Lab 7, Act 1: Running a Chain by Hand

How common is chytrid in this pond, and how sure can we be?

NRES 746, name: _____ group: _____

The scenario

Chytrid fungus (*Batrachochytrium dendrobatidis*, or Bd) has driven steep declines of the Sierra Nevada yellow-legged frog. Field crews swab frogs' skin and test the swabs for Bd. At one high-elevation pond, **2 of 20 frogs tested positive** (same numbers from Week 7's WOD). Model the number of positives as binomial with unknown prevalence p.

You also have prior information: earlier surveys of nearby ponds found Bd in about 20% of frogs. To use it without letting it overwhelm this pond's data, let's give p a weak **Beta(2, 8)** prior: mean 0.2, and worth about as much information as 10 frogs (2 positive, 8 negative), half as many as you swabbed here.

The usual Wald (Hessian-based) interval, $\hat{p} \pm 1.96\sqrt{\hat{p}(1-\hat{p})/n}$, gives 0.10 ± 0.13, or -0.03 to 0.23. A prevalence below zero makes no sense. In this Act you run Monday's Metropolis algorithm by hand, then judge some computer-run chains, to get a more sensible answer.

What is the prevalence of Bd in this pond, and how uncertain is it?

Please keep laptops closed for this Act.

Part 1: Six steps of Metropolis (individual, ~20 min)

The log posterior is the log likelihood plus the log prior, which works out to this equation:

$$\log P(p \mid y) = \underbrace{2\log p + 18\log(1-p)}_{\text{log likelihood}} + \underbrace{\log p + 7\log(1-p)}_{\text{log prior}} + \text{constant}$$

for $0 < p < 1$; outside that range the posterior is 0 (log posterior $-\infty$). The table lists the log likelihood and the log prior on a grid of p values (each shifted by a constant so its largest value is 0). Add them to get the log posterior.

p: log lik., log prior	p: log lik., log prior	p: log lik., log prior	p: log lik., log prior
0.01: -2.89, -1.66	0.11: -0.01, -0.01	0.21: -0.86, -0.20	0.31: -2.52, -0.75
0.02: -1.69, -1.04	0.12: -0.04, 0.00	0.22: -1.00, -0.24	0.32: -2.72, -0.82
0.03: -1.06, -0.71	0.13: -0.09, 0.00	0.23: -1.14, -0.29	0.33: -2.92, -0.90
0.04: -0.67, -0.49	0.14: -0.15, -0.01	0.24: -1.29, -0.33	0.34: -3.14, -0.97
0.05: -0.41, -0.34	0.15: -0.22, -0.02	0.25: -1.45, -0.39	0.35: -3.35, -1.05
0.06: -0.24, -0.23	0.16: -0.30, -0.04	0.26: -1.61, -0.44	0.36: -3.57, -1.13
0.07: -0.12, -0.15	0.17: -0.40, -0.06	0.27: -1.78, -0.50	0.37: -3.80, -1.21
0.08: -0.05, -0.10	0.18: -0.50, -0.09	0.28: -1.96, -0.56	0.38: -4.04, -1.30
0.09: -0.01, -0.05	0.19: -0.61, -0.12	0.29: -2.14, -0.62	0.39: -4.28, -1.39
0.10: 0.00, -0.03	0.20: -0.73, -0.16	0.30: -2.33, -0.69	0.40: -4.53, -1.48

The chain starts at $p = 0.30$. In principle any starting value works, since a well-behaved chain eventually forgets where it began. In practice, pick a value that is plausible under the prior, so the chain doesn't waste many steps reaching the posterior. Here 0.30 is a plausible value on the high side of the prior (about 80% of the prior lies below it), which lets you watch the chain move toward the data. At each step, the proposal is the current value plus a random **jump**. The second column of random numbers is $\log u$, where u is a uniform random number between 0 and 1. Both are already drawn for you.

1a. Run six steps. For each, find the proposal, find the log posterior, lp = log likelihood + log prior, of the current value and of the proposal, compute the difference (proposal minus current), decide whether to accept, and write down the value the chain records at that step.

