Regression Model with a Diffuse Prior

                        The MCMC Procedure

                        Posterior Summaries

                              Standard                 Percentiles
  Parameter         N        Mean   Deviation       25%        50%         75%

  alpha          10000   -11.7707      2.0997   -13.1243   -11.6683    -10.3003
  beta           10000     0.2920      0.0542     0.2537     0.2889      0.3268

                        Posterior Intervals

     Parameter    Alpha      Equal-Tail Interval         HPD Interval

     alpha        0.050     -16.3332     -7.9675     -15.8822     -7.6673
     beta         0.050       0.1951      0.4087       0.1901      0.4027

              Logistic Regression Model with a Diffuse Prior

                        The MCMC Procedure

                   Monte Carlo Standard Errors

                                    Standard
           Parameter       MCSE     Deviation     MCSE/SD

           alpha         0.0422        2.0997      0.0201
           beta          0.00110       0.0542      0.0203

                      Effective Sample Sizes

                                   Correlation
           Parameter       ESS           Time     Efficiency

           alpha        2470.1         4.0484       0.2470
           beta         2435.4         4.1060       0.2435
```

The summary statistics table shows that the sample mean of the output chain for the parameter alpha is −11.7707. This is an estimate of the mean of the marginal posterior distribution for the intercept parameter alpha. The estimated posterior standard deviation for alpha is 2.0997. The two 95% credible intervals for alpha are both negative, which indicates with very high probability that the intercept term is negative. On the other hand, you observe a positive effect on the regression coefficient beta. Exposure to the environment contaminant increases the probability of death.

The Monte Carlo standard errors of each parameter are significantly small relative to the posterior standard deviations. A small MCSE/SD ratio indicates that the Markov chain has stabilized and the mean estimates do not vary much over time. Note that the precision in the parameter estimates increases with the square of the MCMC sample size, so if you want to double the precision, you must quadruple the MCMC sample size.

MCMC chains do not produce independent samples. Each sample point depends on the point before it. In this case, the correlation time estimate, read from the effective sample sizes table, is roughly 4. This means that it takes four observations from the MCMC output to make inferences about alpha with the same precision that you would get from using an independent sample. The effective sample size of 2470 reflects this loss of efficiency. The coefficient beta has similar efficiency. You can often observe that some parameters have significantly better mixing (better efficiency) than others, even in a single Markov chain run.

Output 52.3.2 Plots for Parameters in the Logistic Regression Example

Output 52.3.2 *continued*

Diagnostics for beta

Trace plots and autocorrelation plots of the posterior samples are shown in Output 52.3.2. Convergence looks good in both parameters; there is good mixing in the trace plot and quick drop-off in the ACF plot.

One advantage of Bayesian methods is the ability to directly answer scientific questions. In this example, you might want to find out the posterior probability that the environmental contaminant increases the probability of death—that is, $Pr(\beta > 0|\mathbf{y})$. This can be estimated using the following steps:

```
proc format;
   value betafmt low-0 = 'beta <= 0' 0<-high = 'beta > 0';
run;

proc freq data=beetleout;
   tables beta /nocum;
   format beta betafmt.;
run;
```

Output 52.3.3 Frequency Counts

```
            Logistic Regression Model with a Diffuse Prior

                         The FREQ Procedure

              beta      Frequency      Percent
           ---------------------------------------
           beta > 0         10000       100.00
```

All of the simulated values for β are greater than zero, so the sample estimate of the posterior probability that $\beta > 0$ is 100%. The evidence overwhelmingly supports the hypothesis that increased levels of the environmental contaminant increase the probability of death.

If you are interested in making inference based on any quantities that are transformations of the random variables, you can either do it directly in PROC MCMC or by using the DATA step after you run the simulation. Transformations sometimes can make parameter inference quite formidable using direct analytical methods, but with simulated chains, it is easy to compute chains for any set of parameters. Suppose that you are interested in the lethal dose and want to estimate the level of the covariate x that corresponds to a probability of death, p. Abbreviate this quantity as ldp. In other words, you want to solve the logit transformation with a fixed value p. The lethal dose is as follows:

$$\mathsf{ldp} = \frac{\log\left(\frac{p}{1-p}\right) - \alpha}{\beta}$$

You can obtain an estimate of any ldp by using the posterior mean estimates for α and β. For example, lp95, which corresponds to $p = 0.95$, is calculated as follows:

$$\mathsf{lp95} = \frac{\log\left(\frac{0.95}{1-0.95}\right) + 11.77}{0.29} = 50.79$$

where -11.77 and 0.29 are the posterior mean estimates of α and β, respectively, and 50.79 is the estimated lethal dose that leads to a 95% death rate.

While it is easy to obtain the point estimates, it is harder to estimate other posterior quantities, such as the standard deviation directly. However, with PROC MCMC, you can trivially get estimates of any posterior quantities of lp95. Consider the following program in PROC MCMC:

```
proc mcmc data=beetles ntu=1000 nmc=20000 nthin=2 propcov=quanew
          outpost=beetleout seed=246810 plot=density
          monitor=(pi30 ld05 ld50 ld95);
   ods select PostSummaries PostIntervals densitypanel;
   parms (alpha  beta) 0;
   begincnst;
      c1 = log(0.05 / 0.95);
      c2 = -c1;
   endcnst;

   beginprior;
   prior alpha beta ~ normal(0, var = 10000);
   pi30 = logistic(alpha + beta*30);
   ld05 = (c1 - alpha) / beta;
```

```
    ld50 = - alpha / beta;
    ld95 = (c2 - alpha) / beta;
    endprior;
    pi = logistic(alpha + beta*x);
    model y ~ binomial(n,pi);
run;
ods graphics off;
```

The program estimates four additional posterior quantities. The three lpd quantities, ld05, ld50, and ld95, are the three levels of the covariate that kills 5%, 50%, and 95% of the population, respectively. The predicted probability when the covariate x takes the value of 30 is pi30. The MONITOR= option selects the quantities of interest. The PLOTS= option selects kernel density plots as the only ODS graphical output, excluding the trace plot and autocorrelation plot.

Programming statements between the BEGINCNST and ENDCNST statements define two constants. These statements are executed once at the beginning of the simulation. The programming statements between the BEGINPRIOR and ENDPRIOR statements evaluate the quantities of interest. The symbols, pi30, ld05, ld50, and ld95, are functions of the parameters alpha and beta only. Hence, they should not be processed at the observation level and should be included in the BEGINPRIOR and ENDPRIOR statements. Output 52.3.4 lists the posterior summary and Output 52.3.5 shows the density plots of these posterior quantities.

Output 52.3.4 PROC MCMC Results

```
                  Logistic Regression Model with a Diffuse Prior

                            The MCMC Procedure

                            Posterior Summaries

                                Standard              Percentiles
    Parameter        N      Mean    Deviation      25%       50%       75%

    pi30          10000    0.0524    0.0253     0.0340    0.0477    0.0662
    ld05          10000   29.9281    1.8814    28.8430   30.1727   31.2563
    ld50          10000   40.3745    0.9377    39.7271   40.3165   40.9612
    ld95          10000   50.8210    2.5353    49.0372   50.5157   52.3100


                            Posterior Intervals

        Parameter    Alpha     Equal-Tail Interval        HPD Interval

         pi30        0.050       0.0161     0.1133      0.0109     0.1008
         ld05        0.050      25.6409    32.9660     26.2193    33.2774
         ld50        0.050      38.6706    42.3718     38.6194    42.2811
         ld95        0.050      46.7180    56.7667     46.3221    55.8774
```

The posterior mean estimate of lp95 is 50.82, which is close to the estimate of 50.79 by using the posterior mean estimates of the parameters. With PROC MCMC, in addition to the mean estimate, you can get the standard deviation, quantiles, and interval estimates at any level of significance.

From the density plots, you can see, for example, that the sample distribution for π_{30} is skewed to the right, and almost all of your posterior belief concerning π_{30} is concentrated in the region between zero and 0.15.

Output 52.3.5 Density Plots of Quantities of Interest in the Logistic Regression Example

It is easy to use the DATA step to calculate these quantities of interest. The following DATA step uses the simulated values of α and β to create simulated values from the posterior distributions of ld05, ld50, ld95, and π_{30}:

```
data transout;
   set beetleout;
   pi30 = logistic(alpha + beta*30);
   ld05 = (log(0.05 / 0.95) - alpha) / beta;
   ld50 = (log(0.50 / 0.50) - alpha) / beta;
   ld95 = (log(0.95 / 0.05) - alpha) / beta;
run;
```

Subsequently, you can use SAS/INSIGHT, or the UNIVARIATE, CAPABILITY, or KDE procedures to analyze the posterior sample. If you want to regenerate the default ODS graphs from PROC MCMC, see "Regenerating Diagnostics Plots" on page 3552.

Logistic Regression Model with Jeffreys' Prior

A controlled experiment was run to study the effect of the rate and volume of air inspired on a transient reflex vasoconstriction in the skin of the fingers. Thirty-nine tests under various combinations

of rate and volume of air inspired were obtained (Finney 1947). The result of each test is whether or not vasoconstriction occurred. Pregibon (1981) uses this set of data to illustrate the diagnostic measures he proposes for detecting influential observations and to quantify their effects on various aspects of the maximum likelihood fit. The following statements create the data set vaso:

```
title 'Logistic Regression Model with Jeffreys Prior';
data vaso;
   input vol rate resp @@;
   lvol = log(vol);
   lrate = log(rate);
   ind = _n_;
   datalines;
3.7  0.825  1   3.5  1.09   1   1.25  2.5    1   0.75  1.5   1
0.8  3.2    1   0.7  3.5    1   0.6   0.75   0   1.1   1.7   0
0.9  0.75   0   0.9  0.45   0   0.8   0.57   0   0.55  2.75  0
0.6  3.0    0   1.4  2.33   1   0.75  3.75   1   2.3   1.64  1
3.2  1.6    1   0.85 1.415  1   1.7   1.06   0   1.8   1.8   1
0.4  2.0    0   0.95 1.36   0   1.35  1.35   0   1.5   1.36  0
1.6  1.78   1   0.6  1.5    0   1.8   1.5    1   0.95  1.9   0
1.9  0.95   1   1.6  0.4    0   2.7   0.75   1   2.35  0.03  0
1.1  1.83   0   1.1  2.2    1   1.2   2.0    1   0.8   3.33  1
0.95 1.9    0   0.75 1.9    0   1.3   1.625  1
;
```

The variable resp represents the outcome of a test. The variable lvol represents the log of the volume of air intake, and the variable lrate represents the log of the rate of air intake. You can model the data by using logistic regression. You can model the response with a binary likelihood:

$$\text{resp}_i \sim \text{binary}(p_i)$$

with

$$p_i = \frac{1}{1 + \exp(-(\beta_0 + \beta_1 \text{lvol}_i + \beta_2 \text{lrate}_i))}$$

Let X be the design matrix in the regression. Jeffreys' prior for this model is

$$p(\beta) \propto |X^\top M X|^{1/2}$$

where M is a 39 by 39 matrix with off-diagonal elements being 0 and diagonal elements being $p_i(1 - p_i)$. For details on Jeffreys' prior, see "Jeffreys' Prior" on page 146. You can use a number of matrix functions, such as the determinant function, in PROC MCMC to construct Jeffreys' prior. The following statements illustrate how to fit a logistic regression with Jeffreys' prior:

```
/* fitting a logistic regression with Jeffreys' prior */
%let n = 39;
proc mcmc data=vaso nmc=10000 outpost=mcmcout seed=17;
   ods select PostSummaries PostIntervals;

   array beta[3] beta0 beta1 beta2;
   array m[&n, &n];
   array x[&n, 3];
   array xt[3, &n];
```

```
      array xtm[3, &n];
      array xmx[3, 3];
      array p[&n];

      parms beta0 1 beta1 1 beta2 1;

      begincnst;
         x[ind, 1] = 1;
         x[ind, 2] = lvol;
         x[ind, 3] = lrate;
         if (ind eq &n) then do;
            call transpose(x, xt);
            call zeromatrix(m);
         end;
      endcnst;

      beginprior;
      call mult(x, beta, p);                /* p = x * beta */
      do i = 1 to &n;
         p[i] = 1 / (1 + exp(-p[i]));        /* p[i] = 1/(1+exp(-x*beta)) */
         m[i,i] = p[i] * (1-p[i]);
      end;
      call mult (xt, m, xtm);               /* xtm = xt * m       */
      call mult (xtm, x, xmx);              /* xmx = xtm * x      */
      call det (xmx, lp);                   /* lp = det(xmx)      */
      lp = 0.5 * log(lp);                   /* lp = -0.5 * log(lp) */
      prior beta: ~ general(lp);
      endprior;

      model resp ~ bern(p[ind]);
   run;
```

The first ARRAY statement defines an array beta with three elements: beta0, beta1, and beta2.
The subsequent statements define arrays that are used in the construction of Jeffreys' prior. These
include m (the M matrix), x (the design matrix), xt (the transpose of x), and some additional work
spaces.

The explanatory variables lvol and lrate are saved in the array x in the BEGINCNST and ENDCNST
statements. See "BEGINCNST/ENDCNST Statement" on page 3507 for details. After all the
variables are read into x, you transpose the x matrix and store it to xt. The ZEROMATRIX function
call assigns all elements in matrix m the value zero. To avoid redundant calculation, it is best to
perform these calculations as the last observation of the data set is processed—that is, when ind is
39.

You calculate Jeffreys' prior in the BEGINPRIOR and ENDPRIOR statements. The probability
vector p is the product of the design matrix x and parameter vector beta. The diagonal elements
in the matrix m are $p_i(1 - p_i)$. The expression lp is the logarithm of Jeffreys' prior. The PRIOR
statement assigns lp as the prior for the β regression coefficients. The MODEL statement assigns a
binary likelihood to resp, with probability p[ind]. The p array is calculated earlier using the matrix
function MULT. You use the ind variable to pick out the right probability value for each resp.

Posterior summary statistics are displayed in Output 52.3.6.

Output 52.3.6 PROC MCMC Results, Jeffreys' prior

```
                  Logistic Regression Model with Jeffreys Prior

                             The MCMC Procedure

                             Posterior Summaries

                                    Standard                 Percentiles
     Parameter         N        Mean   Deviation      25%         50%         75%

     beta0          10000    -2.9587      1.3258   -3.8117     -2.7938     -2.0007
     beta1          10000     5.2905      1.8193    3.9861      5.1155      6.4145
     beta2          10000     4.6889      1.8189    3.3570      4.4914      5.8547


                             Posterior Intervals

        Parameter     Alpha      Equal-Tail Interval          HPD Interval

        beta0         0.050      -5.8247      -0.7435      -5.5936      -0.6027
        beta1         0.050       2.3001       9.3789       1.8590       8.7222
        beta2         0.050       1.6788       8.6643       1.3611       8.2490
```

You can also use PROC GENMOD to fit the same model by using the following statements:

```
proc genmod data=vaso descending;
   ods select PostSummaries PostIntervals;
   model resp = lvol lrate / d=bin link=logit;
   bayes seed=17 coeffprior=jeffreys nmc=20000 thin=2;
run;
```

The MODEL statement indicates that resp is the response variable and lvol and lrate are the covariates. The options in the MODEL statement specify a binary likelihood and a logit link function. The BAYES statement requests Bayesian capability. The SEED=, NMC=, and THIN= arguments work in the same way as in PROC MCMC. The COEFFPRIOR=JEFFREYS option requests Jeffreys' prior in this analysis.

The PROC GENMOD statements produce Output 52.3.7, with estimates very similar to those reported in Output 52.3.6. Note that you should not expect to see identical output from PROC GENMOD and PROC MCMC, even with the simulation setup and identical random number seed. The two procedures use different sampling algorithms. PROC GENMOD uses the adaptive rejection metropolis algorithm (ARMS) (Gilks and Wild 1992; Gilks 2003) while PROC MCMC uses a random walk Metropolis algorithm. The asymptotic answers, which means that you let both procedures run an very long time, would be the same as they both generate samples from the same posterior distribution.

Output 52.3.7 PROC GENMOD Results

```
              Logistic Regression Model with Jeffreys Prior

                         The GENMOD Procedure

                           Bayesian Analysis

                          Posterior Summaries

                                      Standard              Percentiles
     Parameter         N       Mean   Deviation       25%        50%         75%

     Intercept     10000    -2.8731      1.3088    -3.6754    -2.7248    -1.9253
     lvol          10000     5.1639      1.8087     3.8451     4.9475     6.2613
     lrate         10000     4.5501      1.8071     3.2250     4.3564     5.6810


                          Posterior Intervals

        Parameter     Alpha     Equal-Tail Interval        HPD Interval

        Intercept     0.050     -5.8246     -0.7271     -5.5774     -0.6060
        lvol          0.050      2.1844      9.2297      2.0112      8.9149
        lrate         0.050      1.5666      8.6145      1.3155      8.1922
```

Poisson Regression

You can use the Poisson distribution to model the distribution of cell counts in a multiway contingency table. Aitkin et al. (1989) have used this method to model insurance claims data. Suppose the following hypothetical insurance claims data are classified by two factors: age group (with two levels) and car type (with three levels). The following statements create the data set:

```
title 'Poisson Regression';
data insure;
   input n c car $ age;
   ln = log(n);
   if car = 'large' then
      do car_dummy1=1;
         car_dummy2=0;
      end;
   else if car = 'medium' then
      do car_dummy1=0;
         car_dummy2=1;
      end;
   else
      do car_dummy1=0;
         car_dummy2=0;
      end;

   datalines;
   500    42   small  0
   1200   37   medium 0
```

```
100     1   large  0
400   101   small  1
500    73   medium 1
300    14   large  1
;
```

The variable n represents the number of insurance policy holders and the variable c represents the number of insurance claims. The variable car is the type of car involved (classified into three groups), and it is coded into two levels. The variable age is the age group of a policy holder (classified into two groups).

Assume that the number of claims c has a Poisson probability distribution and that its mean, μ_i, is related to the factors car and age for observation i by

$$
\begin{aligned}
\log(\mu_i) &= \log(n_i) + \mathbf{x}'\boldsymbol{\beta} \\
&= \log(n_i) + \beta_0 + \\
&\quad \text{car}_i(1)\beta_1 + \text{car}_i(2)\beta_2 + \text{car}_i(3)\beta_3 + \\
&\quad \text{age}_i(1)\beta_4 + \text{age}_i(2)\beta_5
\end{aligned}
$$

The indicator variables $\text{car}_i(j)$ is associated with the jth level of the variable car for observation i in the following way:

$$
\text{car}_i(j) = \begin{cases} 1 & \text{if car} = j \\ 0 & \text{if car} \neq j \end{cases}
$$

A similar coding applies to age. The β's are parameters. The logarithm of the variable n is used as an offset—that is, a regression variable with a constant coefficient of 1 for each observation. Having the offset constant in the model is equivalent to fitting an expanded data set with 3000 observations, each with response variable y observed on an individual level. The log link is used to relate the mean and the factors car and age.

The following statements run PROC MCMC:

```
proc mcmc data=insure outpost=insureout nmc=5000 propcov=quanew
        maxtune=0 seed=7;
   ods select PostSummaries PostIntervals;
   parms alpha 0 beta_car1 0 beta_car2 0 beta_age 0;
   prior alpha beta: ~ normal(0, prec = 1e-6);

   mu = ln + alpha + beta_car1 * car_dummy1
      + beta_car2 * car_dummy2 + beta_age * age;
   model c ~ poisson(exp(mu));
run;
```

The analysis uses a relatively flat prior on all the regression coefficients, with mean at 0 and precision at 10^{-6}. The option MAXTUNE=0 skips the tuning phase because the optimization routine (PROPCOV=QUANEW) provides good initial values and proposal covariance matrix.

There are four parameters in the model: alpha is the intercept; beta_car1 and beta_car2 are coefficients for the class variable car, which has three levels; and beta_age is the coefficient for age. The symbol mu connects the regression model and the Poisson mean by using the log link. The MODEL statement specifies a Poisson likelihood for the response variable c.

Posterior summary and interval statistics are shown in Output 52.3.8.

Output 52.3.8 MCMC Results

```
                         Poisson Regression

                         The MCMC Procedure

                        Posterior Summaries

                                  Standard              Percentiles
       Parameter        N      Mean   Deviation      25%        50%        75%

       alpha          5000   -2.6403    0.1344    -2.7261    -2.6387    -2.5531
       beta_car1      5000   -1.8335    0.2917    -2.0243    -1.8179    -1.6302
       beta_car2      5000   -0.6931    0.1255    -0.7775    -0.6867    -0.6118
       beta_age       5000    1.3151    0.1386     1.2153     1.3146     1.4094


                        Posterior Intervals

          Parameter    Alpha    Equal-Tail Interval       HPD Interval

          alpha        0.050    -2.9201    -2.3837    -2.9133    -2.3831
          beta_car1    0.050    -2.4579    -1.3036    -2.4692    -1.3336
          beta_car2    0.050    -0.9462    -0.4498    -0.9485    -0.4589
          beta_age     0.050     1.0442     1.5898     1.0387     1.5812
```

To fit the same model by using PROC GENMOD, you can do the following. Note that the default normal prior on the coefficients β is $N(0, \text{prec} = 1e - 6)$, the same as used in the PROC MCMC. The following statements run PROC GENMOD and create Output 52.3.9:

```
proc genmod data=insure;
   ods select PostSummaries PostIntervals;
   class car age(descending);
   model c = car age / dist=poisson link=log offset=ln;
   bayes seed=17 nmc=5000 coeffprior=normal;
run;
```

To compare, posterior summary and interval statistics from PROC GENMOD are reported in Output 52.3.9, and they are very similar to PROC MCMC results in Output 52.3.8.

Output 52.3.9 PROC GENMOD Results

```
                          Poisson Regression

                         The GENMOD Procedure

                           Bayesian Analysis

                         Posterior Summaries

                                 Standard              Percentiles
  Parameter          N       Mean   Deviation        25%        50%        75%

  Intercept       5000    -2.6353      0.1299    -2.7243    -2.6312    -2.5455
  carlarge        5000    -1.7996      0.2752    -1.9824    -1.7865    -1.6139
  carmedium       5000    -0.6977      0.1269    -0.7845    -0.6970    -0.6141
  age1            5000     1.3148      0.1348     1.2237     1.3138     1.4067

                         Posterior Intervals

      Parameter    Alpha     Equal-Tail Interval        HPD Interval

      Intercept    0.050     -2.8952    -2.3867     -2.8755    -2.3730
      carlarge     0.050     -2.3538    -1.2789     -2.3424    -1.2691
      carmedium    0.050     -0.9494    -0.4487     -0.9317    -0.4337
      age1         0.050      1.0521     1.5794      1.0624     1.5863
```

Note that the descending option in the CLASS statement reverses the sorting order of the class variable age so that the results agree with PROC MCMC. If this option is not used, the estimate for age has a reversed sign as compared to .

Example 52.4: Nonlinear Poisson Regression Models

This example illustrates how to fit a nonlinear Poisson regression with PROC MCMC. In addition, it shows how you can improve the mixing of the Markov chain by selecting a different proposal distribution or by sampling on the transformed scale of a parameter. This example shows how to analyze count data for calls to a technical support help line in the weeks immediately following a product release. This information could be used to decide upon the allocation of technical support resources for new products. You can model the number of daily calls as a Poisson random variable, with the average number of calls modeled as a nonlinear function of the number of weeks that have elapsed since the product's release. The data are input into a SAS data set as follows:

```
title 'Nonlinear Poisson Regression';
data calls;
   input weeks calls @@;
   datalines;
1   0   1   2   2   2   2   1   3   1   3   3
4   5   4   8   5   5   5   9   6  17   6   9
7  24   7  16   8  23   8  27
;
```

During the first several weeks after a new product is released, the number of questions that technical support receives concerning the product increases in a sigmoidal fashion. The expression for the mean value in the classic Poisson regression involves the log link. There is some theoretical justification for this link, but with MCMC methodologies, you are not constrained to exploring only models that are computationally convenient. The number of calls to technical support tapers off after the initial release, so in this example you can use a logistic-type function to model the mean number of calls received weekly for the time period immediately following the initial release. The mean function $\lambda(t)$ is modeled as follows:

$$\lambda_i = \frac{\gamma}{1 + \exp\left[-(\alpha + \beta t_i)\right]}$$

The likelihood for every observation calls_i is

$$\text{calls}_i \sim \text{Poisson}\left(\lambda_i\right)$$

Past experience with technical support data for similar products suggests using a gamma distribution with shape and scale parameters 3.5 and 12 as the prior distribution for γ, a normal distribution with mean -5 and variance 0.25 as the prior for α, and a normal distribution with mean 0.75 and variance 0.5 as the prior for β.

The following PROC MCMC statements fit this model:

```
ods graphics on;
proc mcmc data=calls outpost=callout seed=53197 ntu=1000 nmc=20000
          propcov=quanew;
   ods select TADpanel;
   parms alpha -4 beta 1 gamma 2;
   prior alpha ~ normal(-5, sd=0.25);
   prior beta  ~ normal(0.75, sd=0.5);
   prior gamma ~ gamma(3.5, scale=12);
   lambda = gamma*logistic(alpha+beta*weeks);
   model calls ~ poisson(lambda);
run;
```

The one PARMS statement defines a block of all parameters and sets their initial values individually. The PRIOR statements specify the informative prior distributions for the three parameters. The assignment statement defines λ, the mean number of calls. Instead of using the SAS function LOGISTIC, you can use the following statement to calculate λ and get the same result:

