Short answers to selected exercises from the practice midterm	Statistics
VU	Fall 2019

3 (a) The Fisher information in a single observation is

$$i_\lambda = \mathbb{V}\mathrm{ar}_\lambda\Big(\frac{\partial}{\partial\lambda}\log p_\lambda(X_1)\Big) = \mathbb{V}\mathrm{ar}_\lambda X_1 = \frac{1}{\lambda^2}.$$

Hence, the Fisher information in the whole vector $(X_1, \dots, X_n)$ is $I_\lambda = n i_\lambda = n/\lambda^2$.

(b) The Cramér-Rao lower bound for the variance of an unbiased estimator of $g(\lambda) = 1/\lambda$ is

$$\frac{(g'(\lambda))^2}{I_\lambda} = \frac{(-1/\lambda^2)^2}{n/\lambda^2} = \frac{1}{n\lambda^2}.$$

(c) The likelihood for λ is $\lambda^n \exp(-n\sum_{i=1}^n x_i)$. By taking logarithms and differentiating with respect to λ we find that the score function is $n/\lambda - \sum_{i=1}^n x_i$. Setting this equal to 0 (and checking that we have a maximum), we see that the MLE for λ is $1/\bar{X}$. Hence, the MLE for $1/\lambda$ is $\bar{X}$.

We have that $\mathbb{E}_\lambda X_1 = 1/\lambda$ and $\mathbb{V}\mathrm{ar}_\lambda X_1 = 1/\lambda^2$. The central limit theorem then implies that

$$\sqrt{n}(\bar{X} - 1/\lambda) \xrightarrow{\mathrm{d}} N(0, 1/\lambda^2)$$

as $n \to \infty$. Hence, for large n, $\sqrt{n}(\bar{X} - 1/\lambda)$ approximately has a $N(0, 1/\lambda^2)$-distribution. This implies that $\bar{X}$ is approximately $N(1/\lambda, 1/(n\lambda^2))$-distributed.

(d) Part (c) shows that for large n, $\bar{X}$ approximately has mean $1/\lambda$ and variance $1/(n\lambda^2)$. Hence the estimator is approximately unbiased for $1/\lambda$ and by (b) its variance approximately equals the Cramér-Rao lower bound.

5 (a) The hypotheses are $H_0 : \mu \leq 1$ and $H_1 : \mu > 1$. This is the situation of the Gauss test. So we use the test statistic $T = \sqrt{n}(\bar{X} - 1)/\sigma$.

(b)
- *Form of the critical region*: Large values of T indicate that H_1 is true. Hence, we use a test of the form "reject H_0 if $T \geq c$" for some appropriately chosen constant $c \in \mathbb{R}$.
- *Exact critical region*: We want a test of level α, i.e. we want that

  $$\sup_{\mu \leq 1} \mathbb{P}_\mu(T \geq c) \leq \alpha.$$

  If μ grows, the probability that T is large grows, hence $\mathbb{P}_\mu(T \geq c)$ is an increasing function of μ. It follows that the supremum is attained at $\mu = 1$, so the requirement reduces to

  $$\mathbb{P}_1(T \geq c) \leq \alpha.$$

  But under $\mathbb{P}_1$, i.e. if $\mu = 1$, the statistic T has a standard normal distribution. It follows that the requirement is fulfilled if $c \geq \xi_{1-\alpha}$. But we want the critical region as large as possible, so we take $c = \xi_{1-\alpha}$.
- *Final test*: Reject H_0 if $T \geq \xi_{1-\alpha}$.