STA 6352, Report 6.6

Carson Slater *Baylor University*

Problem Setup

Use Stan to reproduce the results in Example 6.9 in Section 7.2.2.

Example 6.9

Suppose people at eight sites around the country have been given a treatment and are monitored for adverse events that are assumed to be distributed Poisson.

Specifically, we have counts corresponding to person-time (weeks, months, etc.) for the ith site with Poisson rate θ_i, $i = 1, \ldots, 8$. We are interested in estimating the θ_i's. The data are given below:

Table 1: Adverse events reported for count y_i and person-time t_i

Site	1	2	3	4	5	6	7	8
y_i	47.6	54.1	44.1	41.1	51.1	55.3	41.7	43.7
t_i	10	9	8	11	9	8	10	11

Suppose we wish to fit the following model to this data, where we take t_i to be a known constant:

$$\begin{aligned} f(y_i|\theta_i) &= \frac{e^{-\theta_i t_i}(\theta_i t_i)^{y_i}}{y_i!}, \quad y_i > 0, \theta_i > 0; \\ \pi(\theta_i|\alpha,\beta) &= \frac{\beta^\alpha}{\Gamma(\alpha)}\theta_i^{\alpha-1}e^{-\beta\theta_i}, \quad \alpha > 0, \beta > 0; \\ \pi(\beta) &= \frac{d^c}{\Gamma(c)}\beta^{c-1}e^{-d\beta}, \quad c > 0, d > 0; \\ \pi(\alpha) &= \mu e^{-\mu\alpha}, \quad \mu > 0. \end{aligned}$$

We want the marginal posterior distributions for the θ_i's. Despite the conjugate relationships (gamma for Poisson, inverse-gamma for gamma given α) no closed form expression exists for the desired marginal posteriors. However, the full conditionals are readily available. To derive them, start with the complete Bayesian model:

$$\left[\prod_{i=1}^{k} f(y_i|\theta_i)\pi(\theta_i|\alpha,\beta)\right]\pi(\beta)\pi(\alpha)$$

We can get the full conditional distributions by dropping appropriate terms and normalizing. It is not hard to prove the following results:

The full conditional for θ_i given all other parameters and the data $y = (y_1, \ldots, y_k)$:

$$\begin{aligned}
\pi(\theta_i|\theta_{-i}, \alpha, \beta, y) &\propto \left[\prod_{i=1}^{k} f(y_i|\theta_i)\pi(\theta_i|\alpha, \beta)\right] \pi(\beta)\pi(\alpha) \\
&\propto f(y_i|\theta_i)\pi(\theta_i|\alpha, \beta) \\
&\propto \theta_i^{y_i} e^{-\theta_i t_i} \theta_i^{\alpha-1} e^{-\beta\theta_i} \\
&\propto \theta_i^{y_i+\alpha-1} e^{-(\beta+t_i)\theta_i} \\
&\sim \text{Gamma}(y_i + \alpha, \beta + t_i)
\end{aligned}$$

The full conditional for β is:

$$\begin{aligned}
\pi(\beta|\theta, \alpha, y) &\propto \left[\prod_{i=1}^{k} f(y_i|\theta_i)\pi(\theta_i|\alpha, \beta)\right] \pi(\beta)\pi(\alpha) \\
&\propto \left[\prod_{i=1}^{k} \pi(\theta_i|\alpha, \beta)\right] \pi(\beta) \\
&\propto \prod_{i=1}^{k} \frac{\beta^\alpha}{\Gamma(\alpha)} \theta_i^{\alpha-1} e^{-\beta\theta_i} \cdot \frac{d^c}{\Gamma(c)} \beta^{c-1} e^{-d\beta} \\
&\propto \beta^{k\alpha+c-1} e^{-\beta\left(\sum_{i=1}^{k} \theta_i + d\right)} \\
&\sim \text{Gamma}\left(k\alpha + c, \sum_{i=1}^{k} \theta_i + d\right)
\end{aligned}$$

Similarly, the full conditional for α is:

$$\begin{aligned}
\pi(\alpha|\theta, \beta, y) &\propto \left[\prod_{i=1}^{k} f(y_i|\theta_i)\pi(\theta_i|\alpha, \beta)\right] \pi(\beta)\pi(\alpha) \\
&\propto \left[\prod_{i=1}^{k} \frac{\theta_i^{\alpha-1}}{\Gamma(\alpha)\beta^\alpha}\right] e^{-\alpha/\mu}.
\end{aligned}$$

This last full-conditional is not the kernel of a standard family. Indeed, there is no choice for $\pi(\alpha)$ that would yield a familiar density. However, the full-conditional for α is log-concave in α, the necessary condition for application of the adaptive rejection sampling (ARS) algorithm.

