Solution to Assignment #10

9.1 We have

$$P_\theta\left(L(\boldsymbol{X}) \leq \theta \leq U(\boldsymbol{X})\right) = 1 - P_\theta\left(\{L(\boldsymbol{X}) > \theta\} \cup \{U(\boldsymbol{X}) < \theta\}\right).$$

Because $L(\boldsymbol{x}) \leq U(\boldsymbol{x})$ for all $\boldsymbol{x}$, the events $\{L(\boldsymbol{X}) > \theta\}$ and $\{U(\boldsymbol{X}) < \theta\}$ are mutually exclusive. Therefore,

$$P_\theta\left(\{L(\boldsymbol{X}) > \theta\} \cup \{U(\boldsymbol{X}) < \theta\}\right) = P_\theta\left(L(\boldsymbol{X}) > \theta\right) + P_\theta\left(U(\boldsymbol{X}) < \theta\right) = \alpha_1 + \alpha_2,$$

and the result follows.

9.4 (a) The respective unrestricted MLEs for σ_X^2 and σ_Y^2 are

$$s_X^2 = \frac{1}{n}\sum_{i=1}^n x_i^2 \quad \text{and} \quad s_Y^2 = \frac{1}{m}\sum_{j=1}^m y_j^2.$$

By the invariance principle, the unrestricted MLE of λ is $\hat{\lambda} = s_Y^2/s_X^2$.

When $\lambda = \lambda_0$, the MLE of σ_X^2 is

$$\hat{\sigma}_X^2 = \frac{\lambda_0 n s_X^2 + m s_Y^2}{\lambda_0(n+m)}.$$

Therefore,

$$\begin{aligned}
-2\log\lambda(\boldsymbol{x}) &= (n+m)\log\hat{\sigma}_X^2 + \frac{ns_X^2}{\hat{\sigma}_X^2} + m\log\lambda_0 + \frac{ms_Y^2}{\lambda_0\hat{\sigma}_X^2} - \\
&\quad (n+m)\log s_X^2 - (n+m) - m\log\hat{\lambda} \\
&= (n+m)\log\left(ns_X^2\left[\frac{\lambda_0 + (m/n)\hat{\lambda}}{\lambda_0(n+m)}\right]\right) + m\log\lambda_0 + (n+m) \\
&\quad -(n+m)\log s_X^2 - (n+m) - m\log\hat{\lambda} \\
&= (n+m)\log\left(\frac{n}{n+m}\right) + (n+m)\log\left(1 + \frac{m}{n}\cdot\frac{\hat{\lambda}}{\lambda_0}\right) - m\log(\hat{\lambda}/\lambda_0) \\
&= C_{n,m} + g(\hat{\lambda}/\lambda_0),
\end{aligned}$$

where

$$g(y) = (n+m)\log\left(1 + (m/n)y\right) - m\log y.$$

It follows that the LRT has the form

$$\text{“Reject } H_0 : \lambda = \lambda_0 \text{ iff } g(\hat{\lambda}/\lambda_0) \geq k.\text{”}$$

Now, it's easy to verify that, for $y > 0$, $g(y)$ decreases to a minimum at $y = 1$, and then begins to increase. Therefore, the rejection region of the LRT has the form

$$\frac{\hat{\lambda}}{\lambda_0} \leq c_1 \quad \text{or} \quad \frac{\hat{\lambda}}{\lambda_0} \geq c_2,$$

where $g(c_1) = g(c_2)$ and the rejection region has probability α when H_0 is true.
(b) The statistic $\hat{\lambda}$ is such that

$$\frac{\hat{\lambda}}{\lambda} = \frac{m^{-1}\sum_{j=1}^{m} Y_j^2/\sigma_Y^2}{n^{-1}\sum_{i=1}^{n} X_i^2/\sigma_X^2},$$

which has the F distribution with degrees of freedom m and n. So, when H_0 is true $\hat{\lambda}/\lambda_0$ is distributed $F_{m,n}$.
(c) Let $F_{p,m,n}$ be the pth quantile of the $F_{m,n}$ distribution. Then if $\alpha_1 + \alpha_2 = \alpha$,

$$P\left(F_{\alpha_1,m,n} \le \frac{\hat{\lambda}}{\lambda} \le F_{1-\alpha_2,m,n}\right) = 1 - \alpha.$$

It follows that $[\hat{\lambda}/F_{1-\alpha_2,m,n}, \hat{\lambda}/F_{\alpha_1,m,n}]$ is a $(1-\alpha)100\%$ confidence interval for λ.

9.12 There is no shortage of pivotal quantities in this problem. It is easy to show that each of the following random variables is pivotal:

$$\frac{\bar{X} - \theta}{S/\sqrt{n}}, \quad \frac{\bar{X} - \theta}{\sqrt{\theta}}, \quad \frac{\sum_{i=1}^{n}(X_i - \theta)^2}{\theta} \quad \text{and} \quad \frac{(n-1)S^2}{\theta} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})^2}{\theta}.$$

The first three of these are problematic, however. They sometimes yield what my former SMU professor Campbell Read calls "incredible confidence intervals." They are incredible in the sense that, for certain sets of data, the confidence interval produced is clearly wrong. Take, for example, the pivot $\sqrt{n}(\bar{X} - \theta)/S$. The shortest c.i. produced by this quantity is

$$\bar{X} \pm t_{n-1,\alpha/2}\frac{S}{\sqrt{n}}.$$

Now, even though $\theta > 0$, the event $\{\bar{X} + t_{n-1,\alpha/2}(S/\sqrt{n}) < 0\}$ occurs with positive probability. When this event occurs, the above c.i. contains only negative values, which is incredible since $\Theta = (0, \infty)$.