Step	Current	Jump	Proposal	lp(curr.)	lp(prop.)	Difference	log u	Accept?	Record
1	0.30	−0.08					−0.42		
2		+0.09					−0.85		
3		−0.07					−1.61		
4		−0.17					−0.30		
5		+0.04					−1.20		
6		−0.11					−2.05		

1b. A classmate's version of the algorithm writes down only the accepted proposals and skips the rows where the proposal was rejected. Write out their chain for your six steps. Why does their chain give the wrong picture of the posterior?

1c. The lookup table drops the binomial coefficient, the Beta prior's normalizing constant, and the denominator of Bayes' rule. Would any of your accept/reject decisions change if those were put back in? Why or why not?

Part 2: Which chains can you trust? (individual, ~15 min)

A computer ran four Metropolis chains on the same posterior, 2,000 steps each, with different starting values and jump sizes. Chain D is two chains run side by side from different starting values.

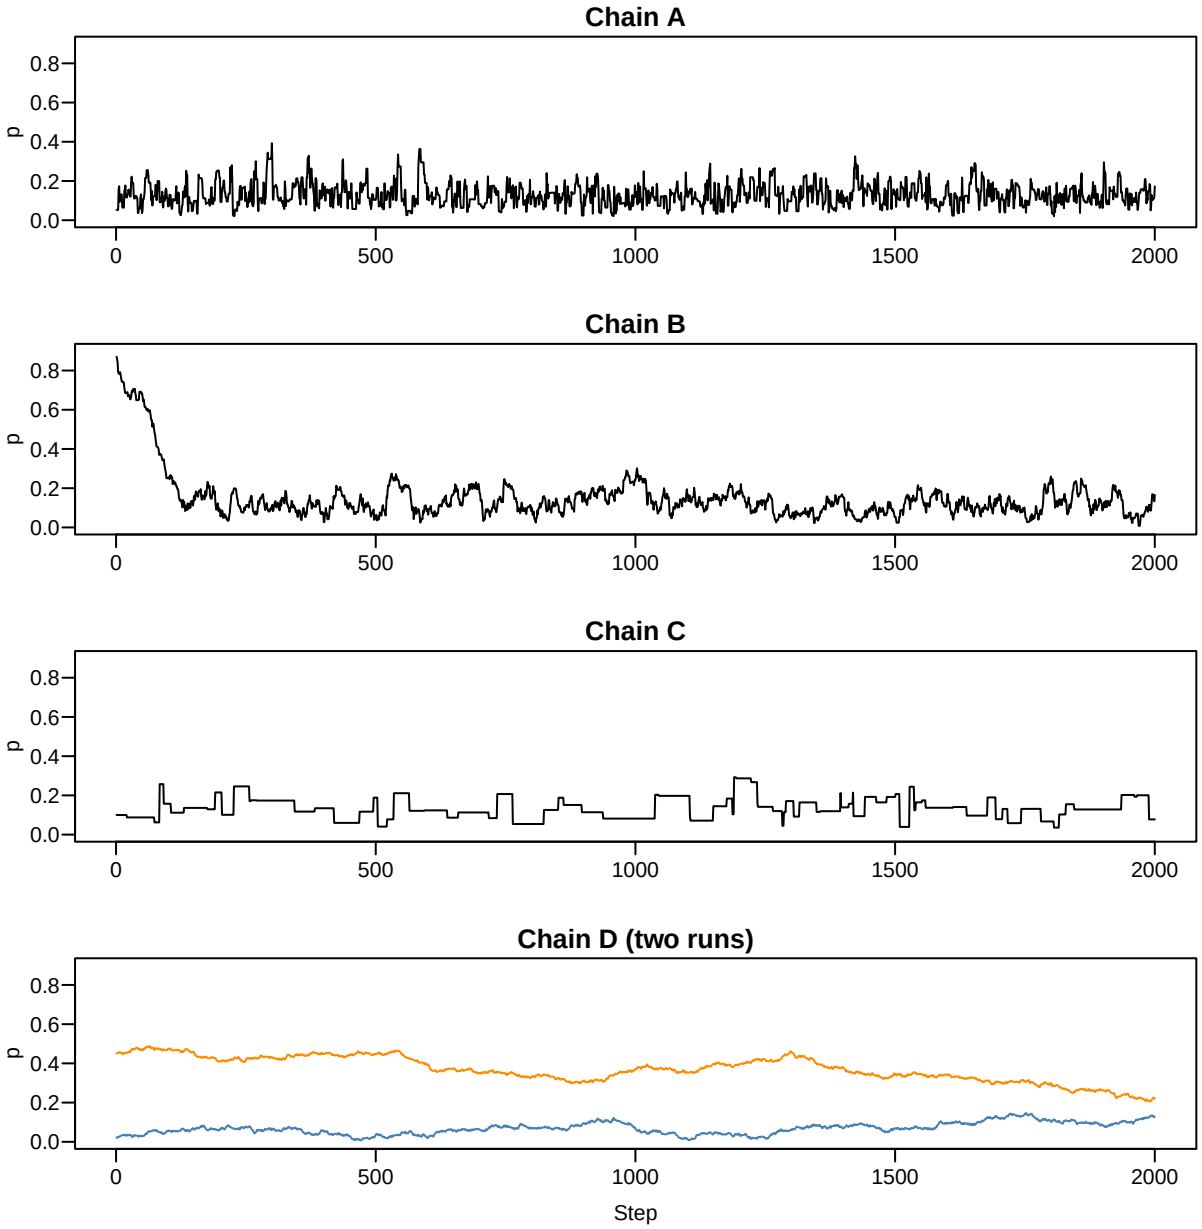