```
lambda = gamma / (1 + exp(-(alpha+beta*weeks)));
```

Mixing is not particularly good with this run of PROC MCMC. The `ods select` statement displays only the diagnostic graphs while excluding all other output. The graphical output is shown in Output 52.4.1.

Output 52.4.1 Plots for Parameters
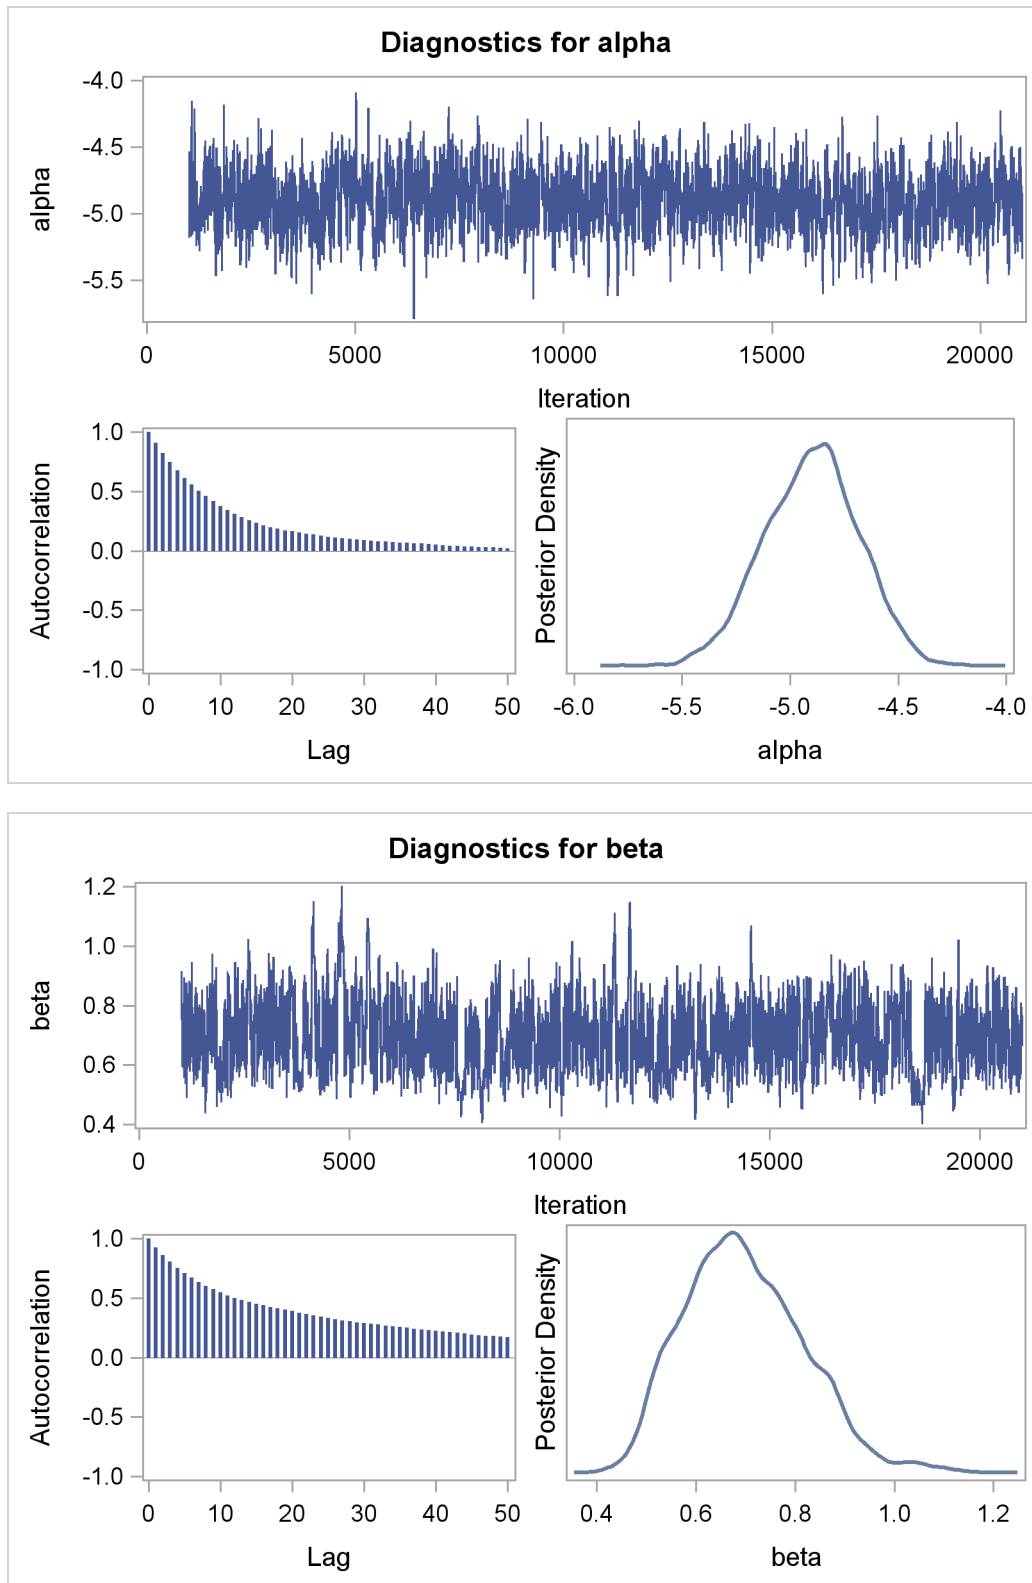
Output 52.4.1 *continued*

By examining the trace plot of the gamma parameter, you see that the Markov chain sometimes gets stuck in the far right tail and does not travel back to the high density area quickly. This effect can be seen around the simulations number 8000 and 18000. One possible explanation for this is that the random walk Metropolis is taking too small of steps in its proposal; therefore it takes more iterations for the Markov chain to explore the parameter space effectively. The step size in the random walk is controlled by the normal proposal distribution (with a multiplicative scale). A (good) proposal distribution is roughly an approximation to the joint posterior distribution at the mode. The curvature of the normal proposal distribution (the variance) does not take into account the thickness of the tail areas. As a result, a random walk Metropolis with normal proposal can have a hard time exploring distributions that have thick tails. This appears to be the case with the posterior distribution of the parameter gamma. You can improve the mixing by using a thicker-tailed proposal distribution, the *t*-distribution. The option PROPDIST controls the proposal distribution. PROPDIST=T(3) changes the proposal from a normal distribution to a *t*-distribution with three degrees of freedom.

The following statements run PROC MCMC and produce Output 52.4.2:

```
proc mcmc data=calls outpost=callout seed=53197 ntu=1000 nmc=20000
          propcov=quanew stats=none propdist=t(3);
   ods select TADpanel;
   parms alpha -4 beta 1 gamma 2;
   prior alpha ~ normal(-5, sd=0.25);
   prior beta  ~ normal(0.75, sd=0.5);
   prior gamma ~ gamma(3.5, scale=12);
   lambda = gamma*logistic(alpha+beta*weeks);
   model calls ~ poisson(lambda);
run;
```

Output 52.4.2 displays the graphical output.

Output 52.4.2 Plots for Parameters, Using a t(3) Proposal Distribution
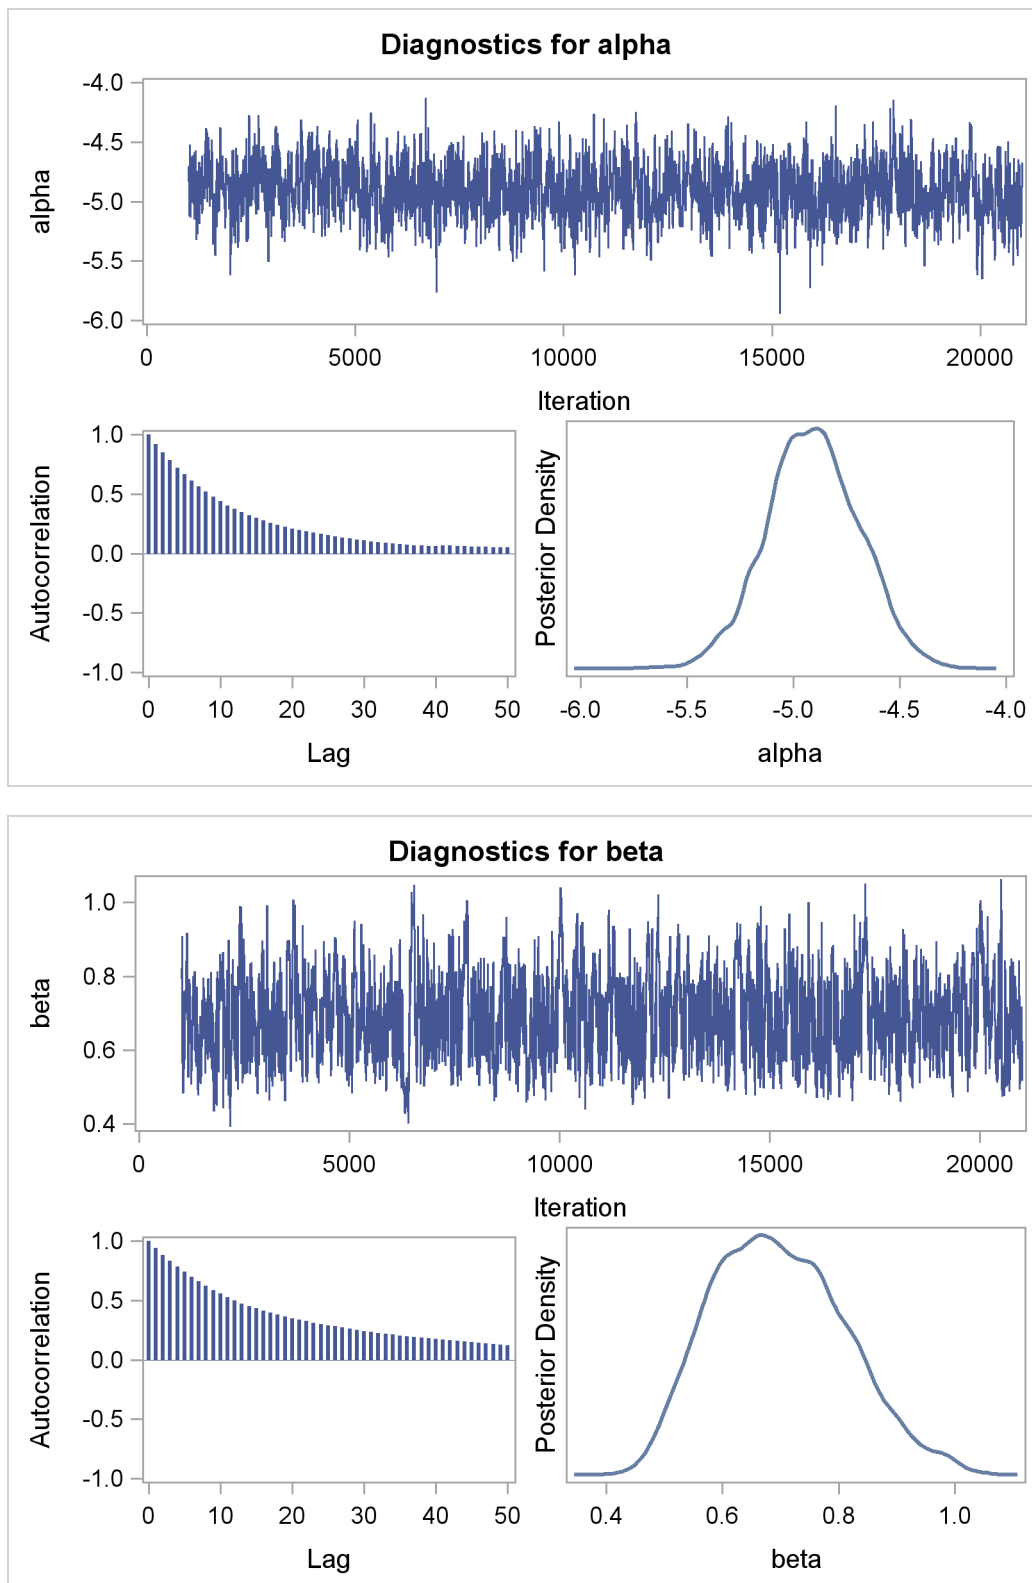
Output 52.4.2 *continued*

Diagnostics for gamma
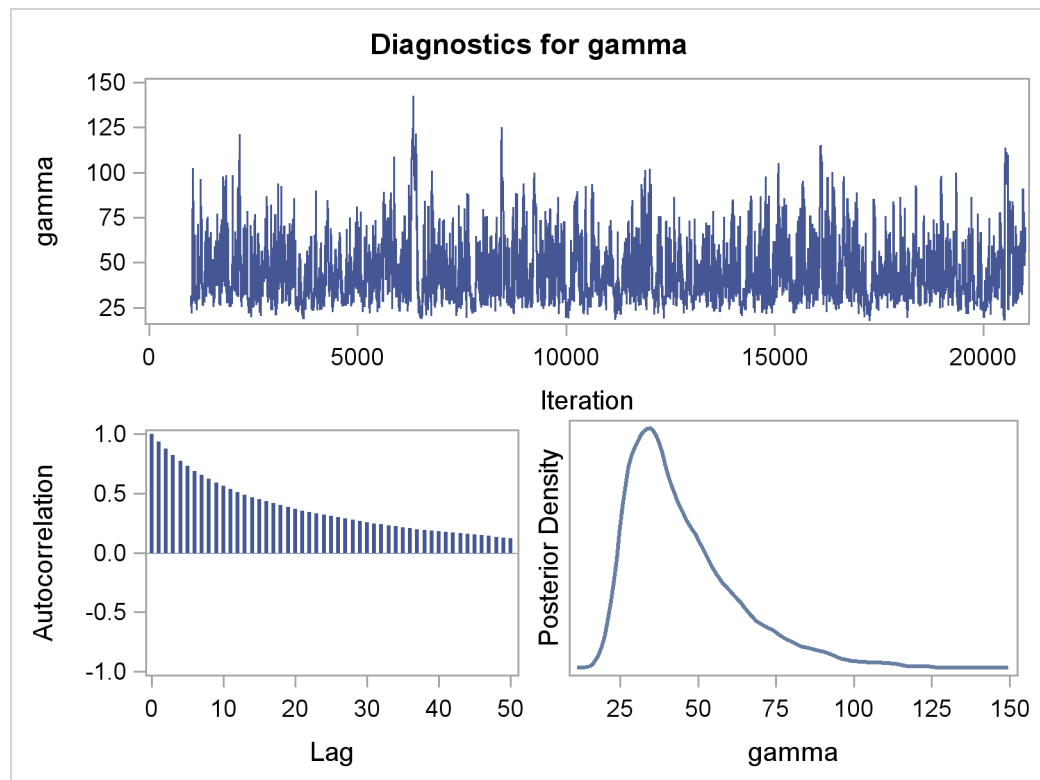
The trace plots are more dense and the ACF plots have faster drop-offs, and you see improved mixing by using a thicker-tailed proposal distribution. If you want to further improve the Markov chain, you can choose to sample the log transformation of the parameter gamma:

$$\text{lg} \sim \text{egamma}(3.5, \text{scale} = 12) \quad \text{is equivalent to} \quad gamma = \exp(\text{lg}) \sim \text{gamma}(3.5, \text{scale} = 12)$$

The parameter gamma has a positive support. Often in this case, it has right-skewed posterior. By taking the log transformation, you can sample on a parameter space that does not have a lower boundary and is more symmetric. This can lead to better mixing.

The following statements produce Output 52.4.4 and Output 52.4.3:

```
proc mcmc data=calls outpost=callout seed=53197 ntu=1000 nmc=20000
          propcov=quanew propdist=t(3)
          monitor=(alpha beta lgamma gamma);
   ods select PostSummaries PostIntervals TADpanel;
   parms alpha -4 beta 1 lgamma 2;
   prior alpha ~ normal(-5, sd=0.25);
   prior beta  ~ normal(0.75, sd=0.5);
   prior lgamma ~ egamma(3.5, scale=12);
   gamma = exp(lgamma);
   lambda = gamma*logistic(alpha+beta*weeks);
   model calls ~ poisson(lambda);
run;
ods graphics off;
```

In the PARMS statement, instead of gamma, you have lgamma. Its prior distribution is egamma, as opposed to the gamma distribution. Note that the following two priors are equivalent to each other:

```
prior lgamma ~ egamma(3.5, scale=12);
prior gamma ~ gamma(3.5, scale=12);
```

The gamma assignment statement transforms lgamma to gamma. The lambda assignment statement calculates the mean for the Poisson by using the gamma parameter. The MODEL statement specifies a Poisson likelihood for the calls response.

The trace plots and ACF plots in Output 52.4.3 show the best mixing seen so far in this example.

Output 52.4.3 Plots for Parameters, Sampling on the Log Scale of Gamma

Output 52.4.3 *continued*

Diagnostics for beta

Diagnostics for lgamma
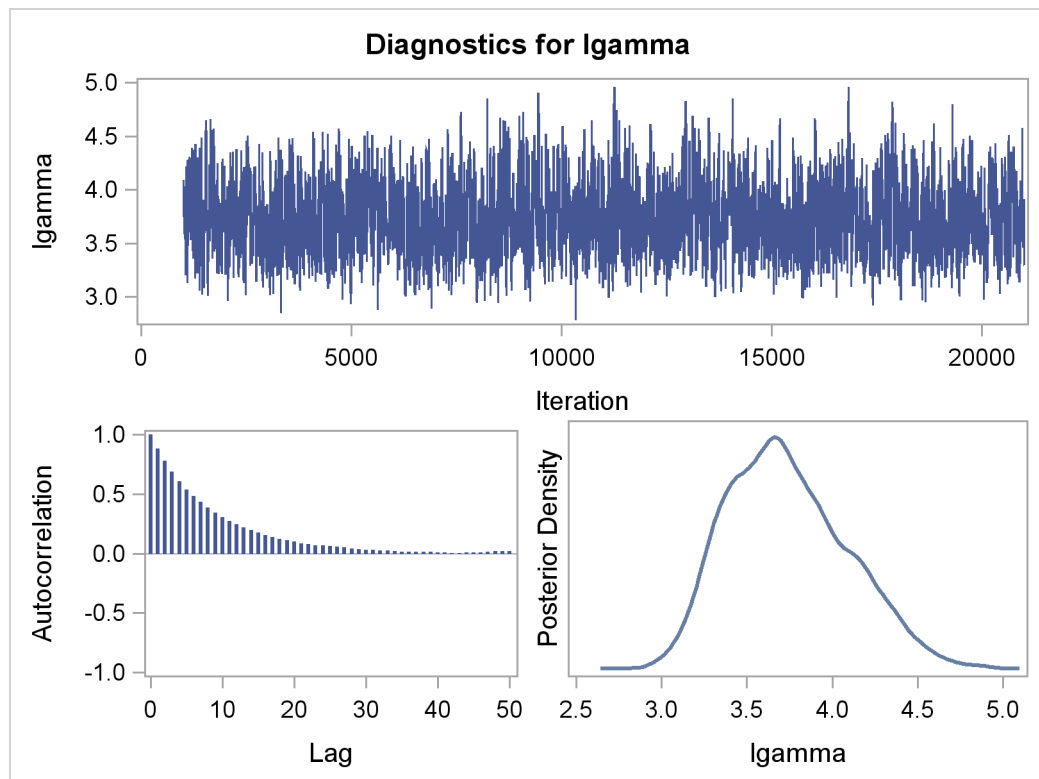
Output 52.4.3 *continued*
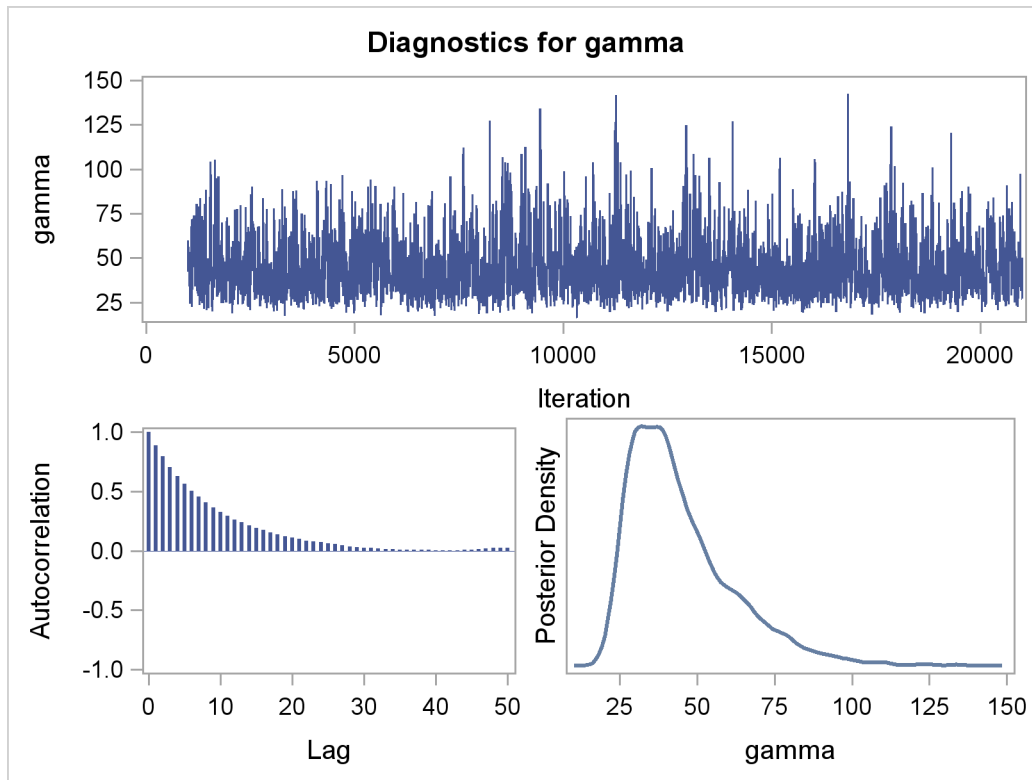
Output 52.4.4 shows the posterior summary statistics of the nonlinear Poisson regression. Note that the lgamma parameter has a more symmetric density than the skewed gamma parameter. The Metropolis algorithm always works better if the target distribution is approximately normal.

Output 52.4.4 MCMC Results, Sampling on the Log Scale of Gamma

```
                        Nonlinear Poisson Regression

                            The MCMC Procedure

                            Posterior Summaries
```

| | | | Standard | | Percentiles | |
Parameter	N	Mean	Deviation	25%	50%	75%
alpha	20000	-4.8907	0.2160	-5.0435	-4.8872	-4.7461
beta	20000	0.6957	0.1089	0.6163	0.6881	0.7698
lgamma	20000	3.7391	0.3487	3.4728	3.7023	3.9696
gamma	20000	44.8136	17.0430	32.2263	40.5415	52.9647

Output 52.4.4 *continued*

```
                       Posterior Intervals

       Parameter    Alpha     Equal-Tail Interval        HPD Interval

       alpha        0.050     -5.3138    -4.4667     -5.3276    -4.4953
       beta         0.050      0.5066     0.9253      0.4868     0.8996
       lgamma       0.050      3.1580     4.4705      3.1222     4.4127
       gamma        0.050     23.5225    87.3972     20.9005    79.4712
```

This example illustrates that PROC MCMC can fit Bayesian nonlinear models just as easily as Bayesian linear models. More importantly, transformations can sometimes improve the efficiency of the Markov chain, and that is something to always keep in mind. Also see "Example 52.12: Using a Transformation to Improve Mixing" on page 3677 for another example of how transformations can improve mixing of the Markov chains.

Example 52.5: Random-Effects Models

This example illustrates how you can use PROC MCMC to fit random effects models. In the example "Mixed-Effects Model" on page 3490 in "Getting Started: MCMC Procedure" on page 3477, you already saw PROC MCMC fit a linear random effects model. There are two more examples in this section. One is a logistic random effects model, and the second one is a nonlinear Poisson regression random effects model. In addition, this section illustrates how to construct prior distributions that depend on input data set variables. Such prior distributions appear frequently in random effects model, especially in cases of hierarchical centering. Although you can use PROC MCMC to analyze random effects models, you might want to first consider some other SAS procedures. For example, you can use PROC MIXED (see Chapter 56, "The MIXED Procedure") to analyze linear mixed effects models, PROC NLMIXED (see Chapter 61, "The NLMIXED Procedure") for nonlinear mixed effects models, and PROC GLIMMIX (see Chapter 38, "The GLIMMIX Procedure") for generalized linear mixed effects models. In addition, a sampling-based Bayesian analysis is available in the MIXED procedure through the PRIOR statement (see "PRIOR Statement" on page 3939).

Logistic Regression Random-Effects Model

This example shows how to fit a logistic random-effects model in PROC MCMC. The data are taken from Crowder (1978). The seeds data set is a 2 × 2 factorial layout, with two types of seeds, *O. aegyptiaca 75* and *O. aegyptiaca 73*, and two root extracts, *bean* and *cucumber*. You observe r, which is the number of germinated seeds, and n, which is the total number of seeds. The independent variables are seed and extract.

The following statements create the data set:

```
title 'Logistic Regression Random-Effects Model';
data seeds;
   input r n seed extract @@;
   ind = _N_;
   datalines;
10   39  0  0     23  62  0  0     23  81  0  0     26  51  0  0
17   39  0  0      5   6  0  1     53  74  0  1     55  72  0  1
32   51  0  1     46  79  0  1     10  13  0  1      8  16  1  0
10   30  1  0      8  28  1  0     23  45  1  0      0   4  1  0
 3   12  1  1     22  41  1  1     15  30  1  1     32  51  1  1
 3    7  1  1
;
```

You can model each observation r_i as having its own probability of success p_i, and the likelihood is as follows:

$$r_i \sim \text{binomial}(n_i, p_i)$$

You can use the logit link function to link the covariates of each observation, seed and extract, to the probability of success:

$$\mu_i = \beta_0 + \beta_1 \cdot \text{seed}_i + \beta_2 \cdot \text{extract}_i + \beta_3 \cdot \text{seed}_i \cdot \text{extract}_i$$
$$p_i = \text{logistic}(\mu_i + \eta_i)$$

where η_i is assumed to be as i.i.d. random effect with a normal prior:

$$\eta_i \sim \text{normal}(0, \text{var} = \sigma^2)$$

The four β regression coefficients and the standard deviation σ^2 in the random effects are model parameters; they are given noninformative priors as follows:

$$\pi(\beta_0, \beta_1, \beta_2, \beta_3) \propto 1$$
$$\pi(\sigma^2) \propto 1/\sigma^2$$

Another way of expressing the same model is as follows:

$$p_i = \text{logistic}(\delta_i)$$

where

$$\delta_i \sim \text{normal}(\beta_0 + \beta_1 \cdot \text{seed}_i + \beta_2 \cdot \text{extract}_i + \beta_3 \cdot \text{seed}_i \cdot \text{extract}_i, \sigma^2)$$

The two models are equivalent. In the first model, the random effects η_i centers at 0 in the normal distribution, and in the second model, δ_i centers at the regression mean. This hierarchical centering can sometimes improve mixing.

From a programming point of view, the second parameterization of the model is more difficult because the prior distribution on δ_i involves the data set variables seed and extract. Each prior distribution depends on a different set of observations in the input data set. Intuitively, you might think that the following statements would specify such a prior:

```
   mu = beta0 + beta1*seed + beta2*extract + beta3*seed*extract;
   prior delta ~ normal(mu, var = v);
```

However, this will not work. This is because the procedure is not able to match the observational level calculation (mu) with elements of a parameter array (there are 21 random effects in delta). Thus, the procedure cannot calculate the log of the prior density correctly. The solution is to cumulatively calculate the joint prior distribution for all $\delta_i, i = 1 \cdots 21$, and assign the prior distribution to all δ by using the GENERAL function.

The following statements generate Output 52.5.1:

```
   proc mcmc data=seeds outpost=postout seed=332786 nmc=100000 thin=10
            ntu=3000 monitor=(beta0-beta3 v);
      ods select PostSummaries ess;
      array delta[21];

      parms delta: 0;
      parms beta0 0 beta1 0 beta2 0 beta3 0 ;
      parms v 1;

      beginprior;
      sigma = sqrt(v);
      endprior;
      w = beta0 + beta1*seed + beta2*extract + beta3*seed*extract;
      if ind eq 1 then
         lp = lpdfnorm(delta[ind], w, sigma);
      else
         lp = lp + lpdfnorm(delta[ind], w, sigma);

      prior v   ~ general(-log(v));
      prior beta: ~ general(1);
      prior delta: ~ general(lp);

      pi = logistic(delta[ind]);
      model r ~ binomial(n = n, p = pi);
   run;
```

PROC MCMC statement specifies the input and output data sets, sets a seed for the random number generator, requests a very large simulation number, thins the Markov chain by 10, and specifies a tuning sample size of 3000. The MONITOR= option selects the parameters of interest. The **ods select** statement displays the summary statistics and effective sample size tables.

The ARRAY statement allocates an array of size 21 for the random effects parameter δ. There are three PARMS statements that place δ, β and σ^2 into three sampling blocks. Calculation of sigma does not involve any observations; hence, it is enclosed in the BEGINPRIOR and ENDPRIOR statements.

The next few lines of statements construct a joint prior distribution for all the δ parameters. The symbol w is the regression mean, whose value changes for every observation. The IF-ELSE statements add the log of the normal density to the symbol lp as PROC MCMC steps through the data set. When ind is 1, lp is the log of the normal density for delta[1] evaluated at the first regression

mean w. As ind gradually increases to 21, lp becomes

$$\sum_i \log(\phi(\delta_i | \beta X_i, \sigma))$$

which is the joint prior distribution for all δ.

The PRIOR statements assign three priors to these parameters, with noninformative priors on σ^2 and β. All of the delta parameters share a joint prior, which is defined by lp. Recall that PROC MCMC adds the log of the prior density to the log of the posterior density at the last observation at every simulation, so the expression lp will have the correct value.

CAUTION: You must define the expression lp before the PRIOR statement for the delta parameters. Switching the order of the PRIOR statement and the programming statements that define lp leads to an incorrect prior distribution for delta. The following statements are wrong because the expression lp has not completed its calculation when lp is added to the log of the posterior density at the last observation of the input data set.

