Exercise sheet 1

Exercise 1. You believe a priori that the mean level of contamination of a certain substance, represented by a random variable $\boldsymbol{\theta}$, follows a normal distribution centred at 10 and with variance 4. Now, you perform two measurements of this substance yielding $y_1 = 5$ and $y_2 = 4$ and your model tells you that these measurements are conditionally independent and normally distributed with mean $\boldsymbol{\theta}$ and variance 9.

Calculate explicitly the posterior distribution of $\boldsymbol{\theta}$. Comment on your results.

Exercise 2. We consider a random variable $\boldsymbol{\theta}$ defined on $(0, \infty)$. Observations are obtained through random experiments $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{n+1}$ on $[0, \infty)$ which are i.i.d. given $\boldsymbol{\theta}$ and follow an exponential distribution defined as

$$p_{\boldsymbol{y}|\boldsymbol{\theta}}(y|\theta) = \mathrm{Ex}(y; \theta) = \theta \exp(-\theta y).$$

Note that $\mathbb{E}(\boldsymbol{y} \,|\, \boldsymbol{\theta} = \theta) = \theta^{-1}$ and $\mathrm{var}(\boldsymbol{y} \,|\, \boldsymbol{\theta} = \theta) = \theta^{-2}$. A priori, you believe that the parameter follows a gamma distribution with parameters $\alpha, \beta > 0$, that is

$$p_{\boldsymbol{\theta}}(\theta) = \mathrm{Ga}(\theta; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} \exp(-\beta\theta).$$

1. Find the posterior distribution of the parameter given the first n observations $y_1, \ldots, y_n$ and discuss the relationship between the posterior mean and the prior and sample means. Assume that, instead of $\boldsymbol{\theta}$, we are interested in making inference about the random variable $\boldsymbol{\psi} = \boldsymbol{\theta}^{-1}$. Derive its posterior distribution and relate its posterior mean to the prior and sample means.
2. Find the predictive distribution of $\boldsymbol{y}_{n+1}$ given $\boldsymbol{y}_{1:n} = (y_1, ..., y_n)$.

[Hint: the mean of a gamma distribution with parameters (α, β) is α/β]

Exercise 3. We consider a random variable $\boldsymbol{\theta}$ defined on $(0, \infty)$. Observations are obtained through random experiments $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{n+1}$ on $\mathbb{N}_0$ which are i.i.d. given $\boldsymbol{\theta}$ and follow a Poisson distribution defined as

$$p_{\boldsymbol{y}|\boldsymbol{\theta}}(y \,|\, \theta) = \mathrm{Po}(y; \theta) = \frac{1}{y!} \theta^y \exp(-\theta).$$

Note that $\mathbb{E}(\boldsymbol{y} \,|\, \boldsymbol{\theta} = \theta) = \theta$ and $\mathrm{var}(\boldsymbol{y} \,|\, \boldsymbol{\theta} = \theta) = \theta$. A priori, you believe that the parameter follows a gamma distribution with parameters $\alpha, \beta > 0$, that is

$$p_{\boldsymbol{\theta}}(\theta) = \mathrm{Ga}(\theta; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} \exp(-\beta\theta).$$

1. Find the posterior distribution of the parameter given the first n observations $y_1, \ldots, y_n$ and discuss the relationship of the posterior mean with the prior and sample means.
2. Find the predictive distribution of $\boldsymbol{y}_{n+1}$ given $\boldsymbol{y}_{1:n} = (y_1, ..., y_n)$.

Exercise 4. We consider a joint random variable $\boldsymbol{\theta} = (\boldsymbol{\mu}, \boldsymbol{\tau})$ defined on $\mathbb{R} \times (0, \infty)$. Observations are obtained through random experiments $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{n+1}$ on $\mathbb{R}$ which are i.i.d. given the random variables $\boldsymbol{\mu}$ and $\boldsymbol{\tau}$ and follow a normal distribution defined as

$$p_{\boldsymbol{y}|\boldsymbol{\mu},\boldsymbol{\tau}}(y \,|\, \mu, \tau) = \mathrm{N}(y; \mu, \tau^{-1}).$$

Suppose also that a priori $(\boldsymbol{\mu}, \boldsymbol{\tau})$ follows a normal-gamma distribution, that is

$$p_{\boldsymbol{\mu}|\boldsymbol{\tau}}(\mu \,|\, \tau) = \mathrm{N}(\mu; \mu_0, (k\tau)^{-1}) \qquad \text{and} \qquad p_{\boldsymbol{\tau}}(\tau) = \mathrm{Ga}(\tau; \alpha, \beta),$$

for some $\mu_0 \in \mathbb{R}$, $k \in \mathbb{N}$ and $\alpha, \beta > 0$.

Find explicitly the predictive distribution of $\boldsymbol{y}_{n+1}$ given $\boldsymbol{y}_{1:n} = (y_1, ..., y_n)$.