STA 6360, Report 4.4

Carson Slater *Baylor University*

1

Prove (4.2.19). Suppose we have n IID observations, $\mathbf{y} \equiv (y_1, y_2, \dots, y_n)$, with $y_i \sim \text{Poisson}(\lambda)$, $\lambda > 0$. We can show that a $\text{Gamma}(\alpha, \beta)$ prior is a conjugate prior.

The prior for λ as $\text{Gamma}(\alpha, \beta)$. The joint distribution of the data and the parameter is given by:

$$P(\mathbf{y}, \lambda) = \prod_{i=1}^{n} \frac{e^{-\lambda}\lambda^{y_i}}{y_i!} \times \frac{\beta^\alpha}{\Gamma(\alpha)} e^{-\beta\lambda}\lambda^{\alpha-1}$$

The likelihood function $P(\mathbf{y}, \lambda)$ is already given. To find the evidence $P(\mathbf{y})$, we need to integrate over all possible values of λ:

$$P(\mathbf{y}) = \int_0^\infty \prod_{i=1}^{n} \frac{e^{-\lambda}\lambda^{y_i}}{y_i!} \times \frac{\beta^\alpha}{\Gamma(\alpha)} e^{-\beta\lambda}\lambda^{\alpha-1} d\lambda$$

We can simplify this integral using the property of the gamma function: $\Gamma(\alpha+1) = \alpha\Gamma(\alpha)$. The result is:

$$P(\mathbf{y}) = \frac{\beta^\alpha \Gamma(\alpha + \sum_{i=1}^n y_i)}{\Gamma(\alpha)} \frac{\prod_{i=1}^n y_i!}{\beta^{\sum_{i=1}^n y_i}}$$

Now, we can write the posterior distribution:

$$P(\lambda|\mathbf{y}) \propto \prod_{i=1}^{n} \frac{e^{-\lambda}\lambda^{y_i}}{y_i!} \times \frac{\beta^\alpha e^{-\beta\lambda}\lambda^{\alpha-1}}{\beta^{\sum_{i=1}^n y_i}}$$

This simplifies to:

$$P(\lambda|\mathbf{y}) \propto \frac{e^{-\lambda(\beta+1)}\lambda^{\sum_{i=1}^n y_i + \alpha - 1}}{\prod_{i=1}^n y_i!}$$

This is a Gamma distribution with parameters $\alpha' = \alpha + \sum_{i=1}^n y_i$ and $\beta' = \beta + 1$. Therefore, the posterior distribution is a Gamma distribution, which proves that the Gamma prior is a conjugate prior for the Poisson distribution.

2

Exercise 2.13 in Gelman's *BDA3*.

(a)

Let y_i be the number of fatal accidents in year i, for $i = 1, \dots, 10$, and λ be the expected number of accidents in a year. The model for the data is $y_i|\lambda \sim \text{Poisson}(\lambda)$.

We use the conjugate family of distributions for convenience. If the prior distribution for λ is $\text{Gamma}(\alpha, \beta)$, then the posterior distribution is $\text{Gamma}(\alpha + 10\bar{y}, \beta + 10)$. Assume a noninformative prior distribution: $(\alpha, \beta) = (1/2, 0)$, this should be ok since we have enough information here: $n = 10$. Since $\beta > 0$ we can use the smallest possible float in R, which we will denote ϵ. Then the posterior distribution is $\lambda|\mathbf{y} \sim \text{Gamma}(238.5, 10 + \epsilon)$.