```
prior delta: ~ general(lp);
w = beta0 + beta1*seed + beta2*extract + beta3*seed*extract;
if ind eq 1 then
   lp = lpdfnorm(delta[ind], w, sigma);
else
   lp = lp + lpdfnorm(delta[ind], w, sigma);
```

The prior you specify in this case is:

$$\sum_{i=1}^{n-1} \log(\phi(\delta_i | \beta X_i, \sigma))$$

The correct log density is the following:

$$\sum_{i=1}^{n} \log(\phi(\delta_i | \beta X_i, \sigma))$$

The symbol pi is the logit transformation. The MODEL specifies the response variable r as a binomial distribution with parameters n and pi.

The mixing is poor in this example. You can see from the effective sample size table (Output 52.5.1) that the efficiency for all parameters is relatively low, even after a substantial amount of thinning. One possible solution is to break the random effects block of parameters (b) into multiple blocks with a smaller number of parameters.

Output 52.5.1 Logistic Regression Random-Effects Model

```
                 Logistic Regression Random-Effects Model

                         The MCMC Procedure

                         Posterior Summaries

                                  Standard             Percentiles
  Parameter        N       Mean   Deviation       25%        50%        75%

  beta0          10000   -0.5503    0.2025    -0.6784    -0.5522    -0.4193
  beta1          10000    0.0626    0.3292    -0.1512     0.0653     0.2760
  beta2          10000    1.3546    0.2876     1.1732     1.3391     1.5349
  beta3          10000   -0.8257    0.4498    -1.1044    -0.8255    -0.5344
  v              10000    0.1145    0.1019     0.0472     0.0875     0.1503


                 Logistic Regression Random-Effects Model

                         The MCMC Procedure

                        Effective Sample Sizes

                                   Correlation
           Parameter        ESS           Time     Efficiency

           beta0          885.3       11.2952         0.0885
           beta1          603.2       16.5771         0.0603
           beta2          854.9       11.6970         0.0855
           beta3          591.6       16.9021         0.0592
           v              273.1       36.6182         0.0273
```

To fit the same model in PROC GLIMMIX, you can use the following statements, which produce Output 52.5.2:

```
proc glimmix data=seeds method=quad;
   ods select covparms parameterestimates;
   ods output covparms=cp parameterestimates=ps;
  class ind;
  model r/n = seed extract seed*extract/ dist=binomial link=logit solution;
  random intercept / subject=ind;
run;
```

Output 52.5.2 Estimates by PROC GLMMIX

```
                    Logistic Regression Random-Effects Model

                            The GLIMMIX Procedure

                        Covariance Parameter Estimates

                                                 Standard
                Cov Parm        Subject     Estimate       Error

                Intercept       ind          0.05577     0.05196


                        Solutions for Fixed Effects

                               Standard
        Effect          Estimate        Error        DF     t Value     Pr > |t|

        Intercept        -0.5484       0.1666        17      -3.29       0.0043
        seed              0.09701      0.2780         0       0.35         .
        extract           1.3370       0.2369         0       5.64         .
        seed*extract     -0.8104       0.3851         0      -2.10         .
```

It is hard to compare point estimates from these two procedures. However, you can visually compare the results by plotting a kernel density plot (by using the posterior sample from PROC MCMC output) on top of a normal approximation plot (by using the mean and standard error estimates from PROC GLIMMIX, for each parameter). This kernel comparison plot is shown in Output 52.5.3.

However, it takes some work to produce the kernel comparison plot. First, you must use PROC KDE to estimate the kernel density for each parameter from MCMC. Next, you want to get the point estimates from the PROC GLIMMIX output. Then, you generate a SAS data set that contains both the kernel density estimates and the gridded estimates based on normal approximations. Finally, you use PROC TEMPLATE (see Chapter 21, "Statistical Graphics Using ODS") to define an appropriate graphical template and produce the comparison plot by using PROC SGRENDER (see the SGRENDER Procedure in the *SAS/GRAPH: Statistical Graphics Procedures Guide*).

The following statements use PROC KDE on the posterior sample data set postout and estimate a kernel density for each parameter, saving the estimates to a SAS data set m1:

```
proc kde data=postout;
   univar beta0 beta1 beta2 beta3 v / out=m1 (drop=count);
run;
```

The following SAS statements take the estimates of all the parameters from the PROC GLIMMIX output, data sets ps and cp, and assign them to macro variables:

```
data gmxest(keep = parm mean sd);
   set ps cp;
   mean = estimate;
   sd = stderr;
   i = _n_-1;
   if(_n_ ne 5) then
      parm = "beta" || put(i, z1.);
```

```
      else
         parm = "var";
   run;

   data msd (keep=mean sd);
      set gmxest;
      do j = 1 to 401;
         output;
      end;
   run;

   data _null_;
      set ps;
      call symputx(compress(effect,'*'), estimate);
      call symputx(compress('s' || effect,'*'), stderr);
   run;

   data _null_;
      set cp;
      call symputx("var", estimate);
      call symputx("var_sd", stderr);
   run;

   %put &intercept &seed &extract &seedextract &var;
   %put &sintercept &sseed &sextract &sseedextract &var_sd;
```

Specifically, the mean estimate of β_0 is assigned to intercept, and the standard error of β_0 is assigned to sintercept. The macro variables seed, extract, seedextract are the mean estimates for β_1, β_2 and β_3, respectively.

To create a SAS data set that contains both the kernel density estimates and the corresponding normal approximation, you can use the %REN and %RESHAPE macros. The %REN macro renames the variables of a SAS data set by appending the suffix name to each variable name, to avoid redundant variable names. The %RESHAPE macro takes an output data set from a PROC KDE run, and transposes it to the right format so that PROC SGRENDER can generate the right graph. The following statements define the %REN and %RESHAPE macros:

```
   /* define macros */
   %macro ren(in=, out=, suffix=);
      %local s;

      proc contents data=&in noprint out=__temp__(keep=name);
      run;

      data _null_;
         length s $ 32000;
         retain s;
         set __temp__ end=eof;
         s = trim(s)||' '||trim(name)||'='||compress(name||"&suffix");
         if eof then call symput('s', trim(s));
      run;

      proc datasets nolist;
```

```
      delete __temp__;
   run; quit;

   data &out;
      set &in(rename=(&s));
   run;
%mend;

%macro reshape(input, output, suffix1=, suffix2=);
   proc sort data=&input;
      by var;
   run;

   data tmp&input;
      set &input;
      by var;
      _n + 1;
      if first.var then _n = 0;
   run;

   proc sort;
      by _n var;
   run;

   proc transpose data=tmp&input out=_by_value_(drop=_n _name_ _label_);
      var value;
      by _n;
      id var;
   run;

   %ren(in=_by_value_, out=_by_value_, suffix=&suffix1)

   proc transpose data=tmp&input out=_by_den_(drop=_n _name_ _label_);
      var density;
      by _n;
      id var;
   run;

   %ren(in=_by_den_, out=_by_den_, suffix=&suffix2)

   data &output;
      merge _by_value_ _by_den_;
   run;

   proc datasets library=work;
      ods exclude all;
      delete tmp&input _by_value_ _by_den_;
   run;
   ods exclude none;
%mend;
```

When you apply the %RESHAPE macro to the data set m1, you create a SAS data set mcmc that has grid values of the β parameters and their corresponding kernel density estimates. Next, you evaluate these parameter grid values in a normal density with the macro variables taken from the

PROC GLIMMIX output:

```
/* create data set mcmc */
%reshape(m1, mcmc, suffix1=, suffix2=_kde);
data all;
   set mcmc;
   beta0_gmx = pdf('normal', beta0, &intercept, &sintercept);
   beta1_gmx = pdf('normal', beta1, &seed, &sseed);
   beta2_gmx = pdf('normal', beta2, &extract, &sextract);
   beta3_gmx = pdf('normal', beta3, &seedextract, &sseedextract);
   v_gmx = pdf('normal', v, &var, &var_sd);
run;
```

In the data set all, you have grid values on β and σ^2, their kernel density estimates from PROC
MCMC, and the normal density evaluated by using estimates from PROC GLIMMIX. To create an
overlaid plot, you first use PROC TEMPLATE to create a 2×3 template as demonstrated by the
following statements:

```
proc template;
   define statgraph twobythree;
      %macro plot;
         begingraph;
            layout lattice / rows=2 columns=3;
               %do i = 0 %to 3;
                  layout overlay /yaxisopts=(label=" ");
                     seriesplot y=beta&i._kde x=beta&i
                        / connectorder=xaxis
                           lineattrs=(pattern=mediumdash color=blue)
                        legendlabel = "MCMC Kernel" name="MCMC";
                     seriesplot y=beta&i._gmx x=beta&i
                        / connectorder=xaxis lineattrs=(color=red)
                        legendlabel="GLIMMIX Approximation" name="GLIMMIX";
                  endlayout;
               %end;
               layout overlay /yaxisopts=(label=" ")
                              xaxisopts=(linearopts=(viewmin=0 viewmax=0.6));
                  seriesplot y=v_kde x=v
                     / connectorder=xaxis
                        lineattrs=(pattern=mediumdash color=blue)
                     legendlabel = "MCMC Kernel" name="MCMC";
                  seriesplot y=v_gmx x=v
                     / connectorder=xaxis lineattrs=(color=red)
                     legendlabel="GLIMMIX Approximation" name="GLIMMIX";
               endlayout;
               Sidebar / align = bottom;
                  discretelegend "MCMC" "GLIMMIX";
               endsidebar;
            endlayout;
         endgraph;
      %mend; %plot;
   end;
run;
```

The kernel density comparison plot is produced by calling PROC SGRENDER (see the SGREN-

DER Procedure in the *SAS/GRAPH: Statistical Graphics Procedures Guide*):

```
proc sgrender data=all template=twobythree;
run;
```

Output 52.5.3 Comparing Estimates from PROC MCMC and PROC GLIMMIX.
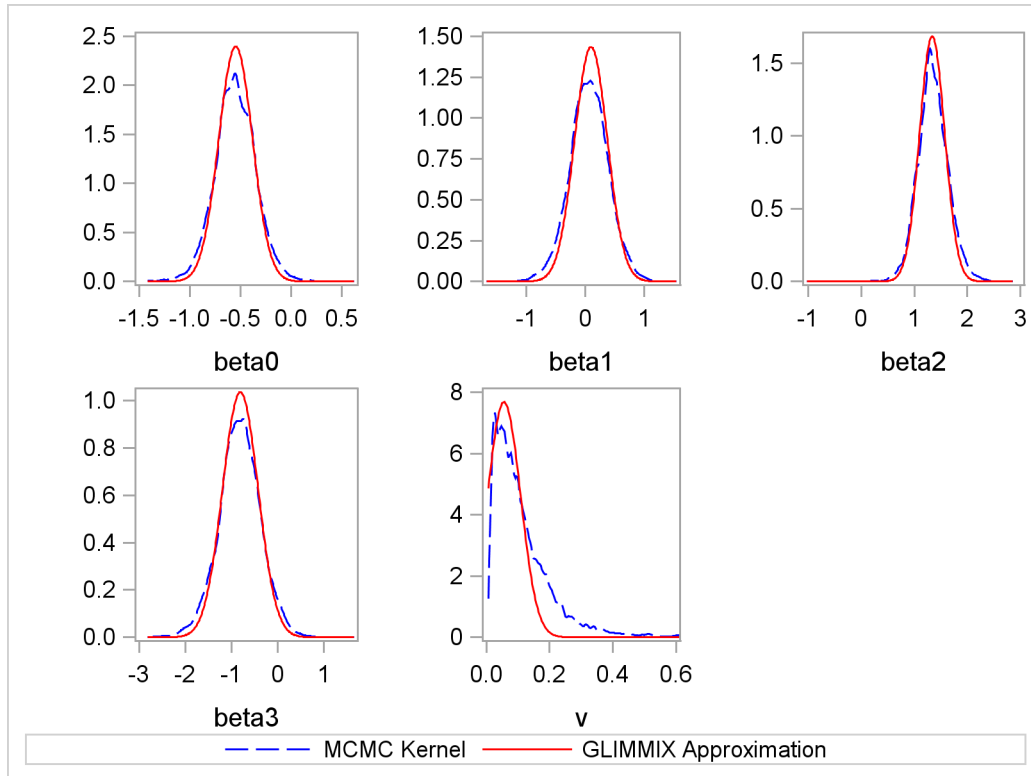
The kernel densities are very similar to each other. Kernel densities from PROC MCMC are not as smooth, possibly due to bad mixing of the Markov chains.

Nonlinear Poisson Regression Random-Effects Model

This example uses the pump failure data of Gaver and O'Muircheartaigh (1987). The number of failures and the time of operation are recorded for 10 pumps. Each of the pumps is classified into one of two groups corresponding to either continuous or intermittent operation. The following statements generate the data set:

```
title 'Nonlinear Poisson Regression Random Effects Model';
data pump;
   input y t group @@;
   pump = _n_;
   logtstd = log(t) - 2.4564900;
   datalines;
 5  94.320 1    1  15.720 2    5  62.880 1
14 125.760 1    3   5.240 2   19  31.440 1
 1   1.048 2    1   1.048 2    4   2.096 2
22  10.480 2
;
```

Each row denotes data for a single pump, and the variable logtstd contains the centered operation times. Letting y_{ij} denote the number of failures for the jth pump in the ith group, Draper (1996) considers the following hierarchical model for these data:

$$
\begin{aligned}
y_{ij}|\lambda_{ij} &\sim \text{Poisson}(\lambda_{ij}) \\
\log \lambda_{ij} &= \alpha_i + \beta_i (\log t_{ij} - \overline{\log t}) + e_{ij} \\
e_{ij}|\sigma^2 &\sim \text{normal}(0, \sigma^2)
\end{aligned}
$$

The model specifies different intercepts and slopes for each group, and the random effect is a mechanism for accounting for over-dispersion. You can use noninformative priors on the parameters α_i, β_i, and σ^2.

$$
\begin{aligned}
\pi(\alpha_1, \alpha_2, \beta_1, \beta_2) &\propto 1 \\
\pi(\sigma^2) &\propto 1/\sigma^2
\end{aligned}
$$

The following statements fit this nonlinear hierarchical model and produce Output 52.5.4:

```
proc mcmc data=pump outpost=postout seed=248601 nmc=100000
          ntu=2000 thin=10
          monitor=(logsig beta1 beta2 alpha1 alpha2 s2 adif bdif);
   ods select PostSummaries;
   array alpha[2];
   array beta[2];
   array llambda[10];

   parms (alpha: beta:) 1;
   parms llambda: 1;
   parms s2 1;

   beginprior;
   sd = sqrt(s2);
   logsig = log(s2)/2;
   adif = alpha1 - alpha2;
   bdif = beta1 - beta2;
   endprior;
   w = alpha[group] + beta[group] * logtstd;
   if pump eq 1 then
      lp = lpdfnorm(llambda[pump], w, sd);
   else
      lp = lp + lpdfnorm(llambda[pump], w, sd);
   prior alpha: beta: ~ general(1);
   prior s2 ~ general(-log(s2));
   prior llambda: ~ general(lp);

   lambda = exp(llambda[pump]);
   model y ~ poisson(lambda);
run;
```

The PROC MCMC statement specifies the input data set (pump), the output data set (postout), a seed for the random number generator, and an MCMC sample of 100000. It also requests a tuning

sample size of 2000 and a thinning rate of 10. The MONITOR= option keeps track of a number of parameters and symbols in the model. The five parameters are beta1, beta2, alpha1, alpha2, and s2. The symbol logsig is the log of the standard deviation, adif measures the difference between alpha1 and alpha2, and bdif measures the difference between beta1 and beta2. The **ods select** statement displays the summary statistics table.

Modeling the random effects e_{ij} with a normal distribution with mean 0 and variance σ^2 is equivalent to modeling $\log \lambda_{ij}$ with a normal distribution with mean $\alpha_i + \beta_i (\log t_{ij} - \overline{\log t})$ and variance σ^2. Here again, the prior distribution on $\log \lambda_{ij}$ depends on the data set variable logstd; hence, the construction of the prior has to take place before the PRIOR statement for $\log \lambda_{ij}$. The symbol lp keeps track of the cumulative log prior density for $\log \lambda_{ij}$.

The symbol lambda is the exponential of the corresponding $\log \lambda_{ij}$, and the MODEL statement gives the response variable y a Poisson likelihood with a mean parameter lambda.

The posterior summary statistics table is shown in Output 52.5.4.

Output 52.5.4 Summary Statistics for the Nonlinear Poisson Regression

```
            Nonlinear Poisson Regression Random Effects Model

                          The MCMC Procedure

                         Posterior Summaries
```

| | | | Standard | | Percentiles | |
Parameter	N	Mean	Deviation	25%	50%	75%
logsig	10000	0.1045	0.3862	-0.1563	0.0883	0.3496
beta1	10000	-0.4467	1.2818	-1.1641	-0.4421	0.2832
beta2	10000	0.5858	0.5808	0.2376	0.5796	0.9385
alpha1	10000	2.9719	2.3658	1.6225	2.9612	4.3137
alpha2	10000	1.6406	0.8674	1.1429	1.6782	2.1673
s2	10000	1.7004	1.7995	0.7316	1.1931	2.0121
adif	10000	1.3313	2.4934	-0.1670	1.2669	2.8032
bdif	10000	-1.0325	1.4186	-1.8379	-1.0284	-0.2189

Draper (1996) reports a posterior mean and standard deviation as follows: $\log \sigma = (0.28, 0.42)$, $\beta_1 = (-0.45, 1.5)$, $\beta_2 = (0.63, 0.68)$, and $\alpha 1 - \alpha 2 = (1.3, 3.0)$. Most estimates from Output 52.5.4 agree with Draper's estimates, with the exception of $\log \sigma$. The difference might be attributed to the different set of prior distributions on α_i, β_i, and σ that are used in this analysis.

You can also use PROC NLMIXED to fit the same model. The following statements run PROC NLMIXED and produce Output 52.5.5:

```
proc nlmixed data=pump;
    ods select parameterestimates additionalestimates;
    ods output additionalestimates=cp parameterestimates=ps;
    parms logsig 0 beta1 1 beta2 1 alpha1 1 alpha2 1;
    if (group = 1) then eta = alpha1 + beta1*logtstd + e;
    else eta = alpha2 + beta2*logtstd + e;
    lambda = exp(eta);
    model y ~ poisson(lambda);
    random e ~ normal(0,exp(2*logsig)) subject=pump;
    estimate 'adif' alpha1-alpha2;
    estimate 'bdif' beta1-beta2;
    estimate 's2' exp(2*logsig);
run;
```

Output 52.5.5 Estimates by PROC NLMIXED

```
               Nonlinear Poisson Regression Random Effects Model

                          The NLMIXED Procedure

                          Parameter Estimates

                   Standard
Parameter   Estimate   Error     DF   t Value   Pr > |t|   Alpha    Lower

logsig      -0.3161    0.3213    9    -0.98     0.3508     0.05    -1.0429
beta1       -0.4256    0.7473    9    -0.57     0.5829     0.05    -2.1162
beta2        0.6097    0.3814    9     1.60     0.1443     0.05    -0.2530
alpha1       2.9644    1.3826    9     2.14     0.0606     0.05    -0.1632
alpha2       1.7992    0.5492    9     3.28     0.0096     0.05     0.5568

                          Parameter Estimates

                   Parameter      Upper     Gradient

                   logsig        0.4107    -0.00002
                   beta1         1.2649    -0.00002
                   beta2         1.4724    -1.61E-6
                   alpha1        6.0921    -5.25E-6
                   alpha2        3.0415    -5.73E-6

                          Additional Estimates

                Standard
 Label  Estimate   Error     DF  t Value  Pr > |t|   Alpha    Lower    Upper

 adif    1.1653    1.4855    9    0.78     0.4529    0.05   -2.1952   4.5257
 bdif   -1.0354    0.8389    9   -1.23     0.2484    0.05   -2.9331   0.8623
 s2      0.5314    0.3415    9    1.56     0.1541    0.05   -0.2410   1.3038
```

Again, the point estimates from PROC NLMIXED for the mean parameters agree relatively closely with the Bayesian posterior means. You can note that there are differences in the likelihood-based

standard errors. This is most likely due to the fact that the Bayesian standard deviations account for the uncertainty in estimating σ^2, whereas the likelihood approach plugs in its estimated value.

You can do a similar kernel density plot that compares the PROC MCMC results, the PROC NLMIXED results and those reported by Draper. The following statements generate Output 52.5.6:

```
data nlmest(keep = parm mean sd);
   set ps cp;
   mean = estimate;
   sd = standarderror;
   if _n_ <= 5 then
      parm = parameter;
   else
      parm = label;
run;

data msd (keep=mean sd);
   set nlmest;
   do j = 1 to 401;
      output;
   end;
run;

data _null_;
   set ps;
   call symputx(compress('m' || parameter,'*'), estimate);
   call symputx(compress('s' || parameter,'*'), standarderror);
run;

data _null_;
   set cp;
   call symputx(compress('m' || label,'*'), estimate);
   call symputx(compress('s' || label,'*'), standarderror);
run;

%put &mlogsig &mbeta1 &mbeta2 &malpha1 &malpha2 &madif &mbdif &ms2;
%put &slogsig &sbeta1 &sbeta2 &salpha1 &salpha2 &sadif &sbdif &ss2;

proc kde data=postout;
   univar logsig beta1 beta2 alpha1 alpha2 adif bdif s2 / out=m1 (drop=count);
run;

%reshape(m1, mcmc, suffix1=, suffix2=_kde);

data all;
   set mcmc;
   logsig_nlm = pdf('normal', logsig, &mlogsig, &slogsig);
   alpha1_nlm = pdf('normal', alpha1, &malpha1, &salpha1);
   alpha2_nlm = pdf('normal', alpha2, &malpha2, &salpha2);
   beta1_nlm = pdf('normal', beta1, &mbeta1, &sbeta1);
   beta2_nlm = pdf('normal', beta2, &mbeta2, &sbeta2);
   adif_nlm = pdf('normal', adif, &madif, &sadif);
   bdif_nlm = pdf('normal', bdif, &mbdif, &sbdif);
```

```
       s2_nlm = pdf('normal', s2, &ms2, &ss2);
       logsig_draper = pdf('normal', logsig, 0.28, 0.42);
       beta1_draper = pdf('normal', beta1, -0.45, 1.5);
       beta2_draper = pdf('normal', beta2, 0.63, 0.68);
       adif_draper = pdf('normal', adif, 1.3, 3.0);
    run;


    proc template;
       define statgraph threebythree;
          %macro plot;
             begingraph;
                layout lattice / rows=3 columns=3;
                   layout overlay /yaxisopts=(label=" ");
                      seriesplot y=logsig_kde x=logsig
                         / connectorder=xaxis
                           lineattrs=(pattern=mediumdash color=blue)
                         legendlabel = "MCMC Kernel" name="MCMC";
                      seriesplot y=logsig_nlm x=logsig
                         / connectorder=xaxis lineattrs=(color=red)
                         legendlabel = "NLMIXED Approximation" name="NLMIXED";
                      seriesplot y=logsig_draper x=logsig
                         / connectorder=xaxis
                           lineattrs=(pattern=shortdash color=green)
                         legendlabel = "Draper (1996) Approximation" name="Draper";
                   endlayout;
                   %do i = 1 %to 2;
                      layout overlay /yaxisopts=(label=" ");
                         seriesplot y=alpha&i._kde x=alpha&i
                            / connectorder=xaxis
                              lineattrs=(pattern=mediumdash color=blue)
                            legendlabel = "MCMC Kernel" name="MCMC";
                         seriesplot y=alpha&i._nlm x=alpha&i
                            / connectorder=xaxis lineattrs=(color=red)
                            legendlabel = "NLMIXED Approximation" name="NLMIXED";
                      endlayout;
                   %end;
                   %do i = 1 %to 2;
                      layout overlay /yaxisopts=(label=" ");
                         seriesplot y=beta&i._kde x=beta&i
                            / connectorder=xaxis
                              lineattrs=(pattern=mediumdash color=blue)
                            legendlabel = "MCMC Kernel" name="MCMC";
                         seriesplot y=beta&i._nlm x=beta&i
                            / connectorder=xaxis lineattrs=(color=red)
                            legendlabel = "NLMIXED Approximation" name="NLMIXED";
                         seriesplot y=beta&i._draper x=beta&i
                            / connectorder=xaxis
                              lineattrs=(pattern=shortdash color=green)
                            legendlabel = "Draper (1996) Approximation" name="Draper";
                      endlayout;
                   %end;
                   layout overlay /yaxisopts=(label=" ");
                      seriesplot y=adif_kde x=adif
                         / connectorder=xaxis
```

```
                              lineattrs=(pattern=mediumdash color=blue)
                           legendlabel = "MCMC Kernel" name="MCMC";
                        seriesplot y=adif_nlm x=adif
                           / connectorder=xaxis lineattrs=(color=red)
                           legendlabel = "NLMIXED Approximation" name="NLMIXED";
                        seriesplot y=adif_draper x=adif
                           / connectorder=xaxis
                             lineattrs=(pattern=shortdash color=green)
                           legendlabel = "Draper (1996) Approximation" name="Draper";
                     endlayout;
                     layout overlay /yaxisopts=(label=" ");
                        seriesplot y=bdif_kde x=bdif
                           / connectorder=xaxis
                             lineattrs=(pattern=mediumdash color=blue)
                           legendlabel = "MCMC Kernel" name="MCMC";
                        seriesplot y=bdif_nlm x=bdif
                           / connectorder=xaxis lineattrs=(color=red)
                           legendlabel = "NLMIXED Approximation" name="NLMIXED";
                     endlayout;
                     layout overlay /yaxisopts=(label=" ")
                                     xaxisopts=(linearopts=(viewmin=0 viewmax=5));
                        seriesplot y=s2_kde x=s2
                           / connectorder=xaxis
                             lineattrs=(pattern=mediumdash color=blue)
                           legendlabel = "MCMC Kernel" name="MCMC";
                        seriesplot y=s2_nlm x=s2
                           / connectorder=xaxis lineattrs=(color=red)
                           legendlabel = "NLMIXED Approximation" name="NLMIXED";
                     endlayout;
                     Sidebar / align = bottom;
                        discretelegend "MCMC" "NLMIXED" "Draper";
                     endsidebar;
                  endlayout;
               endgraph;
            %mend; %plot;
         end;
      run;

      proc sgrender data=all template=threebythree;
      run;
```

The macro %RESHAPE is defined in the example "Logistic Regression Random-Effects Model" on page 3608.

Output 52.5.6 Comparing Estimates from PROC MCMC (dashed blue), PROC NLMIXED (solid red) and Draper (dotted green)
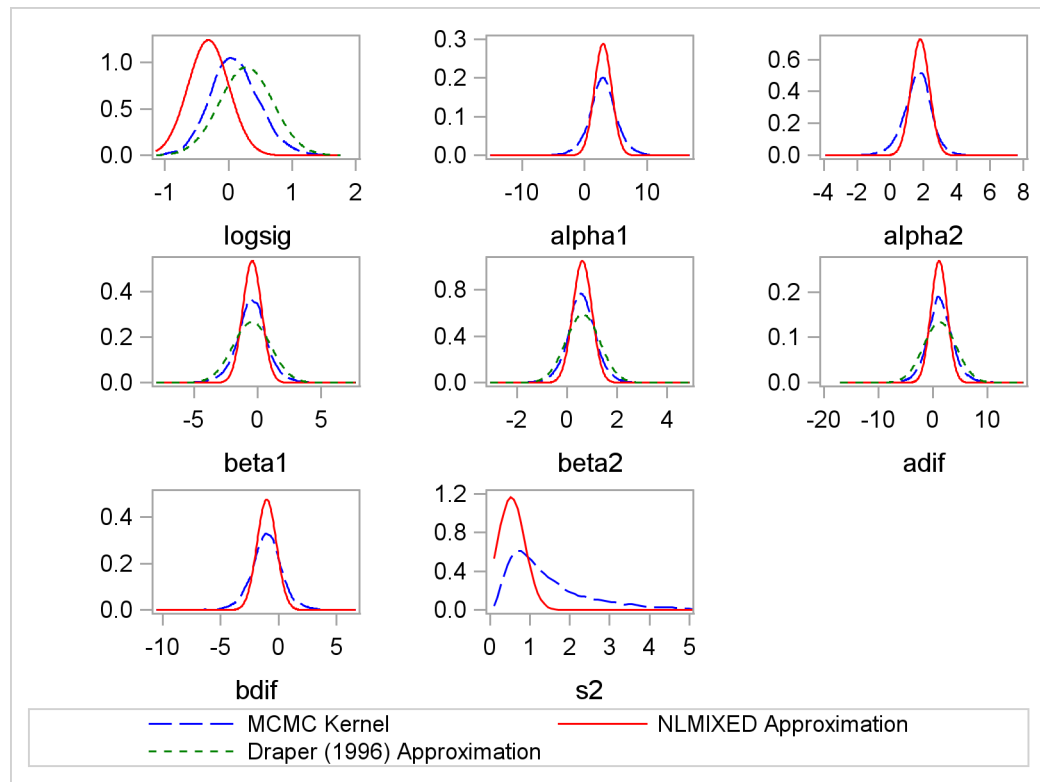
Example 52.6: Change Point Models

Consider the data set from Bacon and Watts (1971), where y_i is the logarithm of the height of the stagnant surface layer and the covariate x_i is the logarithm of the flow rate of water. The following statements create the data set:

