As in the example with non-linear least squares, it will be convenient to define a named parameter set (PARMSET) so we can "poke" a test set of values into it for evaluating (6.14).

CVMODEL has three ways for handling the scale factors (the Λ). The one described above is DMATRIX=MARGINALIZED. The δ is input using the PDF option, and the k with the DFC option.

There are three methods for doing inference on θ– two of which we've already seen, one which is new.

1. Random walk MH
2. Independence chain MH
3. Importance sampling

Importance sampling is discussed in Koop in section 4.3.3 and in the RATS *User's Guide* in section 13.6. It's very similar to independence chain MH. The difference is that independence chain MH overweights "desirable" θ (ones where the posterior is high relative to the probability that we draw it) by staying on them for multiple sweeps, while importance sampling keeps all draws but assigns weights to them. Importance sampling has the advantage that the statistical properties are easier to determine since they're just weighted independent random variables, while MH draws are correlated. However, it can only be used if there isn't a need to do Gibbs sampling. Here, we don't need that since we have a marginal density for θ and a conditional density for $\Lambda|\theta$. If we had only $\theta|\Lambda$ and $\Lambda|\theta$, we would need to do Gibbs sampling and couldn't use importance sampling.

If you can do importance sampling, you should prefer it over independence chain MH. There may, however, be situations in which random walk MH is preferred over importance sampling. This is when it's hard to come up with a good importance function (the direct sampling density for θ). If there are regions where the posterior is high and the probability of hitting them with a draw is low, such points would get very high weights, giving you an weighted average that might be dominated by just a few points. There's a simple measure of whether you have a problem like this: $\left(\sum_{i=1}^{S} w_i\right)^2 / \sum_{i=1}^{S} w_i^2$ should ideally be a high percentage of the number of draws. If it's a small percentage of that (below 10%), your sampler is probably not working very well. Correcting this might require fattening the tails on your candidate density. If no change like that works, then you'll have to switch methods.

Because importance sampling creates records of draws that need weighting, you need to do slightly different calculations in the post processing. For instance, the **DENSITY**instruction requires a WEIGHT option, as does **STATISTICS**.