557: MATHEMATICAL STATISTICS II
HYPOTHESIS TESTING: EXAMPLES

Example 1 Suppose that $X_1, \ldots, X_n \sim N(\theta, 1)$. To test

$$\begin{aligned} H_0 &: \theta = 0 \\ H_1 &: \theta = 1 \end{aligned}$$

the most powerful test at level α is based on the statistic

$$\lambda(\underset{\sim}{x}) = \frac{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|1)}{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|0)} = \frac{(2\pi)^{-n/2}\exp\left\{-\frac{1}{2}\sum_{i=1}^{n}(x_i-1)^2\right\}}{(2\pi)^{-n/2}\exp\left\{-\frac{1}{2}\sum_{i=1}^{n}x_i^2\right\}} = \exp\left\{\sum_{i=1}^{n}x_i - n/2\right\}$$

with critical region $\mathcal{R}$ given by $\underset{\sim}{x} \in \mathcal{R}$ if

$$\sum_{i=1}^{n} x_i - \frac{n}{2} > \log k$$

where k is defined by $\Pr[\underset{\sim}{X} \in \mathcal{R}|\theta = 0] = \alpha$. We can convert this to a rejection region of the form $\overline{X} > c_n$. Now, given $\theta = 0$, $\overline{X} \sim N(0, 1/n)$, so

$$\Pr[\underset{\sim}{X} \in \mathcal{R}|\theta = 0] = \Pr[\overline{X} > c_n|\theta = 0] = 1 - \Phi(\sqrt{n}\,c_n) = \alpha \qquad \therefore \qquad c_n = \frac{\Phi^{-1}(1-\alpha)}{\sqrt{n}}$$

For $\alpha = 0.05$, $\Phi^{-1}(1-\alpha) = 1.645$. Hence we reject H_0 in favour of H_1 if

$$\overline{X} > \frac{1.645}{\sqrt{n}}$$

For example, for $n = 25$, $c_n = 0.329$. The power function $\beta(\theta)$ is given by

$$\beta(\theta) = \Pr[\underset{\sim}{X} \in \mathcal{R}|\theta] = \Pr[\overline{X} > c_n|\theta] = 1 - \Phi(\sqrt{n}(c_n - \theta))$$

which we evaluate specifically at $\theta = 1$. Note that $\beta(\theta)$ is an increasing function of θ so that as θ increases, the power to reject H_0 in favour of H_1 increases.

Example 2 Suppose that $X_1, \ldots, X_n \sim Exp(1/\theta)$. To test

$$\begin{aligned} H_0 &: \theta = 2 \\ H_1 &: \theta > 2 \end{aligned}$$

Let $\theta_0 = 2$, $\theta_1 \in \Theta_1 \equiv (2, \infty)$. The most powerful test of the hypotheses

$$\begin{aligned} H_0 &: \theta = \theta_0 \\ H_1 &: \theta = \theta_1 \end{aligned}$$

is given by the Neyman-Pearson Lemma to be

$$\lambda(\underset{\sim}{x}) = \frac{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1)}{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0)} = \left(\frac{\theta_0}{\theta_1}\right)^n \frac{\exp\left\{-\sum_{i=1}^{n}x_i/\theta_1\right\}}{\exp\left\{-\sum_{i=1}^{n}x_i/\theta_0\right\}} = \left(\frac{2}{\theta_1}\right)^n \exp\left\{-\sum_{i=1}^{n}x_i\left[\frac{1}{\theta_1} - \frac{1}{2}\right]\right\} > k.$$

so that, in terms of the sufficient statistic,

$$T(\underset{\sim}{X}) = \sum_{i=1}^{n} X_i > \frac{\log k - n\log(2/\theta_1)}{\dfrac{1}{2} - \dfrac{1}{\theta_1}}$$

say. Hence the critical region is of the form $T(\underset{\sim}{X}) > c_n$, and as under H_0, $T(\underset{\sim}{X}) \sim Gamma(n, 1/2)$, we require that

$$\Pr[T(\underset{\sim}{X}) > c_n | \theta = 2] = \alpha \qquad \therefore \qquad c_n = q_{n,1/2}(1-\alpha)$$

where $q_{a,b}(1-\alpha)$ is the inverse cdf for the $Gamma(a, b)$ distribution evaluated at $1-\alpha$. Consider tests where $\mathcal{R}_T \equiv \{t : t > c\}$; this test has power function

$$\beta(\theta) = \Pr[T(\underset{\sim}{X}) > c \mid \theta] = \int_c^{\infty} \frac{1}{\theta^n \Gamma(n)} t^{n-1} e^{-t/\theta}\, dt \tag{1}$$

which can be computed numerically. Now, note from equation (1) that $\beta(\theta)$ is a decreasing function of c, so therefore the most powerful test across all possible values of $\theta_1 \in \Theta_1$ that attain size/level α is the one with $c = c_n$. Below is a table of $\beta(\theta)$ for different values of n and θ, when $\alpha = 0.05$ and $c = c_n$:

	θ									
n	2.2	2.4	2.6	2.8	3.0	3.2	3.4	3.6	3.8	4.0
2	0.071	0.095	0.121	0.148	0.176	0.204	0.233	0.261	0.288	0.315
3	0.076	0.105	0.139	0.174	0.211	0.248	0.285	0.321	0.357	0.391
4	0.079	0.115	0.154	0.197	0.242	0.287	0.332	0.376	0.418	0.458
5	0.083	0.123	0.169	0.219	0.272	0.324	0.376	0.426	0.473	0.518
10	0.097	0.160	0.235	0.317	0.401	0.481	0.556	0.624	0.683	0.735
20	0.120	0.223	0.348	0.478	0.598	0.701	0.783	0.846	0.893	0.926

Example 3 Suppose that $X_1, \ldots, X_n \sim Bernoulli(\theta)$. A test of

$$\begin{aligned} H_0 &: \theta \leq \theta_0 \\ H_1 &: \theta > \theta_0 \end{aligned}$$

is required. The likelihood ratio for $\theta_1 < \theta_2$ for this model is

$$\lambda(\underset{\sim}{x}) = \frac{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_2)}{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1)} = \frac{\theta_2^{T(\underset{\sim}{x})}(1-\theta_2)^{n-T(\underset{\sim}{x})}}{\theta_1^{T(\underset{\sim}{x})}(1-\theta_1)^{n-T(\underset{\sim}{x})}} = \left(\frac{\theta_2/(1-\theta_2)}{\theta_1/(1-\theta_1)}\right)^{T(\underset{\sim}{x})} \left(\frac{1-\theta_2}{1-\theta_1}\right)^n$$

where $T(\underset{\sim}{X}) = \sum\limits_{i=1}^{n} X_i$. Thus $\lambda(\underset{\sim}{x})$ is a **monotone increasing** function of $T(\underset{\sim}{x})$ as for $\theta_1 < \theta_2$

$$\frac{\theta_2}{(1-\theta_2)} > \frac{\theta_1}{(1-\theta_1)}$$

and by the Karlin-Rubin theorem, the UMP test at level α is based on the critical region

$$\mathcal{R} \equiv \left\{ \underset{\sim}{x} : T(\underset{\sim}{x}) = \sum_{i=1}^{n} x_i > t_0 \right\}$$

To find t_0, we need to solve

$$\Pr[T(\underset{\sim}{X}) > t_0 \mid \theta_0] = \alpha. \tag{2}$$

Now if $\theta = \theta_0$, then $T(\underset{\sim}{X}) \sim Binomial(n, \theta_0)$, so t_0 need only take integer values on $\{0, \ldots, n\}$. Note that the equation (2) can not be solved for all α, as $T(\underset{\sim}{X})$ has a discrete distribution.

Example 4 Consider the likelihood arising from a random sample $X_1, \ldots, X_n$ following a one-parameter Exponential Family model:

$$L(\theta|\underset{\sim}{x}) = h(\underset{\sim}{x})\{c(\theta)\}^n \exp\{w(\theta)T(\underset{\sim}{x})\}$$

where $T(\underset{\sim}{X}) = \sum\limits_{i=1}^{n} t(X_i)$ is a sufficient statistic. For $\theta_1 < \theta_2$

$$\frac{L(\theta_2|\underset{\sim}{x})}{L(\theta_1|\underset{\sim}{x})} = \left(\frac{c(\theta_2)}{c(\theta_1)}\right)^n \exp\{(w(\theta_2) - w(\theta_1))T(\underset{\sim}{x})\}.$$

This is a monotone function of $T(\underset{\sim}{x})$ if $w(\theta)$ is a monotone function; if $w(\theta)$ is non-decreasing, then the test of the hypothesis

$$\begin{aligned} H_0 &: \theta \leq \theta_0 \\ H_1 &: \theta > \theta_0 \end{aligned}$$

that uses the rejection region $\mathcal{R} \equiv \left\{\underset{\sim}{x} : T(\underset{\sim}{x}) \geq t_0\right\}$, where $\Pr[T(\underset{\sim}{X}) \geq t_0 \mid \theta = \theta_0] = \alpha$, is the UMP α level test.

Example 5 Suppose that $X_1 \ldots, X_{n_1} \sim N(\theta_1, \sigma^2)$ and $Y_1, \ldots, Y_{n_2} \sim N(\theta_2, \sigma^2)$ are independent random samples. To test

$$\begin{aligned} H_0 &: \theta_1 = \theta_2 = \theta,\ \sigma^2 \text{ unspecified} \\ H_1 &: \theta_1 \neq \theta_2,\ \sigma^2 \text{ unspecified} \end{aligned}$$

the likelihood ratio statistic is

$$\lambda(\underset{\sim}{x}, \underset{\sim}{y}) = \frac{\sup\limits_{(\theta,\sigma^2)\in\,\Theta_0} f_{\underset{\sim}{X},\underset{\sim}{Y}|\underset{\sim}{\theta}}(\underset{\sim}{x}, \underset{\sim}{y}|\theta, \sigma^2)}{\sup\limits_{(\theta_1,\theta_2,\sigma^2)\in\,\Theta_1} f_{\underset{\sim}{X},\underset{\sim}{Y}|\theta}(\underset{\sim}{x}, \underset{\sim}{y}|\theta_1, \theta_2, \sigma^2)} = \frac{L_0(\widehat{\theta}, \widehat{\sigma}_0 \mid \underset{\sim}{x}, \underset{\sim}{y})}{L_1(\widehat{\theta}_1, \widehat{\theta}_2, \widehat{\sigma}_1 \mid \underset{\sim}{x}, \underset{\sim}{y})}$$

Under H_0, the maximum likelihood estimators of θ and σ^2 are

$$\begin{aligned} \widehat{\theta} &= \frac{\sum\limits_{i=1}^{n_1} X_i + \sum\limits_{i=1}^{n_2} Y_i}{n_1 + n_2} = \frac{n_1\overline{X} + n_2\overline{Y}}{n_1 + n_2} \\ \widehat{\sigma}_0^2 &= \frac{1}{n_1 + n_2}\left[\sum_{i=1}^{n_1}(X_i - \widehat{\theta})^2 + \sum_{i=1}^{n_2}(Y_i - \widehat{\theta})^2\right] \end{aligned}$$

whereas under H_1, the maximum likelihood estimators of θ_1, θ_2 and σ^2 are

$$\widehat{\theta}_1 = \overline{X} \qquad\qquad \widehat{\theta}_2 = \overline{Y}$$

$$\widehat{\sigma}_1^2 = \frac{1}{n_1 + n_2}\left[\sum_{i=1}^{n_1}(X_i - \widehat{\theta}_1)^2 + \sum_{i=1}^{n_2}(Y_i - \widehat{\theta}_2)^2\right]$$

Therefore

$$\lambda(\underset{\sim}{x}, \underset{\sim}{y}) = \left(\frac{\widehat{\sigma}_1^2}{\widehat{\sigma}_0^2}\right)^{(n_1+n_2)/2}$$

Now $\lambda(\underset{\sim}{x}, \underset{\sim}{y}) \leq k$ is equivalent to

$$\frac{\widehat{\sigma}_1^2}{\widehat{\sigma}_0^2} = \frac{\sum\limits_{i=1}^{n_1}(X_i - \widehat{\theta}_1)^2 + \sum\limits_{i=1}^{n_2}(Y_i - \widehat{\theta}_2)^2}{\sum\limits_{i=1}^{n_1}(X_i - \widehat{\theta})^2 + \sum\limits_{i=1}^{n_2}(Y_i - \widehat{\theta})^2} \leq c_1$$

say. In the denominator

$$\sum_{i=1}^{n_1}(X_i-\widehat{\theta})^2=\sum_{i=1}^{n_1}(X_i-\widehat{\theta}_1+\widehat{\theta}_1-\widehat{\theta})^2 = \sum_{i=1}^{n_1}(X_i-\overline{X})^2+n_1\left(\overline{X}-\frac{n_1\overline{X}+n_2\overline{Y}}{n_1+n_2}\right)^2$$

$$= \sum_{i=1}^{n_1}(X_i-\overline{X})^2+\frac{n_1 n_2^2}{(n_1+n_2)^2}\left(\overline{X}-\overline{Y}\right)^2$$

with an equivalent expression for

$$\sum_{i=1}^{n_2}(Y_i-\widehat{\theta})^2=\sum_{i=1}^{n_2}(Y_i-\overline{Y})^2+\frac{n_1^2 n_2}{(n_1+n_2)^2}\left(\overline{X}-\overline{Y}\right)^2$$

Therefore, after substitution into the inequality above, we have

$$\frac{\sum\limits_{i=1}^{n_1}(X_i-\overline{X})^2+\sum\limits_{i=1}^{n_2}(Y_i-\overline{Y})^2}{\sum\limits_{i=1}^{n_1}(X_i-\overline{X})^2+\sum\limits_{i=1}^{n_2}(Y_i-\overline{Y})^2+\dfrac{n_1 n_2}{n_1+n_2}(\overline{X}-\overline{Y})^2}\le c_1$$

which is equivalent to the inequality

$$\frac{\dfrac{n_1 n_2}{n_1+n_2}(\overline{X}-\overline{Y})^2}{\sum\limits_{i=1}^{n_1}(X_i-\overline{X})^2+\sum\limits_{i=1}^{n_2}(Y_i-\overline{Y})^2}\ge c_2$$

or more familiarly

$$T(\underset{\sim}{X},\underset{\sim}{Y})^2=\frac{(\overline{X}-\overline{Y})^2}{s_P^2\left(\dfrac{1}{n_1}+\dfrac{1}{n_2}\right)}\ge(n_1+n_2-2)c_2=c^2 \tag{3}$$

say, where

$$s_P^2=\frac{1}{n_1+n_2-2}\left[\sum_{i=1}^{n_1}(X_i-\overline{X})^2+\sum_{i=1}^{n_2}(Y_i-\overline{Y})^2\right]$$

is the unbiased estimator of σ^2 under H_1. The statistic on the left hand side of equation (3) is, under H_0, the square of a Student-t random variable with n_1+n_2-2 degrees of freedom, and thus the likelihood ratio test is equivalent to the traditional two-sample t-test for the equality of means. The appropriate value of c can be computed using tables of that distribution; we have for a level α test

$$c=\mathrm{St}_{n_1+n_2-2}^{-1}(1-\alpha/2)$$

where $\mathrm{St}_n^{-1}(p)$ is the inverse cdf of the Student-t density with n degrees of freedom evaluated at probability p. Thus the rejection region $\mathcal{R}_T$ is defined by $\mathcal{R}_T\equiv\{t:(t\le -c)\cup(t\ge c)\}$.

Power Function: The power function β can be formed in terms of the difference $\delta=\theta_1-\theta_2$, and a specific σ. We have

$$\beta(\delta,\sigma)=\Pr[T(\underset{\sim}{X},\underset{\sim}{Y})\in\mathcal{R}_T|\delta,\sigma]=\Pr[T(\underset{\sim}{X},\underset{\sim}{Y})\le -c|\delta,\sigma]+\Pr[T(\underset{\sim}{X},\underset{\sim}{Y})\ge c|\delta,\sigma].$$

To compute these probabilities, we need to compute the distribution of $T(\underset{\sim}{X},\underset{\sim}{Y})$ when the difference between the means is δ. It turns out that this distribution is the **non-central Student-t distribution**: if $Z\sim N(\mu,1)$ and $V\sim\chi^2_\nu$ are independent random variables, then

$$T=\frac{Z}{\sqrt{V/\nu}}\sim\text{Student}(\nu,\mu)$$

for which the pdf can be computed using standard methods from MATH 556. The statistic $T(\underset{\sim}{X}, \underset{\sim}{Y})$ from equation (3) can be written in this fashion, with

$$Z = \sqrt{\frac{n_1 n_2}{n_1+n_2}} \frac{(\overline{X} - \overline{Y})}{\sigma} \qquad V = \frac{\sum_{i=1}^{n_1}(X_i - \overline{X})^2 + \sum_{i=1}^{n_2}(Y_i - \overline{Y})^2}{\sigma} \qquad \mu = \sqrt{\frac{n_1 n_2}{n_1+n_2}} \frac{(\theta_1 - \theta_2)}{\sigma} = \sqrt{\frac{n_1 n_2}{n_1+n_2}} \frac{\delta}{\sigma}$$

The term δ/σ is the **standardized difference** between θ_1 and θ_2, and the form of μ indicates that we can look at power on this standardized scale for different sample sizes. In R, the functions pt and qt compute, respectively, the cdf and inverse cdf for both the Student-t and non-central Student-t distributions; for the probabilities required to compute $\beta(\theta, \sigma)$ the R commands are

```
n<-n1+n2
alpha<-0.05
sigma<-1
delta<-seq(-2,2,by=0.01)
cval<-qt(1-alpha/2,n-2)
mu<-sqrt((n1*n2/(n1+n2)))*(delta/sigma)
beta.power<-pt(-cval,df=n-2,ncp=mu)+1-pt(cval,df=n-2,ncp=mu)
```

The plot below depicts $\beta(\delta/\sigma)$ for $\alpha = 0.05$; note that the power is **higher** as $n = n_1 + n_2$ increases, but that the power for $n = 20$ is also **higher** if $n_1 = n_2 = 10$ than if $n_1 = 5$ and $n_2 = 15$.

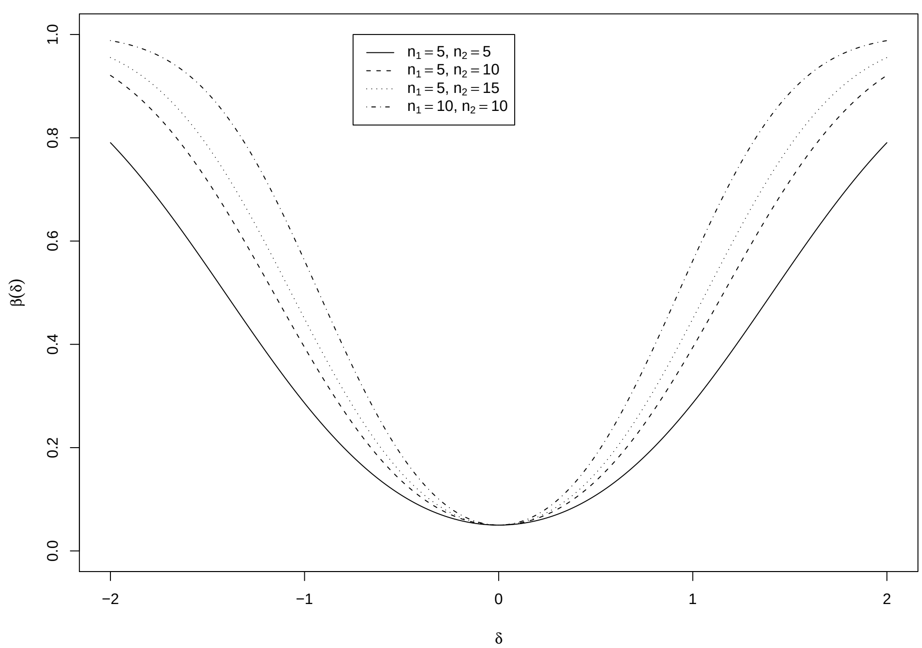

Example 6 Randomized Tests

A test $\mathcal{T}$ with test function $\phi_{\mathcal{R}}(T(\underset{\sim}{x}))$ taking values in $\{0, 1\}$ (with probability one) is termed a *non-randomized* test; given the observed value of statistic $T(\underset{\sim}{x})$, the null hypothesis is (deterministically) **rejected** if $\phi_{\mathcal{R}}(T(\underset{\sim}{x})) = 1$, and is not rejected otherwise. For such a test

$$E_{f_{T|\theta}}[\phi_{\mathcal{R}}(T(\underset{\sim}{x}))|\theta] = \Pr[\phi_{\mathcal{R}}(T(\underset{\sim}{x})) = 1|\theta] = \Pr[T(\underset{\sim}{X}) \in \mathcal{R}|\theta] = \beta(\theta).$$

In the Neyman-Pearson Lemma, for testing parametric models $f_{X|\theta}$ and two possible values θ_0 and θ_1, at level α, the critical region $\mathcal{R}$ is defined by

$$f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) > k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0) \qquad \Longrightarrow \qquad \underset{\sim}{x} \in \mathcal{R}$$

$$f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) < k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0) \qquad \Longrightarrow \qquad \underset{\sim}{x} \in \mathcal{R}'$$

where k is defined by noting the requirement $\Pr[\underset{\sim}{X} \in \mathcal{R}|\theta_0] = \alpha$. However, it may occur that

$$f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) = k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0)$$

in which case the result of the test is ambiguous. A potential resolution of the ambiguity is to construct a *randomized* test, $\mathcal{T}^\star$, where the decision to reject H_0 is potentially **randomly** chosen, but that matches the power of $\mathcal{T}$. Consider the test function $\phi^\star_{\mathcal{R}}(\underset{\sim}{x})$ defined by

$$\phi^\star_{\mathcal{R}}(\underset{\sim}{x}) = \begin{cases} 1 & f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) > k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0) \\ \gamma & f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) = k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0) \\ 0 & f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) < k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0) \end{cases}$$

for $0 \leq \gamma \leq 1$, so that, with a non-zero probability, $\phi^\star_{\mathcal{R}}(\underset{\sim}{x})$ takes a value not equal to zero or one. In this randomized test, the constant γ represents the probability with which H_0 is rejected in the case that

$$f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) = k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0).$$

Note that the requirement $\Pr[\underset{\sim}{X} \in \mathcal{R}|\theta_0] = \alpha$ implies that we much choose γ so that

$$E_{f_{T|\theta}}[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X}))|\theta_0] = \Pr[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X})) = 1|\theta_0] + \gamma \Pr[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X})) = \gamma|\theta_0]$$

The final term needs some explanation; it is equal to the probability of the set

$$A \equiv \{\underset{\sim}{x} : f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1) = k f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0)\}$$

under the model that assumes $\theta = \theta_0$.

For example, suppose that $X_1, \ldots, X_n \sim$ Bernoulli(θ) and consider a test of the simple hypotheses with values $\theta_0 < \theta_1$. Let $T(\underset{\sim}{X})$ be defined by $T(\underset{\sim}{X}) = \sum\limits_{i=1}^{n} X_i$. If

$$\lambda_T(\underset{\sim}{x}) = \frac{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1)}{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0)} = \left(\frac{\theta_1}{\theta_0}\right)^{T(\underset{\sim}{x})} \left(\frac{1-\theta_1}{1-\theta_0}\right)^{n-T(\underset{\sim}{x})}$$

then $\lambda_T(\underset{\sim}{x})$ is an increasing function of $T(\underset{\sim}{x})$. Therefore, there exist constants c and γ such that a test $\mathcal{T}^\star$ can be constructed with test function

$$\phi^\star_{\mathcal{R}}(\underset{\sim}{x}) = \begin{cases} 1 & T(\underset{\sim}{x}) > c \\ \gamma & T(\underset{\sim}{x}) = c \\ 0 & T(\underset{\sim}{x}) \leq c \end{cases}$$

such that

$$\begin{aligned} \alpha = E_{f_{T|\theta}}[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X}))|\theta_0] &= \Pr[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X})) = 1|\theta_0] + \gamma \Pr[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X})) = \gamma|\theta_0] \\ &= \Pr[T(\underset{\sim}{X}) > c|\theta_0] + \gamma \Pr[T(\underset{\sim}{X}) = c|\theta_0] \\ &= \sum_{j=c+1}^{n} \binom{n}{j} \theta_0^j (1-\theta_0)^j + \gamma \binom{n}{c} \theta_0^c (1-\theta_0)^c. \end{aligned}$$

The introduction of the random element allows this equation to be solved exactly, whatever the value of α; this was not possible under the non-randomized rule.

For a specific numerical example, let $n = 20$, $\theta_0 = 0.3$ and $\theta_1 = 0.5$. For $\alpha = 0.05$, the probability distribution of $T(\underset{\sim}{X})$ is Binomial(n, θ), so that the probability $\Pr[T(\underset{\sim}{x}) > c|\theta = 0.3]$ can be computed:

c	5	6	7	8	9	10	11	12	13
$\Pr[T(\underset{\sim}{x}) = c \mid \theta = 0.3]$	0.179	0.192	0.164	0.114	0.065	0.031	0.012	0.004	0.001
$\Pr[T(\underset{\sim}{x}) > c \mid \theta = 0.3]$	0.584	0.392	0.228	0.113	0.048	0.017	0.005	0.001	0.000

Hence choosing c equal to 8 or 9 gives $\Pr[T(\underset{\sim}{x}) > c|\theta = 0.3]$ equal to 0.113 and 0.048 respectively, so that $\alpha = 0.05$ cannot be matched exactly in a non-randomized test (that is, if $\gamma = 0$). However choosing $c = 9$ and $\gamma = 0.308$ in the randomized test yields

$$\Pr[T(\underset{\sim}{X}) > c|\theta_0] + \gamma \Pr[T(\underset{\sim}{X}) = c|\theta_0] = 0.048 + 0.308 \times 0.065 = 0.05 = \alpha$$

so the randomized test that specifies

$$\sum_{i=1}^{n} x_i > 9 \implies \text{Reject } H_0$$

$$\sum_{i=1}^{n} x_i = 9 \implies \text{Reject } H_0 \text{ with probability } \gamma = 0.308$$

$$\sum_{i=1}^{n} x_i < 9 \implies \text{Do Not Reject } H_0$$

has size/level precisely α. The power function is

$$\beta(\theta) = \sum_{j=c+1}^{n} \binom{n}{j} \theta^j (1-\theta)^j + \gamma \binom{n}{c} \theta^c (1-\theta)^c$$

Example 7 Suppose that $X_1, \ldots, X_n \sim$ Uniform$(0, \theta)$. To test

$$\begin{aligned} H_0 &: \theta \leq \theta_0 \\ H_1 &: \theta > \theta_0. \end{aligned}$$

The likelihood ratio for $\theta_1 < \theta_2$ for this model is

$$\lambda(\underset{\sim}{x}) = \frac{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_2)}{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_1)} = \begin{cases} \left(\dfrac{\theta_2}{\theta_1}\right)^n & T(\underset{\sim}{X}) \leq \theta_1 \\ \infty & \theta_1 \leq T(\underset{\sim}{X}) \leq \theta_2 \end{cases}$$

where $T(\underset{\sim}{X}) = X_{(n)} = \min\{X_1, \ldots, X_n\}$. Thus $\lambda(\underset{\sim}{x})$ is a **non decreasing** function of $T(\underset{\sim}{x})$ as for $\theta_1 < \theta_2$, and by the Karlin-Rubin theorem, the UMP test at level α is based on the critical region

$$\mathcal{R} \equiv \left\{ \underset{\sim}{x} : T(\underset{\sim}{x}) = x_{(n)} > t_0 \right\}.$$

To find t_0, we need to solve

$$\Pr[X_{(n)} > t_0 \mid \theta_0] = 1 - \left(\frac{t_0}{\theta_0}\right)^n = \alpha \qquad \therefore \qquad t_0 = \theta_0(1-\alpha)^{1/n}$$

with power function (for $\theta > \theta_0$)

$$\beta(\theta) = 1 - \left(\frac{\theta_0}{\theta}\right)^n (1-\alpha).$$

Now consider the **randomized** test $\mathcal{T}^\star$ with test function

$$\phi^\star_{\mathcal{R}}(\underset{\sim}{x}) = \begin{cases} 1 & x_{(n)} > \theta_0 \\ \alpha & x_{(n)} \leq \theta_0 \end{cases}$$

We have for $\theta > \theta_0$ that

$$\begin{aligned} \beta^\star(\theta) = E_{f_{T|\theta}}[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X}))|\theta] &= \Pr[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X})) = 1|\theta] + \alpha \Pr[\phi^\star_{\mathcal{R}}(T(\underset{\sim}{X})) = \alpha|\theta] \\ &= \Pr[X_{(n)} > \theta_0|\theta] + \alpha \Pr[X_{(n)} \leq \theta_0|\theta] \\ &= 1 - \left(\frac{\theta_0}{\theta}\right)^n + \alpha\left(\frac{\theta_0}{\theta}\right)^n \\ &= 1 - \left(\frac{\theta_0}{\theta}\right)^n (1-\alpha) \end{aligned}$$

thus matching the power of the UMP test described above. Therefore the UMP test is not unique.

Note that for the hypotheses

$$\begin{aligned} H_0 &: \theta = \theta_0 \\ H_1 &: \theta \neq \theta_0 \end{aligned}$$

the likelihood ratio test statistic is

$$\lambda(\underset{\sim}{x}) = \frac{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\theta_0)}{f_{\underset{\sim}{X}|\theta}(\underset{\sim}{x}|\widehat{\theta})} = \begin{cases} \left(\dfrac{x_{(n)}}{\theta_0}\right)^n & x_{(n)} \leq \theta_0 \\ 0 & x_{(n)} > \theta_0 \end{cases}$$

Therefore the likelihood ratio test $\lambda(\underset{\sim}{x}) \leq k$ is has rejection region

$$(X_{(n)} > \theta_0) \cup (X_{(n)}/\theta_0 \leq k^{1/n})$$

To choose k, we require that the size/level is α; as

$$\Pr[(X_{(n)} > \theta_0) \cup (X_{(n)}/\theta_0 \leq k^{1/n})|\theta = \theta_0] = \Pr[X_{(n)} \leq k^{1/n}\theta_0|\theta = \theta_0] = \frac{k\theta_0^n}{\theta_0^n} = k$$

we choose $k = \alpha$. The power function $\beta(\theta)$ is

$$\Pr[(X_{(n)} > \theta_0) \cup (X_{(n)}/\theta_0 < \alpha^{1/n})|\theta] = \begin{cases} \alpha\left(\dfrac{\theta_0}{\theta}\right)^n & 0 < \theta < \theta_0 \\ 1 - (1-\alpha)\left(\dfrac{\theta_0}{\theta}\right)^n & \theta > \theta_0 \end{cases}$$