I prefer the pivot $(n-1)S^2/\theta$ because it is guaranteed to produce intervals containing only positive values. The interval has the form

$$\left[\frac{(n-1)S^2}{\chi^2_{n-1,1-\alpha_2}}, \frac{(n-1)S^2}{\chi^2_{n-1,\alpha_1}}\right],$$

where $\chi^2_{k,p}$ is the pth quantile of the chi-squared distribution with k degrees of freedom and $\alpha_1 + \alpha_2 = \alpha$.

9.25 Let $X_{(1)}$ be the smallest order statistic. The LRT in this problem has the form

$$\text{“Reject } H_0 : \mu = \mu_0 \text{ iff } x_{(1)} < \mu_0 \quad \text{or} \quad x_{(1)} \ge \mu_0 + c,\text{”}$$

where $c > 0$ is chosen to make the test have size α. Using the distributional result stated in the exercise, it is easy to show that $c = \log(\alpha^{-1})/n$. Inverting the test yields the following $(1-\alpha)100\%$ c.i. for μ:

$$\left[X_{(1)} - \log(\alpha^{-1})/n, X_{(1)}\right].$$

An obvious pivotal quantity is $X_{(1)} - \mu$. This leads to an interval of the same form as in Example 9.2.13. The minimum length interval of this type is exactly the same as the one obtained by inverting the LRT.

9.28 (a) The posterior is

$$\pi(\theta, \sigma^2 | \boldsymbol{x}) \propto \sigma^{-(n+2a+3)} e^{-1/(b\sigma^2)} \exp\left(-\frac{1}{2\sigma^2}\left[\sum_{i=1}^{n}(x_i - \theta)^2 + \frac{(\theta-\mu)^2}{\tau^2}\right]\right).$$

The trick now is to complete the square in the exponent. Doing so shows that the posterior is in the same family of joint densities as the prior, and hence the conjugate property follows. Defining $t_n = \sum_{i=1}^n x_i^2$, the quantities a, b, μ and τ^2 in the prior become

$$a^* = a + n/2, \qquad b^* = \left\{\frac{1}{b} + \frac{1}{2}\left[t_n + \frac{\mu^2}{\tau^2} - \frac{(n\bar{x} + \mu/\tau^2)^2}{n + 1/\tau^2}\right]\right\}^{-1},$$

$$\mu^* = \frac{n\bar{x} + \mu/\tau^2}{n + 1/\tau^2} \quad \text{and} \quad \tau^{2*} = \frac{1}{n + 1/\tau^2},$$

respectively, in the posterior.

(b) To find the posterior of θ, we integrate out σ^2 in the joint posterior. Since the prior and posterior are in the same family of densities, we may find the posterior of θ by first deriving the prior marginal of θ, and then replacing a, b, μ and τ^2 by a^*, b^*, μ^* and τ^{2*}, respectively. It is easy to verify that

$$\pi(\theta) = \frac{\Gamma(a+1/2)}{\Gamma(a)} \cdot \sqrt{\frac{b}{2\pi\tau^2}} \left[1 + \frac{(\theta-\mu)^2}{2\tau^2/b}\right]^{-(a+1/2)} \qquad \forall\ \theta.$$

We see that this distribution is symmetric about μ, and hence the HPD region for θ is of the form $\mu^* \pm c$, where c is chosen so that the interval has posterior probability $1-\alpha$.

(c) The answer is "almost, but not quite." Suppose we fix μ and let $a \to 0$, $b \to \infty$ and $\tau^2 \to \infty$. Then

$$\mu^* \to \bar{x}, \quad \tau^{2*} \to \frac{1}{n}, \quad a^* \to \frac{n}{2} \quad \text{and} \quad b^* \to \frac{2}{ns^2},$$

where s^2 is the MLE of σ^2. This yields a posterior of

$$\pi(\theta|\boldsymbol{x}) = \frac{\sqrt{n}}{s} \cdot \frac{\Gamma((n+1)/2)}{\Gamma(n/2)} \frac{1}{\sqrt{\pi n}} \left[1 + \frac{1}{n}\left(\frac{\theta-\mu}{s/\sqrt{n}}\right)^2\right]^{-(n+1)/2},$$

and corresponding HPD region

$$\bar{x} \pm t_{n,\alpha/2} \frac{s}{\sqrt{n}}.$$

This differs from the usual confidence interval in that the latter uses the "unbiased version" of the sample variance and $t_{n-1,\alpha/2}$ in place of $t_{n,\alpha/2}$.

9.36 It's easy to show that $T - \theta$ has the exponential$(2/[n(n+1)])$ distribution. Using Theorem 13, the shortest length $1-\alpha$ confidence interval of the form $[T+a, T+b]$ is

$$\left[T - \frac{2\log\alpha^{-1}}{n(n+1)}, T\right].$$