STA 6352, Report 9.4

Carson Slater *Baylor University*

9.1.4 Example: Repeated measures

- This example is an extension of the Example 2.2.3. Here we also include a time-invariant treatment variable. Thus for the jth observation from the ith subject, the data model is

$$y_{ij} = \mu + a_i + \beta x_i + \varepsilon_{ij} \tag{9.1.5}$$

where $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$, $a_i \sim \mathcal{N}(0, \sigma_a^2)$, and x_i is a binary variable. Here $x_i = 1$ indicates the ith subject received the treatment and $x_i = 0$ indicates receipt of the control. There are $i = 1, 2, \ldots, n$ subjects and $j = 1, 2, \ldots, J$ repeated measurements. We use diffuse normal priors for μ and β, specifically, $\mu \sim \mathcal{N}(0, 1000)$ and $\beta \sim \mathcal{N}(0, 1000)$. For the standard deviations, instead of the uniform prior discussed in Example 2.2.3, we consider a half-normal prior (i.e., a normal truncated at 0). For small numbers of replications, the half-normal is recommended.[1]

- Suppose that we are interested in the Type I error probability ($H_0 : \beta = 0$) and the power at $\beta = 5$ for a sample of size $n = 40$ (20 with treatment and 20 with control) and $J = 4$ repeated measures.
- **Algorithm:**[2]

1. Select parameter values and MCMC specifics: $\mu = 80$, $\sigma^2 = 20$, $\sigma_a^2 = 1$, $\beta = 0$ (null case), and $\beta = 5$ (for power), $M = 2000$, iterations $= 15000$, and burn-in iterations $= 5000$.
2. Generate data using parameters in step 1.
3. Call JAGS from R, specifying parameters to monitor.
4. For each data set determine if $\Pr(\beta > 0 \mid \text{data}) > 1 - \alpha$. If so, this iteration is marked as a rejection. Also keep track of the posterior means and the coverage for credible intervals.
5. Compute empirical average of rejections. This is the Type I error probability when $\beta = 0$, otherwise, it is the power.

- The figure below provides the power for sample sizes from 30 to 100. For a power of 80% a sample of 70 subjects is required. To verify the Type I error rate is controlled we also simulated from the null for a sample of 70 and found the probability of a Type I error to be 0.047. Thus a sample of 70 yields frequentist power of 0.8 and Type I error rate of 0.05.

Report 9.4.

[1] Recall Section 7.3.2.
[2] See the WinBUGS and R code starting on p. 295.

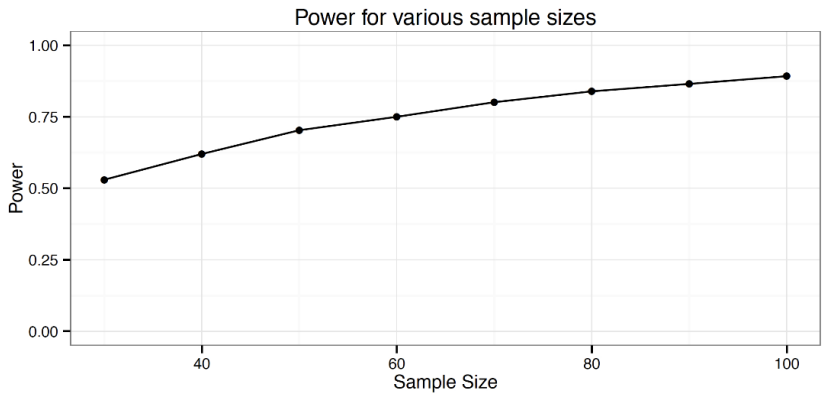

Figure 9.6: Power for $30 \leq n \leq 100$. (15,000 iterations).

Figure 9.6: Power for $30 \leq n \leq 100$. (15,000 iterations).

(a) Write an R script to produce a simulation like that in the last example, using JAGS. Report the same simulation summaries. See the code in Section O.16 of the course notes.

