

STA 6352, Report 6.4

Carson Slater *Baylor University*

Introduction

In this report, we implement a Gibbs sampler in R to generate posterior samples from the beta-binomial model described in Example 6.8. The full conditionals are given as:

- $v|\alpha, \beta \sim \text{Beta}(\alpha, \beta)$
- $u|v \sim \text{Binomial}(n, v)$

We generate posterior samples, summarize results, and compare the empirical histogram with the theoretical beta-binomial distribution.

Gibbs Sampler Implementation

```
# Set parameters
set.seed(42)
N <- 1e6    # Number of iterations
n <- 50     # Binomial sample size
alpha <- 2  # Beta shape parameter
beta <- 5   # Beta shape parameter

# Initialize storage
u <- numeric(N)
v <- numeric(N)
v[1] <- rbeta(1, alpha, beta) # Initialize v

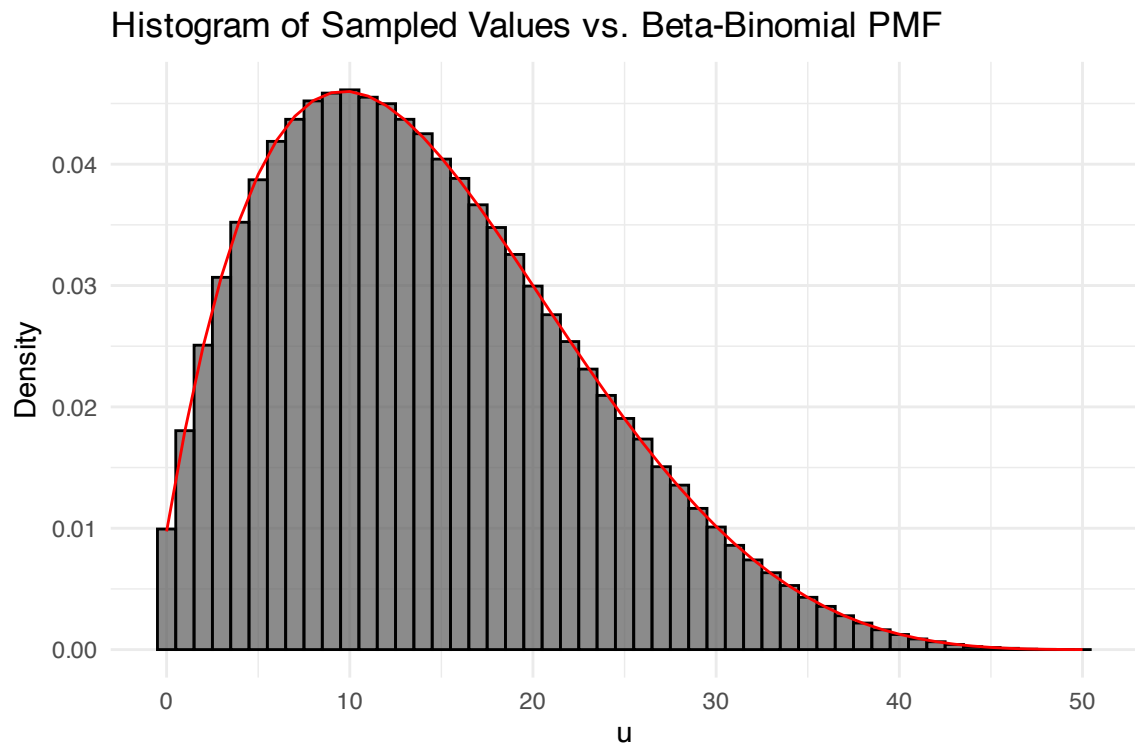
# Gibbs sampling
for (t in 2:N) {
  u[t] <- rbinom(1, n, v[t-1])
  v[t] <- rbeta(1, alpha + u[t], beta + n - u[t])
}

# Posterior mean approximation
posterior_mean <- mean(u / (u + 1))
posterior_mean
```

```
## [1] 0.8927734
```

Results and Visualization

We compare the histogram of the Gibbs samples to the beta-binomial distribution.



Conclusion

The Gibbs sampler successfully generates posterior samples, and our histogram closely aligns with the beta-binomial distribution. The posterior mean approximation is computed, demonstrating the effectiveness of Gibbs sampling for hierarchical models.