STA 6360, Report 1.7

Carson Slater *Baylor University*

Reproducing Example 1.13

Suppose we have a manufacturing process for which the probability of a defective part is θ. Then let $y \sim \text{Bin}(n, \theta)$ be the number of defective parts in a sample of size n, and let $\theta \sim \text{Beta}(u, v)$. After observing the data, we have that $\theta | y \sim \text{Beta}(u + y, v + n - y)$.

Suppose past batches yielded an average proportion of defective parts $\mu = \bar{x}$ and corresponding sample variance $\sigma = S^2$. We can select a $\text{Beta}(u, v)$ prior by solving the ststem of equations simultaneously.

$$\mathbb{E}(\theta) = \frac{u}{u+v} = \mu \text{ and } \mathbb{V}(\theta) = \frac{uv}{(u+v)^2(u+v+1)} = S^2$$

In past batches, the percent defective was 0.2 with a variance of 0.01. This would yield $u = 3$ and $v = 12$. This beta pior is able to be seen below.

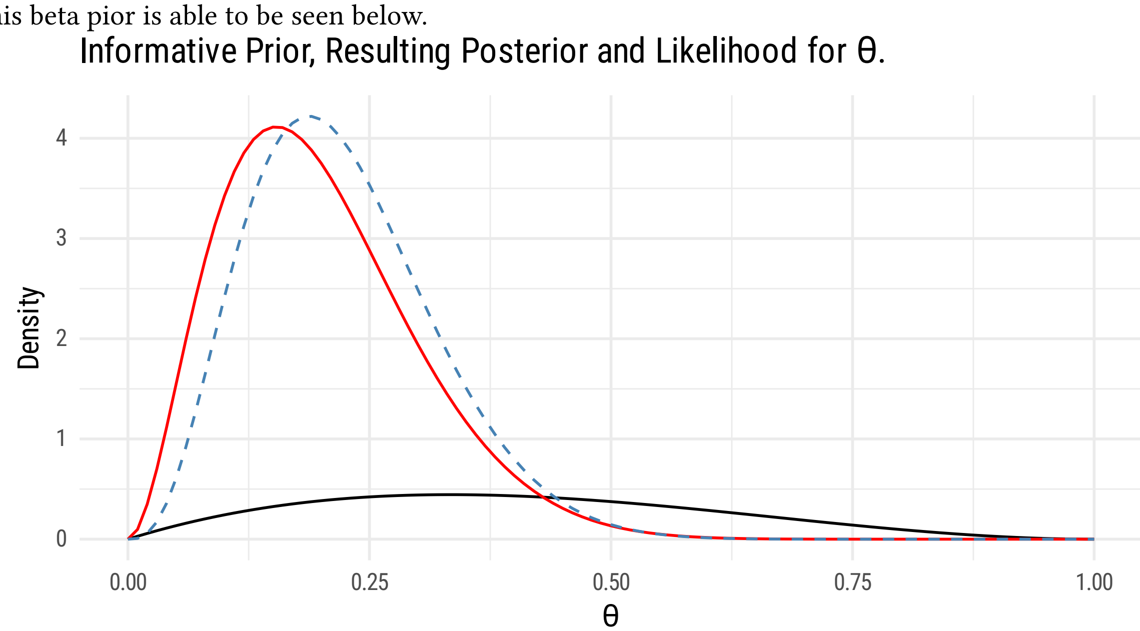

Prior (Red), Posterior (Blue), and Likelihood (Black).

Now suppose we take a sample size of $n = 3$ and find $y = 1$; then the posterior is

$$\pi(\theta \,|\, y = 1) = \text{Beta}(1 + 3, 12 + 3 - 1) = \text{Beta}(4, 14).$$

This would yield a posterior mean and variance of 0.22 and 0.0091 respectively. We have that the 'conventional' estimate of θ would be $1/3$, which is larger than our posterior mean. This is the impact of the prior. To find the HPD credible set, we employ equation 1.9.3 from the notes,

$$F_{\theta|\mathbf{y}}(a + h) - F_{\theta|\mathbf{y}}(a) = 1 - \alpha.$$

yielding

$$\begin{aligned}(a+h)^{u^*-1}(1-a+h)^{v^*-1} &= (a+h)^{u^*}(1-a)^{v^*-1} \\ \implies 1 &= \left(\frac{a+h}{a}\right)^{u^*-1}\left(\frac{1-a-h}{1-a}\right)^{v^*-1},\end{aligned}$$

where u^* and v^* are posterior parameters. Using such equality with the following equation,

$$\pi(a + h|\alpha^*, \beta^*) - \pi(a|\alpha^*, \beta^*) = 1 - \alpha,$$

a and h may be solved for numerically. yielding $a = 0.053$ and $h = 0.357$.

