Durham University

# EXAMINATION PAPER

| Examination Session: | Year: | Exam Code: |
|---|---|---|
| May/June | 2025 | MATH3421-WE01 |

**Title:** Bayesian Computation and Modelling III

| | | |
|---|---|---|
| Time: | 2 hours | |
| Additional Material provided: | Formula sheet | |
| Materials Permitted: | | |
| Calculators Permitted: | Yes | Models Permitted: Casio FX83 series or FX85 series. |

| | |
|---|---|
| Instructions to Candidates: | Answer all questions.<br>Section A is worth 40% and Section B is worth 60%. Within each section, all questions carry equal marks.<br>Write your answer in the white-covered answer booklet with barcodes.<br>Begin your answer to each question on a new page. |

| | |
|---|---|
| **Revision:** | |

ED01/2025

## SECTION A

**Q1** Consider the simulation of values from a $N(0,1)$ distribution, truncated to the left at 0 and to the right at 1. The target density is $\pi(\theta) \propto \exp\{-\theta^2/2\}$, $0 < \theta < 1$, and we wish to generate values of $\theta$ using a *Metropolis-Hastings independence sampler* based on Uniform $U(0,1)$ proposals.

(a) Write down and simplify the acceptance probability for a move from $\theta$ to $\phi$.

(b) For what values of $\phi$ will the acceptance probability in (a) be 1? You should leave your answer in terms of $\theta$.

(c) Given that the current value of the Markov chain simulated by this sampler is $\theta$, find the probability that the chain moves (marginalised over the distribution of proposed values). You should leave your answer in terms of $\theta$ and the standard normal distribution function $\Phi(\cdot)$.

(d) The overall acceptance rate of this sampler (assuming that the chain is in equilibrium) is 0.92. Comment on this acceptance rate.

**Q2** The *tobit model* is widely used for random variables $Y$ that are either exactly equal to zero or continuous and positive. Let $Y^*|\mu, \sigma^2 \sim N(\mu, \sigma^2)$ be a partially latent random variable; if $Y^* > 0$, we observe a value for $Y = Y^*$ but when $Y^* \leq 0$, we observe $Y = 0$. A random sample, $y_1, \ldots, y_N$, from the tobit model is therefore regarded as a set of $N$ censored observations from the $N(\mu, \sigma^2)$ distribution, where zeros are treated as observations that are missing but known to be negative.

(a) Write down the joint likelihood for the model parameters $\mu$ and $\sigma^2$ with respect to $N = n + m$ observations, assuming $y_1, \ldots, y_n$ are strictly positive and the remaining $m$ observations are equal to zero. You may use the standard normal distribution function $\Phi(\cdot)$ in your answer.

(b) Are both $\mu$ and $\sigma^2$ identifiable? Give your reasoning.

(c) An expert decides to model daily rainfall (in millimetres) in Durham as a random sample from the tobit model. They posit the following prior distribution for its parameters:

$$\mu \sim N(a, 4),$$
$$\sigma^2 \sim \text{InvGamma}(2.5, 30),$$

where $a$ is a constant to be chosen. Before seeing the data, the expert believes that $\Pr(Y = 0)$, i.e. the probability of a dry day, should be 0.6. Write down an integral equation that could be solved for $a$ in order to capture the expert's judgement. Note that you need not solve the equation.

(d) Do you think this is a good model for daily rainfall? Give your reasoning.

## SECTION B

**Q3** Consider a first order, homogeneous Markov chain $\underline{\theta}^{(n)}$, $n = 0, 1, \ldots$, where $\underline{\theta}^{(n)} = (\theta_1^{(n)}, \theta_2^{(n)})^T$, with continuous state-space $\mathcal{S}$. Let $\pi(\theta_1, \theta_2)$ be a density that is intended to be a stationary density of the Markov chain and let $p(\underline{\phi}|\underline{\theta})$ denote the transition density of the chain.

(a) Write down an integral equation satisfied by $\pi(\theta_1, \theta_2)$.

(b) Consider a *deterministic scan Gibbs sampler* with target density $\pi(\theta_1, \theta_2)$.

(i) Write down the form of the transition density $p(\underline{\phi}|\underline{\theta})$ of the resulting Markov chain, giving your answer in terms of full conditional densities of the form $\pi(\cdot|\cdot)$.

(ii) Using your answer to (a), show that the Markov chain defined by the deterministic scan Gibbs sampler has stationary density $\pi(\theta_1, \theta_2)$.

(c) Suppose that

$$\pi(\theta_1, \theta_2) \;\propto\; \theta_1^{-1} \exp\left\{2\theta_2 \log\theta_1 \;-\; 2\theta_2^2 \;-\; (\log\theta_1)^2\right\}, \quad \theta_1 \;>\; 0, \;\; \theta_2 \;\in\; \mathbb{R}.$$

(i) Derive and identify all full conditional distributions for this bivariate target.

(ii) The correlation between $\theta_1$ and $\theta_2$ is 0.54. Very briefly comment on the likely impact (if any) of this correlation on the mixing of the deterministic scan Gibbs sampler for this bivariate target.

(iii) The *marginal* distribution of $\theta_2$ is $\mathrm{N}(0, 1/2)$. Suggest a mechanism for *direct* sampling from $\pi(\theta_1, \theta_2)$.

**Q4** The following model is proposed for the time to failure, $Y$, in units of thousands of hours, of light bulbs produced by a particular manufacturer:

$$\begin{aligned} Y|\alpha,\beta &\sim \text{Gamma}(\alpha,\beta), \\ \alpha &\sim \text{Gamma}(a_1,a_2), \\ \beta &\sim \text{Gamma}(b_1,b_2). \end{aligned}$$

(a) The Law of Total Expectation is given by $\mathbb{E}(\mathbf{Y}) = \mathbb{E}\{\mathbb{E}(\mathbf{Y}|\mathbf{X})\}$ for random vectors $\mathbf{X} = (X_1,\ldots,X_p)^T$ and $\mathbf{Y} = (Y_1,\ldots,Y_q)^T$. Use this to show that the mean of the preposterior distribution for $Y$ is given by

$$\mathbb{E}(Y) = \frac{a_1 b_2}{a_2(b_1-1)},$$

for $b_1 > 1$, and undefined otherwise.

(b) Suppose that we decide to focus on the special case where $\alpha$ and $\beta$ have exponential distributions (i.e. $a_1 = b_1 = 1$). Derive the preposterior density for $Y$, $\pi(y)$. Note that $\Gamma(a+1) = a\Gamma(a)$.

(c) We will soon have a random sample of 10 failure times.

(i) Draw a directed acyclic graph that represents the model for $Y_1,\ldots,Y_{10}$.

(ii) Hence or otherwise, explain whether you could calculate the evidence for this model using

$$\pi(\mathbf{y}) = \prod_{i=1}^{10} \pi(y_i),$$

where $\mathbf{y} = (y_1,\ldots,y_{10})^T$ denotes the random sample and $\pi(y_i)$ is an instance of the preposterior density for a single observation, derived in (b).

(d) Would you recommend the use of exponential priors for $\alpha$ and $\beta$ in this model? Explain your reasoning.