Power	Mean_Post	Coverage	Avg_Length
1	5.015953	0.96	3.517635

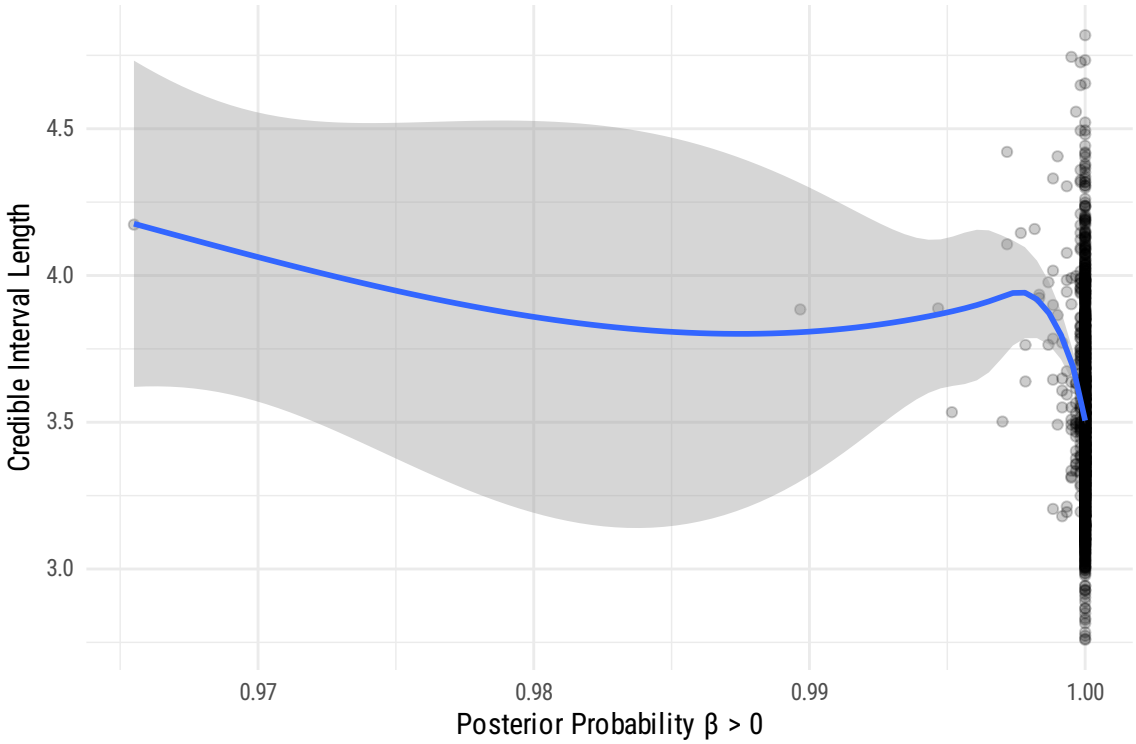

(b) Use the simulation summary graph you coded in the last report to summarize the performance of the credible set for β.

