Question The following data are taken from Bishop, Fienberg and Holland (*Discrete Multivariate Analysis*, MIT Press, 1975, Section 5.2.8, Exercise 4).

Active	Passive Participant					
Participant	R	S	T	U	V	W
R	—	1	5	8	9	0
S	29	—	14	46	4	0
T	0	0	—	0	0	0
U	2	3	1	—	38	2
V	0	0	0	0	—	1
W	9	25	4	6	13	—

The data involve displays between squirrel monkeys (there are six monkeys labeled R through W). Each display has an active and a passive participant giving the two-way classification of the table. Since a monkey cannot display toward itself, the diagonal cells of the table are impossible (structural zeros in the jargon of categorical data analysis).

We wish to fit the following Bayesian model to these data (the same model for which Bishop, Fienberg, and Holland discuss frequentist analysis in their Section 5.2), called the *quasi-independence* model. We assume Poisson sampling, so the likelihood is

$$\prod_{i=1}^{6} \prod_{\substack{j=1 \\ j \neq i}}^{6} \mu_{ij}^{x_{ij}} e^{-\mu_{ij}} \tag{1}$$

where x_{ij} is the data (given in the table above) and the parameters μ_{ij} have the log-linear additive form

$$\log(\mu_{ij}) = \alpha + \beta_i + \gamma_j$$

where in order to obtain identifiability we must have the constraints

$$\sum_{i=1}^{6} \beta_i = 0 \tag{2}$$

and

$$\sum_{j=1}^{6} \gamma_j = 0. \tag{3}$$

Because the third row of the data is all zero, an improper prior would give an improper posterior. So we can't use a flat prior. The conjugate prior we use looks like a gamma with density proportional to

$$\prod_{i=1}^{6} \prod_{\substack{j=1 \\ j \neq i}}^{6} \mu_{ij}^{\epsilon} e^{-\delta \mu_{ij}} \tag{4}$$

but it isn't gamma because we do not consider the μ_{ij} the fundamental random variables but the α, β_i, and γ_j.

Thus the unnormalized posterior density is the product of (1) and (3) considered as a function of the α, β_i, and γ_j. The specific values of the hyperparameters we use are $\epsilon = 0.2$ and $\delta = 0.3$.

Estimate the posterior mean of each of the 13 parameters, α, β_i, and γ_j, giving a Monte Carlo standard error for each estimated quantity. Try to get at least two significant figures (Monte Carlo standard error a hundredth the size of the estimated quantity).

Describe you algorithm in sufficient detail so that it could be reimplemented by and expert in MCMC.

To avoid typing mistakes, the data are in the file

`http://www.stat.umn.edu/~charlie/monkey.dat`

Also provide a link to a machine readable form of any computer code you used for the graders to examine if necessary.