CHAPTER 15

Section 15.1

1.

a. From the information provided, θ can only equal .50 (fair) or .75. The prior probability assignments are $P(\theta = .50) = .80$ and $P(\theta = .75) = 1 - .80 = .20$. That is, $\pi(.50) = .80$ and $\pi(.75) = .20$.

b. As a function of θ, $P(\text{HHHTH})$ equals $\theta^4(1-\theta)$. Apply Bayes' theorem:

$$P(\theta = .5 \mid \text{HHHTH}) = \frac{P(\theta = .5)P(\text{HHHTH} \mid \theta = .5)}{P(\theta = .5)P(\text{HHHTH} \mid \theta = .5) + P(\theta = .75)P(\text{HHHTH} \mid \theta = .75)}$$

$$= \frac{(.80).5^4(1-.5)}{(.80).5^4(1-.5) + (.20).75^4(1-.75)} = .6124$$

That is, the posterior probability that $\theta = .5$ is .6124. Since .75 is the only other possibility, the posterior probability is $P(\theta = .75 \mid \text{HHHTH}) = 1 - .6124 = .3876$. I.e., $\pi(.50|\text{HHHTH}) = .6124$ and $\pi(.75|\text{HHHTH}) = .3876$.

3.

a. For a gamma distribution with parameters α and β, the mean is $\alpha\beta$ and the standard deviation is $\sqrt{\alpha}\beta$. Here, we want the prior distribution of μ to satisfy $\alpha\beta = 15$ and $\sqrt{\alpha}\beta = 5$. Divide the two equations to get $\sqrt{\alpha} = 15/5 = 3$, so $\alpha = 9$; then, $\beta = 15/\alpha = 15/9 = 5/3$. So, the prior for μ will be Gamma(9, 5/3).

b. The prior for μ is Gamma(9, 5/3); conditional on μ, the observations $X_1, \ldots, X_n$ are assumed to be a random sample from a Poisson(μ) distribution. Hence, the <u>numerator</u> of the posterior distribution of μ is

$$\pi(\mu) f(x_1, \ldots, x_n; \mu) = \frac{1}{\Gamma(9)(5/3)^9}\mu^{9-1}e^{-\mu/(5/3)} \cdot \frac{e^{-\mu}\mu^{x_1}}{x_1!} \cdots \frac{e^{-\mu}\mu^{x_n}}{x_n!}$$

$$= C\mu^{8+\Sigma x_i}e^{-(3/5)\mu - n\mu} = C\mu^{\alpha_1 - 1}e^{-\mu/\beta_1},$$

where $\alpha_1 = 9 + \Sigma x_i$ and $\beta_1 = \dfrac{1}{3/5 + n}$. We recognize the last expression above as the "kernel" of a gamma distribution (i.e., the pdf without the constant in front). Therefore, we conclude that the posterior distribution of μ is also gamma, but with the parameters α_1 and β_1 specified above.
With the specific values provided, $n = 10$ and $\Sigma x_i = 136$, so the posterior distribution of μ given these observed data is Gamma(α_1, β_1) = Gamma(145, 5/53).

5. Using (15.1), the numerator of the posterior distribution of μ is

$$\pi(\mu) f(x_1, \ldots, x_n; \mu) = \frac{1}{\Gamma(\alpha_0)\beta_0^{\alpha_0}}\mu^{\alpha_0 - 1}e^{-\mu/\beta_0}\frac{e^{-\mu}\mu^{x_1}}{x_1!} \cdots \frac{e^{-\mu}\mu^{x_n}}{x_n!} = C\mu^{\alpha_0 + \Sigma x_i - 1}e^{-\mu[1/\beta_0 + n]}.$$ As a function of μ, this is the kernel of a gamma pdf, specifically with first parameter $\alpha_0 + \Sigma x_i$ and second parameter $1/[1/\beta_0 + n]$. Therefore, the posterior distribution of μ is Gamma($\alpha_0 + \Sigma x_i$, $1/(n + 1/\beta_0)$).

7. If the prior distribution of p is $\text{Beta}(\alpha_0, \beta_0)$, its pdf is proportional to $p^{\alpha_0 - 1}(1-p)^{\beta_0 - 1}$. The joint pmf of a random sample from a negative binomial distribution is $\binom{x_1 - 1}{r-1} p^r (1-p)^{x_1 - r} \cdots \binom{x_n - 1}{r-1} p^r (1-p)^{x_n - r} \propto p^{nr}(1-p)^{\Sigma x_i - nr}$. Apply Equation (15.1): $\pi(p \mid x_1, \ldots, x_n) \propto \pi(p) f(x_1, \ldots, x_n \mid p) \propto p^{\alpha_0 - 1}(1-p)^{\beta_0 - 1} p^{nr}(1-p)^{\Sigma x_i - nr} = p^{\alpha_0 + nr - 1}(1-p)^{\beta_0 + \Sigma x_i - nr - 1}$. As a function of p, this is the kernel of a beta pdf, but with parameters $\alpha_1 = \alpha_0 + nr$ and $\beta_1 = \beta_0 + \Sigma x_i - nr$. That is, the posterior distribution of p is Beta($\alpha_1 = \alpha_0 + nr$, $\beta_1 = \beta_0 + \Sigma x_i - nr$).

