Lecture 16: UMPU tests in exponential families

Neyman structure

Let $U(X)$ be a sufficient statistic for $P \in \bar{\mathscr{P}}$ and let $\bar{\mathscr{P}}_U$ be the family of distributions of U as P ranges over $\bar{\mathscr{P}}$.
A test is said to have *Neyman structure* w.r.t. U if

$$E[T(X)|U] = \alpha \qquad \text{a.s. } \bar{\mathscr{P}}_U,$$

Clearly, if T has Neyman structure, then

$$E[T(X)] = E\{E[T(X)|U]\} = \alpha \qquad P \in \bar{\mathscr{P}},$$

i.e., T is similar on $\bar{\Theta}_{01}$.
If all tests similar on $\bar{\Theta}_{01}$ have Neyman structure w.r.t. U, then working with tests having Neyman structure is the same as working with tests similar on $\bar{\Theta}_{01}$.

Lemma 6.6

Let $U(X)$ be a sufficient statistic for $P \in \bar{\mathscr{P}}$.
A necessary and sufficient condition for all tests similar on $\bar{\Theta}_{01}$ to have Neyman structure w.r.t. U is that U is boundedly complete for $P \in \bar{\mathscr{P}}$.

Proof

(i) Suppose first that U is boundedly complete for $P \in \bar{\mathscr{P}}$.
Let $T(X)$ be a test similar on $\bar{\Theta}_{01}$.
Then $E[T(X) - \alpha] = 0$ for all $P \in \bar{\mathscr{P}}$.
From the boundedness of $T(X)$, $E[T(X)|U]$ is bounded.
Since $E\{E[T(X)|U] - \alpha\} = E[T(X) - \alpha] = 0$ for all $P \in \bar{\mathscr{P}}$ and U is boundedly complete, $E[T(X)|U] = \alpha$ a.s. $\bar{\mathscr{P}}_U$, i.e., T has Neyman structure.
(ii) Suppose now that all tests similar on $\bar{\Theta}_{01}$ have Neyman structure w.r.t. U.
Suppose also that U is not boundedly complete for $P \in \bar{\mathscr{P}}$.
Then there is a function h such that $|h(u)| \leq C$, $E[h(U)] = 0$ for all $P \in \bar{\mathscr{P}}$, and $h(U) \neq 0$ with positive probability for some $P \in \bar{\mathscr{P}}$.
Let $T(X) = \alpha + ch(U)$, where $c = \min\{\alpha, 1-\alpha\}/C$.
Then T is a test similar on $\bar{\Theta}_{01}$ but T does not have Neyman structure w.r.t. U (because $h(U) \neq 0$).
Thus, U must be boundedly complete for $P \in \bar{\mathscr{P}}$.
This proves the result.

Theorem 6.4 (UMPU tests in multiparameter exponential families)

Suppose that X has the following p.d.f. w.r.t. a σ-finite measure:

$$f_{\theta,\varphi}(x) = \exp\{\theta Y(x) + \varphi^\tau U(x) - \zeta(\theta,\varphi)\},$$

where θ is a real-valued parameter, φ is a vector-valued parameter, and Y (real-valued) and U (vector-valued) are statistics.

(i) For testing $H_0 : \theta \le \theta_0$ versus $H_1 : \theta > \theta_0$, a UMPU test of size α is

$$T_*(Y,U) = \begin{cases} 1 & Y > c(U) \\ \gamma(U) & Y = c(U) \\ 0 & Y < c(U), \end{cases}$$

where $c(u)$ and $\gamma(u)$ are Borel functions determined by

$$E_{\theta_0}[T_*(Y,U)|U=u] = \alpha \quad \text{for every } u$$

and E_{θ_0} is the expectation w.r.t. $f_{\theta_0,\varphi}$.

(ii) For testing $H_0 : \theta \le \theta_1$ or $\theta \ge \theta_2$ versus $H_1 : \theta_1 < \theta < \theta_2$, a UMPU test of size α is

$$T_*(Y,U) = \begin{cases} 1 & c_1(U) < Y < c_2(U) \\ \gamma_i(U) & Y = c_i(U),\ i = 1,2, \\ 0 & Y < c_1(U) \text{ or } Y > c_2(U), \end{cases}$$

Theorem 6.4 (continued)

where $c_i(u)$'s and $\gamma_i(u)$'s are Borel functions determined by

$$E_{\theta_1}[T_*(Y,U)|U=u] = E_{\theta_2}[T_*(Y,U)|U=u] = \alpha \text{ for every } u.$$

