

APTS Coursework for Statistical Asymptotics, 2016

The work provided here is intended to take students no more than a few hours to complete. Students should talk to their supervisors to find out whether or not their department requires this work as part of any accreditation process. It is anticipated that departments will decide on the appropriate level of assessment locally, and may choose to drop some (or indeed all) of the parts, accordingly.

A random variable X has the *inverse Gaussian distribution*, $IG(\mu, \lambda)$, if its probability density function is

$$f(x; \mu, \lambda) = \frac{\sqrt{\lambda}}{\sqrt{2\pi}} x^{-3/2} \exp\left\{-\frac{\lambda(x - \mu)^2}{2\mu^2 x}\right\}, \quad x > 0, \lambda > 0, \mu > 0.$$

(i) Show that this distribution constitutes an exponential family model. What are the natural parameters? What is the minimal sufficient statistic based on an independent sample of size n from this distribution? Find the cumulant generating function of X , and hence calculate the mean and variance of X .

(ii) Suppose X_1, \dots, X_n are independent, identically distributed $IG(\mu, \lambda)$. What is the exact distribution of $\bar{X} = n^{-1} \sum_{i=1}^n X_i$?

(iii) Construct the saddlepoint approximation to the density of \bar{X} and comment on whether or not it is exact.

(iv) Verify that the distribution function of X has the form

$$P(X \leq x) = \Phi\left\{\sqrt{\frac{\lambda}{x}}\left(\frac{x}{\mu} - 1\right)\right\} + \exp\left(\frac{2\lambda}{\mu}\right) \Phi\left\{-\sqrt{\frac{\lambda}{x}}\left(\frac{x}{\mu} + 1\right)\right\},$$

in terms of the standard normal distribution function Φ .

Investigate numerically the accuracy of a Lugannani-Rice (LR) approximation to the distribution function of \bar{X} , for small n , say $n = 5, 10, 15$, and a range of values of μ, λ . Is the asymptotically equivalent form of the LR approximation (obtained by calculating $\Phi(r_s^*)$ where Φ is the standard normal distribution function and r_s^* is introduced in Chapter 3 just after the LN approximation) more, or less, accurate? How bad is normal approximation to the distribution function based directly on the CLT?

(v) Find expressions for the maximum likelihood estimators $\hat{\mu}, \hat{\lambda}$, based on a sample X_1, \dots, X_n . Using the distributional result that $\lambda(\sum X_i^{-1} - n\bar{X}^{-1})$ is distributed as χ_{n-1}^2 , independently of \bar{X} , examine accuracy of the p^* -formula for the joint density of $(\hat{\mu}, \hat{\lambda})$.