Statistical Inference: HW 3

Find the Jeffreys Prior for a location-scale family, i.e. one where

$$f(x|a,b) = b * f((x-a)*b)$$

for $-\infty < a < \infty$, $0 < b < \infty$ for some pdf $f(z)$. (Hint: Write $f(z) = e^{-\phi(z)}$ and do everything in terms of $\phi(z) \equiv -\log f(z)$; change variables in the expectation step to $z = b(x-a)$.)

Let X_j be a sequence of independent $\mathsf{Po}(\lambda)$ random variables.

1. Find the maximum likelihood estimator $\hat{\lambda}_n$ on the basis of the first n observations.
2. Show that $\hat{\lambda}_n$ is **sufficient** for λ.
3. Find the Fisher information $I(\lambda)$.
4. On the basis of the $n = 6$ observations $\mathbf{x} = \{1, 0, 2, 4, 3, 0\}$, find the 10% Likelihoodist Interval for λ (i.e., the set of points λ where the LH function attains at least 10% of its maximum value).
5. On the basis of the same $n = 6$ observations as before, find the equal-tail 90% Confidence Interval for λ, correct to four decimal places. Show the `Splus` code needed for your answer.
6. Find the Jeffrey's prior distribution for λ, and, on the basis of the same $n = 6$ observations as before, find the Bayesian posterior distribution and the posterior mean $\bar{\lambda}_n = \mathsf{E}[\lambda \mid \mathbf{x}]$.
7. Find the Bayesian equal-tail 90% Credible Interval for λ, using the Jeffreys prior.
8. Find the risk functions $R(\lambda, T)$ for both $T = \hat{\lambda}_n$ and $T = \bar{\lambda}_n$. Which one has lower risk?