(iii) For testing $H_0 : \theta_1 \leq \theta \leq \theta_2$ versus $H_1 : \theta < \theta_1$ or $\theta > \theta_2$, a UMPU test of size α is

$$T_*(Y,U) = \begin{cases} 1 & Y < c_1(U) \text{ or } Y > c_2(U) \\ \gamma_i(U) & Y = c_i(U),\ i = 1,2, \\ 0 & c_1(U) < Y < c_2(U), \end{cases}$$

where $c_i(u)$'s and $\gamma_i(u)$'s are Borel functions determined by

$$E_{\theta_1}[T_*(Y,U)|U=u] = E_{\theta_2}[T_*(Y,U)|U=u] = \alpha \text{ for every } u.$$

(iv) For testing $H_0 : \theta = \theta_0$ versus $H_1 : \theta \neq \theta_0$, a UMPU test of size α is given by $T_*(Y,U)$ in (iii), where $c_i(u)$'s and $\gamma_i(u)$'s are Borel functions determined by

$$E_{\theta_0}[T_*(Y,U)|U=u] = \alpha \text{ for every } u$$

and

$$E_{\theta_0}[T_*(Y,U)Y|U=u] = \alpha E_{\theta_0}(Y|U=u) \text{ for every } u.$$

Proof

By sufficiency, we only need to consider tests that are functions of (Y,U).

It follows from Theorem 2.1(i) that the p.d.f. of (Y,U) (w.r.t. a σ-finite measure) is in a natural exponential family of the form $\exp\{\theta y + \varphi^\tau u - \zeta(\theta,\varphi)\}$ and, given $U = u$, the p.d.f. of the conditional distribution of Y (w.r.t. a σ-finite measure ν_u) is in a natural exponential family of the form $\exp\{\theta y - \zeta_u(\theta)\}$.

Hypotheses in (i)-(iv) are of the form $H_0 : \theta \in \Theta_0$ vs $H_1 : \theta \in \Theta_1$ with $\bar{\Theta}_{01} = \{(\theta,\varphi) : \theta = \theta_0\}$ or = $\{(\theta,\varphi) : \theta = \theta_i, i = 1,2\}$.

In case (i) or (iv), U is sufficient and complete for $P \in \bar{\mathscr{P}}$ and, hence, Lemma 6.6 applies.

In case (ii) or (iii), applying Lemma 6.6 to each $\{(\theta,\varphi) : \theta = \theta_i\}$ also shows that working with tests having Neyman structure is the same as working with tests similar on $\bar{\Theta}_{01}$.

By Theorem 2.1, the power functions of all tests are continuous and, hence, Lemma 6.5 applies.

Thus, for (i), it suffices to show T_* is UMP among all tests T satisfying

$$E_{\theta_0}[T(Y,U)|U=u] = \alpha \text{ for every } u \quad (1)$$

and for part (ii) or (iii)), it suffices show T_* is UMP among all tests T satisfying

$$E_{\theta_1}[T(Y,U)|U=u] = E_{\theta_2}[T(Y,U)|U=u] = \alpha \text{ for every } u.$$

For (iv), any unbiased T should satisfy (1) and

$$\frac{\partial}{\partial\theta}E_{\theta,\varphi}[T(Y,U)] = 0, \qquad \theta \in \bar{\Theta}_{01}. \quad (2)$$

One can show (exercise) that (2) is equivalent to

$$E_{\theta,\varphi}[T(Y,U)Y - \alpha Y] = 0, \qquad \theta \in \bar{\Theta}_{01}. \quad (3)$$

Using the argument in the proof of Lemma 6.6, one can show (exercise) that (3) is equivalent to

$$E_{\theta_0}[T(Y,U)Y|U=u] = \alpha E_{\theta_0}(Y|U=u) \text{ for every } u. \quad (4)$$

Hence, for (iv), it suffices to show T_* is UMP among all tests T satisfying (1) and (4).

Note that the power function of any test $T(Y,U)$ is

$$\beta_T(\theta,\varphi) = \int \left[\int T(y,u) dP_{Y|U=u}(y) \right] dP_U(u).$$

Thus, it suffices to show that for every fixed u and $\theta \in \Theta_1$, T_* maximizes

$$\int T(y,u) dP_{Y|U=u}(y)$$

over all T subject to the given side conditions.

Since $P_{Y|U=u}$ is in a one-parameter exponential family, the results in (i) and (ii) follow from Corollary 6.1 and Theorem 6.3, respectively.