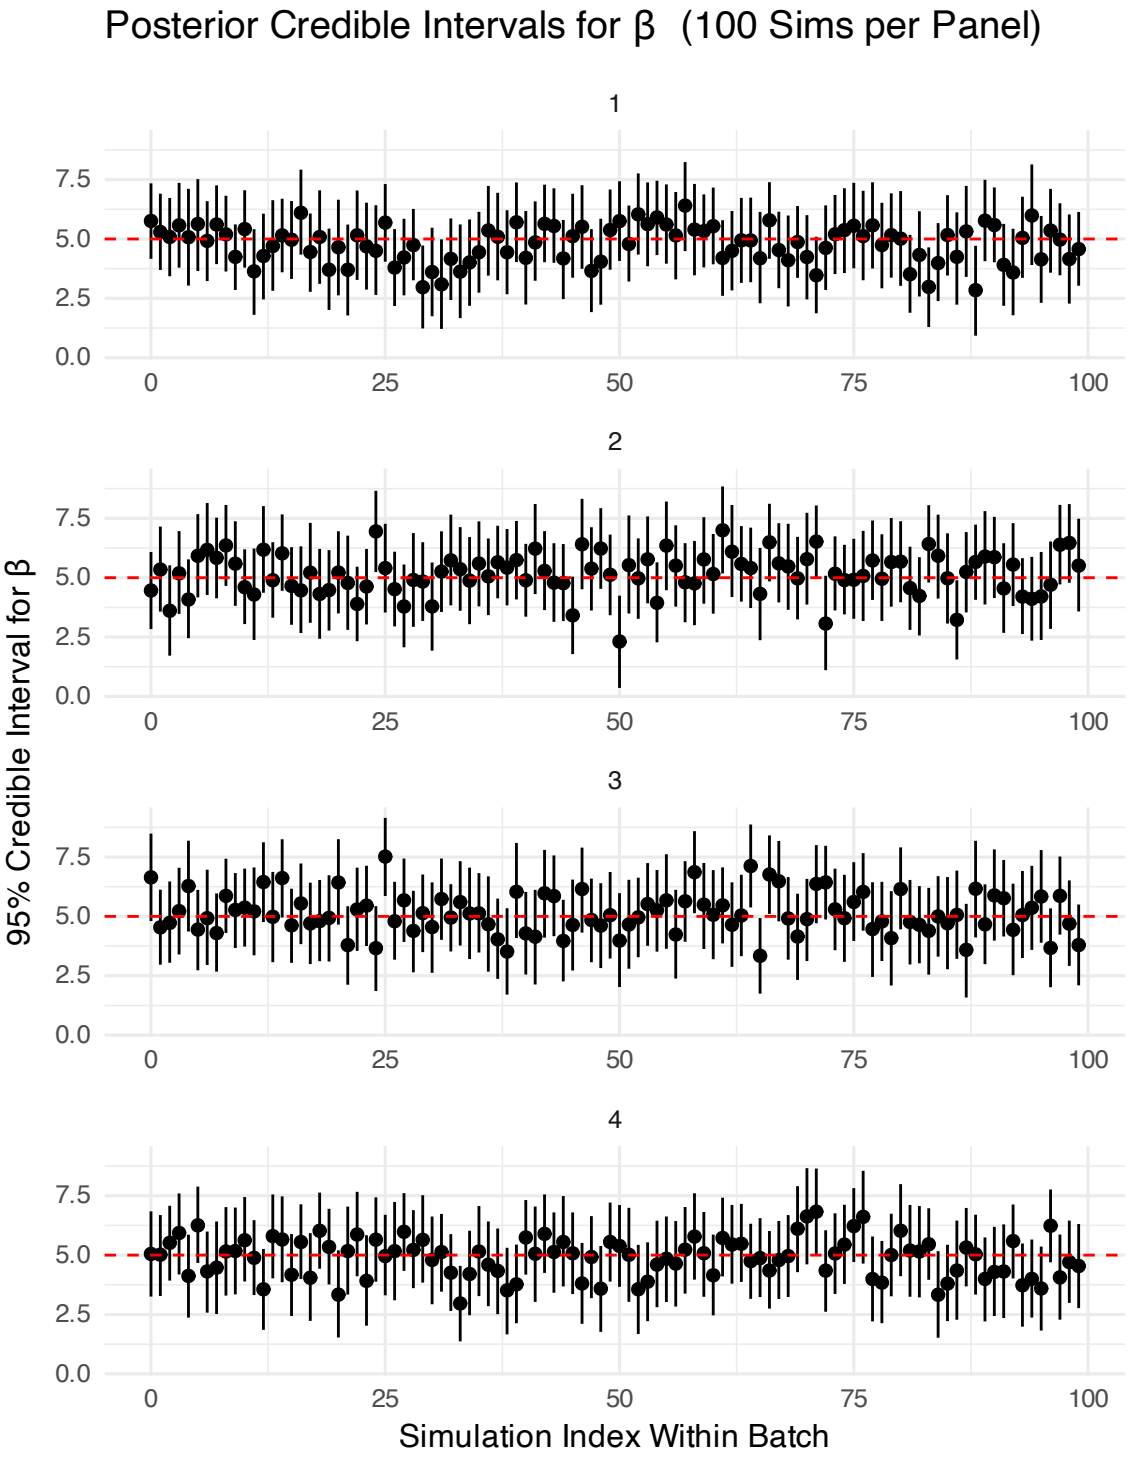

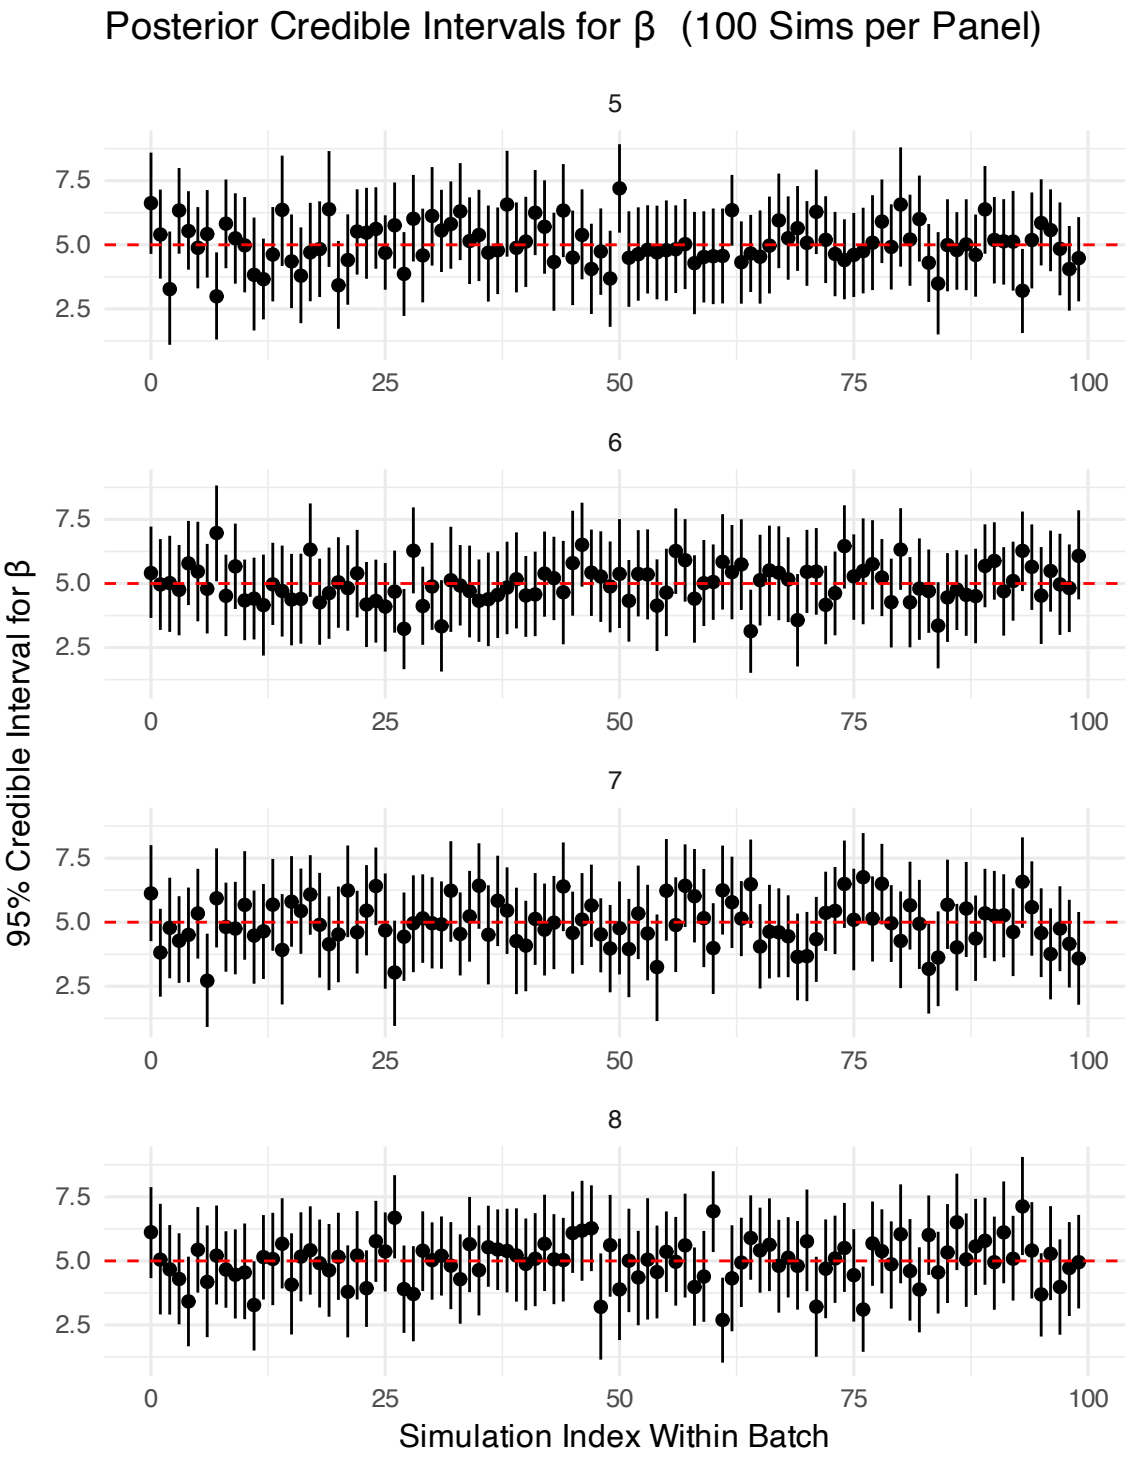
Posterior Credible Intervals for β (100 Sims per Panel)
5
6
7
8
7.5
5.0
2.5
0
25
50
75
100
95% Credible Interval for β
Simulation Index Within Batch

