Homework 7

Due Tuesday, October 27, in class

STAT5710 Applied Stochastic Processes, Fall 2026

Where a part calls for computation, you may either attach or transcribe relevant output. You may include codes, but codes are not required. The sampler from `scripts/mh-inclass.R` will run everything below.

1. Exercises 5.7 and 5.19. Treat these as one problem: derive the acceptance ratio in 5.7, then implement it as asked in 5.19.

2. **Matching the tails.** Sample from a standard Cauchy target using an *independence sampler*, that is, a proposal $q(y \mid x) = q(y)$ that ignores the current state.

   (a) Write the acceptance ratio for an independence sampler in terms of q.

   (b) Implement it three times, with proposals from $N(0, 2^2)$, t_2, and Cauchy$(0, 3)$. Provide a table that reports, for each proposal distribution, the acceptance rate, the lag-one autocorrelation, the longest run of consecutive repeated values, the effective sample size, and an estimate of $P(|X| > 10)$.

   (c) Which independence sampler works best, and why?

3. **How many were made?** Equipment is stamped with consecutive serial numbers $1, 2, \ldots, N$, and N is unknown. You observe four items, with serial numbers 17, 42, 58, and 60.

   (a) Assuming the four were drawn without replacement and that any $N \geq 60$ is equally likely a priori, show that the posterior satisfies $f(N \mid \text{data}) \propto N^{-4}$ for $N \geq 60$.

   (b) Build a Metropolis-Hastings sampler for this posterior using a proposal that moves to $N \pm k$, with k drawn uniformly from $\{1, \ldots, m\}$.

   (c) Tune m, and report the acceptance rate and effective sample size for the values you tried. For your final choice, fix the number of samples and report the acceptance rate, effective sample size, lag-one autocorrelation, and attach the trace plot.

   (d) Estimate the posterior median and the posterior mean. Which one is harder to estimate, and why? (Try running your sampler with several seeds.)

4. Exercise 5.17.