```
title 'Change Point Model';
data stagnant;
   input y x @@;
   ind = _n_;
   datalines;
 1.12  -1.39   1.12  -1.39   0.99  -1.08   1.03  -1.08
 0.92  -0.94   0.90  -0.80   0.81  -0.63   0.83  -0.63
 0.65  -0.25   0.67  -0.25   0.60  -0.12   0.59  -0.12
 0.51   0.01   0.44   0.11   0.43   0.11   0.43   0.11
 0.33   0.25   0.30   0.25   0.25   0.34   0.24   0.34
 0.13   0.44  -0.01   0.59  -0.13   0.70  -0.14   0.70
-0.30   0.85  -0.33   0.85  -0.46   0.99  -0.43   0.99
-0.65   1.19
;
```

A scatter plot (Output 52.6.1) shows the presence of a nonconstant slope in the data. This suggests a change point regression model (Carlin, Gelfand, and Smith 1992). The following statements

generate the scatter plot in Output 52.6.1:

```
proc sgplot data=stagnant;
   scatter x=x y=y;
run;
```

Output 52.6.1 Scatter Plot of the Stagnant Data Set

Let the change point be cp. Following formulation by Spiegelhalter et al. (1996), the regression model is as follows:

$$y_i \sim \begin{cases} \text{normal}(\alpha + \beta_1(x_i - \text{cp}), \sigma^2) & \text{if } x_i < \text{cp} \\ \text{normal}(\alpha + \beta_2(x_i - \text{cp}), \sigma^2) & \text{if } x_i \geq \text{cp} \end{cases}$$

You might consider the following diffuse prior distributions:

$$\begin{aligned} \pi(\text{cp}) &\sim \text{uniform}(-1.3, 1.1) \\ \pi(\alpha, \beta_1, \beta_2) &\sim \text{normal}(0, \text{var} = 1e6) \\ \pi(\sigma^2) &\sim \text{uniform}(0, 5) \end{aligned}$$

The following statements generate Output 52.6.2:

```
proc mcmc data=stagnant outpost=postout seed=24860 ntu=1000
          nmc=20000;
   ods select PostSummaries;
   ods output PostSummaries=ds;

   array beta[2];
```

```
      parms alpha cp beta1 beta2;
      parms s2;

      prior cp ~ unif(-1.3, 1.1);
      prior s2  ~ uniform(0, 5);
      prior alpha beta:  ~ normal(0, v = 1e6);

      j = 1 + (x >= cp);
      mu = alpha + beta[j] * (x - cp);
      model y ~ normal(mu, var=s2);
   run;
```

The PROC MCMC statement specifies the input data set (stagnant), the output data set (postout), a random number seed, a tuning sample of 1000, and an MCMC sample of 20000. The `ods select` statement displays only the summary statistics table. The `ods output` statement saves the summary statistics table to the data set ds.

The ARRAY statement allocates an array of size 2 for the beta parameters. You can use beta1 and beta2 as parameter names without allocating an array, but having the array makes it easier to construct the likelihood function. The two PARMS statements put the five model parameters in two blocks. The three PRIOR statements specify the prior distributions for these parameters.

The symbol j indicates the segment component of the regression. When x is less than the change point, (x >= cp) returns 0 and j is assigned the value 1; if x is greater than or equal to the change point, (x >= cp) returns 1 and j is 2. The symbol mu is the mean for the jth segment, and beta[j] changes between the two regression coefficients depending on the segment component. The MODEL statement assigns the normal model to the response variable y.

Posterior summary statistics are shown in Output 52.6.2.

Output 52.6.2 MCMC Estimates of the Change Point Regression Model

```
                        Change Point Model

                        The MCMC Procedure

                        Posterior Summaries
```

			Standard		Percentiles	
Parameter	N	Mean	Deviation	25%	50%	75%
alpha	20000	0.5349	0.0249	0.5188	0.5341	0.5509
cp	20000	0.0283	0.0314	0.00728	0.0303	0.0493
beta1	20000	-0.4200	0.0146	-0.4293	-0.4198	-0.4111
beta2	20000	-1.0136	0.0167	-1.0248	-1.0136	-1.0023
s2	20000	0.000451	0.000145	0.000348	0.000425	0.000522

You can use PROC SGPLOT to visualize the model fit. Output 52.6.3 shows the fitted regression lines over the original data. In addition, on the bottom of the plot is the kernel density of the posterior marginal distribution of cp, the change point. The kernel density plot shows the relative variability of the posterior distribution on the data plot. You can use the following statements to create the plot:

```
data _null_;
   set ds;
   call symputx(parameter, mean);
run;

data b;
   missing A;
   input x1 @@;
   if x1 eq .A then x1 = &cp;
   if _n_ <= 2 then y1 = &alpha + &beta1 * (x1 - &cp);
   else y1 = &alpha + &beta2 * (x1 - &cp);
   datalines;
   -1.5 A 1.2
   ;

proc kde data=postout;
   univar cp / out=m1 (drop=count);
run;

data m1;
   set m1;
   density = (density / 25) - 0.653;
run;

data all;
   set stagnant b m1;
run;

proc sgplot data=all noautolegend;
   scatter x=x y=y;
   series x=x1 y=y1 / lineattrs = graphdata2;
   series x=value y=density / lineattrs = graphdata1;
run;
```

The macro variables &alpha, &beta1, &beta2, and &cp store the posterior mean estimates from the data set ds. The data set predicted contains three predicted values, at the minimum and maximum values of x and the estimated change point &cp. These input values give you fitted values from the regression model. Data set m1 contains the kernel density estimates of the parameter cp. The density is scaled down so the curve would fit in the plot. Finally, you use PROC SGPLOT to overlay the scatter plot, regression line and kernel density plots in the same graph.

Output 52.6.3 Estimated Fit to the Stagnant Data Set

Example 52.7: Exponential and Weibull Survival Analysis

This example covers two commonly used survival analysis models: the exponential model and the Weibull model. The deviance information criterion (DIC) is used to do model selections, and you can also find programs that visualize posterior quantities. Exponential and Weibull models are widely used for survival analysis. This example shows you how to use PROC MCMC to analyze the treatment effect for the E1684 melanoma clinical trial data. These data were collected to assess the effectiveness of using interferon alpha-2b in chemotherapeutic treatment of melanoma. The following statements create the data set:

```
data e1684;
   input t t_cen treatment @@;
   if t = . then do;
     t = t_cen;
     v = 0;
   end;
   else
     v = 1;
   ifn = treatment - 1;
   et = exp(t);
   lt = log(t);
   drop t_cen;
   datalines;
```

```
1.57808  0.00000  2   1.48219  0.00000  2     .           7.33425  1
2.23288  0.00000  1     .        9.38356  2   3.27671  0.00000  1
   .       9.64384  1   1.66575  0.00000  2   0.94247  0.00000  1

... more lines ...

3.39178  0.00000  1     .        4.36164  2     .        4.81918  2
;
```

The data set e1684 contains the following variables: t is the failure time that equals the censoring time whether the observation was censored, v indicates whether the observation is an actual failure time or a censoring time, treatment indicates two levels of treatments, and ifn indicates the use of interferon as a treatment. The variables et and lt are the exponential and logarithm transformation of the time t. The published data contains other potential covariates that are not listed here. This example concentrates on the effectiveness of the interferon treatment.

Exponential Survival Model

The density function for exponentially distributed survival times is as follows:

$$p(t_i | \lambda_i) = \lambda_i \exp(-\lambda_i t_i)$$

Note that this formulation of the exponential distribution is different from what is used in the SAS probability function PDF. The definition used in PDF for the exponential distributions is as follows:

$$p(t_i | \nu_i) = \frac{1}{\nu_i} \exp(-\frac{t_i}{\nu_i})$$

The relationship between λ and ν is as follows:

$$\lambda_i = \frac{1}{\nu_i}$$

The corresponding survival function, using the λ_i formulation, is as follows:

$$S(t_i | \lambda_i) = \exp(-\lambda_i t_i)$$

If you have a sample $\{t_i\}$ of n independent exponential survival times, each with mean λ_i, then the likelihood function in terms of λ is as follows:

$$
\begin{aligned}
L(\lambda | t) &= \Pi_{i=1}^{n} p(t_i | \lambda_i)^{\nu_i} S(t_i | \lambda_i)^{1-\nu_i} \\
&= \Pi_{i=1}^{n} (\lambda_i \exp(-\lambda_i t_i))^{\nu_i} (\exp(-\lambda_i t_i))^{1-\nu_i} \\
&= \Pi_{i=1}^{n} \lambda_i^{\nu_i} \exp(-\lambda_i t_i)
\end{aligned}
$$

If you link the covariates to λ with $\lambda_i = \exp x_i' \beta$, where x_i is the vector of covariates corresponding to the ith observation and β is a vector of regression coefficients, then the log-likelihood function is as follows:

$$l(\beta | t, x) = \sum_{i=1}^{n} \nu_i x_i' \beta - t_i \exp(x_i' \beta)$$

In the absence of prior information about the parameters in this model, you can choose diffuse normal priors for the β:

$$\beta \sim \text{normal}(0, \text{sd} = 10000)$$

There are two ways to program the log-likelihood function in PROC MCMC. You can use the SAS functions LOGPDF and LOGSDF. Alternatively, you can use the simplified log-likelihood function, which is more computationally efficient. You get identical results by using either approaches.

The following PROC MCMC statements fit an exponential model with simplified log-likelihood function:

```
title 'Exponential Survival Model';
ods graphics on;
proc mcmc data=e1684 outpost=expsurvout nmc=10000 seed=4861;
    ods select PostSummaries PostIntervals TADpanel
              ess mcse;
   parms (beta0 beta1) 0;
   prior beta: ~ normal(0, sd = 10000);
   /***************************************************/
   /* (1) the logpdf and logsdf functions are not used  */
   /***************************************************/
   /*      nu = 1/exp(beta0 + beta1*ifn);
          llike = v*logpdf("exponential", t, nu) +
                  (1-v)*logsdf("exponential", t, nu);
   */
   /***************************************************/
   /* (2) the simplified likelihood formula is used    */
   /***************************************************/
   l_h = beta0 + beta1*ifn;
   llike = v*(l_h) -  t*exp(l_h);
   model general(llike);
run;
ods graphics off;
```

The two assignment statements that are commented out calculate the log-likelihood function by using the SAS functions LOGPDF and LOGSDF for the exponential distribution. The next two assignment statements calculate the log likelihood by using the simplified formula. The first approach is slower because of the redundant calculation involved in calling both LOGPDF and LOGSDF.

An examination of the trace plots for β_0 and β_1 (see Output 52.7.1) reveals that the sampling has gone well with no particular concerns about the convergence or mixing of the chains.

Output 52.7.1 Posterior Plots for β_0 and β_1 in the Exponential Survival Analysis
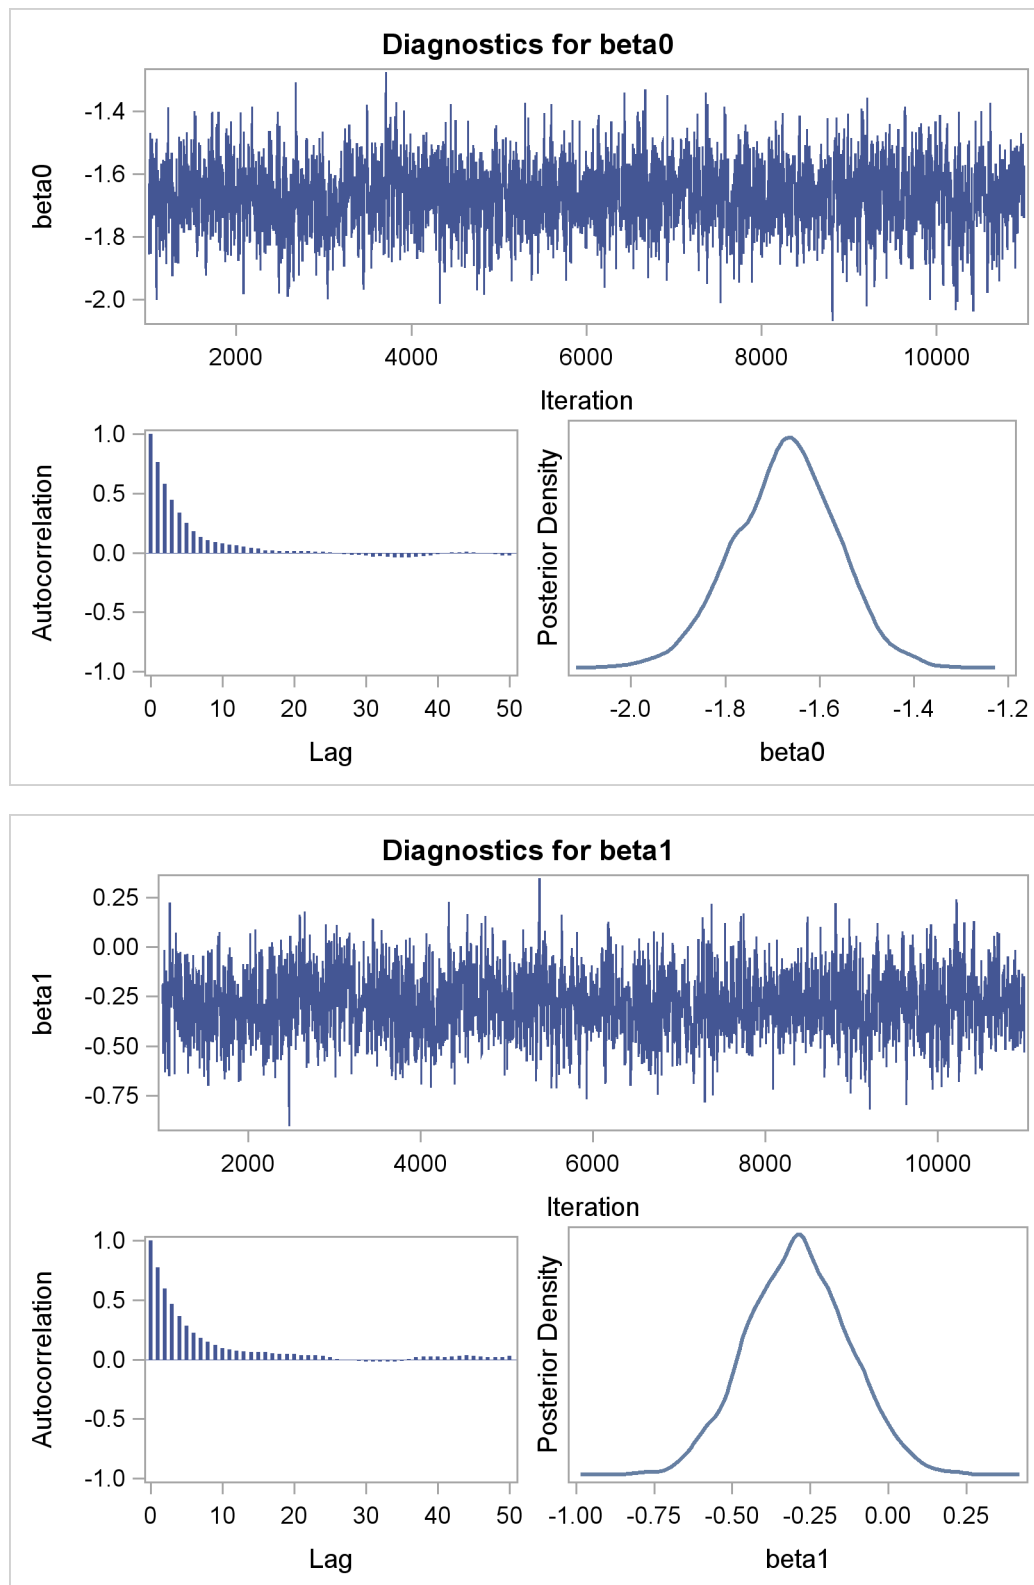
The MCMC results are shown in Output 52.7.2.

Output 52.7.2 Posterior Summary and Interval Statistics

```
                    Exponential Survival Model

                        The MCMC Procedure

                        Posterior Summaries

                                   Standard           Percentiles
   Parameter        N      Mean    Deviation     25%        50%        75%

   beta0         10000   -1.6715    0.1091    -1.7426    -1.6684    -1.5964
   beta1         10000   -0.2879    0.1615    -0.4001    -0.2892    -0.1803

                        Posterior Intervals

      Parameter    Alpha     Equal-Tail Interval       HPD Interval

      beta0        0.050      -1.8907    -1.4639    -1.8930    -1.4673
      beta1        0.050      -0.5985     0.0300    -0.6104     0.0169
```

The Monte Carlo standard errors and effective sample sizes are shown in Output 52.7.3. The posterior means for β_0 and β_1 are estimated with high precision, with small standard errors with respect to the standard deviation. This indicates that the mean estimates have stabilized and do not vary greatly in the course of the simulation. The effective sample sizes are roughly the same for both parameters.

Output 52.7.3 MCSE and ESS

```
                    Exponential Survival Model

                        The MCMC Procedure

                    Monte Carlo Standard Errors

                                     Standard
          Parameter        MCSE      Deviation     MCSE/SD

          beta0          0.00302      0.1091       0.0277
          beta1          0.00485      0.1615       0.0301

                    Effective Sample Sizes

                                   Correlation
          Parameter        ESS          Time      Efficiency

          beta0         1304.1        7.6682       0.1304
          beta1         1107.2        9.0319       0.1107
```

The next part of this example shows fitting a Weibull regression to the data and then comparing the two models with DIC to see which one provides a better fit to the data.

Weibull Survival Model

The density function for Weibull distributed survival times is as follows:

$$p(t_i | \alpha, \lambda_i) = \alpha t_i^{\alpha-1} \exp(\lambda_i - \exp(\lambda_i) t_i^{\alpha})$$

Note that this formulation of the Weibull distribution is different from what is used in the SAS probability function PDF. The definition used in PDF is as follows:

$$p(t_i | \alpha, \gamma_i) = \exp\left(-\left(\frac{t_i}{\gamma_i}\right)^{\alpha}\right) \frac{\alpha}{\gamma_i} \left(\frac{t_i}{\gamma_i}\right)^{\alpha-1}$$

The relationship between λ and γ in these two parameterizations is as follows:

$$\lambda_i = -\alpha \log \gamma_i$$

The corresponding survival function, using the λ_i formulation, is as follows:

$$S(t_i | \alpha, \lambda_i) = \exp(-\exp(\lambda_i) t_i^{\alpha})$$

If you have a sample $\{t_i\}$ of n independent Weibull survival times, with parameters α, and λ_i, then the likelihood function in terms of α and λ is as follows:

$$
\begin{aligned}
L(\alpha, \lambda | t) &= \Pi_{i=1}^{n} p(t_i | \alpha, \lambda_i)^{\nu_i} S(t_i | \alpha, \lambda_i)^{1-\nu_i} \\
&= \Pi_{i=1}^{n} (\alpha t_i^{\alpha-1} \exp(\lambda_i - \exp(\lambda_i) t_i^{\alpha}))^{\nu_i} (\exp(-\exp(\lambda_i) t_i^{\alpha}))^{1-\nu_i} \\
&= \Pi_{i=1}^{n} (\alpha t_i^{\alpha-1} \exp(\lambda_i))^{\nu_i} (\exp(-\exp(\lambda_i) t_i^{\alpha}))
\end{aligned}
$$

If you link the covariates to λ with $\lambda_i = x_i'\beta$, where x_i is the vector of covariates corresponding to the ith observation and β is a vector of regression coefficients, the log-likelihood function becomes this:

$$l(\alpha, \beta | t, x) = \sum_{i=1}^{n} \nu_i (\log(\alpha) + (\alpha - 1) \log(t_i) + x_i'\beta) - \exp(x_i'\beta) t_i^{\alpha})$$

As with the exponential model, in the absence of prior information about the parameters in this model, you can use diffuse normal priors on β. You might wish to choose a diffuse gamma distribution for α. Note that when $\alpha = 1$, the Weibull survival likelihood reduces to the exponential survival likelihood. Equivalently, by looking at the posterior distribution of α, you can conclude whether fitting an exponential survival model would be more appropriate than the Weibull model.

PROC MCMC also allows you to make inference on any functions of the parameters. Quantities of interest in survival analysis include the value of the survival function at specific times for specific treatments and the relationship between the survival curves for different treatments. With PROC MCMC, you can compute a sample from the posterior distribution of the interested survival functions at any number of points. The data in this example range from about 0 to 10 years, and the treatment of interest is the use of interferon.

Like in the previous exponential model example, there are two ways to fit this model: using the SAS functions LOGPDF and LOGSDF, or using the simplified log likelihood functions. The example uses the latter method. The following statements run PROC MCMC and produce Output 52.7.4:

```
title 'Weibull Survival Model';
proc mcmc data=e1684 outpost=weisurvout nmc=10000 seed=1234
          monitor=(_parms_ surv_ifn surv_noifn);
   ods select PostSummaries;
   ods output PostSummaries=ds PostIntervals=is;
   array surv_ifn[10];
   array surv_noifn[10];
   parms alpha 1 (beta0 beta1) 0;
   prior beta: ~ normal(0, var=10000);
   prior alpha ~ gamma(0.001,is=0.001);

   begincnst;
      surv_ifn1 = 1;
      surv_noifn1 = 1;
   endcnst;

   beginprior;
      do t = 1 to 10;
        surv_ifn[t] = exp(-exp(beta0+beta1)*t**alpha);
        surv_noifn[t] = exp(-exp(beta0)*t**alpha);
      end;
   endprior;

   lambda = beta0 + beta1*ifn;
   /***************************************************/
   /* (1) the logpdf and logsdf functions are not used  */
   /***************************************************/
   /*      gamma =  exp(-lambda /alpha);
           llike = v*logpdf('weibull', t, alpha, gamma) +
                   (1-v)*logsdf('weibull', t, alpha, gamma);
   */
   /***************************************************/
   /* (2) the simplified likelihood formula is used    */
   /***************************************************/
   llike  = v*(log(alpha) + (alpha-1)*log(t) + lambda) -
             exp(lambda)*(t**alpha);
   model general(llike);
run;
```

The MONITOR= option indicates the parameters and quantities of interest that PROC MCMC tracks. The symbol _PARMS_ specifies all model parameters. The array surv_ifn stores the expected survival probabilities for patients who received interferon over a period of 10 years. Similarly, surv_noifn stores the expected survival probabilities for patients who did not received interferon.

In the BEGINCNST and ENDCNST statements, two elements, one each from the allocated arrays of surv_inf and surv_noifn, are referenced once. This is needed for PROC MCMC to monitor the desired array elements correctly. If none of the array elements are referenced in the program, these arrays become temporary arrays with the variable name information stripped. As a result, PROC MCMC is no longer able to monitor the indicated symbols.

The BEGINPRIOR and ENDPRIOR statements enclose the calculations for the survival probabilities. The assignment statements proceeding the MODEL statement calculate the log likelihood for the Weibull survival model. The MODEL statement specifies the log likelihood that you pro-

grammed.

An examination of the trace plots for α, β_0, and β_1 (not displayed here) reveals that the sampling has gone well, with no particular concerns about the convergence or mixing of the chains.

Output 52.7.4 displays the posterior summary statistics.

Output 52.7.4 Posterior Summary Statistics

```
                    Weibull Survival Model

                      The MCMC Procedure

                     Posterior Summaries

                            Standard              Percentiles
Parameter          N      Mean    Deviation      25%       50%       75%

alpha           10000    0.7856    0.0533      0.7488    0.7849    0.8225
beta0           10000   -1.3414    0.1389     -1.4321   -1.3424   -1.2463
beta1           10000   -0.2918    0.1683     -0.4050   -0.2919   -0.1711
surv_ifn1       10000    0.8212    0.0237      0.8054    0.8228    0.8374
surv_ifn2       10000    0.7128    0.0308      0.6919    0.7138    0.7337
surv_ifn3       10000    0.6283    0.0352      0.6039    0.6282    0.6522
surv_ifn4       10000    0.5588    0.0383      0.5326    0.5589    0.5852
surv_ifn5       10000    0.5001    0.0405      0.4728    0.4997    0.5276
surv_ifn6       10000    0.4497    0.0420      0.4204    0.4489    0.4786
surv_ifn7       10000    0.4060    0.0431      0.3760    0.4051    0.4350
surv_ifn8       10000    0.3677    0.0437      0.3375    0.3664    0.3972
surv_ifn9       10000    0.3340    0.0440      0.3035    0.3325    0.3638
surv_ifn10      10000    0.3041    0.0440      0.2736    0.3024    0.3337
surv_noifn1     10000    0.7685    0.0280      0.7501    0.7701    0.7876
surv_noifn2     10000    0.6360    0.0349      0.6131    0.6376    0.6598
surv_noifn3     10000    0.5372    0.0386      0.5119    0.5384    0.5637
surv_noifn4     10000    0.4593    0.0407      0.4330    0.4599    0.4878
surv_noifn5     10000    0.3960    0.0417      0.3686    0.3959    0.4256
surv_noifn6     10000    0.3437    0.0421      0.3159    0.3432    0.3727
surv_noifn7     10000    0.2999    0.0419      0.2715    0.2993    0.3282
surv_noifn8     10000    0.2629    0.0412      0.2349    0.2624    0.2909
surv_noifn9     10000    0.2313    0.0403      0.2037    0.2302    0.2581
surv_noifn10    10000    0.2041    0.0392      0.1767    0.2033    0.2302
```

An examination of the α parameter reveals that the exponential model might not be inappropriate here. The estimated posterior mean of α is 0.7856 with a posterior standard deviation of 0.0533. As noted previously, if $\alpha = 1$, then the Weibull survival distribution is the exponential survival distribution. With these data, you can see that the evidence is in favor of $\alpha < 1$. The value 1 is almost 4 posterior standard deviations away from the posterior mean. The following statements compute the posterior probability of the hypothesis that $\alpha < 1$.:

```
proc format;
   value alphafmt low-<1 = 'alpha < 1' 1-high = 'alpha >= 1';
run;

proc freq data=weisurvout;
   tables alpha /nocum;
   format alpha alphafmt.;
run;
```

The PROC FREQ results are shown in Output 52.7.5.

Output 52.7.5 Frequency Analysis of α