Solution

We implement the model in Stan as follows:

```
model_string <- "
data {
  int<lower=1> K;          // Number of sites
```

```
  int<lower=0> y[K];       // Adverse event counts
  real<lower=0> t[K];      // Person-time for each site
  real<lower=0> c;
  real<lower=0> d;
  real<lower=0> mu;
}

parameters {
  real<lower=0> alpha;
  real<lower=0> beta;
  real<lower=0> theta[K];
}

model {
  // Priors
  alpha ~ exponential(mu);
  beta ~ gamma(c, d);

  for (i in 1:K)
    theta[i] ~ gamma(alpha, beta);

  // Likelihood
  for (i in 1:K)
    y[i] ~ poisson(theta[i] * t[i]);
}

// stan_poisson_model.stan
"

write_lines(model_string, "stan_poisson_model.stan")
```

The results can be seem as follows:

	mean	se_mean	sd	2.5%	25%	50%	75%	97.5%
alpha	4.1601	0.0266	1.5836	1.6822	3.0236	3.9415	5.1077	7.8441
beta	0.8298	0.0058	0.3345	0.3132	0.5833	0.7822	1.0305	1.5920
theta[1]	4.8172	0.0081	0.6809	3.6123	4.3328	4.7892	5.2644	6.2287
theta[2]	5.9234	0.0113	0.7963	4.4963	5.3696	5.8829	6.4280	7.5725
theta[3]	5.4556	0.0099	0.8036	3.9919	4.8931	5.4213	5.9646	7.1790
theta[4]	3.8035	0.0076	0.5813	2.7349	3.4147	3.7770	4.1631	5.0526
theta[5]	5.6192	0.0092	0.7492	4.2655	5.1041	5.5841	6.0971	7.1818
theta[6]	6.7060	0.0106	0.8802	5.0927	6.1079	6.6740	7.2347	8.5554
theta[7]	4.2642	0.0080	0.6321	3.1062	3.8289	4.2296	4.6770	5.6062
theta[8]	4.0699	0.0074	0.5783	2.9851	3.6756	4.0451	4.4441	5.2921

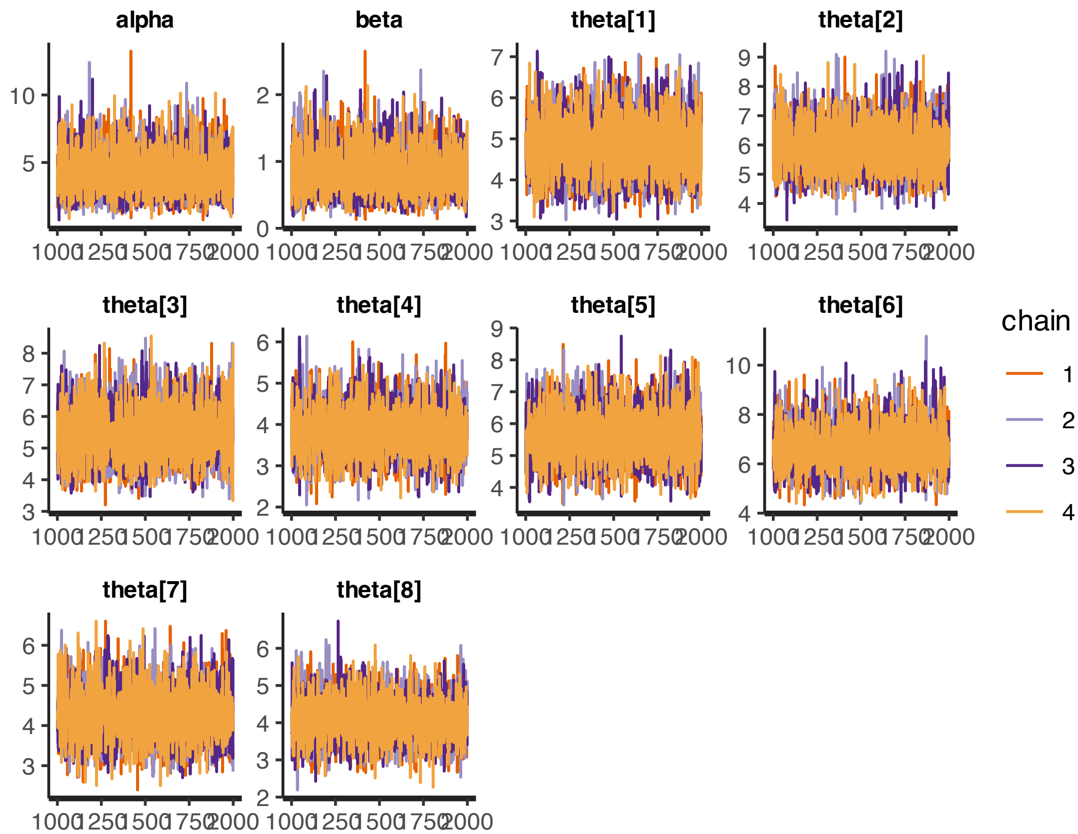
alpha
beta
theta[1]
theta[2]
theta[3]
theta[4]
theta[5]
theta[6]
theta[7]
theta[8]
chain
1
2
3
4