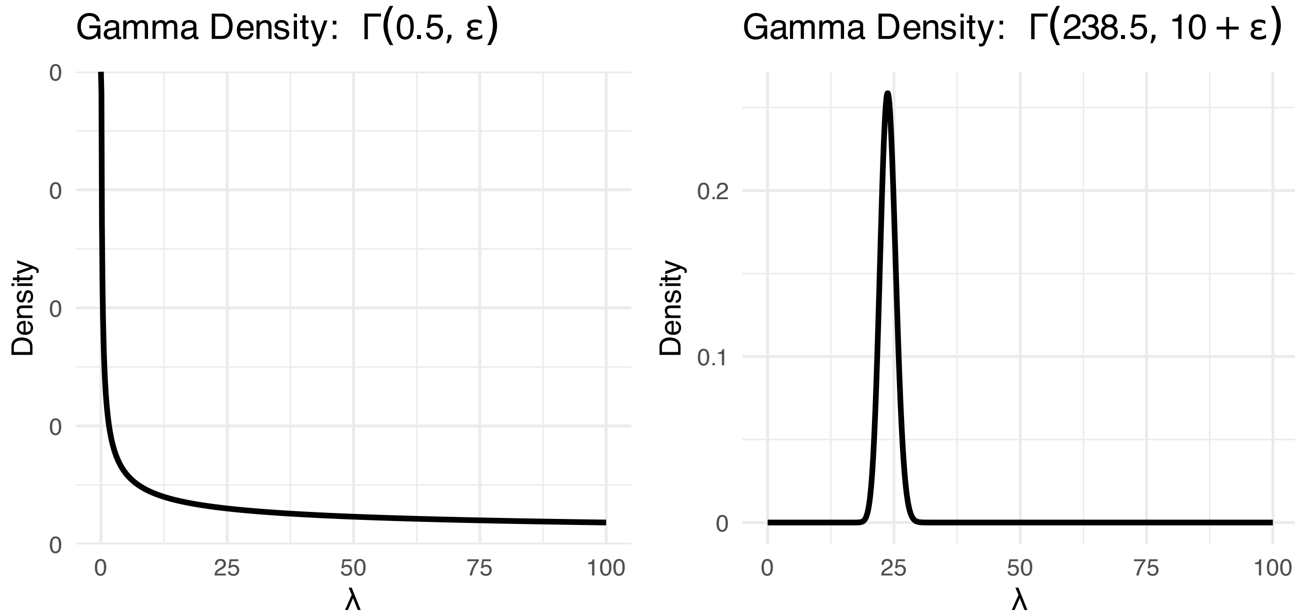

To obtain a 95% posterior predictive interval, we draw θ from its posterior, then draw y from the corresponding Poisson distribution. With these draws, we can obtain the necessary quantiles.

Table 1: Posterior predictive interval.

2.5%	97.5%
14.975	34

(b)

Table 2: Data for airline accidents.

Year	Fatal Accidents	Passenger Deaths	Death Rate
1976	24	734	0.19
1977	25	516	0.12
1978	31	754	0.15
1979	31	877	0.16
1980	22	814	0.14
1981	21	362	0.06
1982	26	764	0.13
1983	20	809	0.13
1984	16	223	0.03
1985	22	1066	0.15

In part (a), we ignored how many flights there are. We can incorporate this information into our model by using passenger_miles as a measure of exposure. The parameter θ is now the rate of fatal accidents per year per 100 million passenger miles. Note that this rate is over an order of magnitude smaller than the death rate in the table because the number of fatal accidents is an order of magnitude smaller than the number of passenger deaths. The posterior is $\theta \mid y \, \mathrm{gamma}(238.5,\ 57158.69)$.

Table 3: Posterior predictive interval.

5%	95%
24	44

(c)

Repeat (a) above, replacing 'fatal accidents' with 'passenger deaths.'
Here we use the same model as in part a but for the number of passenger deaths instead of fatal accidents.