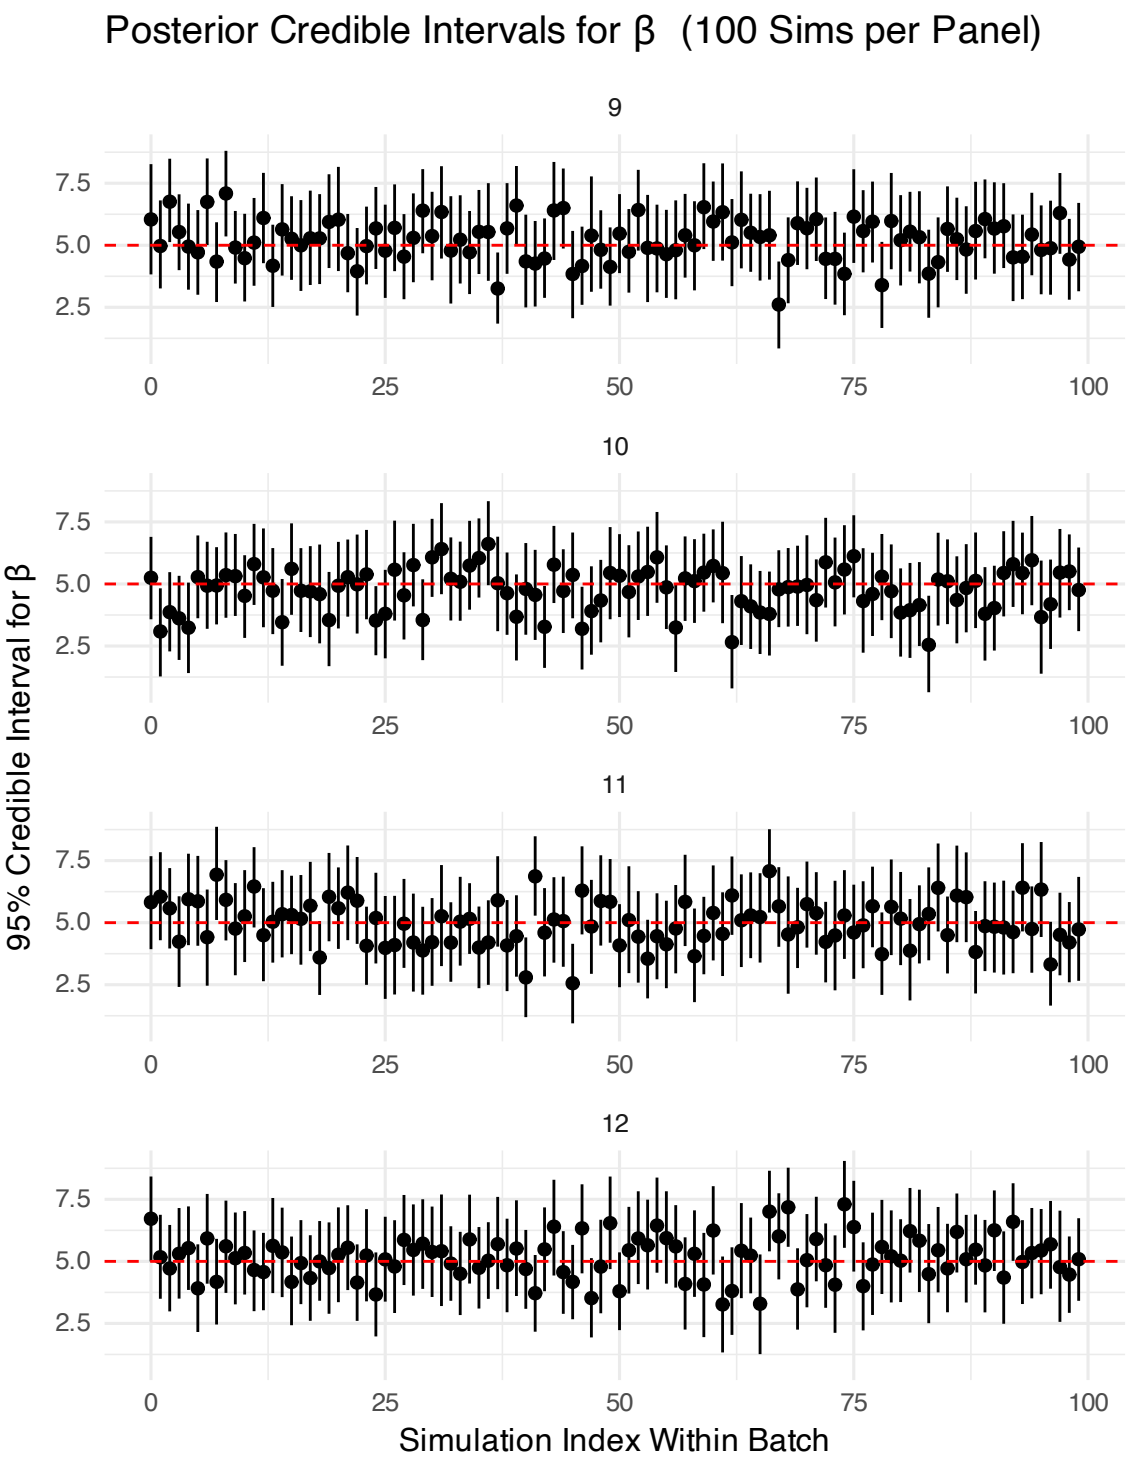
Posterior Credible Intervals for β (100 Sims per Panel)
9
10
11
12
7.5
5.0
2.5
0
25
50
75
100
95% Credible Interval for β
Simulation Index Within Batch