```
                    Weibull Survival Model

                      The FREQ Procedure

               alpha     Frequency     Percent
          ---------------------------------------
          alpha < 1        10000        100.00
```

The output from PROC FREQ shows that 100% of the 10000 simulated values for α are less than 1. This is a very strong indication that the exponential model is too restrictive to model these data well.

You can examine the estimated survival probabilities over time individually, either through the posterior summary statistics or by looking at the kernel density plots. Alternatively, you might find it more informative to examine these quantities in relation with each other. For example, you can use a side-by-side box plot to display these posterior distributions by using PROC SGPLOT ("Statistical Graphics Using ODS" on page 497). First you need to take the posterior output data set weisurvout and stack variables that you want to plot. For example, to plot all the survival times for patients who received interferon, you want to stack surv_inf1–surv_inf10. The macro %Stackdata takes an input data set dataset, stacks the wanted variables vars, and outputs them into the output data set.

The following statements define the macro stackdata:

```
/* define macro stackdata */
%macro StackData(dataset,output,vars);
   data &output;
      length var $ 32;
      if 0 then set &dataset nobs=nnn;
      array lll[*] &vars;
      do jjj=1 to dim(lll);
         do iii=1 to nnn;
            set &dataset point=iii;
            value = lll[jjj];
            call vname(lll[jjj],var);
            output;
         end;
      end;
```

```
        stop;
        keep var value;
    run;
%mend;

/* stack the surv_ifn variables and saved them to survifn. */
%StackData(weisurvout, survifn, surv_ifn1-surv_ifn10);
```

Once you stack the data, use PROC SGPLOT to create the side-by-side box plots. The following statements generate Output 52.7.6:

```
proc sgplot data=survifn;
    yaxis label='Survival Probability' values=(0 to 1 by 0.2);
    xaxis label='Time' discreteorder=data;
    vbox value / category=var;
run;
```

Output 52.7.6 Side-by-Side Box Plots of Estimated Survival Probabilities

There is a clear decreasing trend over time of the survival probabilities for patients who receive the treatment. You might ask how does this group compare to those who did not receive the treatment? In this case, you want to overlay the two predicted curves for the two groups of patients and add the corresponding credible interval. See Output 52.7.7. To generate the graph, you first take the posterior mean estimates from the ODS output table ds and the lower and upper HPD interval estimates is, store them in the data set surv, and draw the figure by using PROC SGPLOT.

The following statements generate data set surv:

```
data surv;
   set ds;
   if _n_ >= 4 then do;
      set is point=_n_;
      group = 'with interferon   ';
      time = _n_ - 3;
      if time > 10 then do;
         time = time - 10;
         group = 'without interferon';
      end;
      output;
   end;
   keep time group mean hpdlower hpdupper;
run;
```

The following SGPLOT statements generate Output 52.7.7:

```
proc sgplot data=surv;
   yaxis label="Survival Probability" values=(0 to 1 by 0.2);
   series x=time y=mean / group = group name='i';
   band x=time lower=hpdlower upper=hpdupper / group = group transparency=0.7;
   keylegend 'i';
run;
```

In Output 52.7.7, the solid line is the survival curve for patients who received interferon; the shaded region centers at the solid line is the 95% HPD intervals; the medium-dashed line is the survival curve for patients who did not receive interferon; and the shaded region around the dashed line is the corresponding 95% HPD intervals.

Output 52.7.7 Predicted Survival Probability Curves with 95% HPD Intervals

The plot suggests that there is an effect of using interferon because patients who received interferon have sustained better survival probabilities than those who did not. However, the effect might not be very significant, as the 95% credible intervals of the two groups do overlap. For more on these interferon studies, refer to Ibrahim, Chen, and Sinha (2001).

Weibull or Exponential?

Although the evidence from the Weibull model fit shows that the posterior distribution of α has a significant amount of density mass less than 1, suggesting that the Weibull model is a better fit to the data than the exponential model, you might still be interested in comparing the two models more formally. You can use the Bayesian model selection criterion (see the section "Deviance Information Criterion (DIC)" on page 172) to determine which model fits the data better.

The PROC MCMC DIC option requests the calculation of DIC, and the procedure displays the ODS output table DIC. The table includes the posterior mean of the deviation, $\overline{D(\theta)}$, deviation at the estimate, $D(\overline{\theta})$, effective number of parameters, p_D, and DIC. It is important to remember that the standardizing term, $p(\mathbf{y})$, which is a function of the data alone, is not taken into account in calculating the DIC. This term is irrelevant only if you compare two models that have the *same* likelihood function. If you do not have identical likelihood functions, using DIC for model selection purposes without taking this standardizing term into account can produce incorrect results. In addition, you want to be careful in interpreting the DIC whenever you use the GENERAL function to construct the log-likelihood, as the case in this example. Using the GENERAL function, you can

obtain identical posterior samples with two log-likelihood functions that differ only by a constant. This difference translates to a difference in the DIC calculation, which could be very misleading.

If $\alpha = 1$, the Weibull likelihood is identical to the exponential likelihood. It is safe in this case to directly compare DICs from these two models. However, if you do not want to work out the mathematical detail or you are uncertain of the equivalence, a better way of comparing the DICs is to run the Weibull model twice: once with α being a parameter and once with $\alpha = 1$. This ensures that the likelihood functions are the same, and the DIC comparison is meaningful.

The following statements fit a Weibull model:

```
title 'Model Comparison between Weibull and Exponential';
proc mcmc data=e1684 outpost=weisurvout nmc=10000 seed=4861 dic;
   ods select dic;
   parms alpha 1 (beta0 beta1) 0;
   prior beta: ~ normal(0, var=10000);
   prior alpha ~ gamma(0.001,is=0.001);

   lambda = beta0 + beta1*ifn;
   llike  = v*(log(alpha) + (alpha-1)*log(t) + lambda) -
            exp(lambda)*(t**alpha);
   model general(llike);
run;
```

The DIC option requests the calculation of DIC, and the table is displayed is displayed in Output 52.7.8:

Output 52.7.8 DIC Table from the Weibull Model

```
           Model Comparison between Weibull and Exponential

                      The MCMC Procedure

                 Deviance Information Criterion

     Dbar (posterior mean of deviance)              858.623
     Dmean (deviance evaluated at posterior mean)   855.633
     pD (effective number of parameters)              2.990
     DIC (smaller is better)                        861.614

     The [D]GENERAL function is used in this program. To make
     meaningful comparisons, you must ensure that all
     [D]GENERAL functions include appropriate normalizing
     constants. Otherwise, DIC comparisons can be misleading.
```

The note in Output 52.7.8 reminds you of the importance of ensuring identical likelihood functions when you use the GENERAL function. The DIC value is 861.6.

Based on the same set of code, the following statements fit an exponential model by setting $\alpha = 1$:

```
proc mcmc data=e1684 outpost=expsurvout nmc=10000 seed=4861 dic;
   ods select dic;
   parms beta0 beta1 0;
   prior beta: ~ normal(0, var=10000);
   begincnst;
      alpha = 1;
   endcnst;

   lambda = beta0 + beta1*ifn;
   llike  = v*(log(alpha) + (alpha-1)*log(t) + lambda) -
            exp(lambda)*(t**alpha);
   model general(llike);
run;
```

Output 52.7.9 displays the DIC table.

Output 52.7.9 DIC Table from the Exponential Model

```
          Model Comparison between Weibull and Exponential

                        The MCMC Procedure

                   Deviance Information Criterion

      Dbar (posterior mean of deviance)            870.133
      Dmean (deviance evaluated at posterior mean)  868.190
      pD (effective number of parameters)            1.943
      DIC (smaller is better)                       872.075


      The [D]GENERAL function is used in this program. To make
      meaningful comparisons, you must ensure that all
      [D]GENERAL functions include appropriate normalizing
      constants. Otherwise, DIC comparisons can be misleading.
```

The DIC value of 872.075 is greater than 861. A smaller DIC indicates a better fit to the data; hence, you can conclude that the Weibull model is more appropriate for this data set. You can see the equivalencing of the exponential model you fitted in "Exponential Survival Model" on page 3629 by running the following comparison.

The following statements are taken from the section "Exponential Survival Model" on page 3629, and they fit the same exponential model:

```
proc mcmc data=e1684 outpost=expsurvout1 nmc=10000 seed=4861 dic;
   ods select none;
   parms (beta0 beta1) 0;
   prior beta: ~ normal(0, sd = 10000);
   l_h = beta0 + beta1*ifn;
   llike = v*(l_h) -  t*exp(l_h);
   model general(llike);
run;

proc compare data=expsurvout compare=expsurvout1;
   var beta0 beta1;
run;
```

The posterior samples of beta0 and beta1 in the data set expsurvout1 are identical to those in the data set expsurvout. The comparison results are not shown here.

Example 52.8: Cox Models

This example has two purposes. One is to illustrate how to use PROC MCMC to fit a Cox proportional hazard model. Specifically, two models are considered: time independent and time dependent models. However, note that it is much easier to fit a Bayesian Cox model by specifying the BAYES statement in PROC PHREG (see Chapter 64, "The PHREG Procedure"). If you are interested only in fitting a Cox regression survival model, you should use PROC PHREG.

The main purpose of this example is to demonstrate how to model data that are not independent. That is the case where the likelihood for observation i depends on other observations in the data set. In other words, if you work with a likelihood function that cannot be broken down simply as $L(\mathbf{y}) = \prod_i^n L(y_i)$, you can use this example for illustrative purposes. By default, PROC MCMC assumes that the programming statements and model specification is intended for a single row of observations in the data set. The Cox model is chosen because the complexity in the data structure requires more elaborate coding.

The Cox proportional hazard model is widely used in the analysis of survival time, failure time, or other duration data to explain the effect of exogenous explanatory variables. The data set used in this example is taken from Krall, Uthoff, and Harley (1975), who analyzed data from a study on myeloma in which researchers treated 65 patients with alkylating agents. Of those patients, 48 died during the study and 17 survived. The following statements generate the data set that is used in this example:

```
data Myeloma;
   input Time Vstatus LogBUN HGB Platelet Age LogWBC Frac
         LogPBM Protein SCalc;
   label Time='survival time'
         VStatus='0=alive 1=dead';
   datalines;
 1.25  1  2.2175   9.4  1  67  3.6628  1  1.9542  12  10
```

```
    1.25  1  1.9395   12.0  1  38  3.9868  1  1.9542  20  18
    2.00  1  1.5185    9.8  1  81  3.8751  1  2.0000   2  15

    ... more lines ...

    77.00  0  1.0792   14.0  1  60  3.6812  0  0.9542   0  12
    ;
proc sort data = Myeloma;
    by descending time;
run;

data _null_;
    set Myeloma nobs=_n;
    call symputx('N', _n);
    stop;
run;
```

The variable Time represents the survival time in months from diagnosis. The variable VStatus consists of two values, 0 and 1, indicating whether the patient was alive or dead, respectively, at the end of the study. If the value of VStatus is 0, the corresponding value of Time is censored. The variables thought to be related to survival are LogBUN (log(BUN) at diagnosis), HGB (hemoglobin at diagnosis), Platelet (platelets at diagnosis: 0=abnormal, 1=normal), Age (age at diagnosis in years), LogWBC (log(WBC) at diagnosis), Frac (fractures at diagnosis: 0=none, 1=present), LogPBM (log percentage of plasma cells in bone marrow), Protein (proteinuria at diagnosis), and SCalc (serum calcium at diagnosis). Interest lies in identifying important prognostic factors from these explanatory variables. In addition, there are 65 (&n) observations in the data set Myeloma. The likelihood used in these examples is the Brewslow likelihood:

$$L(\beta) = \prod_{i=1}^{n} \left[\prod_{j=1}^{d_i} \frac{\exp(\beta' Z_j(t_i))}{\sum_{l \in \mathcal{R}_i} \exp(\beta' Z_l(t_i))} \right]^{v_i}$$

where

- β is the vector parameters

- n is the total number of observations in the data set

- t_i is the ith time, which can be either event time or censored time

- $Z_l(t)$ is the vector explanatory variables for the lth individual at time t

- d_i is the multiplicity of failures at t_i. If there are no ties in time, d_i is 1 for all i.

- $\mathcal{R}_i$ is the risk set for the ith time t_i, which includes all observations that have survival time greater than or equal to t_i

- v_i indicates whether the patient is censored. The value 0 corresponds to censoring. Note that the censored time t_i enters the likelihood function only through the formation of the risk set $\mathcal{R}_i$.

Priors on the coefficients are independent normal priors with very large variance (1e6). Throughout this example, the symbol bZ represents the regression term $\beta' Z_j(t_i)$ in the likelihood, and the symbol S represents the term $\sum_{l \in \mathcal{R}_i} \exp(\beta' Z_l(t_i))$.

Time Independent Model

The regression model considered in this example uses the following formula:

$$\beta' Z_j \;=\; \beta_1 \text{logbun} + \beta_2 \text{hgb} + \beta_3 \text{platelet} + \beta_4 \text{age} +$$
$$\beta_5 \text{logwbc} + \beta_6 \text{frac} + \beta_7 \text{logpbm} + \beta_8 \text{protein} + \beta_9 \text{scalc}$$

The hard part of coding this in PROC MCMC is the construction of the risk set $\mathcal{R}_i$. $\mathcal{R}_i$ contains all observations that have survival time greater than or equal to t_i. First suppose that there are no ties in time. Sorting the data set by the variable time into descending order gives you $\mathcal{R}_i$ that is in the right order. Observation i's risk set consists of all data points j such that $j \;<= i$ in the data set. You can cumulatively increment S in the SAS statements.

With potential ties in time, at observation i, you need to know whether any subsequent observations, $i + 1$ and so on, have the same survival time as t_i. Suppose that the ith, the $i + 1$th, and the $i + 2$th observations all have the same survival time; all three of them need to be included in the risk set calculation. This means that to calculate the likelihood for some observations, you need to access both the previous and subsequent observations in the data set. There are two ways to do this. One is to use the LAG function; the other is to use the option JOINTMODEL.

The LAG function returns values from a queue (see *SAS Language Reference: Dictionary*). So for the ith observation, you can use LAG1 to access variables from the previous row in the data set. You want to compare the lag1 value of time with the current time value. Depending on whether the two time values are equal, you can add correction terms in the calculation for the risk set S.

The following statements sort the data set by time into descending order, with the largest survival time on top:

```
title 'Cox Model with Time Independent Covariates';
proc freq data=myeloma;
   ods select none;
   tables time / out=freqs;
run;

proc sort data = freqs;
   by descending time;
run;

data myelomaM;
   set myeloma;
   ind = _N_;
run;
```

The following statements run PROC MCMC and produce Output 52.8.1:

```
proc mcmc data=myelomaM outpost=outi nmc=50000 ntu=3000 seed=1;
   ods select PostSummaries PostIntervals;
   array beta[9];
   parms beta: 0;
   prior beta: ~ normal(0, var=1e6);
```

```
   bZ = beta1 * LogBUN + beta2 * HGB + beta3 * Platelet
      + beta4 * Age + beta5 * LogWBC + beta6 * Frac +
      beta7 * LogPBM + beta8 * Protein + beta9 * SCalc;

   if ind = 1 then do;          /* first observation        */
      S = exp(bZ);
      l = vstatus * bZ;
      v = vstatus;
   end;
   else if (1 < ind < &N) then do;
      if (lag1(time) ne time) then do;
         l = vstatus * bZ;
         l = l - v * log(S);  /* correct the loglike value  */
         v = vstatus;         /* reset v count value        */
         S = S + exp(bZ);
      end;
      else do;                 /* still a tie               */
         l = vstatus * bZ;
         S = S + exp(bZ);
         v = v + vstatus;      /* add # of nonsensored values */
      end;
   end;
   else do;                    /* last observation          */
      if (lag1(time) ne time) then do;
         l = - v * log(S);     /* correct the loglike value  */
         S = S + exp(bZ);
         l = l + vstatus * (bZ - log(S));
      end;
      else do;
         S = S + exp(bZ);
         l = vstatus * bZ - (v + vstatus) * log(S);
      end;
   end;
   model general(l);
run;
```

The symbol bZ is the regression term, which is independent of the time variable. The symbol ind indexes observation numbers in the data set. The symbol S keeps track of the risk set term for every observation. The symbol l calculates the log likelihood for each observation. Note that the value of l for observation ind is not necessarily the correct log likelihood value for that observation, especially in cases where the observation ind is in the tied times. Correction terms are added to subsequent values of l when the time variable becomes different in order to make up the difference. The total sum of l calculated over the entire data set is correct. The symbol v keeps track of the sum of vstatus, as censored data do not enter the likelihood and need to be taken out.

You use the function LAG1 to detect if two adjacent time values are different. If they are, you know that the current observation is in a different risk set than the last one. You then need to add a correction term to the log likelihood value of l. The IF-ELSE statements break the observations into three parts: the first observation, the last observation and everything in the middle.

Output 52.8.1 Summary Statistics on Cox Model with Time Independent Explanatory Variables and Ties in the Survival Time, Using PROC MCMC

```
                 Cox Model with Time Independent Covariates

                          The MCMC Procedure

                         Posterior Summaries

                                   Standard                Percentiles
 Parameter         N       Mean    Deviation       25%        50%        75%

 beta1          50000     1.7600     0.6441      1.3275     1.7651     2.1947
 beta2          50000    -0.1308     0.0720     -0.1799    -0.1304    -0.0817
 beta3          50000    -0.2017     0.5148     -0.5505    -0.1965     0.1351
 beta4          50000    -0.0126     0.0194     -0.0257    -0.0128    0.000641
 beta5          50000     0.3373     0.7256     -0.1318     0.3505     0.8236
 beta6          50000     0.3992     0.4337      0.0973     0.3864     0.6804
 beta7          50000     0.3749     0.4861      0.0464     0.3636     0.6989
 beta8          50000     0.0106     0.0271    -0.00723     0.0118     0.0293
 beta9          50000     0.1272     0.1064      0.0579     0.1300     0.1997


                          Posterior Intervals

     Parameter     Alpha     Equal-Tail Interval        HPD Interval

        beta1      0.050      0.4649      3.0214      0.5117      3.0465
        beta2      0.050     -0.2704      0.0114     -0.2746     0.00524
        beta3      0.050     -1.2180      0.8449     -1.2394      0.7984
        beta4      0.050     -0.0501      0.0257     -0.0512      0.0245
        beta5      0.050     -1.1233      1.7232     -1.1124      1.7291
        beta6      0.050     -0.4136      1.2970     -0.4385      1.2575
        beta7      0.050     -0.5551      1.3593     -0.5423      1.3689
        beta8      0.050     -0.0451      0.0618     -0.0451      0.0616
        beta9      0.050     -0.0933      0.3272     -0.0763      0.3406
```

An alternative to using the LAG function is to use the PROC option JOINTMODEL. With this option, the log-likelihood function you specify applies not to a single observation but to the entire data set. See "Modeling Dependent Data" on page 3549 for details on how to properly use this option. The basic idea is that you store all necessary data set variables in arrays and use only the arrays to construct the log likelihood of the entire data set. This approach works here because for every observation i, you can use index to access different values of arrays to construct the risk set S. To use the JOINTMODEL option, you need to do some additional data manipulation. You want to create a stop variable for each observation, which indicates the observation number that should be included in S for that observation. For example, if observations 4, 5, 6 all have the same survival time, the stop value for all of them is 6.

The following statements generate a new data set myelomaM that contains the stop variable:

```
data myelomaM;
   merge myelomaM freqs(drop=percent);
   by descending time;
   retain stop;
   if first.time then do;
      stop = _n_ + count - 1;
   end;
run;
```

The following SAS program fits the same Cox model by using the JOINTMODEL option:

```
proc mcmc data=myelomaM outpost=outa nmc=50000 ntu=3000 seed=1 jointmodel;
   ods select none;
   array beta[9];          array timeA[&n];        array vstatusA[&n];
   array logbunA[&n];      array hgbA[&n];         array plateletA[&n];
   array ageA[&n];         array logwbcA[&n];      array fracA[&n];
   array logpbmA[&n];      array proteinA[&n];     array scalcA[&n];
   array stopA[&n];        array bZ[&n];           array S[&n];

   begincnst;
      timeA[ind]=time;           vstatusA[ind]=vstatus;
      logbunA[ind]=logbun;       hgbA[ind]=hgb;
      plateletA[ind]=platelet;   ageA[ind]=age;
      logwbcA[ind]=logwbc;       fracA[ind]=frac;
      logpbmA[ind]=logpbm;       proteinA[ind]=protein;
      scalcA[ind]=scalc;         stopA[ind]=stop;
   endcnst;

   parms (beta:) 0;
   prior beta: ~ normal(0, var=1e6);

   jl = 0;

   /* calculate each bZ and cumulatively adding S as if there are no ties.*/
   do i = 1 to &n;
      bZ[i] = beta1*logbunA[i] + beta2*hgbA[i] + beta3*plateletA[i] +
              beta4*ageA[i] + beta5*logwbcA[i] + beta6*fracA[i] +
              beta7*logpbmA[i] + beta8*proteinA[i] + beta9*scalcA[i];
      if(i eq 1) then S[i] = exp(bZ[i]);
      else S[i] = S[i-1] + exp(bZ[i]);
   end;

   do i = 1 to &n;
      /* correct the S[i] term, when needed. */
      if(stopA[i] > i) then do;
         do j = (i+1) to stopA[i];
            S[i] = S[i] + exp(bZ[j]);
         end;
      end;
      jl = jl + vstatusA[i] * (bZ[i] - log(S[i]));
   end;
   model general(jl);
run;
```

No output tables were produced as this PROC MCMC run gives identical posterior samples as the previous example.

Multiple ARRAY statements allocate array symbols that are used to store the parameters (beta), the response (timeA), the covariates (vstatusA, logbunA, hgbA, plateletA, ageA, logwbcA, fracA, logpbmA, proteinA, scalcA, and stopA), and the work space (bZ and S). The bZ and S arrays store the regression term and the risk set term for every observation. The BEGINCNST and ENDCNST statements enclose programming statements that read the data set variables into these arrays. This was explained in the section "BEGINCNST/ENDCNST Statement" on page 3507. The rest of the programming statements construct the log likelihood for the entire data set.

You use two DO-loops to loop through the array-based data set and sum up observational level log likelihood. In the first DO-loop, you calculate bZ for each observation and cumulatively add the risk set term S as if there are no ties in time. This underevaluates some of the S elements. For those observations that have a tied time, you make the necessary correction to the corresponding S values. The correction takes place in the second DO-loop. Any observation that has a tied time also has a stopA[i] that is different from i. You add the right terms to S and sum up the joint log likelihood jl. The MODEL statement specifies that the log likelihood takes on the value of jl. Note that there is no data set variable outside of the BEGINCNST and ENDCNST statements in this example. For the JOINTMODEL to work correctly, PROC MCMC does not access the data set during the entire simulation.

To see that you get identical results from these two approaches, use PROC COMPARE to compare the posterior samples from two runs:

```
proc compare data=outi compare=outa;
   ods select comparesummary;
   var beta1-beta9;
run;
```

The output is omitted.

Generally speaking, the JOINTMODEL option can be slightly faster than using the default setup. The savings come from avoiding the overhead cost of accessing the data set repeatedly at every iteration. Note, however, that the speed gain is not guaranteed because it largely depends on the efficiency of your programs.

PROC PHREG fits the same model, and you get very similar results to PROC MCMC. The following statements run PROC PHREG and produce Output 52.8.2:

```
proc phreg data=Myeloma;
   ods select PostSummaries PostIntervals;
   ods output posteriorsample = phout;
   model Time*VStatus(0)=LogBUN HGB Platelet Age LogWBC
                         Frac LogPBM Protein Scalc;
   bayes seed=1 nmc=10000;
run;
```

Output 52.8.2 Summary Statistics for Cox Model with Time Independent Explanatory Variables and Ties in the Survival Time, Using PROC PHREG

```
                   Cox Model with Time Independent Covariates

                            The PHREG Procedure

                            Bayesian Analysis

                            Posterior Summaries

                                  Standard                 Percentiles
   Parameter          N       Mean    Deviation        25%        50%        75%

   LogBUN          10000     1.7610      0.6593     1.3173     1.7686     2.2109
   HGB             10000    -0.1279      0.0727    -0.1767    -0.1287    -0.0789
   Platelet        10000    -0.2179      0.5169    -0.5659    -0.2360     0.1272
   Age             10000    -0.0130      0.0199    -0.0264    -0.0131   0.000492
   LogWBC          10000     0.3150      0.7451    -0.1718     0.3321     0.8253
   Frac            10000     0.3766      0.4152     0.0881     0.3615     0.6471
   LogPBM          10000     0.3792      0.4909     0.0405     0.3766     0.7023
   Protein         10000     0.0102      0.0267   -0.00745     0.0106     0.0283
   SCalc           10000     0.1248      0.1062     0.0545     0.1273     0.1985


                            Posterior Intervals

      Parameter     Alpha     Equal-Tail Interval          HPD Interval

      LogBUN        0.050        0.4418     3.0477      0.4107      2.9958
      HGB           0.050       -0.2718     0.0150     -0.2801     0.00599
      Platelet      0.050       -1.1952     0.8296     -1.1871      0.8341
      Age           0.050       -0.0514     0.0259     -0.0519      0.0251
      LogWBC        0.050       -1.2058     1.7228     -1.1783      1.7483
      Frac          0.050       -0.3995     1.2316     -0.4273      1.2021
      LogPBM        0.050       -0.5652     1.3671     -0.5939      1.3241
      Protein       0.050       -0.0437     0.0611     -0.0405      0.0637
      SCalc         0.050       -0.0935     0.3264     -0.0846      0.3322
```

Output 52.8.3 shows kernel density plots that compare the posterior marginal distributions of all the beta parameters from the PROC MCMC run and the PROC PHREG run. The following statements generate the comparison:

```
proc kde data=outi;
   ods exclude all;
   univar beta1 beta2 beta3 beta4 beta5 beta6 beta7 beta8 beta9
      / out=m1 (drop=count);
run;
ods exclude none;

%reshape(m1, mcmc, suffix1=, suffix2=md);

data phout;
   set phout(drop = LogPost Iteration);
   beta1 = LogBUN;  beta2 = HGB;      beta3 = Platelet;
```

```
    beta4 = Age;      beta5 = LogWBC;   beta6 = Frac;
    beta7 = LogPBM;  beta8 = Protein;  beta9 = SCalc;
    drop LogBUN HGB Platelet Age LogWBC Frac LogPBM Protein SCalc;
run;

proc kde data=phout;
   ods exclude all;
   univar beta1 beta2 beta3 beta4 beta5 beta6 beta7 beta8 beta9
      / out=m2 (drop=count);
run;
ods exclude none;

