Exercise: Sampling from a truncated beta posterior using MH

Consider this example from (?, p55). Suppose we make a new version of a product and ask $N = 20$ people if they prefer it to the old version; $X = 12$ people say yes. Let θ be probability they prefer the new version. We want to compute $\pi(\theta) = p(\theta|X = 12, N = 20)$.

Let $X \sim Bin(N, \theta)$. Suppose the prior is flat but we know that at least half the people will prefer the new version, which we encode as $\theta \sim U(0.5, 1)$. In other words, our probability model is

$$p(X|\theta, N) = \binom{N}{X} \theta^X (1-\theta)^{N-X} \quad (1)$$

$$p(\theta) = \frac{1}{1-0.5} I(0.5 \leq \theta \leq 1) \quad (2)$$

$$\pi(\theta) = p(\theta|X, N) \propto \theta^X (1-\theta)^{N-X} I(0.5 \leq \theta \leq 1) \quad (3)$$

Note that the truncated uniform prior is not conjugate to the binomial likelihood. We can compute the posterior using MH. Although it is possible to use proposals that only propose valid values of $\theta \in [0.5, 1]$, it is common to transform such constrained parameters to unconstrained form. Define

$$\phi = \log \frac{\theta - 0.5}{1 - \theta} \quad (4)$$

so $\phi \in (-\infty, \infty)$, with inverse transform

$$\theta = \frac{0.5 + e^\phi}{1 + e^\phi} \quad (5)$$

Now we can use a Gaussian proposal on ϕ.

1. Show that the transformed target density is

$$p(\phi|X) \propto \frac{(0.5 + e^\phi)^{12} e^\phi}{(1 + e^\phi)^{22}} \quad (6)$$

2. Use the MH algorithm to draw samples from $p(\phi|X)$. Use a Gaussian proposal with variance σ^2. Try $\sigma = 0.5$ and $\sigma = 10$. Use the following code snippet to ensure reproducible results

```
setSeed(1); xinit = rand(1,1); % initial state
```

Draw 40,000 samples, discarding the first 2000 for burnin (these numbers are somewhat arbitrary). Plot a histogram of all the samples of ϕ post burnin, and also a trace of the last 500 samples of ϕ. Finally, convert the samples of ϕ back to the 0:1 scale using

$$\theta = \frac{0.5 + e^\phi}{1 + e^\phi} \quad (7)$$

and plot a histogram of these. Your results should look like Figure 1 and Figure 2.

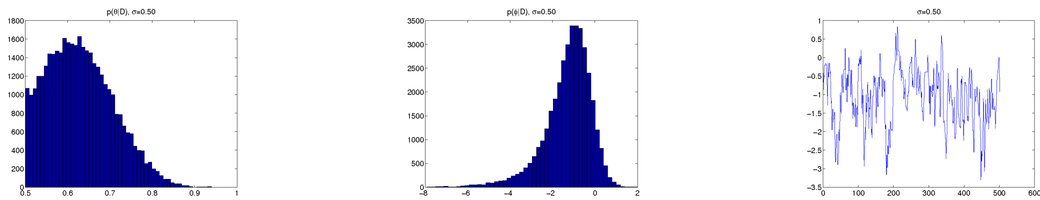

(a) (b) (c)

Figure 1: An example of the Metropolis algorithm for sampling from a binomial distribution with uniform prior using a Gaussian proposal with $\sigma = 0.5$. We used 40,000 samples and a burnin of 2000. Left: samples of the original parameter θ. The peak is near the MLE of $\hat{\theta}^{ML} = 0.6$. Middle: samples of the transformed parameter ϕ. Right: plot of the last 500 samples of ϕ. Figure generated by matlabEx.m.

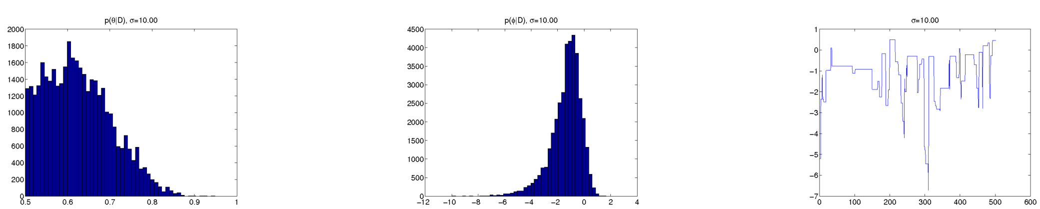

(a) (b) (c)

Figure 2: Same as Figure 1, except the Gaussian proposal has $\sigma = 10$. On the right we see the chain is not mixing is well, so the histograms are narrower and more blocky.