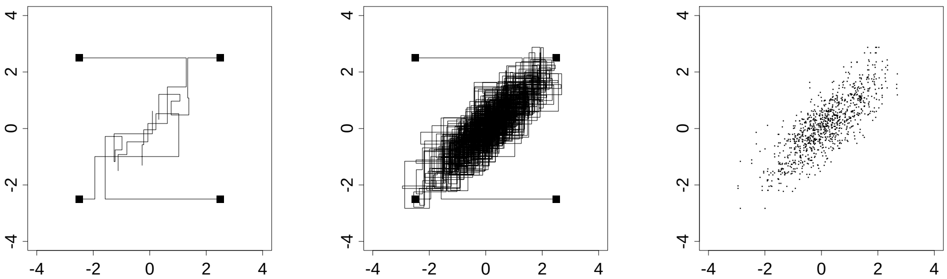

Figure 18.7 *Four independent sequences of the Gibbs sampler for a simple example with two parameters. Initial values of the $n_{\text{chains}} = 4$ sequences are indicated by solid squares. (a) First 10 iterations, showing the component-by-component updating of the Gibbs iterations. (b) After 500 iterations, when the sequences have reached approximate convergence. (c) The points from the second halves of the sequences.*

(b) Choose some number n_{iter} of iterations (typically a somewhat large number such as 1000). For each iteration:

   Update the parameters, or batches of parameters, one at a time. For each parameter or batch, take a random simulation draw given the data and the current estimate of all the other parameters. (We illustrate with some examples below.)

2. Evaluate the mixing of the simulated chains using the $\hat{R}$ summary, which we have already discussed in Section 16.4 in the context of interpreting the output from Bugs models.
3. If convergence is poor, run longer or alter the model, following the advice in Section 16.9.

The key part of this algorithm is the sequential updating step. Bugs performs it automatically, but here we will show how to compute Gibbs updates "manually" in R for multilevel linear regressions. Our purpose is not to set you up to program these yourself but rather to give enough insight that you can understand roughly how Bugs works, and thus better diagnose and fix problems when Bugs is not working so well.

We present in this section the steps of Gibbs sampling for a series of multilevel linear regressions: first a model with no predictors, then including a predictor at the individual level, then adding one at the group level.

The basic Gibbs sampler structure described here works for multilevel regressions, with the new twist that the regression coefficients can be estimated using an adaptation of classical least squares regression. (Model (18.10) can be considered as a special case of regression with only an intercept and no slope parameters, but this case is so simple that least squares matrix computations were not required.)

Gibbs sampler for a multilevel model with no predictors

We first go through the steps of the Gibbs sampler—mathematically and as programmed in R—for the multilevel model (18.10) with data in groups and no predictors.

The Gibbs sampler starts with initial values for all the parameters and then updates the parameters in turn, giving each a random estimate based on the data and the current guess of the other parameters in the model. For the simple model we are considering here, the Gibbs updating steps are: