CHAPTER 5 - HYPOTHESIS TESTING - COMPOUND NULL HYPOTHESES

§5.1 Monotone likelihood ratio tests

Can we find UMP tests when the null and alternative hypotheses are both compound and in the form

$$H_0 : \theta \leq \theta_0 \quad \text{vs} \quad H_1 : \theta > \theta_0?$$

The answer is yes, provided the underlying distribution $f_\theta(\mathbf{x})$ possesses a property called the Monotone Likelihood Ratio.

Definition 5.1

A joint distribution $f_\theta(\mathbf{x})$ has a *Monotone Likelihood Ratio* in a statistic $T(\mathbf{x})$ if for any two values of the parameter, θ_1 and θ_2, with $\theta_1 < \theta_2$, the ratio $\frac{f_{\theta_2}(\mathbf{x})}{f_{\theta_1}(\mathbf{x})}$ depends on $\mathbf{x}$ only through $T(\mathbf{x})$, and this ratio is a non-decreasing function of $T(\mathbf{x})$.

Theorem 5.1

If a joint distribution $f_\theta(\mathbf{x})$ has a Monotone Likelihood Ratio in a statistic $T(\mathbf{x})$, then a uniformly most powerful test at size α of the hypotheses

$$H_0 : \theta \leq \theta_0 \quad \text{versus} \quad H_1 : \theta > \theta_0$$

is to reject H_0 if $T(\mathbf{x}) \geq \mathbf{k}$ where $\mathsf{P}_{\theta_0}(T(\mathbf{X}) \geq k) = \alpha$.

Proof: First consider a pair of simple hypotheses

$$H_0 : \theta = \theta_A \quad \text{versus} \quad H_1 : \theta = \theta_B$$

where $\theta_A < \theta_B$. Applying the Neyman-Pearson lemma, the optimal test of this pair of hypotheses has

$$\begin{aligned} C^* &= \left\{ \mathbf{x} : \frac{f_{\theta_A}(\mathbf{x})}{f_{\theta_B}(\mathbf{x})} < k \right\} \\ &= \left\{ \mathbf{x} : \frac{f_{\theta_B}(\mathbf{x})}{f_{\theta_A}(\mathbf{x})} > \frac{1}{k} \right\} \\ &= \{\mathbf{x} : T(\mathbf{x}) > \mathbf{k}^*\} \text{ because of the MLR property} \\ \text{where } \mathsf{P}_{\theta_A}(\mathbf{X} : T(\mathbf{x}) > \mathbf{k}^*) &= \alpha \end{aligned}$$

Now consider two choices for θ_A or θ_B:

$\theta_A = \theta_0$ **so** $\theta_B > \theta_0$**:** Since the critical region C^* does not depend on θ_B, it is also the uniformly most powerful test of the pair of hypotheses

$$H_0 : \theta = \theta_0 \quad \text{versus} \quad H_1 : \theta > \theta_0$$

Notice that Θ_1 is here also the alternative parameter space for the pair of compound hypotheses of interest. Since the test C^* is uniformly most powerful,

$$\mathsf{P}_\theta(\mathbf{X} \in C^*) \geq \mathsf{P}_\theta(\mathbf{X} \in C) \; \forall \, \theta \in \Theta_1 \text{ and for all size } \alpha \; C \tag{1}$$

Fix k^*_α by setting $\mathsf{P}_{\theta_0}(T(\mathbf{x}) > \mathbf{k}^*_\alpha) = \alpha$.

$\theta_B = \theta_0$ **so** $\theta_A < \theta_0$**:** Using k^*_α defined above, we consider the size of the most powerful test of

$$H_0 : \theta = \theta_A \quad \text{versus} \quad H_1 : \theta = \theta_0$$

that is the size of the test using $C^* = \{\mathbf{x} : T(\mathbf{x}) > \mathbf{k}^*_\alpha\}$.

$$\begin{aligned}
\text{size} &= \mathsf{P}_{\theta_A}(T(\mathbf{X}) > \mathbf{k}^*_\alpha) \\
&\leq \text{ power of the test (by Example 3 on problem sheet 7)} \\
&= \mathsf{P}_{\theta_0}(T(\mathbf{X}) > \mathbf{k}^*_\alpha) \\
&= \alpha \\
\text{i.e. } \mathsf{P}_\theta(T(\mathbf{X}) > \mathbf{k}^*_\alpha) &\leq \mathsf{P}_{\theta_0}(T(\mathbf{X}) > \mathbf{k}^*_\alpha) \; \forall\, \theta < \theta_\mathbf{0} \qquad (2)
\end{aligned}$$

Equations (1) and (2) taken together are the two conditions which must be satisfied for a test to be uniformly most powerful and hence the UMP test is to reject H_0 if $T(\mathbf{x}) \geq \mathbf{k}$ where $\mathsf{P}_{\theta_0}(T(\mathbf{X}) \geq k) = \alpha$.

Examples

1. Exponential Suppose $X_1, \ldots, X_n$ are iid $\exp(\lambda)$ random variables, where λ is the rate parameter:

$$\begin{aligned}
f_\lambda(\mathbf{x}) &= \lambda^n \exp(-\lambda \sum_{i=1}^n x_i) \\
\frac{f_{\lambda_2}(\mathbf{x})}{f_{\lambda_1}(\mathbf{x})} &= \left(\frac{\lambda_2}{\lambda_1}\right)^n \exp(-(\lambda_2 - \lambda_1) \sum_{i=1}^n x_i) \\
&= \left(\frac{\lambda_2}{\lambda_1}\right)^n \exp(-(\lambda_2 - \lambda_1) T(\mathbf{x})) \text{ where } T(\mathbf{x}) = \sum_{i=1}^n x_i
\end{aligned}$$

When $\lambda_2 > \lambda_1$, the likelihood ratio is a decreasing function of $T(\mathbf{x}) = \sum_{i=1}^n x_i$. Therefore there is a Monotone Likelihood Ratio in $-\sum_{i=1}^n X_i$. Theorem 5.1 then tells us that the UMP test of

$$\begin{aligned}
H_0 : \lambda \leq \lambda_0 \qquad &\text{vs} \qquad H_1 : \lambda > \lambda_0 \\
\text{is } C^* &= \left\{\mathbf{x} : -\sum_{i=1}^n x_i \geq k\right\} \\
&= \left\{\mathbf{x} : \sum_{i=1}^n x_i \leq k_\alpha\right\} \\
\text{where } \alpha &= \mathsf{P}_{\lambda_0}(\sum_{i=1}^n X_i \leq k_\alpha)
\end{aligned}$$

We know that adding independent exponential distributions leads to a Gamma distribution, $\sum_{i=1}^n X_i \sim \Gamma am(n, \lambda_0)$ when $\lambda = \lambda_0$, and so k_α is the lower α quantile of $\Gamma am(n, \lambda_0)$.

2. Poisson Suppose $X_1, \ldots, X_n$ are iid $Pois(\mu)$ random variables:

$$\begin{aligned}
f_\mu(\mathbf{x}) &= \frac{\mu^{\sum_{i=1}^n x_i} \exp(-n\mu)}{\prod_{i=1}^n x_i!} \\
\frac{f_{\mu_2}(\mathbf{x})}{f_{\mu_1}(\mathbf{x})} &= \left(\frac{\mu_2}{\mu_1}\right)^{\sum_{i=1}^n x_i} \exp(-n(\mu_2 - \mu_1)) \\
&= \left(\frac{\lambda_2}{\lambda_1}\right)^{T(\mathbf{x})} \exp(-n(\mu_2 - \mu_1)) \text{ where } T(\mathbf{x}) = \sum_{i=1}^n x_i
\end{aligned}$$

When $\mu_2 > \mu_1$, the likelihood ratio is an increasing function of $T(\mathbf{x}) = \sum_{i=1}^n x_i$. Therefore there is a Monotone Likelihood Ratio in $\sum_{i=1}^n X_i$. Theorem 5.1 then tells us that the UMP test of

$$\begin{aligned} H_0 : \mu \leq \mu_0 \quad & \text{vs} \quad H_1 : \mu > \mu_0 \\ \text{is } C^* &= \left\{ \mathbf{x} : \sum_{i=1}^n x_i \geq k_\alpha \right\} \\ \text{where } \alpha &= \mathsf{P}_{\mu_0}(\sum_{i=1}^n X_i \geq k_\alpha) \end{aligned}$$

We know that adding independent Poisson distributions leads to another Poisson distribution, $\sum_{i=1}^n X_i \sim Pois(n\mu_0)$ when $\mu = \mu_0$, and so k_α is the upper α quantile of $Pois(n\mu_0)$.

Two generalisations of Theorem 5.1

1. How do we find the UMP test for the other direction of one-sided test

$$H_0 : \theta \geq \theta_0 \quad \text{vs} \quad H_1 : \theta < \theta_0?$$

A slight modification of Theorem 5.1 tells us that the UMP test is to reject H_0 if $T(\mathbf{x}) \leq \mathbf{k}$ where $\mathsf{P}_{\theta_0}(T(\mathbf{X}) \leq k) = \alpha$.

2. How do we tackle a two-sided alternative

$$H_0 : \theta \in [\theta_1, \theta_2] \quad \text{vs} \quad H_1 : \theta < \theta_1 \text{ or } \theta > \theta_2?$$

We return to the concept of unbiased tests (§4.3) and construct a size α uniformly most powerful unbiased test from the union of two size $\alpha/2$ critical regions corresponding to the one-sided tests related to the two distinct intervals of the alternative parameter space:

$$\begin{aligned} \text{Test 1: } H_0 : \theta \geq \theta_1 \quad & \text{vs} \quad H_1 : \theta < \theta_1 \\ C_1 &= \{\mathbf{x} : T(\mathbf{x}) \leq k_1\} \\ \text{where } \mathsf{P}_{\theta_1}(T(\mathbf{X}) \leq k_1) &= \frac{\alpha}{2} \\ \text{Test 2: } H_0 : \theta \leq \theta_2 \quad & \text{vs} \quad H_1 : \theta > \theta_2 \\ C_2 &= \{\mathbf{x} : T(\mathbf{x}) \geq k_2\} \\ \text{where } \mathsf{P}_{\theta_2}(T(\mathbf{X}) \geq k_2) &= \frac{\alpha}{2} \\ \text{Overall } C^* &= C_1 \cup C_2. \end{aligned}$$

Example Suppose $X_1, \ldots, X_n$ are iid $N(\mu, 1)$ random variables:

$$\begin{aligned} f_\mu(\mathbf{x}) &= \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\sum_{i=1}^n (x_i - \mu)^2\right) \\ \frac{f_{\mu_2}(\mathbf{x})}{f_{\mu_1}(\mathbf{x})} &= \exp\left(-\frac{1}{2}(\sum_{i=1}^n (x_i - \mu_2)^2 - \sum_{i=1}^n (x_i - \mu_1)^2)\right) \\ &= \exp\left(-\frac{1}{2}(n(\mu_2^2 - \mu_1^2) - 2(\mu_2 - \mu_1)T(\mathbf{x}))\right) \text{ where } T(\mathbf{x}) = \sum_{i=1}^n x_i \end{aligned}$$

When $\mu_2 > \mu_1$, the likelihood ratio is an increasing function of $T(\mathbf{x}) = \sum_{i=1}^n x_i$. Therefore there is a Monotone Likelihood Ratio in $\sum_{i=1}^n X_i$.

Suppose we wish to test

$$H_0 : \mu \in [-\gamma, \gamma] \quad \text{vs} \quad H_1 : \mu < -\gamma \text{ or } \mu > \gamma$$

For a size α uniformly most powerful unbiased test, we construct two size $\alpha/2$ tests, one for each tail:

$$\begin{aligned}
\text{Test 1: } H_0 : \mu \geq -\gamma \quad & \text{vs} \quad H_1 : \mu < -\gamma \\
C_1 &= \{\mathbf{x} : T(\mathbf{x}) \leq k_1\} \\
\text{where } \mathsf{P}_{-\gamma}(T(\mathbf{X}) \leq k_1) &= \frac{\alpha}{2} \\
\text{Test 2: } H_0 : \mu \leq \gamma \quad & \text{vs} \quad H_1 : \mu > \gamma \\
C_2 &= \{\mathbf{x} : T(\mathbf{x}) \geq k_2\} \\
\text{where } \mathsf{P}_{\gamma}(T(\mathbf{X}) \geq k_2) &= \frac{\alpha}{2}
\end{aligned}$$

Since $\sum_{i=1}^n X_i \sim N(n\mu, n)$, k_1 is the lower $\alpha/2$ quantile of a $N(-n\gamma, n)$, and k_2 is the upper $\alpha/2$ quantile of a $N(n\gamma, n)$. Overall we reject H_0 if $\sum_{i=1}^n x_i$ is less than k_1 or greater than k_2.

§5.2 Generalised likelihood ratio tests

Generalised likelihood ratio tests are a very general class of test which do not promise any optimality properties, but which are quite intuitive, can deal with a wide range of hypotheses, can be used when there are additional unknown parameters (nuisance parameters), and which sometimes have a useful approximate asymptotic behaviour which can be used to generate critical values.

Definition 5.2

Suppose $X_1, \ldots, X_n$ have joint distribution $f_\theta(\mathbf{x})$ then the *Generalised Likelihood Ratio* test of the hypotheses

$$H_0 : \theta \in \Theta_0 \quad \text{versus} \quad H_1 : \theta \in \Theta_1$$

uses test statistic

$$\begin{aligned}
T(\mathbf{x}) &= \frac{\sup_{\theta \in \Theta_0} f_\theta(\mathbf{x})}{\sup_{\theta \in \Theta} f_\theta(\mathbf{x})} \text{ where } \Theta = \Theta_0 \cup \Theta_1 \\
&= \frac{f_{\hat{\theta}_0}(\mathbf{x})}{f_{\hat{\theta}}(\mathbf{x})}
\end{aligned}$$

where $\hat{\theta}_0$ is the MLE under H_0 and $\hat{\theta}$ is the MLE under $\Theta_0 \cup \Theta_1$. The null hypothesis is rejected for particularly small values of $T(\mathbf{x})$.

Provided that $X_1, \ldots, X_n$ are independent and identically distributed under H_0, the CRLB exists, and the hypotheses are nested, an approximate asymptotic result holds:

$$-2 \ln T(\mathbf{X}) \sim \chi^2_{k-r} \text{ under } H_0$$

where k is the dimension of Θ and r is the dimension of Θ_0. If H_0 is not true, the test statistic will tend to be larger than this.

Examples

1. Bernoulli, two populations Suppose $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$ are two mutually independent sets of variables with $X_i \sim Bern(p_1)$ and $Y_i \sim Bern(p_2)$, and we wish to test

$$H_0 : p_1 = p_2 \quad \text{versus} \quad H_1 : p_1 \neq p_2.$$

To apply a Generalised Likelihood Ratio test, we need to find the MLEs of the common parameter under H_0, which we denote p, and the MLEs of the two parameters under $H_0 \cup H_1$:

Under H_0

$$\begin{aligned}
L(p) &= p^{\sum_{i=1}^n x_i}(1-p)^{n-\sum_{i=1}^n x_i} p^{\sum_{i=1}^n y_i}(1-p)^{n-\sum_{i=1}^n y_i} \\
\ell(p) &= (\sum_{i=1}^n x_i + \sum_{i=1}^n y_i)\ln p + (2n - \sum_{i=1}^n x_i - \sum_{i=1}^n y_i)\ln(1-p) \\
\frac{d\ell(p)}{dp} &= \frac{\sum_{i=1}^n x_i + \sum_{i=1}^n y_i}{p} - \frac{2n - \sum_{i=1}^n x_i - \sum_{i=1}^n y_i}{1-p} \\
\frac{d^2\ell(p)}{dp^2} &= -\frac{\sum_{i=1}^n x_i + \sum_{i=1}^n y_i}{p^2} - \frac{2n - \sum_{i=1}^n x_i - \sum_{i=1}^n y_i}{(1-p)^2} \\
\Rightarrow \hat{p} &= \frac{\sum_{i=1}^n x_i + \sum_{i=1}^n y_i}{2n}.
\end{aligned}$$

Under $H_0 \cup H_1$

$$\begin{aligned}
L(p_1, p_2) &= p_1^{\sum_{i=1}^n x_i}(1-p_1)^{n-\sum_{i=1}^n x_i} p_2^{\sum_{i=1}^n y_i}(1-p_2)^{n-\sum_{i=1}^n y_i} \\
\ell(p_1, p_2) &= (\sum_{i=1}^n x_i)\ln p_1 + (n - \sum_{i=1}^n x_i)\ln(1-p_1) + (\sum_{i=1}^n y_i)\ln p_2 + (n - \sum_{i=1}^n y_i)\ln(1-p_2) \\
\frac{\partial\ell(p_1, p_2)}{\partial p_1} &= \frac{\sum_{i=1}^n x_i}{p_1} - \frac{n - \sum_{i=1}^n x_i}{1-p_1} \\
\frac{\partial\ell(p_1, p_2)}{\partial p_2} &= \frac{\sum_{i=1}^n y_i}{p_2} - \frac{n - \sum_{i=1}^n y_i}{1-p_2} \\
\frac{\partial^2\ell(p_1, p_2)}{\partial p_1^2} &= -\frac{\sum_{i=1}^n x_i}{p_1^2} - \frac{n - \sum_{i=1}^n x_i}{(1-p_1)^2} \\
\frac{\partial^2\ell(p_1, p_2)}{\partial p_2^2} &= -\frac{\sum_{i=1}^n y_i}{p_2^2} - \frac{n - \sum_{i=1}^n y_i}{(1-p_2)^2} \\
\frac{\partial^2\ell(p_1, p_2)}{\partial p_1 p_2} &= 0 \\
\Rightarrow \hat{p}_1 &= \frac{\sum_{i=1}^n x_i}{n} \\
\text{and } \hat{p}_2 &= \frac{\sum_{i=1}^n y_i}{n}
\end{aligned}$$

Test statistic

$$T(\mathbf{x}, \mathbf{y}) = \frac{L(p)|_{\hat{p}}}{L(p_1, p_2)|_{\hat{p}_1, \hat{p}_2}}$$

$$
\begin{aligned}
&= \frac{\hat{p}^{2n\hat{p}}(1-\hat{p})^{2n(1-\hat{p})}}{\hat{p}_1^{n\hat{p}_1}(1-\hat{p}_1)^{n(1-\hat{p}_1)}\hat{p}_2^{n\hat{p}_2}(1-\hat{p}_2)^{n(1-\hat{p}_2)}} \\
&= \left(\frac{\hat{p}^{2\hat{p}}(1-\hat{p})^{2(1-\hat{p})}}{\hat{p}_1^{\hat{p}_1}(1-\hat{p}_1)^{(1-\hat{p}_1)}\hat{p}_2^{\hat{p}_2}(1-\hat{p}_2)^{(1-\hat{p}_2)}}\right)^n
\end{aligned}
$$

We would reject H_0 if this $T(\mathbf{x},\mathbf{y})$ were particularly small. However, as the sampling distribution may be difficult to derive, we can use instead the asymptotic result which says that

$$-2\ln T(\mathbf{X},\mathbf{Y}) \sim \chi_1^2 \text{ as } n \to \infty \text{ under } H_0.$$

We would reject H_0 at size α if $-2\ln T(\mathbf{x},\mathbf{y})$ exceeds the top α quantile of a χ_1^2.

2. Shifted exponential, two populations Suppose $X_1,\ldots,X_n$ and $Y_1,\ldots,Y_n$ are two mutually independent sets of shifted exponential variables, sharing a common known shift η but potentially different mean parameters:

$$
\begin{aligned}
f_{\theta_1}(x) &= \frac{1}{\theta_1}\exp\left(-\frac{x-\eta}{\theta_1}\right),\ x>\eta \\
f_{\theta_2}(y) &= \frac{1}{\theta_2}\exp\left(-\frac{y-\eta}{\theta_2}\right),\ y>\eta
\end{aligned}
$$

We are interested in testing

$$H_0:\theta_1=\theta_2 \quad \text{versus} \quad H_1:\theta_1\neq\theta_2.$$

To apply a Generalised Likelihood Ratio test, we need to find the MLEs of the common parameter under H_0, which we denote θ, and the MLEs of the two parameters under $H_0\cup H_1$:

Under H_0

$$
\begin{aligned}
L(\theta) &= \left(\frac{1}{\theta}\right)^{2n}\exp\left(-\frac{\sum_{i=1}^n(x_i-\eta)+\sum_{i=1}^n(y_i-\eta)}{\theta}\right) \\
\ell(\theta) &= -2n\ln\theta-\frac{\sum_{i=1}^n(x_i-\eta)+\sum_{i=1}^n(y_i-\eta)}{\theta} \\
\frac{d\ell(\theta)}{d\theta} &= \frac{-2n}{\theta}+\frac{\sum_{i=1}^n(x_i-\eta)+\sum_{i=1}^n(y_i-\eta)}{\theta^2} \\
\frac{d^2\ell(\theta)}{d\theta^2} &= \frac{2n}{\theta^2}-2\frac{\sum_{i=1}^n(x_i-\eta)+\sum_{i=1}^n(y_i-\eta)}{\theta^3} \\
\Rightarrow\ \hat{\theta} &= \frac{\sum_{i=1}^n(x_i-\eta)+\sum_{i=1}^n(y_i-\eta)}{2n}.
\end{aligned}
$$

Under $H_0\cup H_1$

$$
\begin{aligned}
L(\theta_1,\theta_2) &= \left(\frac{1}{\theta_1}\right)^{n}\exp\left(-\frac{\sum_{i=1}^n(x_i-\eta)}{\theta_1}\right)\left(\frac{1}{\theta_2}\right)^{n}\exp\left(-\frac{\sum_{i=1}^n(y_i-\eta)}{\theta_2}\right) \\
\ell(\theta_1,\theta_2) &= -n\ln\theta_1-\frac{\sum_{i=1}^n(x_i-\eta)}{\theta_1}-n\ln\theta_2-\frac{\sum_{i=1}^n(y_i-\eta)}{\theta_2} \\
\frac{\partial\ell(\theta_1,\theta_2)}{\partial\theta_1} &= \frac{-n}{\theta_1}+\frac{\sum_{i=1}^n(x_i-\eta)}{\theta_1^2}
\end{aligned}
$$

$$
\begin{aligned}
\frac{\partial \ell(\theta_1,\theta_2)}{\partial \theta_2} &= \frac{-n}{\theta_2} + \frac{\sum_{i=1}^n (y_i-\eta)}{\theta_2^2} \\
\frac{\partial^2 \ell(\theta_1,\theta_2)}{\partial \theta_1^2} &= \frac{n}{\theta_1^2} - 2\frac{\sum_{i=1}^n (x_i-\eta)}{\theta_1^3} \\
\frac{\partial^2 \ell(\theta_1,\theta_2)}{\partial \theta_2^2} &= \frac{n}{\theta_2^2} - 2\frac{\sum_{i=1}^n (y_i-\eta)}{\theta_2^3} \\
\frac{\partial^2 \ell(\theta_1,\theta_2)}{\partial \theta_1 \theta_2} &= 0 \\
\Rightarrow\ \hat{\theta}_1 &= \frac{\sum_{i=1}^n (x_i-\eta)}{n} \\
\text{and}\ \ \hat{\theta}_2 &= \frac{\sum_{i=1}^n (y_i-\eta)}{n}
\end{aligned}
$$

Test statistic

$$
\begin{aligned}
T(\mathbf{x},\mathbf{y}) &= \frac{L(\theta)|_{\hat{\theta}}}{L(\theta_1,\theta_2)|_{\hat{\theta}_1,\hat{\theta}_2}} \\
&= \frac{\left(\frac{1}{\hat{\theta}}\right)^{2n} \exp\left(-\frac{\sum_{i=1}^n (x_i-\eta)+\sum_{i=1}^n (y_i-\eta)}{\hat{\theta}}\right)}{\left(\frac{1}{\hat{\theta}_1}\right)^{n} \exp\left(-\frac{\sum_{i=1}^n (x_i-\eta)}{\hat{\theta}_1}\right)\left(\frac{1}{\hat{\theta}_2}\right)^{n} \exp\left(-\frac{\sum_{i=1}^n (y_i-\eta)}{\hat{\theta}_2}\right)} \\
&= \frac{\left(\frac{1}{\hat{\theta}}\right)^{2n} \exp(-2n)}{\left(\frac{1}{\hat{\theta}_1}\right)^{n} \exp(-n)\left(\frac{1}{\hat{\theta}_2}\right)^{n} \exp(-n)} \\
&= \left(\frac{\hat{\theta}_1\hat{\theta}_2}{\hat{\theta}^2}\right)^n
\end{aligned}
$$

We would reject H_0 if this $T(\mathbf{x},\mathbf{y})$ were particularly small. However, as the sampling distribution may be difficult to derive, we can use instead the asymptotic result which says that

$$-2\ln T(\mathbf{X},\mathbf{Y}) \sim \chi_1^2 \ \text{ as } n \to \infty \text{ under } H_0.$$

We would reject H_0 at size α if $-2\ln T(\mathbf{x},\mathbf{y})$ exceeds the top α quantile of a χ_1^2.

3. Shifted exponential, two populations, with nuisance parameter How would the test change, compared to Example 2, if the shift parameter η was unknown rather than known, i.e. η is a nuisance parameter?

The first point to note is that if η is unknown, then the distributions are no longer part of an exponential family (since the support depends on an unknown parameter). Secondly, the likelihood function, under H_0 or under $H_0 \cup H_1$ is a function of one additional variable (denote this additional variable η_0 under H_0 and just η under $H_0 \cup H_1$ to remind ourselves that the MLEs may not be equal in the two situations). Notice that, under H_0

$$
\begin{aligned}
L(\theta,\eta_0) &= \left(\frac{1}{\theta}\right)^{2n} \exp\left(-\frac{\sum_{i=1}^n (x_i-\eta_0)+\sum_{i=1}^n (y_i-\eta_0)}{\theta}\right) \\
\ell(\theta,\eta_0) &= -2n\ln\theta - \frac{\sum_{i=1}^n (x_i-\eta_0)+\sum_{i=1}^n (y_i-\eta_0)}{\theta} \\
\frac{\partial \ell(\theta,\eta_0)}{\partial \eta_0} &= \frac{2n}{\theta}.
\end{aligned}
$$

From this we can see that $\hat{\eta}_0$ will occur on the boundary of the parameter space where η is as large as possible, that is

$$\begin{aligned}
\hat{\eta}_0 &= \min\{x_1, \ldots, x_n, y_1, \ldots, y_n\} \\
\text{while } \hat{\theta} &= \frac{\sum_{i=1}^n (x_i - \hat{\eta}_0) + \sum_{i=1}^n (y_i - \hat{\eta}_0)}{2n}.
\end{aligned}$$

Similarly, under $H_0 \cup H_1$

$$\begin{aligned}
\hat{\eta} &= \min\{x_1, \ldots, x_n, y_1, \ldots, y_n\} \\
\hat{\theta}_1 &= \frac{\sum_{i=1}^n (x_i - \hat{\eta})}{n} \\
\hat{\theta}_2 &= \frac{\sum_{i=1}^n (y_i - \hat{\eta})}{n}
\end{aligned}$$

Calculating the test statistic,

$$T(\mathbf{x}, \mathbf{y}) = \left(\frac{\hat{\theta}_1 \hat{\theta}_2}{\hat{\theta}^2}\right)^n$$

(noting that these MLEs are as above, and are not the same as those in Example 2). A crucial difference with Example 2 is that the asymptotic result does not hold because the the conditions for the CRLB to exist fail (namely that the support does not depend on unknown parameters).

4. Normal, with nuisance parameter Suppose $X, \ldots, X_n$ are iid $N(\mu, \sigma^2)$ where neither μ nor σ^2 are known, and we wish to test:

$$H_0 : \mu = \mu_0 \quad \text{versus} \quad H_1 : \mu \neq \mu_0$$

The fact that we have the unknown nuisance parameter σ^2 means we cannot use the UMPU test which would follow from an application of the Neyman-Pearson lemma, and so we consider instead a Generalised likelihood Ratio test.

Under H_0

$$\begin{aligned}
\hat{\mu}_0 &= \mu_0 \\
\hat{\sigma}_0^2 &= \frac{1}{n}\sum_{i=1}^n (x_i - \hat{\mu}_0)^2 \\
L(\hat{\mu}_0, \hat{\sigma}_0^2) &= \left(\frac{1}{\sqrt{2\pi\hat{\sigma}_0^2}}\right)^n \exp\left(-\frac{1}{2\hat{\sigma}_0^2}\sum_{i=1}^n (x_i - \hat{\mu}_0)^2\right) \\
&= \left(\frac{1}{\sqrt{2\pi\hat{\sigma}_0^2}}\right)^n \exp(-n/2)
\end{aligned}$$

Under $H_0 \cup H_1$

$$\begin{aligned}
\hat{\mu} &= \frac{1}{n}\sum_{i=1}^n x_i \\
\hat{\sigma}^2 &= \frac{1}{n}\sum_{i=1}^n (x_i - \hat{\mu})^2
\end{aligned}$$

$$\begin{aligned} L(\hat{\mu}, \hat{\sigma}^2) &= \left(\frac{1}{\sqrt{2\pi\hat{\sigma}^2}}\right)^n \exp(-\frac{1}{2\hat{\sigma}^2}\sum_{i=1}^n (x_i - \hat{\mu})^2) \\ &= \left(\frac{1}{\sqrt{2\pi\hat{\sigma}^2}}\right)^n \exp(-n/2) \end{aligned}$$

Test statistic

$$T(\mathbf{x}) = \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_0^2}\right)^{n/2}$$

We would reject H_0 if $T(\mathbf{x})$ is unusually small. Although the asymptotic result would hold in this case, we can actually find an exact result here (which will turn out to be the usual t-test):

$$\begin{aligned} C &= \left\{\mathbf{x} : \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_0^2}\right)^{n/2} < k\right\} \\ &= \left\{\mathbf{x} : \frac{\hat{\sigma}^2}{\hat{\sigma}_0^2} < k_2\right\} \\ &= \left\{\mathbf{x} : \frac{\hat{\sigma}_0^2}{\hat{\sigma}^2} > k_3\right\} \\ &= \left\{\mathbf{x} : \frac{\sum_{i=1}^n (x_i - \mu_0)^2}{\sum_{i=1}^n (x_i - \bar{x})^2} > k_3\right\} \\ &= \left\{\mathbf{x} : \frac{\sum_{i=1}^n (x_i - \bar{x} + \bar{x} - \mu_0)^2}{\sum_{i=1}^n (x_i - \bar{x})^2} > k_3\right\} \\ &= \left\{\mathbf{x} : \frac{\sum_{i=1}^n (x_i - \bar{x})^2 + n(\bar{x} - \mu_0)^2}{\sum_{i=1}^n (x_i - \bar{x})^2} > k_3\right\} \\ &= \left\{\mathbf{x} : 1 + \frac{n(\bar{x} - \mu_0)^2}{\sum_{i=1}^n (x_i - \bar{x})^2} > k_3\right\} \\ &= \left\{\mathbf{x} : \frac{(\bar{x} - \mu_0)^2}{\sum_{i=1}^n (x_i - \bar{x})^2/n} > k_4\right\} \\ &= \left\{\mathbf{x} : \frac{(\bar{x} - \mu_0)^2}{(n-1)S^2/n} > k_4\right\} \text{ where } S^2 = \sum_{i=1}^n (x_i - \bar{x})^2/(n-1) \\ &= \left\{\mathbf{x} : \frac{(\bar{x} - \mu_0)^2}{S^2/n} > k_5\right\} \\ &= \left\{\mathbf{x} : \frac{|\bar{x} - \mu_0|}{\sqrt{S^2/n}} > k_6\right\} \end{aligned}$$

However, we know that under H_0, $\frac{\bar{X} - \mu_0}{\sqrt{S^2/n}} \sim t_{n-1}$ and so this test becomes the usual two-tailed t-test.

What would have changed if we had been interested in the one-tailed alternative instead?

$$H_0 : \mu = \mu_0 \quad \text{versus} \quad H_1 : \mu > \mu_0$$

Crucially for the Generalised Likelihood Ratio test, the parameter space under $H_0 \cup H_1$ changes:

$$\Theta_0 \cup \Theta_1 = \Big\{(\mu, \sigma^2) : \mu \geq \mu_0, \sigma^2 > 0\Big\}$$

This means that when finding the MLEs under $H_0 \cup H_1$, the constraint on the parameter space must be considered:

$$\frac{\partial \ell(\mu, \sigma^2)}{\partial \mu} = \frac{1}{\sigma^2} \sum_{i=1}^{n} (x_i - \mu)$$

If $\bar{x} < \mu_0$, then the turning point of this first derivative of the log likelihood does not occur in the parameter space, and this first derivative is negative over $\Theta_0 \cup \Theta_1$ implying that the MLE of μ occurs at the smallest possible value, i.e. $\hat{\mu} = \mu_0$. If this happens, then both MLEs under H_0 and $H_0 \cup H_1$ are identical. The consequence of this for the test statistic is

$$T(\mathbf{x}) = \begin{cases} 1, & \bar{x} < \mu_0 \\ \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_0^2}\right)^{n/2}, & \bar{x} \geq \mu_0 \end{cases}$$

Knowing that the generalised likelihood ratio is always less than or equal to one in value, this tells us that the smallest values can only occur in the case where $\bar{x} \geq \mu_0$. Intuitively this makes sense as we would be unlikely to reject this H_0 in favour of the alternative that $\mu > \mu_0$ when we observed an $\bar{x} < \mu_0$. In terms of finding a critical region for the test, we can also update the last steps of the previous derivation

$$\begin{aligned} C &= \left\{\mathbf{x} : \left(\frac{\hat{\sigma}^2}{\hat{\sigma}_0^2}\right)^{n/2} < k\right\} \\ &= \left\{\mathbf{x} : \frac{(\bar{x} - \mu_0)^2}{S^2/n} > k_5\right\} \\ &= \left\{\mathbf{x} : \frac{\bar{x} - \mu_0}{\sqrt{S^2/n}} > k_6\right\} \end{aligned}$$

to see that we are only looking at one tail of the t distribution to find the critical value.