Exercise sheet 1

	Poisson	exponential	beta	gamma
parameter	$\lambda > 0$	$\lambda > 0$	$\alpha, \beta > 0$	$\alpha, \beta > 0$
support	$\mathbb{N}_0$	$[0, \infty)$	$[0, 1]$	$(0, \infty)$
$\boldsymbol{x} \sim$	$\mathrm{Po}(k; \lambda) = \frac{\lambda^k}{k!} \exp(-\lambda)$	$\mathrm{Ex}(x; \lambda) = \lambda \exp(-\lambda x)$	$\mathrm{Be}(x; \alpha, \beta) = \frac{1}{\mathrm{B}(\alpha,\beta)} x^{\alpha-1}(1-x)^{\beta-1}$	$\mathrm{Ga}(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp(-\beta x)$
$\mathbb{E}(\boldsymbol{x})$	λ	λ^{-1}	$\alpha/(\alpha+\beta)$	α/β
$\mathrm{var}(\boldsymbol{x})$	λ	λ^{-2}	$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$	α/β^2

Table 1: Standard probability distributions with their mean $\mathbb{E}(\boldsymbol{x})$ and variance $\mathrm{var}(\boldsymbol{x})$ with $\boldsymbol{x}$ a random variable distributed accordingly.

Theorem 1. *Let $\boldsymbol{\theta}$ be a random variable on $\mathbb{R}$ with probability density function $p_{\boldsymbol{\theta}}(\cdot)$ and let $f : \mathbb{R} \to \mathbb{R}$ be a monotonic function, then the probability density function $p_{\boldsymbol{\psi}}(\cdot)$ of the random variable $\boldsymbol{\psi} = f(\boldsymbol{\theta})$ is characterised by*

$$p_{\boldsymbol{\psi}}(\psi) = p_{\boldsymbol{\theta}}(f^{-1}(\psi)) \left| \frac{\mathrm{d}}{\mathrm{d}\psi} f^{-1}(\psi) \right|, \qquad \psi \in \mathbb{R}.$$

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, i.e. $p_{\boldsymbol{y}|\boldsymbol{\theta}}(y|\theta) = \mathrm{Ex}(y; \theta)$.

1. Choose a suitable prior distribution from Table 1 and explain your choice.
2. 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.
3. 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.
4. Find the predictive distribution of $\boldsymbol{y}_{n+1}$ given $\boldsymbol{y}_{1:n} = (y_1, ..., y_n)$.

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, i.e. $p_{\boldsymbol{y}|\boldsymbol{\theta}}(y \mid \theta) = \mathrm{Po}(y; \theta)$.

1. Choose a suitable prior distribution from Table 1 and explain your choice.
2. 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.
3. 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

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

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

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

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 $\boldsymbol{y}_{1:n} = (y_1, ..., y_n)$.