9. The prior distribution of μ is $\pi(\mu) = \sqrt{\frac{\tau_0}{2\pi}} e^{-\tau_0(\mu - \mu_0)^2/2}$. The lognormal pdf with $\sigma = 1$ is $f(x; \mu) = \frac{1}{\sqrt{2\pi}x} e^{-(\ln x - \mu)^2/2}$. Using Equation (15.1), the numerator of the posterior distribution of μ is

$$\pi(\mu) f(x_1, \ldots, x_n; \mu) = \sqrt{\frac{\tau_0}{2\pi}} e^{-\tau_0(\mu - \mu_0)^2/2} \cdot \frac{1}{\sqrt{2\pi}x_1} e^{-(\ln x_1 - \mu)^2/2} \cdots \frac{1}{\sqrt{2\pi}x_n} e^{-(\ln x_n - \mu)^2/2}$$

$$\propto \exp\left(-\frac{\tau_0(\mu - \mu_0)^2}{2} - \frac{\sum(\ln x_i - \mu)^2}{2}\right) = \exp\left(-\frac{1}{2}\left[\tau_0(\mu - \mu_0)^2 + \sum(\mu - \ln x_i)^2\right]\right)$$

Because, as a function of μ, this is the exponential of a quadratic function of μ, we know that the posterior distribution of μ is normal (because, by definition, any pdf of that form is a normal pdf). To identify the parameters, note the following: if Y is normal with mean μ_1 and precision τ_1, then the exponential portion of the normal pdf has the form $-\frac{1}{2}\tau_1(y - \mu_1)^2 = -\frac{1}{2}[\tau_1 y^2 - 2\tau_1\mu_1 y + C]$. So, if we can expand the above expression in this way, we can identify the posterior precision and mean hyperparameters by looking at the coefficients on μ^2 and μ. Let's proceed:

$-\frac{1}{2}\left[\tau_0(\mu - \mu_0)^2 + \sum(\mu - \ln x_i)^2\right] = -\frac{1}{2}\left[(\tau_0 + n)\mu^2 - 2[\tau_0\mu_0 + \Sigma \ln x_i]\mu + C\right]$, where C does not depend on μ. Therefore, the posterior hyperparameters are $\tau_1 = \tau_0 + n$ and $\tau_1\mu_1 = \tau_0\mu_0 + \Sigma \ln x_i \Rightarrow \mu_1 = \frac{\tau_0\mu_0 + \Sigma \ln x_i}{\tau_0 + n}$.

In conclusion, the posterior distribution of μ is normal, with parameters μ_1 and τ_1 as above.

Section 15.2

11. In Exercise 3, the posterior distribution of μ was found to be Gamma(145, 5/53).

a. The Bayes estimate of μ is the mean of the posterior distribution:
$E(\mu \mid x_1, \ldots, x_n) = \alpha_1\beta_1$ = 145(5/53) = 13.68.

b. With the aid of software, we can determine the .025 and .975 quantiles of the posterior distribution. Using R, `qgamma(c(.025,.975),shape=145,scale=5/53)` returns 11.54338 and 15.99373. Thus, after observing the data in Exercise 3, there's a 95% (posterior) chance that μ is between 11.54 and 15.99.

13.

a. From the conjugate prior proposition regarding the beta distribution, the posterior distribution of p is Beta with parameters $\alpha_1 = \alpha_0 + x = 1 + n$ and $\beta_1 = \beta_0 + n - x = 1 + n - n = 1$. Hence, the posterior mean of p is $\alpha_1 / (\alpha_1 + \beta_1) = (1+n)/(1+n+1) = (n+1)/(n+2)$.

b. Imagine two prior trials, one success and one failure. Then we observe n successes in the next n trials, for n + 1 total successes out of these combined n + 2 trials. The relative frequency of successes is then (n+1)/(n+2).

c. This seems like an odd, and arbitrary, application of Laplace's idea. Why start with two days, with the sun rising on only one of them? Then, no matter how many days follow with the sun rising, we still include those two days on which the sun rose only once.

15. The same method as in Example 15.9 applies, but with n = 1 and σ = 3.5355. In particular, the posterior distribution of μ is still normal, with variance $\sigma_1^2 = \left(\frac{1}{3.5355^2} + \frac{1}{7.5^2}\right)^{-1}$ = 10.227 and mean $\mu_1 = \sigma_1^2\left(\frac{1(118.28)}{3.5355^2} + \frac{110}{7.5^2}\right)$ = 116.77. This is identical to the results of Example 15.9.

17. We are interested in the Beta distribution with parameters 804 and 772. This distribution is approximately normal, with mean μ = 804/1576 = .510 and variance $804(772)/(1576)^2(1577)$ = .0001585, or σ = .0126. Thus, a 95% credibility interval for p is .510 ± 1.96(.0126) = (.485,.535).

19.

a. The prior expectation of μ is the mean of the Gamma(α_0, β_0) distribution, $\alpha_0\beta_0$.

b. Exercise 5 established that the posterior distribution of μ is Gamma($\alpha_0 + \Sigma x_i$, $1/(n + 1/\beta_0)$). Hence the Bayes estimator is $\hat{\mu} = \alpha_1\beta_1 = \frac{\alpha_0 + \Sigma X_i}{n + 1/\beta_0}$.

c. By the Law of Large Numbers, $\bar{X} \to \mu^*$ as $n \to \infty$. Divide the numerator and denominator of $\hat{\mu}$ by n and take the limit: $\frac{\alpha_0 + \Sigma X_i}{n + 1/\beta_0} = \frac{(\alpha_0/n) + \bar{X}}{1 + 1/(n\beta_0)} \to \frac{0 + \mu^*}{1 + 0} = \mu^*$.