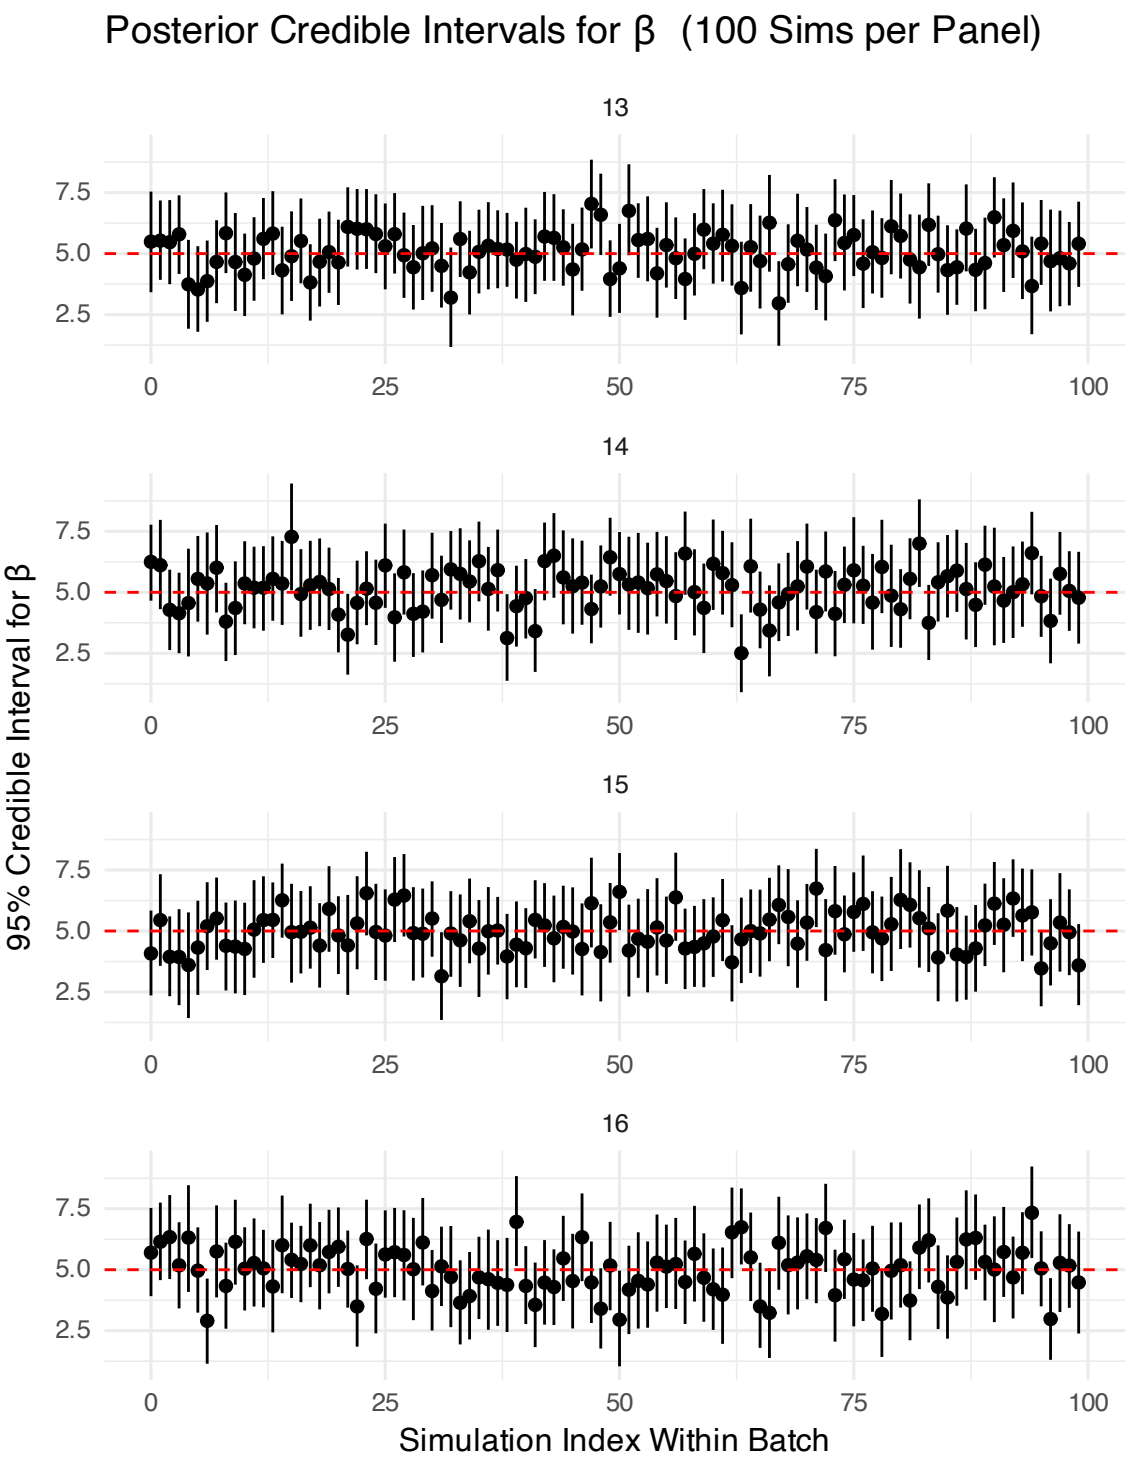
Posterior Credible Intervals for β (100 Sims per Panel)
13
14
15
16
7.5
5.0
2.5
0
25
50
75
100
95% Credible Interval for β
Simulation Index Within Batch