Imposing a Uniform Prior

Using a $\text{Unif}(0, 1)$ prior on the previous example would yield a $\text{Beta}(1 + 1, 1 + 3 - 1) = \text{Beta}(2, 3)$ posterior.

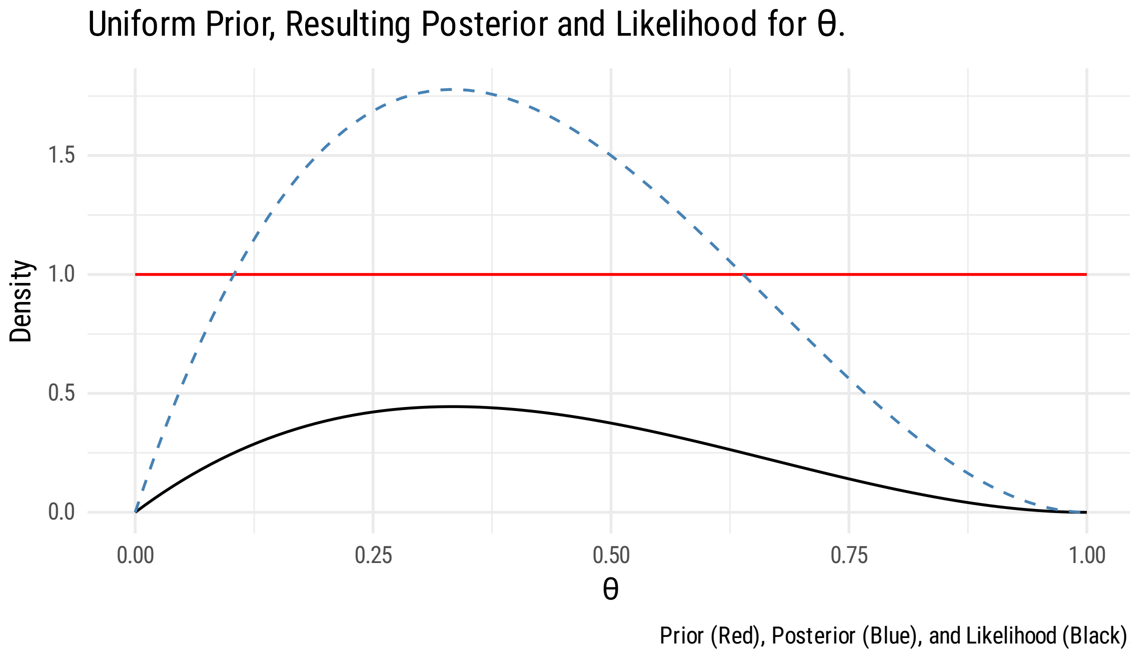

Repeating with $n = 10$.

Assuming we have a sample size of $n = 10$, but only 1 of these parts are defective, then an informative $\text{Beta}(3, 12)$ prior would yield a $\text{Beta}(4, 21)$ posterior, and a $\text{Unif}(0, 1)$ prior would yield a $\text{Beta}(2, 10)$ posterior. We demonstrate these with graphics.

