CSC 2541, Small exercise #5, due in class April 5, worth 5% of the mark

Suppose that we observe a non-negative integer count, x, that we model as coming from a Poisson distribution with a positive real parameter θ, for which the probability of x is $P(x) = (1/x!)\theta^x \exp(-\theta)$. We use an improper prior for θ, with $P(\theta) \propto I(\theta > 0)/\theta$, where $I(\theta > 0)$ is the indicator with value 1 where $\theta > 0$ and value zero otherwise. This improper prior will produce a proper posterior distribution as long as $x \geq 1$.

1. Find the exact posterior mean of θ, assuming that the observed x is at least 1. For the remaining questions, we will find approximations to the posterior mean, pretending that we don't know this exact answer, but we can compare to the exact answer to see how good the approximations are.

2. Approximate the posterior distribution of θ by a normal distribution whose mode coincides with the true mode, and whose variance is such that the the second derivative of the log density at the mode is the same for the true posterior distribution as for the approximation. What do you need to assume about x for this approximation to make sense? What is the mean of θ in this approximation to the posterior?

3. Find an approximation to the mean of the posterior distribution by Laplace's method. Express this mean as the ratio of the integral of θ times the product of prior density and likelihood for θ, divided by the integral of the product of prior density and likelihood for θ. Approximate the numerator and denominator of this ratio separately by Laplace's method in order to get an approximation to the mean. What do you need to assume about x for this approximation to make sense? Among values of x where the approximation makes sense, for what value of x is the error in the approximation largest, and what is the magnitude of the error at that value for x?

4. Suppose that we were to try to approximate the posterior distribution, P, of θ by a normal distribution, Q, such that $KL(Q||P)$ is minimized. What would happen? (Hint: No elaborate calculations are needed to answer this question.)