Background and notation for simple Bayes sampling example

Previously we considered maximum likelihood estimation for the location θ and scale s of a Cauchy model for the data

```
1.2 2.4 1.3 1.3 0.0 1.0 1.8 0.8 4.6 1.4
```

I.e. the model is that $y_1, \ldots, y_{10}$ as above are modelled as iid according to

$$p(y_i|\theta, s) = \frac{1}{\pi s} \cdot \frac{1}{1 + \left(\frac{y_i - \theta}{s}\right)^2}$$

and likelihood

$$L(\underline{y}|\theta, s) = \prod_{i=1^n} p(y_i|\theta, s)$$

We now consider a simple sampling scheme to explore the posterior distribution of the parameters θ and s in this model, when θ and s have independent priors

$$\begin{array}{rcl} \theta & \sim & N(\theta_0, \sigma_0^2) \\ s^2 & \sim & 1/\chi^2_{df} \end{array}$$

First, we consider generating a sample from the posterior $\pi(\theta, s|\underline{y})$ by the simple rejection scheme

Rejection for Bayes.

- *Generate θ from $\pi(\theta)$;*
- *Generate U from $Unif[0, L(y|\hat{\theta})\pi(\theta)]$;*
- *If $U < L(y|\theta)\pi(\theta)$, accept θ as an observation from $p(\theta|y)$;*
  otherwise reject θ and start over.

(in the example we drop θ from steps 2 and 3 and instead generate $U \sim Unif[0, 1]$ and check $U < L(y|\theta)/L(y|\hat{\theta})$).

Second, we can use importance sampling/SIR to generate a sample from a posterior using a different prior, as follows: Given an iid sample

$$\theta_1, \ldots, \theta_M \text{ from } p(\theta|\underline{y}) \sim L(\underline{y}|\theta)\pi(\theta),$$

if we resample

$$\theta_1^*, \ldots, \theta_n^* \text{ as i.i.d. from } \theta_1, \ldots, \theta_M, \text{ with probabilities } q_1, \ldots q_M$$

where $q_i \propto \pi^*(\theta_i)/\pi(\theta_i)$ the result will be

$$\theta_1^*, \ldots, \theta_n^* \text{ iid, } \overset{approx}{\sim} p^*(\theta|\underline{y}) \propto L(\underline{y}|\theta)\pi^*(\theta).$$