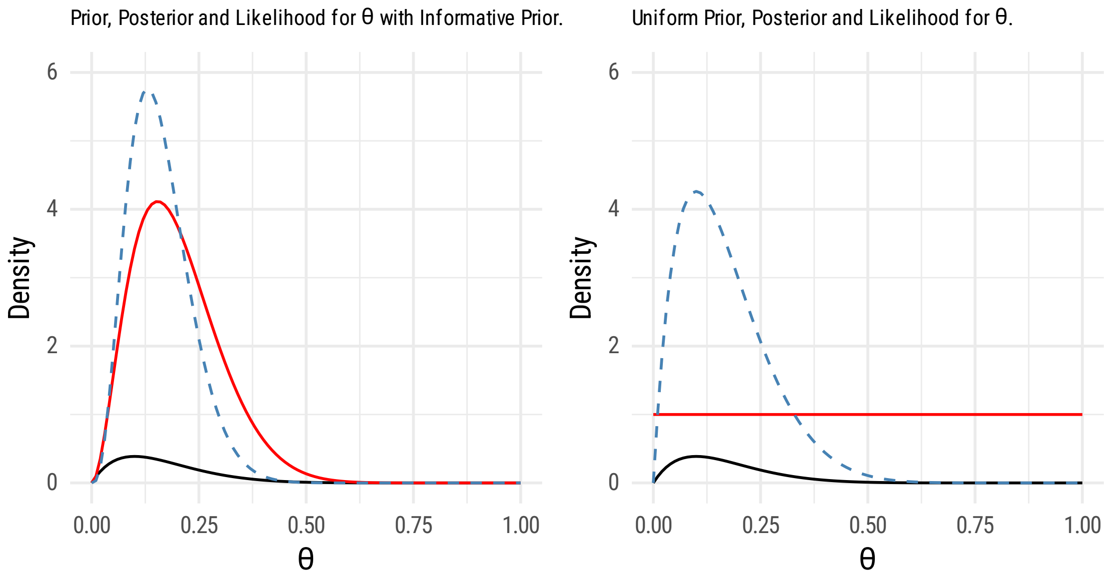

Prior (Red), Posterior (Blue), and Likelihood (Black); n = 10.

Contrasting these two results, the posterior yielded by the informative prior appears to be empowered by prior itself, and is supported by the likelihood. Alternatively, a uniform prior appears to create more 'work' for the likelihood function.

Repeating with $n = 20$.

Assuming we have a sample size of $n = 20$, but only 1 of these parts are defective, then an informative $\text{Beta}(3, 12)$ prior would yield a $\text{Beta}(4, 31)$ posterior, and a $\text{Unif}(0, 1)$ prior would yield a $\text{Beta}(2, 20)$ posterior. We demonstrate these with graphics.

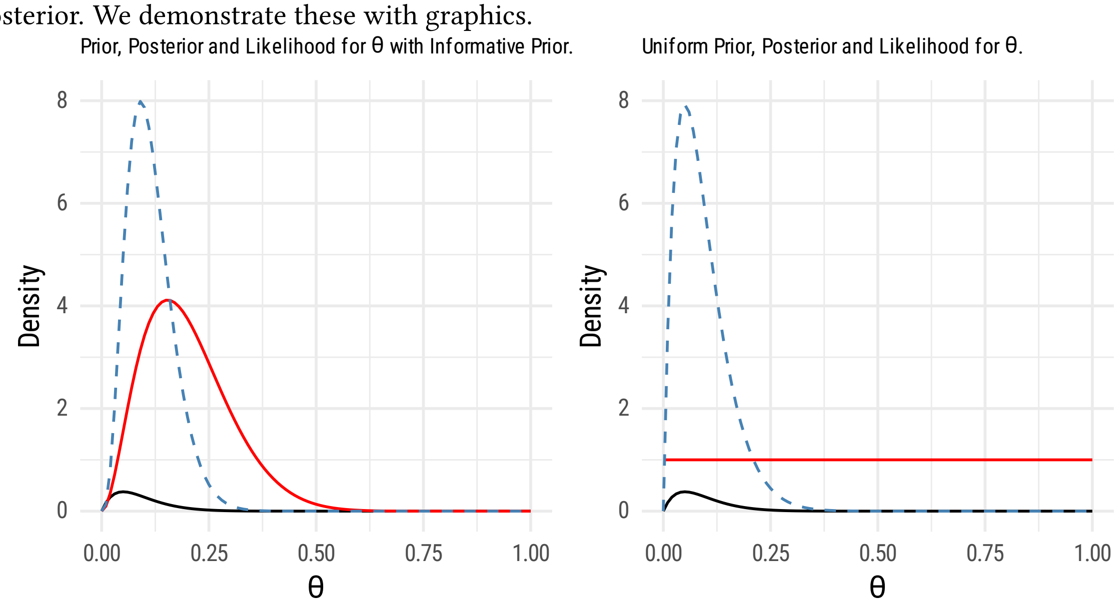

Prior (Red), Posterior (Blue), and Likelihood (Black); n = 10.