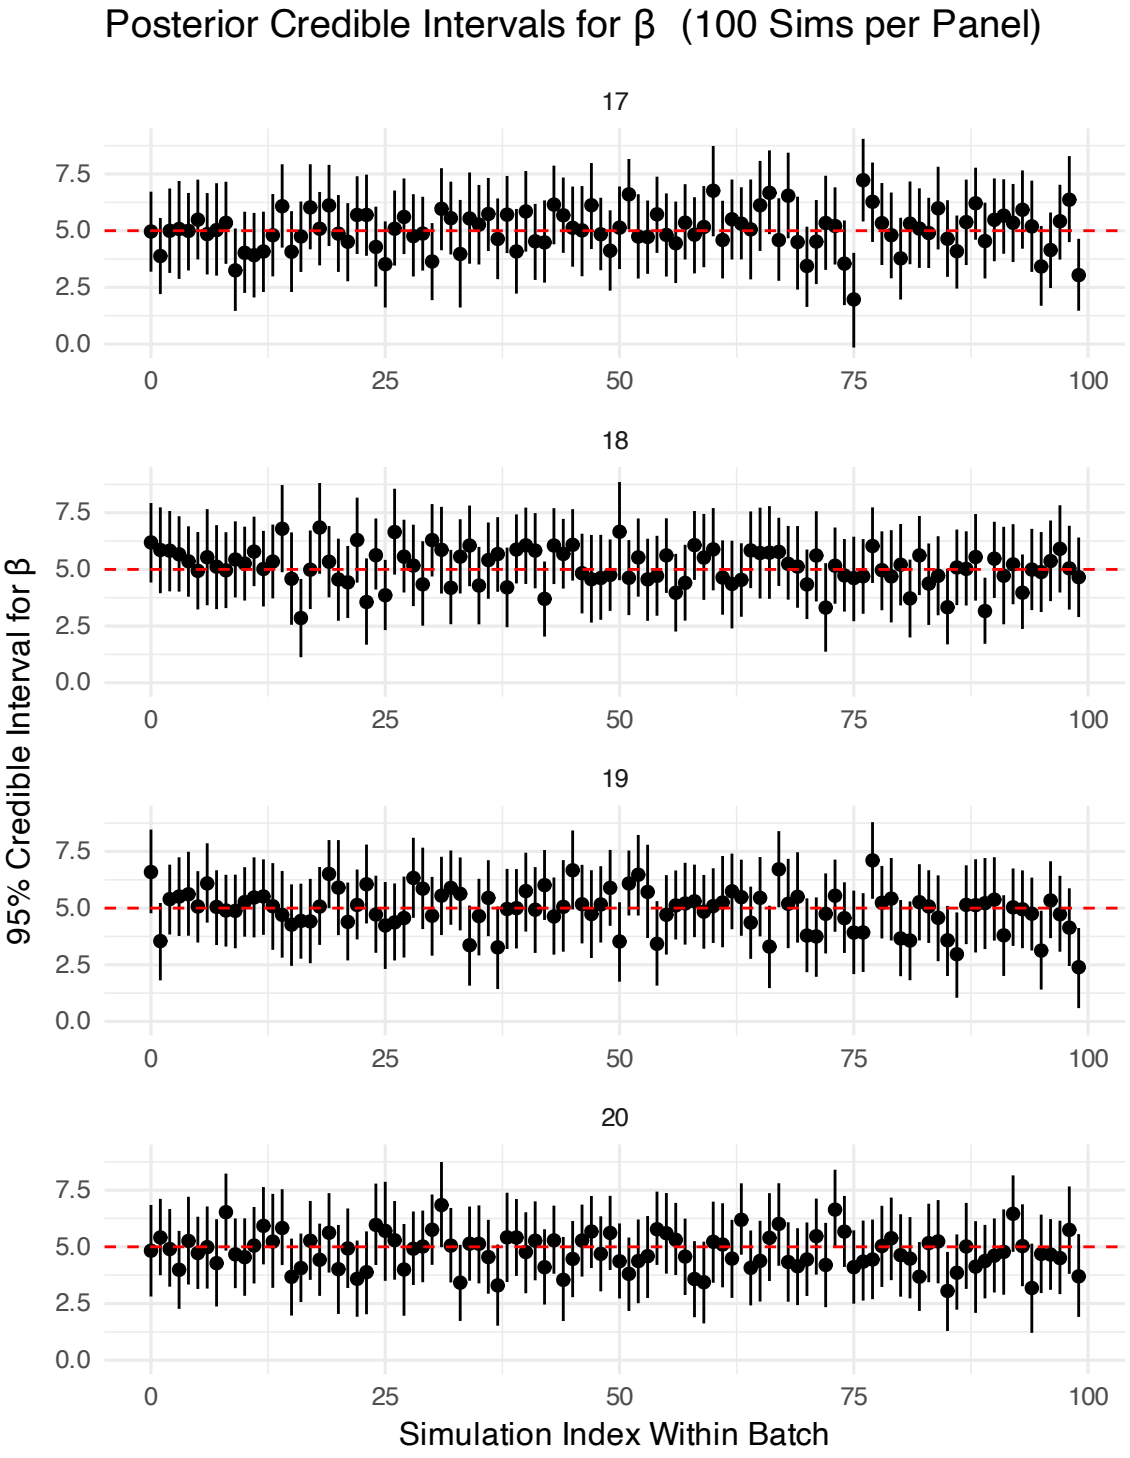
Posterior Credible Intervals for β (100 Sims per Panel)
17
18
19
20
95% Credible Interval for β
Simulation Index Within Batch
0.0
2.5
5.0
7.5
0
25
50
75
100