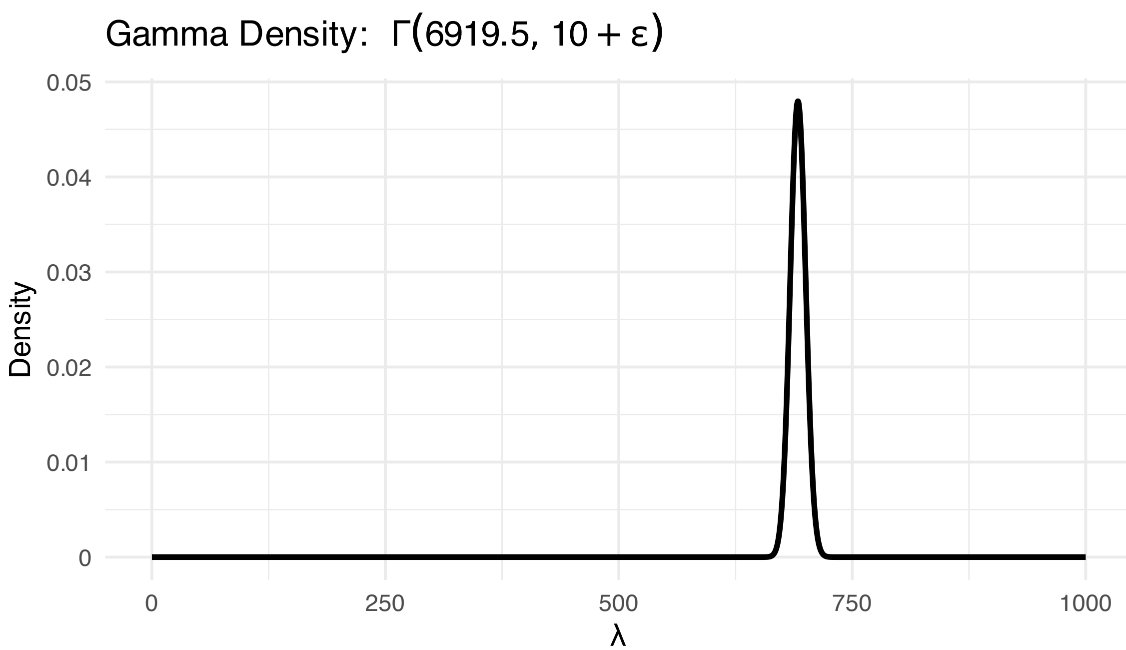

Here we see that only 1 of the 10 observations in the dataset lie within the 95% posterior predictive interval.

Table 4: Posterior predictive interval.

2.5%	97.5%
643	745

(d)

Repeat (b) above, replacing 'fatal accidents' with 'passenger deaths.'
Here we use the same model as in part a but for the number of passenger deaths instead of fatal accidents.

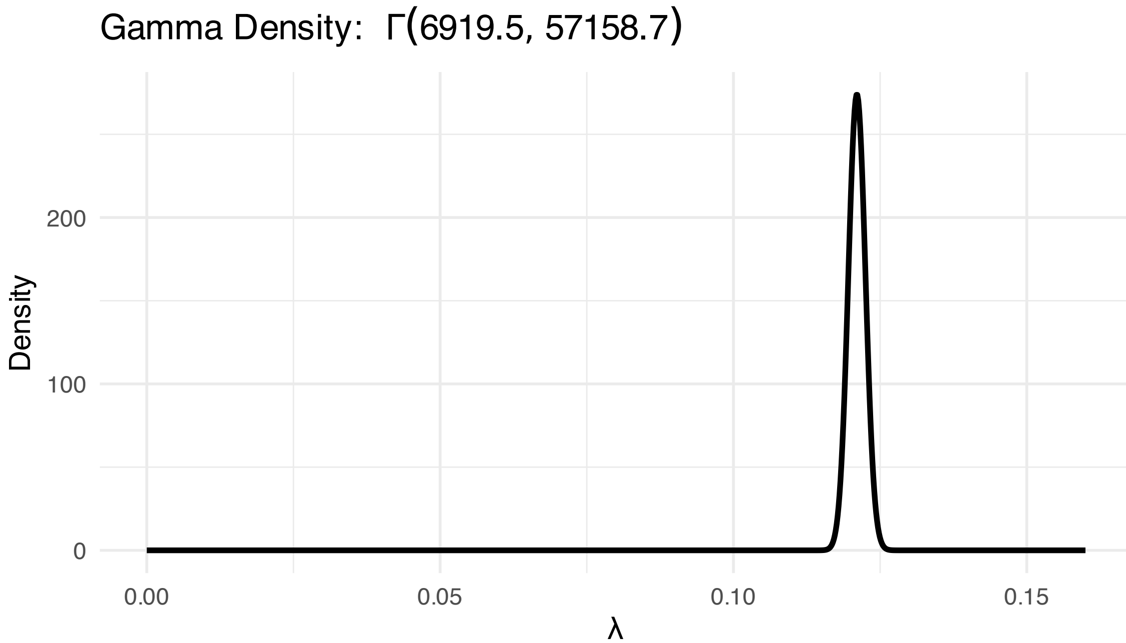

Here we see that none of the observed values falls into the 95% posterior predictive interval.

Table 5: Posterior predictive interval (d).

5%	95%
913.95	1021.05

(e)

There are a number of issues to consider that are not mentioned in the question or suggested by the data. The number of fatal accidents depends on the number of miles flown by airplanes: if there are more flights, there will likely be more accidents. However, the number of flights isn't directly accounted for in the number of passenger miles since the number of passengers per flight can vary from year to year. In any case, the number of passenger deaths per year is not independent because passengers on the same flight will have more similar survival chances.