%reshape(m2, phreg, suffix1=p, suffix2=pd);

data all;
   merge mcmc phreg;
run;

proc template;
   define statgraph threebythree;
      %macro plot;
         begingraph;
            layout lattice / rows=3 columns=3;
               %do i = 1 %to 9;
                  layout overlay /yaxisopts=(label=" ");
                     seriesplot y=beta&i.md x=beta&i
                        / connectorder=xaxis
                          lineattrs=(pattern=mediumdash color=blue)
                        legendlabel = "MCMC" name="MCMC";
                     seriesplot y=beta&i.pd x=beta&i.p
                        / connectorder=xaxis lineattrs=(color=red)
                        legendlabel = "PHREG" name="PHREG";
                  endlayout;
               %end;
               Sidebar / align = bottom;
                  discretelegend "MCMC" "PHREG";
               endsidebar;
            endlayout;
         endgraph;
      %mend; %plot;
   end;
run;

proc sgrender data=all template=threebythree;
   title "Kernel Density Comparison";
run;
```

The macro %RESHAPE is defined in the example "Logistic Regression Random-Effects Model" on page 3608. The posterior densities are almost identical to one another.

Output 52.8.3 Comparing Estimates from PROC MCMC and PROC PHREG
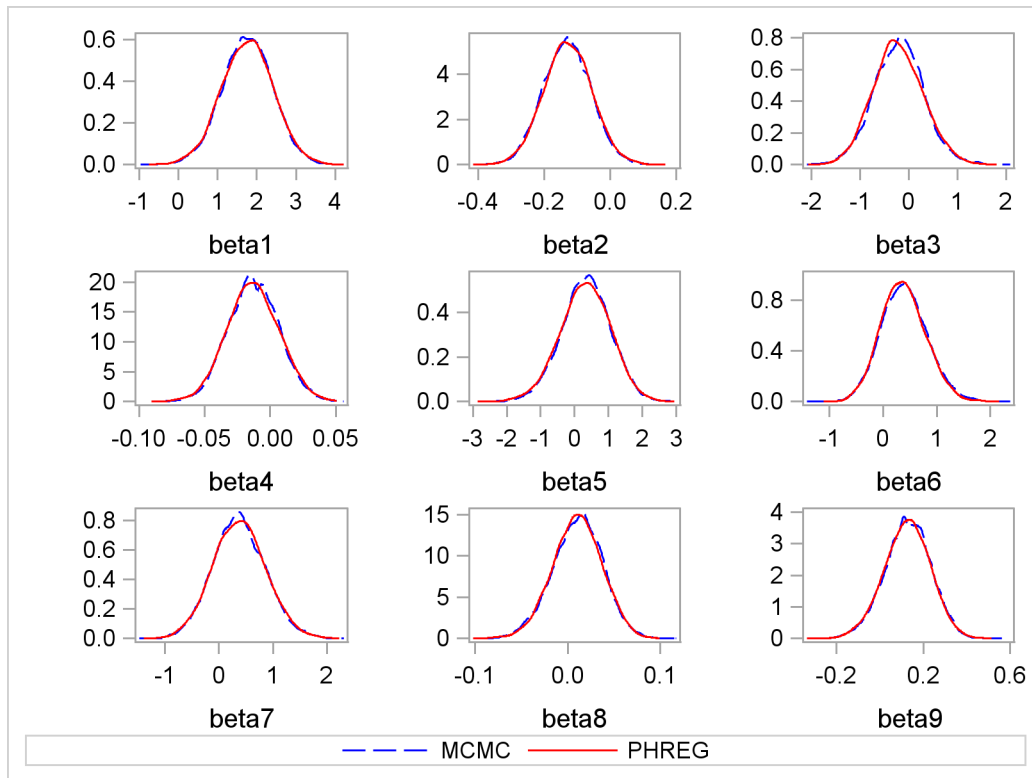
Time Dependent Model

To model $Z_i(t_i)$ as a function of the survival time, you can relate time t_i to covariates by using this formula:

$$\beta'Z_j(t_i) = (\beta_1 + \beta_2 t_i)\text{logbun} + (\beta_3 + \beta_4 t_i)\text{hgb} + (\beta_5 + \beta_6 t_i)\text{platelet}$$

For illustrational purposes, only three explanatory variables, LOGBUN, HBG, and PLATELET, are used in this example.

Since $Z_j(t_i)$ depends on t_i, every term in the summation of $\sum_{l \in \mathcal{R}_i} \exp(\beta'Z_l(t_i))$ is a product of the current time t_i and all observations that are in the risk set. You can use the JOINTMODEL option, as in the last example, or you can modify the input data set such that every row contains not only the current observation but also all observations that are in the corresponding risk set. When you construct the log likelihood for each observation, you have all the relevant data at your disposal.

The following statements illustrate how you can create a new data set with different risk sets at different rows:

```
title 'Cox Model with Time Dependent Covariates';
ods select none;
proc freq data=myeloma;
   tables time / out=freqs;
run;
ods select all;

proc sort data = freqs;
   by descending time;
run;

data myelomaM;
   set myeloma;
   ind = _N_;
run;

data myelomaM;
   merge myelomaM freqs(drop=percent); by descending time;
   retain stop;
   if first.time then do;
      stop = _n_ + count - 1;
   end;
run;

%macro array(list);
   %global mcmcarray;
   %let mcmcarray = ;
   %do i = 1 %to 32000;
      %let v = %scan(&list, &i, %str( ));
      %if %nrbquote(&v) ne %then %do;
         array _&v[&n];
         %let mcmcarray = &mcmcarray array _&v[&n] _&v.1 - _&v.&n%str(;);
            do i = 1 to stop;
               set myelomaM(keep=&v) point=i;
               _&v[i] = &v;
            end;
      %end;
      %else %let i = 32001;
   %end;
%mend;

data z;
   set myelomaM;
   %array(logbun hgb platelet);
   drop vstatus logbun hgb platelet count stop;
run;

data myelomaM;
   merge myelomaM z; by descending time;
run;
```

The data set myelomaM contains 65 observations and 209 variables. For each observation, you see added variables stop, _logbun1 through _logbun65, _hgb1 through _hgb65, and _platelet1 through _platelet65. The variable stop indicates the number of observations that are in the risk set of the current observation. The rest are transposed values of model covariates of the entire data set. The data set contains a number of missing values. This is due to the fact that only the relevant observations are kept, such as _logbun*1* to _logbun*stop*. The rest of the cells are filled in with missing values. For example, the first observation has a unique survival time of 92 and stop is 1, making it a risk set of itself. You see nonmissing values only in _logbun1, _hgb1, and _platelet1.

The following statements fit the Cox model by using PROC MCMC:

```
proc mcmc data=myelomaM outpost=outi nmc=50000 ntu=3000 seed=17
          missing=ac;
   ods select PostSummaries PostIntervals;
   array beta[6];
   &mcmcarray
   parms (beta:) 0;
   prior beta: ~ normal(0, prec=1e-6);

   b = (beta1 + beta2 * time) * logbun +
       (beta3 + beta4 * time) * hgb +
       (beta5 + beta6 * time) * platelet;
   S = 0;
   do i = 1 to stop;
      S = S + exp( (beta1 + beta2 * time) * _logbun[i] +
                   (beta3 + beta4 * time) * _hgb[i] +
                   (beta5 + beta6 * time) * _platelet[i]);
   end;
   loglike = vstatus * (b - log(S));

   model general(loglike);
run;
```

Note that the option MISSING= is set to AC. This is due to missing cells in the input data set. You must use this option so that PROC MCMC retains observations that contain missing values.

The macro variable &mcmcarray is defined in the earlier part in this example. You can use a **%put** statement to print its value:

```
%put &mcmcarray;
```

This statement prints the following:

```
array _logbun[65] _logbun1 - _logbun65; array _hgb[65] _hgb1 - _hgb65; array
_platelet[65] _platelet1 - _platelet65;
```

The macro uses the ARRAY statement to allocate three arrays, each of which links their corresponding data set variables. This makes it easier to reference these data set variables in the program. The PARMS statement puts all the parameters in the same block. The PRIOR statement gives them normal priors with large variance. The symbol b is the regression term, and S is cumulatively added from 1 to stop for each observation in the DO-loop. The symbol loglike completes the construction of log likelihood for each observation and the MODEL statement completes the model specification.

Posterior summary and interval statistics are shown in Output 52.8.4.

Output 52.8.4 Summary Statistics on Cox Model with Time Dependent Explanatory Variables and Ties in the Survival Time, Using PROC MCMC

```
                   Cox Model with Time Dependent Covariates

                           The MCMC Procedure

                           Posterior Summaries

                                Standard              Percentiles
   Parameter        N        Mean   Deviation      25%        50%        75%

   beta1         50000     3.2397      0.8226    2.6835     3.2413     3.7830
   beta2         50000    -0.1411      0.0471   -0.1722    -0.1406    -0.1092
   beta3         50000    -0.0369      0.1017   -0.1064    -0.0373     0.0315
   beta4         50000    -0.00409     0.00360  -0.00656   -0.00408   -0.00167
   beta5         50000     0.3548      0.7359   -0.1634     0.3530     0.8445
   beta6         50000    -0.0417      0.0359   -0.0661    -0.0423    -0.0181


                           Posterior Intervals

      Parameter    Alpha    Equal-Tail Interval        HPD Interval

      beta1        0.050      1.6399      4.8667      1.6664      4.8752
      beta2        0.050     -0.2351     -0.0509     -0.2294     -0.0458
      beta3        0.050     -0.2337      0.1642     -0.2272      0.1685
      beta4        0.050     -0.0111      0.00282    -0.0112      0.00264
      beta5        0.050     -1.0317      1.8202     -1.0394      1.8100
      beta6        0.050     -0.1107      0.0295     -0.1122      0.0269
```

You can also use the option JOINTMODEL to get the same inference and avoid transposing the data for every observation:

```
proc mcmc data=myelomaM outpost=outa nmc=50000 ntu=3000 seed=17 jointmodel;
   ods select none;
   array beta[6];       array timeA[&n];     array vstatusA[&n];
   array logbunA[&n];   array hgbA[&n];      array plateletA[&n];
   array stopA[&n];     array bZ[&n];        array S[&n];

   begincnst;
      timeA[ind]=time;          vstatusA[ind]=vstatus;
      logbunA[ind]=logbun;      hgbA[ind]=hgb;
      plateletA[ind]=platelet;  stopA[ind]=stop;
   endcnst;

   parms (beta:) 0;
   prior beta: ~ normal(0, prec=1e-6);

   jl = 0;
   do i = 1 to &n;
      v1 = beta1 + beta2 * timeA[i];
      v2 = beta3 + beta4 * timeA[i];
```

```
            v3 = beta5 + beta6 * timeA[i];
            bZ[i] = v1 * logbunA[i] + v2 * hgbA[i] + v3 * plateletA[i];

            /* sum over risk set without considering ties in time. */
            S[i] = exp(bZ[i]);
            if (i > 1) then do;
               do j = 1 to (i-1);
                  b1 = v1 * logbunA[j] + v2 * hgbA[j] + v3 * plateletA[j];
                  S[i] = S[i] + exp(b1);
               end;
            end;
         end;

         /* make correction to the risk set due to ties in time. */
         do i = 1 to &n;
            if(stopA[i] > i) then do;
               v1 = beta1 + beta2 * timeA[i];
               v2 = beta3 + beta4 * timeA[i];
               v3 = beta5 + beta6 * timeA[i];
               do j = (i+1) to stopA[i];
                  b1 = v1 * logbunA[j] + v2 * hgbA[j] + v3 * plateletA[j];
                  S[i] = S[i] + exp(b1);
               end;
            end;
            jl = jl + vstatusA[i] * (bZ[i] - log(S[i]));
         end;
         model general(jl);
      run;
```

The multiple ARRAY statements allocate array symbols that are used to store the parameters (beta), the response (timeA), the covariates (vstatusA, logbunA, hgbA, plateletA, and stopA), and work space (bZ and S). The bZ and S arrays store the regression term and the risk set term for every observation. Programming statements in the BEGINCNST and ENDCNST statements input the response and covariates from the data set to the arrays.

Using the same technique shown in the example "Time Independent Model" on page 3644, the next DO-loop calculates the regression term and corresponding S for every observation, pretending that there are no ties in time. This means that the risk set for observation i involves only observation 1 to i. The correction terms are added to the corresponding S[i] in the second DO-loop, conditional on whether the stop variable is greater than the observation count itself. The symbol jl cumulatively adds the log likelihood for the entire data set, and the MODEL statement specifies the joint log-likelihood function.

The following statements run PROC COMPARE and show that the output data set outa contains identical posterior samples as outi:

```
   proc compare data=outi compare=outa;
      ods select comparesummary;
      var beta1-beta6;
   run;
```

The results are not shown here.

The following statements use PROC PHREG to fit the same time dependent Cox model:

```
proc phreg data=Myeloma;
   ods select PostSummaries PostIntervals;
   ods output posteriorsample = phout;
   model Time*VStatus(0)=LogBUN z2 hgb z3 platelet z4;
   z2 = Time*logbun;
   z3 = Time*hgb;
   z4 = Time*platelet;
   bayes seed=1 nmc=10000;
run;
```

Coding is simpler than PROC MCMC. See Output 52.8.5 for posterior summary and interval statistics:

Output 52.8.5 Summary Statistics on Cox Model with Time Dependent Explanatory Variables and Ties in the Survival Time, Using PROC PHREG

```
                   Cox Model with Time Dependent Covariates

                           The PHREG Procedure

                            Bayesian Analysis

                           Posterior Summaries

                                  Standard              Percentiles
    Parameter        N       Mean    Deviation      25%       50%       75%

    LogBUN        10000     3.2423      0.8311     2.6838    3.2445    3.7929
    z2            10000    -0.1401      0.0482    -0.1723   -0.1391   -0.1069
    HGB           10000    -0.0382      0.1009    -0.1067   -0.0385    0.0297
    z3            10000    -0.00407     0.00363   -0.00652  -0.00404  -0.00162
    Platelet      10000     0.3778      0.7524    -0.1500    0.3389    0.8701
    z4            10000    -0.0419      0.0364    -0.0660   -0.0425   -0.0178

                           Posterior Intervals

       Parameter     Alpha      Equal-Tail Interval         HPD Interval

       LogBUN        0.050       1.6059     4.8785       1.5925     4.8582
       z2            0.050      -0.2361    -0.0494      -0.2354    -0.0492
       HGB           0.050      -0.2343     0.1598      -0.2331     0.1603
       z3            0.050      -0.0113     0.00297     -0.0109     0.00322
       Platelet      0.050      -0.9966     1.9464      -1.1342     1.7968
       z4            0.050      -0.1124     0.0296      -0.1142     0.0274
```

Output 52.8.6 shows a kernel density comparison plot that compares posterior marginal distributions of all the beta parameters from the PROC MCMC run and the PROC PHREG run. The following statements generate Output 52.8.6:

```
proc kde data=outi;
   ods exclude all;
   univar beta1 beta2 beta3 beta4 beta5 beta6 / out=m1 (drop=count);
run;
ods exclude none;
```

```
%reshape(m1, mcmc, suffix1=, suffix2=md);

data phout;
   set phout(drop = LogPost Iteration);
   beta1 = LogBUN;  beta2 = z2;        beta3 = HGB;
   beta4 = z3;       beta5 = Platelet; beta6 = z4;
   drop LogBUN HGB Platelet z2-z4;
run;

proc kde data=phout;
   ods exclude all;
   univar beta1 beta2 beta3 beta4 beta5 beta6 / out=m2 (drop=count);
run;
ods exclude none;

%reshape(m2, phreg, suffix1=p, suffix2=pd);

data all;
   merge mcmc phreg;
run;

proc template;
   define statgraph twobythree;
      %macro plot;
         begingraph;
            layout lattice / rows=2 columns=3;
               %do i = 1 %to 6;
                  layout overlay /yaxisopts=(label=" ");
                     seriesplot y=beta&i.md x=beta&i
                        / connectorder=xaxis
                          lineattrs=(pattern=mediumdash color=blue)
                        legendlabel = "MCMC" name="MCMC";
                     seriesplot y=beta&i.pd x=beta&i.p
                        / connectorder=xaxis lineattrs=(color=red)
                        legendlabel = "PHREG" name="PHREG";
                  endlayout;
               %end;
               Sidebar / align = bottom;
                  discretelegend "MCMC" "PHREG";
               endsidebar;
            endlayout;
         endgraph;
      %mend; %plot;
   end;
run;

proc sgrender data=all template=twobythree;
   title "Kernel Density Comparison";
run;
```

The macro %RESHAPE is defined in the example "Logistic Regression Random-Effects Model"
on page 3608.

Output 52.8.6 Comparing Estimates from PROC MCMC and PROC PHREG
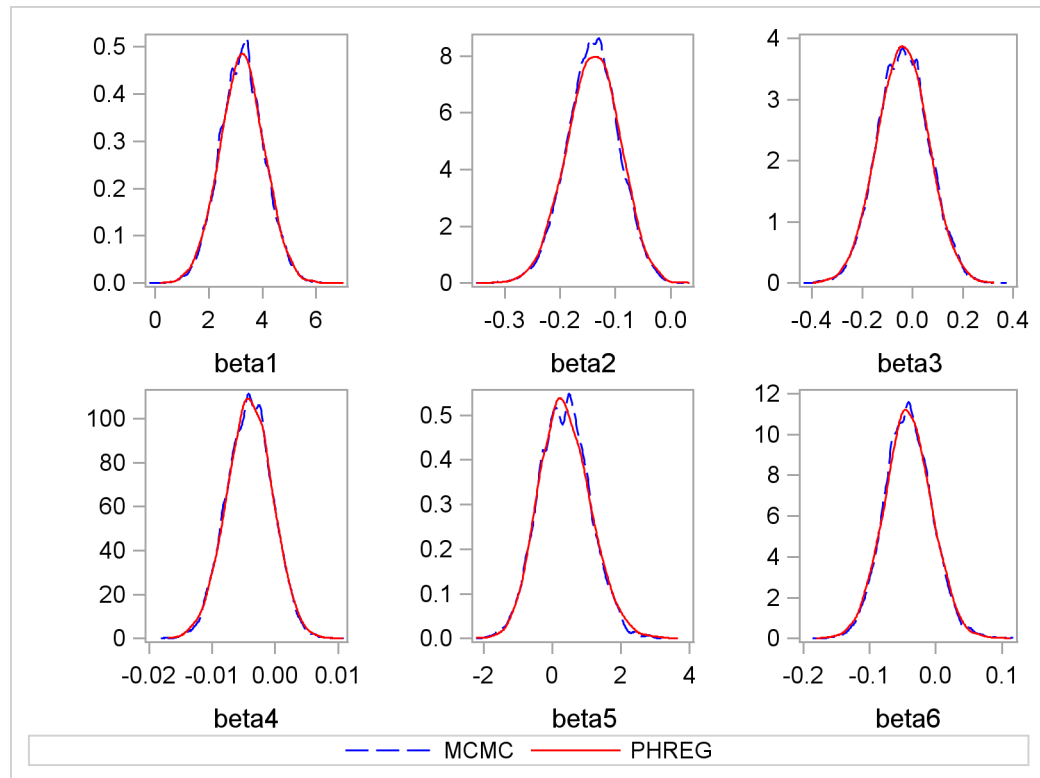
Example 52.9: Normal Regression with Interval Censoring

You can use PROC MCMC to fit failure time data that can be right, left, or interval censored. To illustrate, a normal regression model is used in this example.

Assume that you have the following simple regression model with no covariates:

$$\mathbf{y} = \mu + \sigma\epsilon$$

where $\mathbf{y}$ is a vector of response values (the failure times), μ is the grand mean, σ is an unknown scale parameter, and ϵ are errors from the standard normal distribution. Instead of observing y_i directly, you only observe a truncated value t_i. If the true y_i occurs after the censored time t_i, it is called *right censoring*. If y_i occurs before the censored time, it is called *left censoring*. A failure time y_i can be censored at both ends, and this is called *interval censoring*. The likelihood for y_i is as follows:

$$p(y_i|\mu) = \begin{cases} \phi(y_i|\mu,\sigma) & \text{if } y_i \text{ is uncensored} \\ S(t_{l,i}|\mu) & \text{if } y_i \text{ is right censored by } t_{l,i} \\ 1 - S(t_{r,i}|\mu) & \text{if } y_i \text{ is left censored by } t_{r,i} \\ S(t_{l,i}|\mu) - S(t_{r,i}|\mu) & \text{if } y_i \text{ is interval censored by } t_{l,i} \text{ and } t_{r,i} \end{cases}$$

where $S(\cdot)$ is the survival function, $S(t) = Pr(T > t)$.

Gentleman and Geyer (1994) uses the following data on cosmetic deterioration for early breast cancer patients treated with radiotherapy:

```
title 'Normal Regression with Interval Censoring';
data cosmetic;
   label tl = 'Time to Event (Months)';
   input tl tr @@;
   datalines;
45  .    6 10    .  7 46  .  46  .    7 16  17  .    7 14
37 44    .  8    4 11  15  .  11 15  22  .  46  .  46  .
25 37  46  .   26 40  46  .  27 34  36 44  46  .  36 48
37  .  40  .   17 25  46  .  11 18  38  .   5 12  37  .
 .  5  18  .   24  .  36  .   5 11  19 35  17 25  24  .
32  .  33  .   19 26  37  .  34  .  36  .
;
```

The data consist of time interval endpoints (in months). Nonmissing equal endpoints (tl = tr) indicates uncensoring; a nonmissing lower endpoint (tl ≠ .) and a missing upper endpoint (tr = .) indicates right censoring; a missing lower endpoint (tl = .) and a nonmissing upper endpoint (tr ≠ .) indicates left censoring; and nonmissing unequal endpoints (tl ≠ tr) indicates interval censoring.

With this data set, you can consider using proper but diffuse priors on both μ and σ, for example:

$$\pi(\mu) \quad \propto \quad \phi(0, \text{sd} = 1000)$$
$$\pi(\sigma) \quad \propto \quad f_\Gamma(0.001, \text{iscale} = 0.001)$$

where f_Γ is the gamma density function.

The following SAS statements fit an interval censoring model and generate Output 52.9.1:

```
proc mcmc data=cosmetic outpost=postout seed=1 nmc=20000 missing=AC;
   ods select PostSummaries PostIntervals;
   parms mu 60 sigma 50;

   prior mu ~ normal(0, sd=1000);
   prior sigma ~ gamma(shape=0.001,iscale=0.001);

   if (tl^=. and tr^=. and tl=tr) then
      llike = logpdf('normal',tr,mu,sigma);
   else if (tl^=. and tr=.) then
      llike = logsdf('normal',tl,mu,sigma);
   else if (tl=. and tr^=.) then
      llike = logcdf('normal',tr,mu,sigma);
   else
      llike = log(sdf('normal',tl,mu,sigma) -
         sdf('normal',tr,mu,sigma));

   model general(llike);
run;
```

Because there are missing cells in the input data, you want to use the MISSING=AC option so that PROC MCMC does not delete any observations that contain missing values. The IF-ELSE statements distinguish different censoring cases for y_i, according to the likelihood. The SAS functions

LOGCDF, LOGSDF, LOGPDF, and SDF are useful here. The MODEL statement assigns llike as the log likelihood to the response. The Markov chain appears to have converged in this example (evidence not shown here), and the posterior estimates are shown in Output 52.9.1.

Output 52.9.1 Interval Censoring

```
                  Normal Regression with Interval Censoring

                          The MCMC Procedure

                          Posterior Summaries

                                    Standard              Percentiles
     Parameter        N      Mean   Deviation      25%         50%        75%

     mu            20000   41.7807    5.7882    37.7220     41.3468    45.2249
     sigma         20000   29.1122    6.0503    24.8774     28.2210    32.4250


                          Posterior Intervals

        Parameter     Alpha      Equal-Tail Interval        HPD Interval

          mu          0.050      32.0499     54.6104      31.3604    53.6115
          sigma       0.050      20.0889     43.1335      19.4041    41.6742
```

Example 52.10: Constrained Analysis

Conjoint analysis uses regression techniques to model consumer preferences and to estimate consumer utility functions. A problem with conventional conjoint analysis is that sometimes your estimated utilities do not make sense. Your results might suggest, for example, that the consumers would prefer to spend more on a product than to spend less. With PROC MCMC, you can specify constraints on the part-worth utilities (parameter estimates). Suppose that the consumer product being analyzed is an off-road motorcycle. The relevant attributes are how large each motorcycle is (less than 300cc, 301–550cc, and more than 551cc), how much it costs (less than $5000, $5001–$6000, $6001–$7000, and more than $7000), whether or not it has an electric starter, whether or not the engine is counter-balanced, and whether the bike is from Japan or Europe. The preference variable is a ranking of the bikes. You could perform an ordinary conjoint analysis with PROC TRANSREG (see Chapter 90, "The TRANSREG Procedure") as follows:

```
title 'Constrained Conjoint Analysis';

options validvarname=any;
proc format;
   value sizef  1 = '< 300cc' 2 = '300-550cc' 3 = '> 551cc';
   value pricef 1 = '< $5000' 2 = '$5000 - $6000'
                3 = '$6001 - $7000' 4 = '> $7000';
   value startf 1 = 'Electric Start' 2 = 'Kick Start';
   value balf   1 = 'Counter Balanced' 2 = 'Unbalanced';
   value orif   1 = 'Japanese' 2 = 'European';
run;
```

```
data bikes;
   input Size Price Start Balance Origin Rank @@;
   format size sizef. price pricef. start startf.
          balance balf. origin orif.;
   datalines;
2 1 2 1 2  3  1 4 2 2 2  7  1 2 1 1 2  6
3 3 1 1 2  1  1 3 2 1 1  5  3 4 2 2 2 12
2 3 2 2 1  9  1 1 1 2 1  8  2 2 1 2 2 10
2 4 1 1 1  4  3 1 1 2 1 11  3 2 2 1 1  2
;


proc transreg data=bikes utilities cprefix=0 lprefix=0;
   ods select Utilities;
   model identity(rank / reflect) =
         class(size price start balance origin / zero=sum);
   output out=coded(drop=intercept) replace;
run;
```

The DATA step reads the experimental design and dependent variable Rank and assigns formats to label the factor levels. PROC TRANSREG is run specifying UTILITIES, which requests a conjoint analysis. The rank variable is reflected around its mean ($1 \rightarrow 12$, $2 \rightarrow 11$, ..., $12 \rightarrow 1$) so that in the analysis, larger part-worth utilities correspond to higher preference. The OUT=CODED data set contains the reflected ranks and a binary coding of the factors that can be used in other analyses. Refer to Kuhfeld (2004) for more information about conjoint analysis and coding with PROC TRANSREG.

The Utilities table from the conjoint analysis is shown in Output 52.10.1. Notice the part-worth utilities for price. The part-worth utility for < \$5000 is 0.25. As price increases to the \$5000–\$6000 range, utility decreases to −0.5. Then as price increases to the \$6001–\$7000 range, part-worth utility *increases* to 0.5. Finally, for the most expensive bikes, utility decreases again to −0.25. In cases like this, you might want to impose constraints on the solution so that the part-worth utility for price never increases as prices go up.

Output 52.10.1 Ordinary Conjoint Analysis by PROC TRANSREG

```
                    Constrained Conjoint Analysis

                       The TRANSREG Procedure

             Utilities Table Based on the Usual Degrees of Freedom

                                            Importance
                                 Standard   (% Utility
    Label                Utility    Error      Range)      Variable

    Intercept             6.5000   0.95743                 Intercept

    < 300cc              -0.0000   1.35401      0.000      Class.< 300cc
    300-550cc            -0.0000   1.35401                 Class.300-550cc
    > 551cc               0.0000   1.35401                 Class.> 551cc

    < $5000               0.2500   1.75891     13.333      Class.< $5000
    $5000 - $6000        -0.5000   1.75891                 Class.$5000 - $6000
    $6001 - $7000         0.5000   1.75891                 Class.$6001 - $7000
    > $7000              -0.2500   1.75891                 Class.> $7000

    Electric Start       -0.1250   1.01550      3.333      Class.Electric Start
    Kick Start            0.1250   1.01550                 Class.Kick Start

    Counter Balanced      3.0000   1.01550     80.000      Class.Counter Balanced
    Unbalanced           -3.0000   1.01550                 Class.Unbalanced

    Japanese             -0.1250   1.01550      3.333      Class.Japanese
    European              0.1250   1.01550                 Class.European
```

You could run PROC TRANSREG again, specifying monotonicity constraints on the part-worth utilities for price:

```
proc transreg data=bikes utilities cprefix=0 lprefix=0;
   ods select ConservUtilities;
   model identity(rank / reflect) =
         monotone(price / tstandard=center)
         class(size start balance origin / zero=sum);
run;
```

The output from this PROC TRANSREG step is shown in Output 52.10.2.

Output 52.10.2 Constrained Conjoint Analysis by PROC TRANSREG

```
                       Constrained Conjoint Analysis

                          The TRANSREG Procedure

                Utilities Table Based on Conservative Degrees of Freedom


                                          Importance
                               Standard   (% Utility
      Label            Utility   Error       Range)        Variable

      Intercept         6.5000   0.97658                   Intercept

      Price            -0.1581      .         7.143        Monotone(Price)
      < $5000           0.2500      .
      $5000 - $6000     0.0000      .
      $6001 - $7000     0.0000      .
      > $7000          -0.2500      .