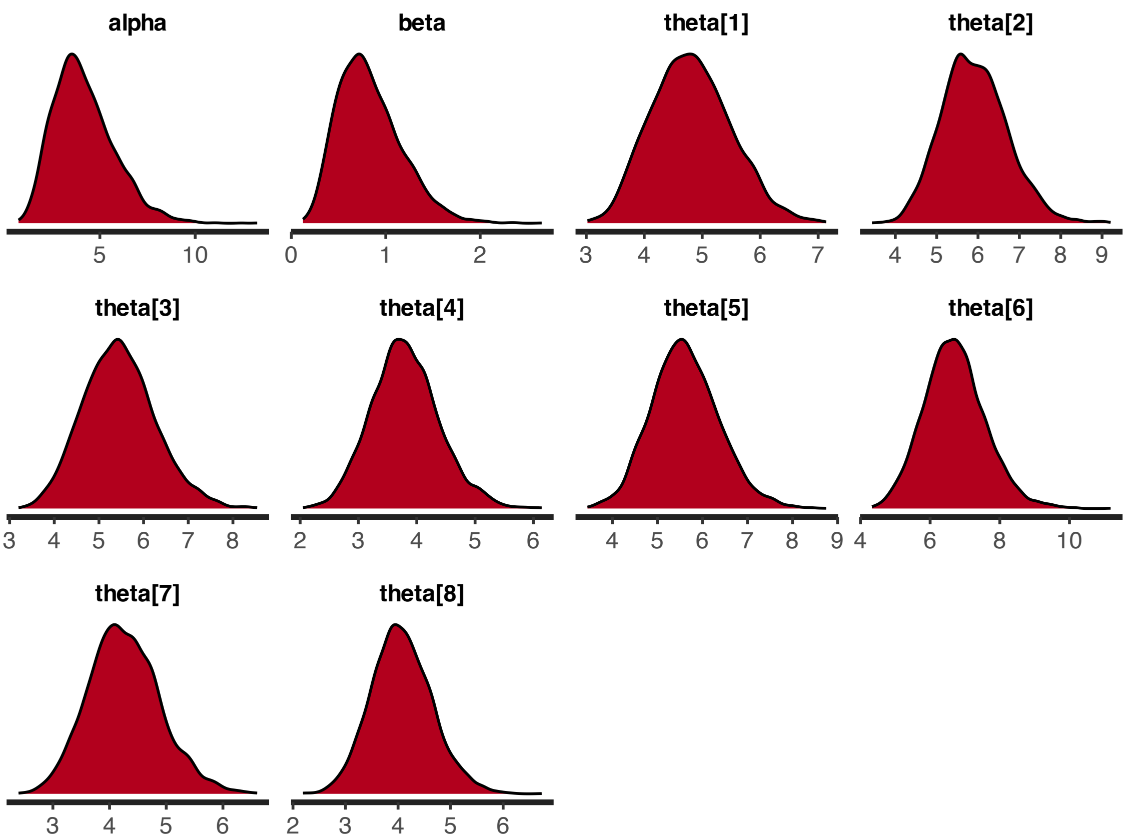
alpha
beta
theta[1]
theta[2]
theta[3]
theta[4]
theta[5]
theta[6]
theta[7]
theta[8]