2a. For each chain, would you use it to describe the posterior? If not, what went wrong, and what would you change?

Chain	Use it?	What went wrong (if anything)?	What would you change?
A			
B			
C			
D			

2b. About how many steps of Chain B would you throw away before using it, and why?

Steps to discard: _____

2c. Here is Chain A after dropping its first 200 steps. The dashed lines mark its 2.5% and 97.5% quantiles.

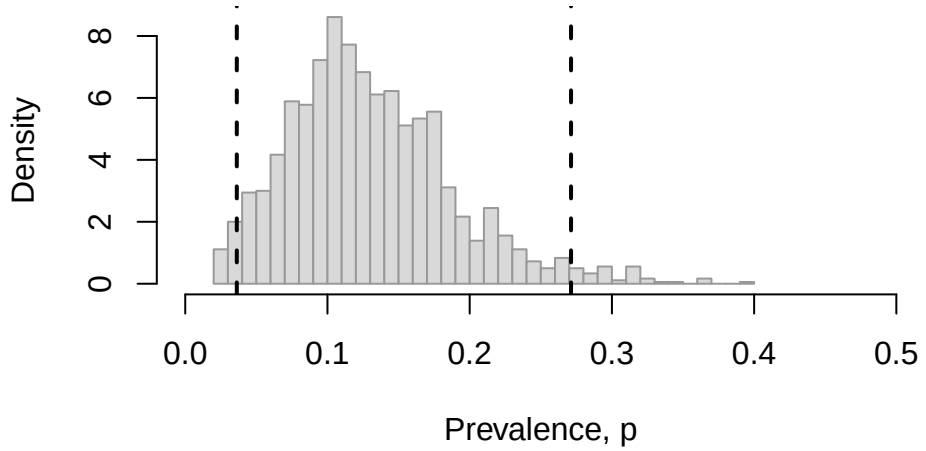

Posterior mean 0.130; 95% credible interval 0.036 to 0.271.

How does this interval differ from the Wald interval (-0.03 to 0.23), and why? With a Beta(2, 8) prior and these data, which named distribution should this histogram match?

Group discussion and revision (~10 min)

Compare your answers with your group, especially your table in 1a and your verdicts in 2a, and revise as needed. Record any changes you made, new insights, and any "muddiest points" or topics you want to review.