      < 300cc          -0.0000   1.38109      0.000        Class.< 300cc
      300-550cc         0.0000   1.38109                   Class.300-550cc
      > 551cc           0.0000   1.38109                   Class.> 551cc

      Electric Start   -0.2083   1.00663      5.952        Class.Electric Start
      Kick Start        0.2083   1.00663                   Class.Kick Start

      Counter Balanced  3.0000   0.97658     85.714        Class.Counter Balanced
      Unbalanced       -3.0000   0.97658                   Class.Unbalanced

      Japanese         -0.0417   1.00663      1.190        Class.Japanese
      European          0.0417   1.00663                   Class.European
```

This monotonicity constraint is one of the few constraints on the part-worth utilities that you can specify in PROC TRANSREG. In contrast, PROC MCMC allows you to specify any constraint that can be written in the DATA step language. You can perform the restricted conjoint analysis with PROC MCMC by using the coded factors that were output from PROC TRANSREG. The data set is coded.

The likelihood is a simple regression model:

$$\text{rank}_i \sim \text{normal}(\mathbf{x}_i' \boldsymbol{\beta}, \sigma)$$

where rank is the response, the covariates are '< 300cc'n, '300-500cc'n, '< $5000'n, '$5000 - $6000'n, '$6001 - $7000'n, 'Electric Start'n, 'Counter Balanced'n, and Japanese. Note that OPTIONS VALID-VARNAME=ANY allows PROC TRANSREG to create names for the coded variables with blanks and special characters. That is why the name-literal notation ('*variable-name*'n) is used for the input data set variables.

Suppose that there are two constraints you want to put on some of the parameters: one is that the parameters for '< $5000'n, '$5000 - $6000'n, and '$6001 - $7000'n decrease in order, and the other is that the parameter for 'Counter Balanced'n is strictly positive. You can consider a truncated multivariate normal prior as follows:

$$\pi \left(\beta_{\text{'< $5000'n}}, \beta_{\text{'$5000 - $6000'n}}, \beta_{\text{'$6001 - $7000'n}}, \beta_{\text{'Counter Balanced'n}} \right) \sim \text{MVN}(0, \sigma \boldsymbol{I})$$

with the following set of constraints:

$$\beta_{\text{'< \$5000'n}} > \beta_{\text{'\$5000 - \$6000'n}} > \beta_{\text{'\$6001 - \$7000'n}} > 0$$

$$\beta_{\text{'Counter Balanced'n}} > 0$$

The condition that $\beta_{\text{'\$6001 - \$7000'n}} > 0$ reflects an implied constraint that, by definition, 0 is the utility for the highest price range, > \$7000, which is the reference level for the binary coded price variable. The following statements fit the desired model:

```
proc mcmc data=coded outpost=bikesout ntu=3000 nmc=50000 thin=10
          seed=448;
   ods select PostSummaries;
   array sigma[4,4] sigma1-sigma16;
   array mu[4] mu1-mu4;

   begincnst;
      call identity(sigma);
      call mult(sigma, 100, sigma);
      call zeromatrix(mu);
      rc = logmpdfsetsq('v', of sigma1-sigma16);
   endcnst;

   parms intercept pw300cc pw300_550cc pwElectricStart pwJapanese ltau 1;
   parms pw5000 0.3 pw5000_6000 0.2 pw6001_7000 0.1 pwCounterBalanced 1;

   beginprior;
   prior intercept pw300: pwE: pwJ: ~ normal(0, var=100);
   if (pw5000      >= pw5000_6000 & pw5000_6000 >= pw6001_7000 &
      pw6001_7000 >= 0             & pwCounterBalanced > 0) then
      lp = logmpdfnormal(of mu1-mu4, pw5000, pw5000_6000,
                         pw6001_7000, pwCounterBalanced, 'v');
   else
      lp = .;
   prior pw5000 pw5000_6000 pw6001_7000 pwC: ~ general(lp);
   prior ltau  ~ egamma(0.001, scale=1000);
   tau = exp(ltau);
   endprior;

   mean = intercept +
         pw300cc             * '< 300cc'n            +
         pw300_550cc         * '300-550cc'n          +
         pw5000              * '< $5000'n            +
         pw5000_6000         * '$5000 - $6000'n      +
         pw6001_7000         * '$6001 - $7000'n      +
         pwElectricStart   * 'Electric Start'n    +
         pwCounterBalanced * 'Counter Balanced'n +
         pwJapanese          * Japanese;
   model rank ~ normal(mean, prec=tau);
run;

data _null_;
   rc = logmpdffree();
run;
```

The two ARRAY statements allocate a 4 × 4 dimensional array for the prior covariance and an array of size 4 for the prior means. In the BEGINCNST and ENDCNST statements, the CALL IDENTITY function sets sigma to be an identity matrix; the CALL MULT function sets sigma's diagonal elements to be 100 (the diagonal variance terms); the CALL ZEROMATRIX function sets mu to be a vector of zeros (the prior means); and the LOGMPDFSETSQ function sets up sigma to be called in a multivariate normal density function later. For matrix functions in PROC MCMC, see the section "Matrix Functions in PROC MCMC" on page 3529. For multivariate density functions, see the section "Multivariate Density Functions" on page 3544. It is important to note that if you used the LOGMPDFSET or the LOGMPDFSETSQ functions to set up covariance matrix, you must free the memory allocated by these functions after you exit PROC MCMC. To free the memory, use the function LOGMPDFFREE.

There are two PARMS statements, with each of them naming a block of parameters. The first PARMS statement blocks the following: the intercept, the two size parameters, the one start-type parameter, the one origin parameter, and the log of the precision. The second PARMS statement blocks the three price parameters and the one balance parameter, parameters that have the constraint multivariate normal prior. The second PARMS statement also specifies initial values for the parameter estimates. The initial values reflect the constraints on these parameters. The initial part-worth utilities all decrease from 0.3 to 0.2 to 0.1 to 0.0 (for the implicit reference level) as the prices increase. Also, the initial part-worth utility for the counter-balanced engine is set to a positive value, 1.

In the PRIOR statements, regression coefficients without constraints are given an independent normal prior with mean at 0 and variance of 100. The next IF-ELSE construction imposes the constraints. When these constraints are met, pw5000, pw5000_6000, pw6001_7000, pwCounterBalanced are jointly distributed as a multivariate normal prior with mean mu and covariance sigma (as defined via the symbol 'v' in the BEGINCNST and ENDCNST statements). Otherwise, the prior is not defined and lp is assigned a missing value.

The parameter ltau is given an egamma prior. It is an equivalent prior to placing a gamma prior, with the same configuration, on tau = exp(ltau). For the definition of the egamma distribution, see the section "Standard Distributions" on page 3529. This transformation often improves mixing (see "Example 52.4: Nonlinear Poisson Regression Models" on page 3599 and "Example 52.12: Using a Transformation to Improve Mixing" on page 3677). The next assignment statement transforms ltau back to tau.

The model specification is linear. The mean is comprised of an intercept and the sum of terms like pw300cc * '< 300cc'n, which is a parameter times an input data set variable. The MODEL statement specifies that the linear model for rank is normally distributed with mean mean and precision tau.

After the PROC MCMC run, you *must* run the memory clean up function LOGMPDFFREE, which should produce the following note in the log file:

```
NOTE: The matrix -  v - has been deleted.
```

The MCMC results are shown in Output 52.10.3.

Output 52.10.3 MCMC Results

```
                      Constrained Conjoint Analysis

                          The MCMC Procedure

                          Posterior Summaries

                                  Standard              Percentiles
    Parameter            N     Mean  Deviation      25%       50%       75%

    intercept          5000   2.2052    2.6285    0.8089    2.3658    3.8732
    pw300cc            5000   0.0780    2.5670   -1.4062    0.0717    1.5850
    pw300_550cc        5000  -0.0173    2.5378   -1.5136   -0.00275   1.4536
    pwElectricStart    5000  -1.2175    2.1805   -2.4933   -1.1041    0.1410
    pwJapanese         5000  -0.4212    2.1485   -1.6575   -0.4102    0.7909
    ltau               5000  -2.4440    0.7293   -2.9024   -2.3787   -1.9177
    pw5000             5000   4.3724    2.4962    2.6418    3.9163    5.5202
    pw5000_6000        5000   2.6649    1.8227    1.3878    2.2894    3.5162
    pw6001_7000        5000   1.4880    1.3303    0.5077    1.1389    2.0849
    pwCounterBalanced  5000   5.9056    2.0591    4.6440    5.9033    7.1036
```

The estimates of the part-worth utility for the price categories are ordered as expected. This agrees with the intuition that there is a higher preference for a less expensive motor bike when all other things are equal, and that is what you see when you look at the estimated posterior means for the price part-worths. The estimated standard deviations of the price part-worths in this model are of approximately the same order of magnitude as the posterior means. This indicates that the part-worth utilities for this subject are not significantly far from each other, and that this subject's ranking of the options was not significantly influenced by the difference in price.

One advantage of Bayesian analysis is that you can incorporate prior information in the data analysis. Constraints on the parameter space are one possible source of information that you might have before you examine the data. This example shows that it can easily be accomplished in PROC MCMC.

Example 52.11: Implement a New Sampling Algorithm

This example illustrates using the UDS statement to implement a new Markov chain sampler. The algorithm demonstrated here is proposed by Holmes and Held (2006), hereafter referred to as HH. They presented a Gibbs sampling algorithm for generating draws from the posterior distribution of the parameters in a probit regression model. The notation follows closely to HH.

The data used here is the remission data set from a PROC LOGISTIC example:

```
title 'Implement a New Sampling Algorithm';
data inputdata;
   input remiss cell smear infil li blast temp;
   ind = _n_;
   label remiss='Complete Remission';
   datalines;
   1  0.8   0.83  0.66  1.9  1.1    0.996

... more lines ...

   0  1     0.73  0.73  0.7  0.398  0.986
;
```

The variable remiss is the cancer remission indicator variable with a value of 1 for remission and a value of 0 for nonremission. There are six explanatory variables: cell, smear, infil, li, blast, and temp. These variables are the risk factors thought to be related to cancer remission. The binary regression model is as follows:

$$\text{remiss}_i \sim \text{binary}(p_i)$$

where the covariates are linked to p_i through a probit transformation:

$$\text{probit}(p_i) = \mathbf{x}'\boldsymbol{\beta}$$

$\boldsymbol{\beta}$ are the regression coefficients and $\mathbf{x}'$ the explanatory variables. Suppose that you want to use independent normal priors on the regression coefficients:

$$\beta_i \sim \text{normal}(0, \text{var} = 25)$$

Fitting a logistic model with PROC MCMC is straightforward. You can use the following statements:

```
proc mcmc data=inputdata nmc=100000 propcov=quanew seed=17
          outpost=mcmcout;
   ods select PostSummaries ess;
   parms beta0-beta6;
   prior beta: ~ normal(0,var=25);
   mu = beta0 + beta1*cell + beta2*smear +
        beta3*infil +  beta4*li + beta5*blast +  beta6*temp;
   p = cdf('normal', mu, 0, 1);
   model remiss ~ bern(p);
run;
```

The expression mu is the regression mean, and the CDF function links mu to the probability of remission p in the binary likelihood.

The summary statistics and effective sample sizes tables are shown in Output 52.11.1. There are high autocorrelations among the posterior samples, and efficiency is relatively low. The correlation time is reduced only after a large amount of thinning.

Output 52.11.1 Random Walk Metropolis

```
                    Implement a New Sampling Algorithm

                           The MCMC Procedure

                          Posterior Summaries

                                     Standard              Percentiles
     Parameter        N       Mean   Deviation      25%          50%         75%

     beta0       100000    -2.0531     3.8299    -4.6418     -2.0354      0.5638
     beta1       100000     2.6300     2.8270     0.6563      2.5272      4.4846
     beta2       100000    -0.8426     3.2108    -3.0270     -0.8263      1.3429
     beta3       100000     1.5933     3.5491    -0.7993      1.6190      3.9695
     beta4       100000     2.0390     0.8796     1.4312      2.0028      2.6194
     beta5       100000    -0.3184     0.9543    -0.9613     -0.3123      0.3418
     beta6       100000    -3.2611     3.7806    -5.8050     -3.2736     -0.7243

                    Implement a New Sampling Algorithm

                           The MCMC Procedure

                         Effective Sample Sizes

                                      Correlation
            Parameter           ESS          Time     Efficiency

            beta0            4280.8       23.3602         0.0428
            beta1            4496.5       22.2398         0.0450
            beta2            3434.1       29.1199         0.0343
            beta3            3856.6       25.9294         0.0386
            beta4            3659.7       27.3245         0.0366
            beta5            3229.9       30.9610         0.0323
            beta6            4430.7       22.5696         0.0443
```

As an alternative to the random walk Metropolis, you can use the Gibbs algorithm to sample from the posterior distribution. The Gibbs algorithm is described in the section "Gibbs Sampler" on page 154. While the Gibbs algorithm generally applies to a wide range of statistical models, the actual implementation can be problem-specific. In this example, performing a Gibbs sampler involves introducing a class of auxiliary variables (also known as latent variables). You first reformulate the model by adding a z_i for each observation in the data set:

$$y_i = \begin{cases} 1 & \text{if } z_i > 0 \\ 0 & \text{otherwise} \end{cases}$$

$$z_i = \mathbf{x}'_i \boldsymbol{\beta} + \epsilon_i$$

$$\epsilon \sim \text{normal}(0, 1)$$

$$\boldsymbol{\beta} \sim \pi(\boldsymbol{\beta})$$

If $\boldsymbol{\beta}$ has a normal prior, such as $\pi(\boldsymbol{\beta}) = N(\mathbf{b}, \mathbf{v})$, you can work out a closed form solution to the full conditional distribution of $\boldsymbol{\beta}$ given the data and the latent variables z_i. The full conditional

distribution is also a multivariate normal, due to the conjugacy of the problem. See the section "Conjugate Priors" on page 146. The formula is shown here:

$$
\begin{aligned}
\boldsymbol{\beta}|\mathbf{z}, \mathbf{x} &\sim \text{normal}(\mathbf{B}, \mathbf{V}) \\
\mathbf{B} &= \mathbf{V}((v)^{-1}\mathbf{b} + \mathbf{x}'\mathbf{z}) \\
\mathbf{V} &= (\mathbf{v}^{-1} + \mathbf{x}'\mathbf{x})^{-1}
\end{aligned}
$$

The advantage of creating the latent variables is that the full conditional distribution of $\mathbf{z}$ is also easy to work with. The distribution is a truncated normal distribution:

$$
z_i|\boldsymbol{\beta}, \mathbf{x}_i, y_i \sim \begin{cases} \text{normal}(\mathbf{x}_i\boldsymbol{\beta}, 1)I(z_i > 0) & \text{if } y_i = 1 \\ \text{normal}(\mathbf{x}_i\boldsymbol{\beta}, 1)I(z_i \le 0) & \text{otherwise} \end{cases}
$$

You can sample $\boldsymbol{\beta}$ and z iteratively, by drawing $\boldsymbol{\beta}$ given z and vice verse. HH point out that a high degree of correlation could exist between $\boldsymbol{\beta}$ and z, and it makes this iterative way of sampling inefficient. As an improvement, HH proposed an algorithm that samples $\boldsymbol{\beta}$ and $\mathbf{z}$ jointly. At each iteration, you sample z_i from the posterior marginal distribution (this is the distribution that is conditional only on the data and not on any parameters) and then sample $\boldsymbol{\beta}$ from the same posterior full conditional distribution as described previously:

1. Sample z_i from its posterior marginal distribution:

$$
\begin{aligned}
z_i|\mathbf{z}_{-i}, y_i &\sim \begin{cases} \text{normal}(m_i, v_i)I(z_i > 0) & \text{if } y_i = 1 \\ \text{normal}(m_i, v_i)I(z_i \le 0) & \text{otherwise} \end{cases} \\
m_i &= \mathbf{x}_i\mathbf{B} - w_i(z_i - \mathbf{x}_i\mathbf{B}) \\
v_i &= 1 + w_i \\
w_i &= h_i/(1 - h_i) \\
h_i &= (\mathbf{H})_i i, \mathbf{H} = \mathbf{x}\mathbf{V}\mathbf{x}'
\end{aligned}
$$

2. Sample $\boldsymbol{\beta}$ from the same posterior full conditional distribution described previously.

For a detailed description of each of the conditional terms, refer to the original paper.

PROC MCMC cannot sample from the probit model by using this sampling scheme but you can implement the algorithm by using the UDS statement. To sample z_i from its marginal, you need a function that draws random variables from a truncated normal distribution. The functions, RLT-NORM and RRTNORM, generate left- and right-truncated normal variates, respectively. The algorithm is taken from Robert (1995).

The functions are written in PROC FCMP (see the FCMP Procedure in the *Base SAS Procedures Guide*):

```
proc fcmp outlib=sasuser.funcs.uds;
   /**************************************/
   /* Generate left-truncated normal variate */
   /**************************************/
   function rltnorm(mu,sig,lwr);
   if lwr<mu then do;
      ans = lwr-1;
      do while(ans<lwr);
         ans = rand('normal',mu,sig);
      end;
   end;
   else do;
      mul = (lwr-mu)/sig;
      alpha = (mul + sqrt(mul**2 + 4))/2;
      accept=0;
      do while(accept=0);
         z = mul + rand('exponential')/alpha;
         lrho = -(z-alpha)**2/2;
         u = rand('uniform');
         lu = log(u);
         if lu <= lrho then accept=1;
      end;
      ans = sig*z + mu;
   end;
   return(ans);
   endsub;

   /**************************************/
   /* Generate right-truncated normal variate */
   /**************************************/
   function rrtnorm(mu,sig,uppr);
   ans = 2*mu - rltnorm(mu,sig, 2*mu-uppr);
   return(ans);
   endsub;
run;
```

The function call to RLTNORM(mu,sig,lwr) generates a random number from the left-truncated normal distribution:

$$\theta \sim \text{normal}(\text{mu}, \text{sd} = \text{sig}) I(\theta > \text{lwr})$$

Similarly, the function call to RRTNORM(mu,sig,uppr) generates a random number from the right-truncated normal distribution:

$$\theta \sim \text{normal}(\text{mu}, \text{sd} = \text{sig}) I(\theta < \text{lwr})$$

These functions are used to generate the latent variables z_i.

Using the algorithm A1 from the HH paper as an example, Output 52.37 lists a line-by-line implementation with the PROC MCMC coding style. The table is broken into three portions: set up the constants, initialize the parameters, and sample one draw from the posterior distribution. The left column of the table is identical to the A1 algorithm stated in the appendix of HH. The right column of the table lists SAS statements.

Table 52.37 Holmes and Held (2006), algorithm A1. Side-by-Side Comparison to SAS

Define Constants	In the BEGINCNST/ENDCNST Statements

$V \leftarrow (X^T X + v^{-1})^{-1}$	```call transpose(x,xt); /* xt = transpose(x) */```
	```call mult(xt,x,xtx);```
	```call inv(v,v); /* v = inverse(v) */```
	```call addmatrix(xtx,v,xtx); /* xtx = xtx+v */```
	```call inv(xtx,v); /* v = inverse(xtx) */```

$L \leftarrow \mathrm{Chol}(V)$	```call chol(v,L);```
$S \leftarrow V X^T$	```call mult(v,xt,S);```

FOR $j = 1$ to n
　$H[j] \leftarrow X[j,]S[,j]$
　$W[j] \leftarrow H[j]/(1 - H[j])$
　$Q[j] \leftarrow W[j] + 1$
END
	```call mult(x,S,HatMat);```
	```do j=1 to &n;```
	``` H = HatMat[j,j];```
	``` W[j] = H/(1-H);```
	``` sQ[j] = sqrt(W[j] + 1); /* use s.d. in SAS */```
	``` end;```

Initial Values	In the BEGINCNST/ENDCNST Statements

$Z \sim \mathrm{normal}(0, I_n)Ind(Y, Z)$	```do j=1 to &n;```
	``` if(y[j]=1) then```
	``` Z[j] = rltnorm(0,1,0);```
	``` else```
	``` Z[j] = rrtnorm(0,1,0);```
	``` end;```

$B \leftarrow SZ$	```call mult(S,Z,B);```

Table 52.37 (continued)

Draw One Sample	Subroutine HH
| | ```
do j=1 to &n;
  zold = Z[j];
  m = 0;
  do k= 1 to &p;
    m = m + X[j,k] * B[k];
  end;
  m = m - W[j]*(Z[j]-m);
  if (y[j]=1) then
    Z[j] = rltnorm(m,sQ[j],0);
  else
    Z[j] = rrtnorm(m,sQ[j],0);
  diff = Z[j] - zold;
  do k= 1 to &p;
    B[k] = B[k] + diff * S[k,j];
  end;
end;
do j = 1 to &p;
  T[j] = rand('normal');
end;
call mult(L,T,T);
call addmatrix(B,T,beta);
``` |

In the left column:

FOR $j = 1$ to n

  $z_{old} \leftarrow Z[j]$

  $m \leftarrow X[j,]B$

  $m \leftarrow m - W[j](Z[j] - m)$

  $Z[j] \sim \text{normal}(m, Q[j])Ind(Y[j], Z[j])$

  $B \leftarrow B + (Z[j] - z_{old})S[, j]$

END

$T \sim \text{normal}(0, I_p)$

$\beta[, i] \leftarrow B + LT$

The following statements define the subroutine HH (algorithm A1) in PROC FCMP and store it in library `sasuser.funcs.uds`:

```
/* define the HH algorithm in PROC FCMP. */
%let n = 27;
%let p = 7;
options cmplib=sasuser.funcs;
proc fcmp outlib=sasuser.funcs.uds;
   subroutine HH(beta[*],Z[*],B[*],x[*,*],y[*],W[*],sQ[*],S[*,*],L[*,*]);
   outargs beta, Z, B;
   array T[&p] / nosym;
   do j=1 to &n;
      zold = Z[j];
      m = 0;
      do k = 1 to &p;
         m = m + X[j,k] * B[k];
      end;
      m = m - W[j]*(Z[j]-m);
      if (y[j]=1) then
         Z[j] = rltnorm(m,sQ[j],0);
      else
         Z[j] = rrtnorm(m,sQ[j],0);
      diff = Z[j] - zold;
      do k = 1 to &p;
```

```
        B[k] = B[k] + diff * S[k,j];
      end;
   end;
   do j=1 to &p;
      T[j] = rand('normal');
   end;
   call mult(L,T,T);
   call addmatrix(B,T,beta);
   endsub;
run;
```

Note that one-dimensional array arguments take the form of *name*[*] and two-dimensional array arguments take the form of *name*[*,*]. Three variables, beta, Z, and B, are OUTARGS variables, making them the only arguments that can be modified in the subroutine. For the UDS statement to work, all OUTARGS variables have to be model parameters. Technically, only beta and Z are model parameters, and B is not. The reason that B is declared as an OUTARGS is because the array has to be updated throughout the simulation, and this is the only way to modify its values. The input array x contains all of the explanatory variables, and the array y stores the response. The rest of the input arrays, W, sQ, S, and L, store constants as detailed in the algorithm. The following statements illustrate how to fit a Bayesian probit model by using the HH algorithm:

```
options cmplib=sasuser.funcs;

proc mcmc data=inputdata nmc=5000 monitor=(beta) outpost=hhout;
   ods select PostSummaries ess;
   array xtx[&p,&p];        /* work space                        */
   array xt[&p,&n];         /* work space                        */
   array v[&p,&p];          /* work space                        */
   array HatMat[&n,&n];     /* work space                        */
   array S[&p,&n];          /* V * Xt                            */
   array W[&n];
   array y[&n];             /* y stores the response variable    */
   array x[&n,&p];          /* x stores the explanatory variables */
   array sQ[&n];            /* sqrt of the diagonal elements of Q */
   array B[&p];             /* conditional mean of beta          */
   array L[&p,&p];          /* Cholesky decomp of conditional cov */
   array Z[&n];             /* latent variables Z                */
   array beta[&p] beta0-beta6;   /* regression coefficients      */

   begincnst;
      y[ind] = remiss;
      x[ind,1] = 1;
      x[ind,2] = cell;
      x[ind,3] = smear;
      x[ind,4] = infil;
      x[ind,5] = li;
      x[ind,6] = blast;
      x[ind,7] = temp;
      if ind=&n then do;
         call identity(v);
         call mult(v, 25, v);
         call transpose(x,xt);
         call mult(xt,x,xtx);
         call inv(v,v);
```

```
                call addmatrix(xtx,v,xtx);
                call inv(xtx,v);
                call chol(v,L);
                call mult(v,xt,S);
                call mult(x,S,HatMat);
                do j=1 to &n;
                    H = HatMat[j,j];
                    W[j] = H/(1-H);
                    sQ[j] = sqrt(W[j] + 1);
                end;

                do j=1 to &n;
                    if(y[j]=1) then
                        Z[j] = rltnorm(0,1,0);
                    else
                        Z[j] = rrtnorm(0,1,0);
                end;
                call mult(S,Z,B);
            end;
            call streaminit(83101);
        endcnst;

        uds HH(beta,Z,B,x,y,W,sQ,S,L);
        parms z: beta: 0 B1-B7 / uds;
        prior z: beta: B1-B7 ~ general(0);

        model general(0);
    run;
```

The **options** statement names the catalog of FCMP subroutines to use. The cmplib library stores the subroutine HH. You do not need to set a random number seed in the PROC MCMC statement because all random numbers are generated from the HH subroutine. The initialization of the rand function is controlled by the streaminit function, which is called in the program with a seed value of 83101.

A number of arrays are allocated. Some of them, such as xtx, xt, v, and HatMat, are work space that is used to construct constant arrays. Other arrays are used in the subroutine sampling. Explanations of the arrays are shown in comments in the statements.

In the BEGINCNST and ENDCNST statements, you read data set variables in the arrays x and y. You want to hold off on any array calculations until both arrays are filled. When the condition ind=&n is satisfied, you calculate all the constant terms and assign initial values to Z and B. For listings of all array functions and their definitions, see the section "Matrix Functions in PROC MCMC" on page 3529.

The UDS statement declares that the subroutine HH is used to sample the parameters beta, Z, and B. You also specify the UDS option in the PARMS statement. Because all parameters are updated through the UDS interface, there is no need to declare the actual form of the prior for any of the parameters. Each parameter is declared to have a prior of general(0). Similarly, it is not necessary to declare the actual form of the likelihood. The MODEL statement also takes a flat likelihood of the form **general(0)**.

Summary statistics and effective sample sizes are shown in Output 52.11.2. The posterior estimates are very close to what was shown in Output 52.11.1. The HH algorithm produces samples that are much less correlated.

Output 52.11.2 Holms and Held

```
                    Implement a New Sampling Algorithm

                           The MCMC Procedure

                          Posterior Summaries

                              Standard                  Percentiles
    Parameter        N       Mean    Deviation      25%       50%        75%

    beta0          5000    -2.0567     3.8260     -4.6537   -2.0777     0.5495
    beta1          5000     2.7254     2.8079      0.7812    2.6678     4.5370
    beta2          5000    -0.8318     3.2017     -2.9987   -0.8626     1.2918
    beta3          5000     1.6319     3.5108     -0.7481    1.6636     4.0302
    beta4          5000     2.0567     0.8800      1.4400    2.0266     2.6229
    beta5          5000    -0.3473     0.9490     -0.9737   -0.3267     0.2752
    beta6          5000    -3.3787     3.7991     -5.9089   -3.3504    -0.7928


