Name: ____________________

Intermediate Bayesian Modeling
Midterm Exam

Instructions: Attempt as many parts of as many problems as possible. Show enough work to convince me that you know what you are doing. Write your answer to each problem on the blank paper provided. Write on only one side of each page. Do well!

1. (40 pts.) Video sharing websites like YouTube are interested in how many views are received by videos on their platform. Let u be the number of views received by a particular video in a one-hour span. Assume that u is well-modeled by a $\text{Pois}(\theta)$ distribution. Further assume that a Bayesian statistician puts a $\text{Gamma}\left(1, {}^{1}\!/\!{}_{150}\right)$ prior on the rate parameter θ for this Poisson distribution.

   (a) Find the marginal distribution for u based on this model and prior. (15 pts.)

   (b) Describe how this marginal distribution can be used to test whether seeing $u = 1280$ views in one hour is consistent with the given model and prior. (10 pts.)

   (c) Consider a collection of exponential random variables $v_i \mid \theta \overset{\text{iid}}{\sim} \text{Exp}(\theta)$, which can be thought of as the waiting times (in hours) between new viewers watching the previously discussed video. Explain the conditions under which identical Bayesian inferences will be made for θ using either the u data or the v_i data. (15 pts.)

2. (40 pts.) Recall that in the development of the Deviance Information Criterion, the penalization term p_D is given by

$$p_D = \mathrm{E}_{\theta|y}[-2\ell(\theta \mid y)] + 2\ell(\hat{\theta} \mid y),$$

   where $\ell(\theta \mid y)$ is the log-likelihood for some parameter θ given observed data y, and where $\hat{\theta}$ is an estimator for θ.

   (a) Prove that if $\hat{\theta}$ is the posterior mean for θ, p_D must be non-negative whenever the likelihood is log-concave ($\frac{d^2}{d\theta^2}\ell(\theta \mid y) < 0$ for all θ). (25 pts.)

   (b) In certain circumstances and for certain choices of $\hat{\theta}$, the quantity p_D may be negative. This would be inconvenient, since p_D is a penalty term. If $\hat{\theta}$ is chosen to be the posterior median, what additional constraint or constraints—beyond log-concavity—could be imposed to guarantee that p_D is positive? (You can explain this in words; you don't need to show it mathematically.) (15 pts.)

3. (40 pts.) Let X_i, $i \in \mathcal{N}$ be an infinite exchangeable sequence of random quantities and define Y_h and Y_k to be the averages of h and of k random quantities from among the X_i's. Assume De Finetti's law of large numbers applies here—that is, for any ε and θ, there exist H and K such that

$$h \geq H, k \geq K \longrightarrow \Pr\left[|Y_h - Y_k| > \varepsilon\right] < \theta.$$

If $\Phi_n(\xi) = \Pr\left[Y_n \leq \xi\right]$, prove that $\lim_{n\to\infty} \Phi_n(\xi)$ exists.

4. (40 pts.) A common tool for thinking about discrete probabilities is the Pólya urn scheme. We imagine an urn containing α red balls and β blue balls. Balls are drawn from the urn and observed. This is analogous to sampling from a population with two options—successes or failures on a Bernoulli trial, say. Different distributions can be modeled in this way based on whether or how balls are replaced in the urn after being drawn.

Imagine you are at the decennial Statistics Carnival, playing a game. In this game, a carnival worker has you draw a ball from an urn. Then you return the ball to the carnival worker. The carnival worker then takes one of three actions (always the same after each draw):

(1) The carnival worker does not return the ball to the urn.

(2) The carnival worker returns the ball to the urn.

(3) The carnival worker returns the ball to the urn, adding another ball of the same color.

You will draw 10 balls, and then you will guess which action the carnival worker has been taking. The game begins with one red ball and one blue ball in the urn.

(a) The carnival worker indicates that the three replacement actions are equally probable. Explain why the carnival worker is obviously lying. (5 pts.)

(b) What probability distribution is followed by Y, the total number of red balls seen in 10 draws, if the carnival worker takes Action 2. (10 pts.)

(c) Action 3 gives rise to a beta-binomial distribution with parameters $n = 10$, $\alpha = 1$, and $\beta = 1$. This is the same as the marginal distribution for Y if Y were distributed $\text{Bin}\,(10, p)$ with a $\text{Beta}\,(1, 1)$ prior on the probability of a red ball, p. Find the probability mass function for this beta-binomial distribution. (10 pts.)

(d) Assume that you've seen $Y = 8$ red balls in 10 draws. Calculate the Bayes factor comparing Model A—the carnival worker replaces balls according to Action 2, and Model B—the carnival worker replaces balls according to Action 3. (15 pts.)

(e) EXTRA CREDIT: Pretend $\alpha = \beta = 10$. What probability distribution would be followed by Y if the carnival worker took Action 1?

Table 1: Common distributions and densities.

Distribution	Notation	Density
Bernoulli	$\text{Bern}\,(\theta)$	$f(y\|\theta) = \theta^y(1-\theta)^{1-y}$; $y = 0, 1$
Binomial	$\text{Bin}\,(n, \theta)$	$f(y\|\theta) = \binom{n}{y}\theta^y(1-\theta)^{n-y}$; $y = 0, 1, \ldots, n$
Negative Binomial	$\text{NegBin}\,(r, \theta)$	$f(y\|\theta) = \binom{y}{r}\theta^{y-r}(1-\theta)^r$; $y = r, r+1, \ldots$
Beta	$\text{Beta}\,(a, b)$	$p(\theta) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)}\theta^{a-1}(1-\theta)^{b-1}I_{(0,1)}(\theta)$
Poisson	$\text{Pois}\,(\theta)$	$f(y\|\theta) = \theta^y e^{-\theta}/y!$; $y = 0, 1, 2, \ldots$
Exponential	$\text{Exp}\,(\theta)$	$f(y\|\theta) = \theta e^{-\theta y}I_{(0,\infty)}(y)$
Gamma / Erlang	$\text{Gamma}\,(a, b)$	$p(\theta) = [b^a/\Gamma(a)]\theta^{a-1}e^{-b\theta}I_{(0,\infty)}(\theta)$
Weibull	$\text{Weibull}\,(\alpha, \lambda)$	$f(y\|\alpha, \lambda) = \lambda\alpha y^{\alpha-1}\exp\left(-\lambda y^\alpha\right)I_{(0,\infty)}(y)$

Table 2: Some conjugate families.

$f(y\|\theta)$	$p(\theta)$	$p(\theta\|y)$
$\text{Bin}\,(n, \theta)$	$\text{Beta}\,(a, b)$	$\text{Beta}\,(a+y, b+n-y)$
$\text{NegBin}\,(r, \theta)$	$\text{Beta}\,(a, b)$	$\text{Beta}\,(a+y-r, b+r)$
$\text{Pois}\,(\theta)$	$\text{Gamma}\,(a, b)$	$\text{Gamma}\,(a+\sum y_i, b+n)$
$\text{Exp}\,(\theta)$	$\text{Gamma}\,(a, b)$	$\text{Gamma}\,(a+n, b+\sum y_i)$
$\text{Gamma}\,(k, \theta)$	$\text{Gamma}\,(a, b)$	$\text{Gamma}\,(a+nk, b+\sum y_i)$