                    Implement a New Sampling Algorithm

                           The MCMC Procedure

                         Effective Sample Sizes

                                  Correlation
          Parameter        ESS         Time     Efficiency

          beta0          3651.3       1.3694      0.7303
          beta1          1563.8       3.1973      0.3128
          beta2          5005.9       0.9988      1.0012
          beta3          4853.2       1.0302      0.9706
          beta4          2611.2       1.9148      0.5222
          beta5          3049.2       1.6398      0.6098
          beta6          3503.2       1.4273      0.7006
```

The following statements generate the kernel density comparison plots shown in Output 52.11.3:

```
proc kde data=mcmcout;
   ods exclude all;
   univar beta0 beta1 beta2 beta3 beta4 beta5 beta6 / out=m1(drop=count);
run;
ods exclude none;

%reshape(m1, mcmc, suffix1=, suffix2=md);

proc kde data=hhout(drop = LogPost logprior loglike Iteration z1-z27 b1-b7);;
   ods exclude all;
   univar beta0 beta1 beta2 beta3 beta4 beta5 beta6
           / out=m2 (drop=count);
run;
ods exclude none;

%reshape(m2, hh, suffix1=p, suffix2=pd);

data all;
   merge mcmc hh;
run;

proc template;
   define statgraph threebythree;
      %macro plot;
         begingraph;
            layout lattice / rows=3 columns=3;
               %do i = 0 %to 6;
                  layout overlay /yaxisopts=(label=" ");
                     seriesplot y=beta&i.md x=beta&i
                        / connectorder=xaxis
                           lineattrs=(pattern=mediumdash color=blue)
                        legendlabel = "PROC MCMC" name="MCMC";
                     seriesplot y=beta&i.pd x=beta&i.p
                        / connectorder=xaxis lineattrs=(color=red)
                        legendlabel = "Holmes and Held" name="HH";
                  endlayout;
               %end;
               Sidebar / align = bottom;
                  discretelegend "MCMC" "HH";
               endsidebar;
            endlayout;
         endgraph;
      %mend; %plot;
   end;
run;

proc sgrender data=all template=threebythree;
   title "Kernel Density Comparison";
run;
```

The macro %RESHAPE is defined in the example "Logistic Regression Random-Effects Model" on page 3608.

Output 52.11.3 Kernel Density Comparison

It is interesting to compare the two approaches of fitting a generalized linear model. The random walk Metropolis on a seven-dimensional parameter space produces autocorrelations that are substantially higher than the HH algorithm. A much longer chain is needed to produce roughly equivalent effective sample sizes. On the other hand, the Metropolis algorithm is faster to run. The running time of these two examples is roughly the same, with the random walk Metropolis with 100000 samples, a 20-fold increase over that in the HH algorithm example. The speed difference can be attributed to a number of factors, ranging from the implementation of the software and the overhead cost of calling PROC FCMP subroutine and functions. In addition, the HH algorithm requires more parameters by creating an equal number of latent variables as the sample size. Sampling more parameters takes time. A larger number of parameters also increases the challenge in convergence diagnostics, because it is imperative to have convergence in all parameters before you make valid posterior inferences. Finally, you might feel that coding in PROC MCMC is easier. However, this really is not a fair comparison to make here. Writing a Metropolis algorithm from scratch would have probably taken just as much, if not more, effort than the HH algorithm.

Example 52.12: Using a Transformation to Improve Mixing

Proper transformations of parameters can often improve the mixing in PROC MCMC. You already saw this in "Example 52.4: Nonlinear Poisson Regression Models" on page 3599, which sampled using the log scale of parameters that priors that are strictly positive, such as the gamma priors.

This example shows another useful transformation: the logit transformation on parameters that take a uniform prior on [0, 1].

The data set is taken from Sharples (1990). It is used in Chaloner and Brant (1988) and Chaloner (1994) to identify outliers in the data set in a two-level hierarchical model. Congdon (2003) also uses this data set to demonstrates the same technique. This example uses the data set to illustrate how mixing can be improved using transformation and does not address the question of outlier detection as in those papers. The following statements create the data set:

```
title 'Using Transformation to Improve Mixing';
data inputdata;
   input nobs grp y @@;
   ind = _n_;
   datalines;
1 1 24.80 2 1 26.90 3 1 26.65
4 1 30.93 5 1 33.77 6 1 63.31
1 2 23.96 2 2 28.92 3 2 28.19
4 2 26.16 5 2 21.34 6 2 29.46
1 3 18.30 2 3 23.67 3 3 14.47
4 3 24.45 5 3 24.89 6 3 28.95
1 4 51.42 2 4 27.97 3 4 24.76
4 4 26.67 5 4 17.58 6 4 24.29
1 5 34.12 2 5 46.87 3 5 58.59
4 5 38.11 5 5 47.59 6 5 44.67
;
```

There are five groups (grp, $j = 1, \cdots, 5$) with six observations (nobs, $i = 1, \cdots, 6$) in each. The two-level hierarchical model is specified as follows:

$$
\begin{aligned}
y_{ij} &\sim \text{normal}(\theta_j, \text{prec} = \tau_w) \\
\theta_j &\sim \text{normal}(\mu, \text{prec} = \tau_b) \\
\mu &\sim \text{normal}(0, \text{prec} = 1e - 6) \\
\tau &\sim \text{gamma}(0.001, \text{iscale} = 0.001) \\
p &\sim \text{uniform}(0, 1)
\end{aligned}
$$

with the precision parameters related to each other in the following way:

$$
\begin{aligned}
\tau_b &= \tau/p \\
\tau_w &= \tau_b - \tau
\end{aligned}
$$

The total number of parameters in this model is eight: $\theta_1, \cdots, \theta_5, \mu, \tau$, and p.

The following statements fit the model:

```
ods graphics on;
proc mcmc data=inputdata nmc=50000 thin=10 outpost=m1 seed=17
          plot=trace;
   ods select ess tracepanel;
   array theta[5];

   parms theta:;
   parms p tau;
   parms mu ;

   beginprior;
   hyper p ~ uniform(0,1);
   hyper tau ~ gamma(shape=0.001,iscale=0.001);
   hyper mu ~ normal(0,prec=0.00000001);
   taub = tau/p;
   prior theta: ~ normal(mu,prec=taub);
   tauw = taub-tau;
   endprior;

   model y ~ normal(theta[grp],prec=tauw);
run;
```

The **ods select** statement displays the effective sample size table and the trace plots. The **ods graphics on** statement requests ODS Graphics. The PROC MCMC statement specifies the usual options for the MCMC run and produces trace plots (PLOTS=TRACE). The ARRAY statement allocates an array of size 5 for theta. The three PARMS statements put parameters in three different blocks. The remaining statements specify the hyperprior, prior, and likelihood for the data, as described previously. The resulting trace plots are shown in Output 52.12.1, and the effective sample sizes table is shown in Output 52.12.2.

Output 52.12.1 Trace Plots

Output 52.12.1 *continued*

Trace Plots

Output 52.12.2 Bad Effective Sample Sizes

```
                 Using Transformation to Improve Mixing

                          The MCMC Procedure

                       Effective Sample Sizes

                                   Correlation
          Parameter          ESS          Time      Efficiency

          theta1          2207.5        2.2650         0.4415
          theta2          1713.5        2.9180         0.3427
          theta3          1458.5        3.4281         0.2917
          theta4          1904.4        2.6255         0.3809
          theta5           585.9        8.5345         0.1172
          p                 77.2       64.7758         0.0154
          tau              140.8       35.5052         0.0282
          mu              3340.3        1.4969         0.6681
```

The trace plots show that most parameters have relatively good mixing. Two exceptions appear to be p and τ. The trace plot of p shows a slow periodic movement. The τ parameter does not have good mixing either. When the values are close to zero, the chain stands there for long periods of time. An inspection of the effective sample sizes table reveals the same conclusion: p and τ have

much smaller ESSs than the rest of the parameters.

A scatter plot of the posterior samples of p and τ reveals why mixing is bad in these two dimensions. The following statements generate the scatter plot in Output 52.12.3:

```
title 'Scatter Plot of Parameters on Original Scales';

proc sgplot data=m1;
   yaxis label = 'p';
   xaxis label = 'tau' values=(0 to 0.4 by 0.1);
   scatter x = tau y = p;
run;
```

Output 52.12.3 Scatter Plot of τ versus p

The two parameters clearly have a nonlinear relationship. It is not surprising that the Metropolis algorithm does not work well here. The algorithm is designed for cases where the parameters are linearly related with each other.

Instead of sampling on τ, you can sample on the log of τ. The formulation

$$
\begin{aligned}
\pi(\tau) &\propto f_\Gamma(0.001, \text{iscale} = 0.001) \\
\pi(\log(\tau)) &\propto f_{e\Gamma}(0.001, \text{iscale} = 0.001)
\end{aligned}
$$

where f_Γ and $f_{e\Gamma}$ are density functions for the gamma and egamma distributions. See the section "Standard Distributions" on page 3529 for the definitions of the distributions. In addition, you can

sample on the logit of p. The formulation

$$\pi(p) \propto f_{\text{uniform}}(0, 1)$$

is equivalent to the following transformation:

$$\begin{aligned} \mathsf{lgp} &= \operatorname{logit}(p) \\ \pi(\mathsf{lgp}) &\propto \exp(-\mathsf{lgp})(1 + \exp(-\mathsf{lgp}))^{-2} \end{aligned}$$

The following statements fit the same model by using transformed parameters:

```
proc mcmc data=inputdata nmc=50000 thin=10 outpost=m2 seed=17
         monitor=(tau p mu theta) plot=trace;
   ods select ess tracepanel;
   array theta[5];

   parms theta:;
   parms lgp 0 ltau ;
   parms mu ;

   beginprior;
   prior ltau ~ egamma(shape=0.001,iscale=0.001);
   lp = -lgp - 2*log(1+exp(-lgp));
   prior lgp ~ general(lp);
   tau = exp(ltau);
   p = (1+exp(-lgp))**-1;
   prior mu ~ normal(0,prec=0.00000001);
   taub = tau/p;
   prior theta: ~ normal(mu,prec=taub);
   tauw = taub-tau;
   endprior;

   model y ~ normal(theta[grp],prec=tauw);
run;
ods graphics off;
```

The variable lgp is the logit transformation of p, and ltau is the log transformation of τ. The prior for ltau is egamma. The lp assignment statement evaluates the log density of $\pi(\mathsf{lgp})$. The tau and p assignment statements transform the parameters back to their original scales. The prior distributions for mu, theta, and the log likelihood in the MODEL statement remain unchanged. Trace plots (Output 52.12.4) and effective sample size (Output 52.12.5) both show significant improvements in the mixing for both p and τ.

Output 52.12.4 Trace Plots after Transformation

Output 52.12.4 *continued*

Trace Plots

Output 52.12.5 Effective Sample Sizes after Transformation

```
            Scatter Plot of Parameters on Original Scales

                        The MCMC Procedure

                     Effective Sample Sizes

                                Correlation
            Parameter          ESS          Time      Efficiency

            tau             1916.5        2.6089          0.3833
            p               2468.7        2.0253          0.4937
            mu              3273.9        1.5272          0.6548
            theta1          2184.5        2.2888          0.4369
            theta2          1938.1        2.5799          0.3876
            theta3          1947.1        2.5679          0.3894
            theta4          2115.8        2.3632          0.4232
            theta5          2152.0        2.3235          0.4304
```

The following statements generate Output 52.12.6 and Output 52.12.7:

```
title 'Scatter Plot of Parameters on Transformed Scales';

proc sgplot data=m2;
   yaxis label = 'logit(p)';
   xaxis label = 'log(tau)';
   scatter x = ltau y = lgp;
run;

title 'Scatter Plot of Parameters on Original Scales';

proc sgplot data=m2;
   yaxis label = 'p';
   xaxis label = 'tau';
   scatter x = tau y = p;
run;
```

Output 52.12.6 Scatter Plot of $\log(\tau)$ versus $\mathrm{logit}(p)$, After Transformation

Output 52.12.7 Scatter Plot of τ versus p, After Transformation

The scatter plot of $\log(\tau)$ versus $\text{logit}(p)$ shows a linear relationship between the two transformed parameters, and this explains the improvement in mixing. In addition, the transformations also help the Markov chain better explore in the original parameter space. Output 52.12.7 shows a scatter plot of τ versus p. The plot is similar to Output 52.12.3. However, note that tau has a far longer tail in Output 52.12.7, extending all the way to 0.4 as opposed to 0.15 in Output 52.12.3. This means that the second Markov chain can explore this dimension of the parameter more efficiently, and as a result, you are able to draw more precise inference with an equal number of simulations.

Example 52.13: Gelman-Rubin Diagnostics

PROC MCMC does not have the Gelman-Rubin test (see the section "Gelman and Rubin Diagnostics" on page 161) as a part of its diagnostics. The Gelman-Rubin diagnostics rely on parallel chains to test whether they all converge to the same posterior distribution. This example demonstrates how you can carry out this convergence test. The regression model from the section "Simple Linear Regression" on page 3478 is used. The model has three parameters: β_0 and β_1 are the regression coefficients, and σ^2 is the variance of the error distribution.

The following statements generate the data set:

```
title 'Simple Linear Regression, Gelman-Rubin Diagnostics';

data Class;
   input Name $ Height Weight @@;
   datalines;
Alfred  69.0 112.5   Alice   56.5  84.0   Barbara 65.3   98.0
Carol   62.8 102.5   Henry   63.5 102.5   James   57.3   83.0
Jane    59.8  84.5   Janet   62.5 112.5   Jeffrey 62.5   84.0
John    59.0  99.5   Joyce   51.3  50.5   Judy    64.3   90.0
Louise  56.3  77.0   Mary    66.5 112.0   Philip  72.0 150.0
Robert  64.8 128.0   Ronald 67.0 133.0    Thomas  57.5   85.0
William 66.5 112.0
;
```

To run a Gelman-Rubin diagnostic test, you want to start Markov chains at different places in the parameter space. Suppose that you want to start β_0 at 10, -15, and 0; β_1 at -5, 10, and 0; and σ^2 at 1, 20, and 50. You can put these starting values in the following init SAS data set:

```
data init;
   input Chain beta0 beta1 sigma2;
   datalines;
   1   10  -5   1
   2  -15  10  20
   3    0   0  50
;
```

The following statements run PROC MCMC three times, each with starting values specified in the data set init:

```
/* define constants */
%let nchain = 3;
%let nparm = 3;
%let nsim = 50000;
%let var = beta0 beta1 sigma2;

%macro gmcmc;
   %do i=1 %to &nchain;
      data _null_;
         set init;
         if Chain=&i;
         %do j = 1 %to &nparm;
            call symputx("init&j", %scan(&var, &j));
         %end;
         stop;
      run;

      proc mcmc data=class outpost=out&i init=reinit nbi=0 nmc=&nsim
                stats=none seed=7;
         parms beta0 &init1 beta1 &init2;
         parms sigma2 &init3;
         prior beta0 beta1 ~ normal(0, var = 1e6);
```

```
            prior sigma2 ~ igamma(3/10, scale = 10/3);
            mu = beta0 + beta1*height;
            model weight ~ normal(mu, var = sigma2);
         run;
      %end;
   %mend;

   ods listing close;
   %gmcmc;
   ods listing;
```

The macro variables nchain, nparm, nsim, and var define the number of chains, the number of parameters, the number of Markov chain simulations, and the parameter names, respectively. The macro GMCMC gets initial values from the data set init, assigns them to the macro variables init1, init2 and init3, starts the Markov chain at these initial values, and stores the posterior draws to three output data sets: out1, out2, and out3.

In the PROC MCMC statement, the INIT=REINIT option restarts the Markov chain after tuning at the assigned initial values. No burn-in is requested.

You can use the autocall macro GELMAN to calculate the Gelman-Rubin statistics by using the three chains. The GELMAN macro has the following arguments:

```
   %macro gelman(dset, nparm, var, nsim, nc=3, alpha=0.05);
```

The argument dset is the name of the data set that stores the posterior samples from all the runs, nparm is the number of parameters, var is the name of the parameters, nsim is the number of simulations, nc is the number of chains with a default value of 3, and alpha is the α significant level in the test with a default value of 0.05. This macro creates two data sets: _Gelman_Ests stores the diagnostic estimates and _Gelman_Parms stores the names of the parameters.

The following statements calculate the Gelman-Rubin diagnostics:

```
   data all;
      set out1(in=in1) out2(in=in2) out3(in=in3);
      if in1 then Chain=1;
      if in2 then Chain=2;
      if in3 then Chain=3;
   run;

   %gelman(all, &nparm, &var, &nsim);

   data GelmanRubin(label='Gelman-Rubin Diagnostics');
      merge _Gelman_Parms _Gelman_Ests;
   run;

   proc print data=GelmanRubin;
   run;
```

The Gelman-Rubin statistics are shown in Output 52.13.1.

Output 52.13.1 Gelman-Rubin Diagnostics of the Regression Example

```
              Simple Linear Regression, Gelman-Rubin Diagnostics

                              Between-      Within-                    Upper
           Obs    Parameter     chain        chain      Estimate      Bound

            1       beta0      5384.76      1168.64       1.0002      1.0001
            2       beta1         1.20         0.30       1.0002      1.0002
            3       sigma2     8034.41      2890.00       1.0010      1.0011
```

The Gelman-Rubin statistics do not reveal any concerns about the convergence or the mixing of the multiple chains. To get a better visual picture of the multiple chains, you can draw overlapping trace plots of these parameters from the three Markov chains runs.

The following statements create Output 52.13.2:

```
/* plot the trace plots of three Markov chains. */
%macro trace;
   %do i = 1 %to &nparm;
        proc sgplot data=all cycleattrs;
           series x=Iteration y=%scan(&var, &i) / group=Chain;
        run;
   %end;
%mend;
%trace;
```

Output 52.13.2 Trace Plots of Three Chains for Each of the Parameters

Output 52.13.2 *continued*

Simple Linear Regression, Gelman-Rubin Diagnostics

Simple Linear Regression, Gelman-Rubin Diagnostics

The trace plots show that three chains all eventually converge to the same regions even though they

started at very different locations. In addition to the trace plots, you can also plot the potential scale reduction factor (PSRF). See the section "Gelman and Rubin Diagnostics" on page 161 for definition and details.

The following statements calculate PSRF for each parameter. They use the GELMAN macro repeatedly and can take a while to run:

```
/* define sliding window size */
%let nwin = 200;
data PSRF;
run;

%macro PSRF(nsim);
   %do k = 1 %to %sysevalf(&nsim/&nwin, floor);
      %gelman(all, &nparm, &var, nsim=%sysevalf(&k*&nwin));
      data GelmanRubin;
           merge _Gelman_Parms _Gelman_Ests;
      run;

      data PSRF;
         set PSRF GelmanRubin;
      run;
   %end;
%mend PSRF;

options nonotes;
%PSRF(&nsim);
options notes;

data PSRF;
   set PSRF;
   if _n_ = 1 then delete;
run;

proc sort data=PSRF;
   by Parameter;
run;

%macro sepPSRF(nparm=&nparm, var=&var, nsim=&nsim);
   %do k = 1 %to &nparm;
      data save&k; set PSRF;
         if _n_ > %sysevalf(&k*&nsim/&nwin, floor) then delete;
         if _n_ < %sysevalf((&k-1)*&nsim/&nwin + 1, floor) then delete;
         Iteration + &nwin;
      run;

      proc sgplot data=save&k(firstobs=10) cycleattrs;
         series x=Iteration y=Estimate;
         series x=Iteration y=upperbound;
         yaxis label="%scan(&var, &k)";
      run;
   %end;
%mend sepPSRF;

%sepPSRF(nparm=&nparm, var=&var, nsim=&nsim);
```

Output 52.13.3 PSRF Plot for Each Parameter

Output 52.13.3 *continued*

PSRF is the square root of the ratio of the between-chain variance and the within-chain variance. A large PSRF indicates that the between-chain variance is substantially greater than the within-chain variance, so that longer simulation is needed. You want the PSRF to converge to 1 eventually, as it appears to be the case in this simulation study.

References

Aitkin, M., Anderson, D., Francis, B., and Hinde, J. (1989), *Statistical Modelling in GLIM*, Oxford: Oxford Science Publications.

Atkinson, A. C. (1979), "The Computer Generation of Poisson Random Variables," *Applied Statistics*, 28, 29–35.

Atkinson, A. C. and Whittaker, J. (1976), "A Switching Algorithm for the Generation of Beta Random Variables with at Least One Parameter Less Than One," *Proceedings of the Royal Society of London, Series A*, 139, 462–467.

Bacon, D. W. and Watts, D. G. (1971), "Estimating the Transition between Two Intersecting Straight Lines," *Biometrika*, 58, 525–534.

Berger, J. O. (1985), *Statistical Decision Theory and Bayesian Analysis*, Second Edition, New York: Springer-Verlag.

Box, G. E. P. and Cox, D. R. (1964), "An Analysis of Transformations," *Journal of the Royal Statistics Society, Series B*, 26, 211–234.

Carlin, B. P., Gelfand, A. E., and Smith, A. F. M. (1992), "Hierarchical Bayesian Analysis of Changepoint Problems," *Applied Statistics*, 41(2), 389–405.

Chaloner, K. (1994), "Residual Analysis and Outliers in Bayesian Hierarchical Models," in *Aspects of Uncertainty: A Tribute to D. V. Lindley*, 149–157, New York: Wiley.

Chaloner, K. and Brant, R. (1988), "A Bayesian Approach to Outlier Detection and Residual Analysis," *Biometrika*, 75(4), 651–659.

Cheng, R. C. H. (1978), "Generating Beta Variates with Non-integral Shape Parameters," *Communications ACM*, 28, 290–295.

Congdon, P. (2003), *Applied Bayesian Modeling*, John Wiley & Sons.

Crowder, M. J. (1978), "Beta-Binomial Anova for Proportions," *Applied Statistics*, 27, 34–37.

Draper, D. (1996), "Discussion of the Paper by Lee and Nelder," *Journal of the Royal Statistical Society, Series B*, 58, 662–663.

Eilers, P. H. C. and Marx, B. D. (1996), "Flexible Smoothing with *B*-Splines and Penalties," *Statistical Science*, 11, 89–121, with discussion.

Finney, D. J. (1947), "The Estimation from Individual Records of the Relationship between Dose and Quantal Response," *Biometrika*, 34, 320–334.

Fisher, R. A. (1935), "The Fiducial Argument in Statistical Inference," *Annals of Eugenics*, 6, 391–398.

Fishman, G. S. (1996), *Monte Carlo: Concepts, Algorithms, and Applications*, New York: John Wiley & Sons.

Gaver, D. P. and O'Muircheartaigh, I. G. (1987), "Robust Empirical Bayes Analysis of Event Rates," *Technometrics*, 29, 1–15.

Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2004), *Bayesian Data Analysis*, Second Edition, London: Chapman & Hall.

Gentleman, R. and Geyer, C. J. (1994), "Maximum Likelihood for Interval Censored Data: Consistency and Computation," *Biometrika*, 81, 618–623.

Gilks, W. (2003), "Adaptive Metropolis Rejection Sampling (ARMS)," software from MRC Biostatistics Unit, Cambridge, UK, http://www.maths.leeds.ac.uk/~wally.gilks/adaptive.rejection/web_page/Welcome.html.

Gilks, W. R. and Wild, P. (1992), "Adaptive Rejection Sampling for Gibbs Sampling," *Applied Statistics*, 41, 337–348.

Holmes, C. C. and Held, L. (2006), "Bayesian Auxiliary Variable Models for Binary and Multinomial Regression," *Bayesian Analysis*, 1(1), 145–168, http://ba.stat.cmu.edu/journal/2006/vol01/issue01/held.pdf.

Ibrahim, J. G., Chen, M. H., and Sinha, D. (2001), *Bayesian Survival Analysis*, New York: Springer-Verlag.

Kass, R. E., Carlin, B. P., Gelman, A., and Neal, R. (1998), "Markov Chain Monte Carlo in Practice: A Roundtable Discussion," *The American Statistician*, 52, 93–100.

Krall, J. M., Uthoff, V. A., and Harley, J. B. (1975), "A Step-up Procedure for Selecting Variables Associated with Survival," *Biometrics*, 31, 49–57.

Kuhfeld, W. F. (2004), *Conjoint Analysis*, Technical report, SAS Institute Inc., `http://support.sas.com/resources/papers/tnote/tnote_marketresearch.html`.

Matsumoto, M. and Kurita, Y. (1992), "Twisted GFSR Generators," *ACM Transactions on Modeling and Computer Simulation*, 2(3), 179–194.

Matsumoto, M. and Kurita, Y. (1994), "Twisted GFSR Generators," *ACM Transactions on Modeling and Computer Simulation*, 4(3), 254–266.

Matsumoto, M. and Nishimura, T. (1998), "Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator," *ACM Transactions on Modeling and Computer Simulation*, 8, 3–30.

McGrath, E. J. and Irving, D. C. (1973), *Techniques for Efficient Monte Carlo Simulation, Volume II: Random Number Generation for Selected Probability Distributions*, Technical report, Science Applications Inc., La Jolla, CA.

Michael, J. R., Schucany, W. R., and Haas, R. W. (1976), "Generating Random Variates Using Transformations with Multiple Roots," *American Statistician*, 30(2), 88–90.

Pregibon, D. (1981), "Logistic Regression Diagnostics," *Annals of Statistics*, 9, 705–724.

Ripley, B. D. (1987), *Stochastic Simulation*, New York: John Wiley & Sons.

Robert, C. (1995), "Simulation of Truncated Normal Variables," *Statistics and Computing*, 5, 121–125.

Roberts, G. O., Gelman, A., and Gilks, W. R. (1997), "Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms," *Annual 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.

Rubin, D. B. (1981), "Estimation in Parallel Randomized Experiments," *Journal of Educational Statistics*, 6, 377–411.

Schervish, M. J. (1995), *Theory of Statistics*, New York: Springer-Verlag.

Sharples, L. (1990), "Identification and Accommodation of Outliers in General Hierarchical Models," *Biometrika*, 77, 445–453.

Spiegelhalter, D. J., Thomas, A., Best, N. G., and Gilks, W. R. (1996), "BUGS Examples, Volume 2, Version 0.5, (version ii)," .

Subject Index

Syntax Index

Your Turn

We welcome your feedback.

- If you have comments about this book, please send them to **yourturn@sas.com**. Include the full title and page numbers (if applicable).

- If you have comments about the software, please send them to **suggest@sas.com**.

SAS® Publishing Delivers!

Whether you are new to the work force or an experienced professional, you need to distinguish yourself in this rapidly changing and competitive job market. SAS® Publishing provides you with a wide range of resources to help you set yourself apart. Visit us online at support.sas.com/bookstore.

SAS® Press

Need to learn the basics? Struggling with a programming problem? You'll find the expert answers that you need in example-rich books from SAS Press. Written by experienced SAS professionals from around the world, SAS Press books deliver real-world insights on a broad range of topics for all skill levels.

support.sas.com/saspress

SAS® Documentation

To successfully implement applications using SAS software, companies in every industry and on every continent all turn to the one source for accurate, timely, and reliable information: SAS documentation. We currently produce the following types of reference documentation to improve your work experience:

- Online help that is built into the software.
- Tutorials that are integrated into the product.
- Reference documentation delivered in HTML and PDF – **free** on the Web.
- Hard-copy books.

support.sas.com/publishing

SAS® Publishing News

Subscribe to SAS Publishing News to receive up-to-date information about all new SAS titles, author podcasts, and new Web site features via e-mail. Complete instructions on how to subscribe, as well as access to past issues, are available at our Web site.

support.sas.com/spn

§sas | THE POWER TO KNOW®