The result in (iii) follows from Theorem 6.3(ii) by considering $1 - T_*$.

To prove the result in (iv), it suffices to show that if Y has the p.d.f. given by $\exp\{\theta y - \zeta_u(\theta)\}$ and if u is treated as a constant in (1) and (4), T_* in (iii) with a fixed u is UMP subject to conditions (1) and (4).

We now omit u in the following proof for (iv), which is very similar to the proof of Theorem 6.3.

First, $(\alpha, \alpha E_{\theta_0}(Y))$ is an interior point of the set of points $(E_{\theta_0}[T(Y)], E_{\theta_0}[T(Y)Y])$ as T ranges over all tests of the form $T(Y)$.

By Lemma 6.2 and Proposition 6.1, for testing $\theta = \theta_0$ versus $\theta = \theta_1$, the UMP test is equal to 1 when

$$(k_1 + k_2 y)e^{\theta_0 y} < C(\theta_0, \theta_1)e^{\theta_1 y},$$

where k_i's and $C(\theta_0, \theta_1)$ are constants.

This inequality is equivalent to

$$a_1 + a_2 y < e^{by}$$

for some constants a_1, a_2, and b.

This region is either one-sided or the outside of an interval.

By Theorem 6.2(ii), a one-sided test has a strictly monotone power function and therefore cannot satisfy (4).

Thus, this test must have the form of T_* in (iii).

Since T_* in (iii) does not depend on θ_1, by Lemma 6.1, it is UMP over all tests satisfying (1) and (4); in particular, the test $\equiv \alpha$.

Thus, T_* is UMPU.

Finally, it can be shown that all the c- and γ-functions in (i)-(iv) are Borel functions of u (see Lehmann (1986, p. 149)).

Example 6.11

A problem arising in many different contexts is the comparison of two treatments.

If the observations are integer-valued, the problem often reduces to testing the equality of two Poisson distributions (e.g., a comparison of the radioactivity of two substances or the car accident rate in two cities) or two binomial distributions (when the observation is the number of successes in a sequence of trials for each treatment).

Consider first the Poisson problem in which X_1 and X_2 are independently distributed as the Poisson distributions $P(\lambda_1)$ and $P(\lambda_2)$.

The p.d.f. of $X = (X_1, X_2)$ is

$$[e^{-(\lambda_1+\lambda_2)}/x_1!x_2!]\exp\{x_2\log(\lambda_2/\lambda_1)+(x_1+x_2)\log\lambda_1\}$$

w.r.t. the counting measure on $\{(i,j) : i = 0,1,2,...,j = 0,1,2,...\}$.

The hypotheses such as $\lambda_1 = \lambda_2$ and $\lambda_1 \geq \lambda_2$ are equivalent to $\theta = 0$ and $\theta \leq 0$, respectively, where $\theta = \log(\lambda_2/\lambda_1)$.

The p.d.f. of X is in a multiparameter exponential family with $\varphi = \log\lambda_1$, $Y = X_2$, and $U = X_1 + X_2$.

Thus, Theorem 6.4 applies.

To obtain various tests in Theorem 6.4, it is enough to derive the conditional distribution of $Y = X_2$ given $U = X_1 + X_2 = u$.

Using the fact that $X_1 + X_2$ has the Poisson distribution $P(\lambda_1 + \lambda_2)$, one can show that

$$P(Y = y|U = u) = \binom{u}{y} p^y (1-p)^{u-y} I_{\{0,1,\ldots,u\}}(y), \quad u = 0, 1, 2, \ldots,$$

where $p = \lambda_2/(\lambda_1 + \lambda_2) = e^\theta/(1 + e^\theta)$.

This is the binomial distribution $Bi(p, u)$.

On the boundary set $\bar{\Theta}_{01}$, $\theta = \theta_j$ (a known value) and the distribution $P_{Y|U=u}$ is known.

Consider next the binomial problem in which X_j, $j = 1, 2$, are independently distributed as the binomial distributions $Bi(p_j, n_j)$, $j = 1, 2$, respectively, where n_j's are known but p_j's are unknown.

The p.d.f. of $X = (X_1, X_2)$ is

$$\binom{n_1}{x_1}\binom{n_2}{x_2}(1-p_1)^{n_1}(1-p_2)^{n_2}\exp\left\{x_2 \log\frac{p_2(1-p_1)}{p_1(1-p_2)} + (x_1 + x_2)\log\frac{p_1}{(1-p_1)}\right\}$$

w.r.t. the counting measure on $\{(i,j) : i = 0,1,...,n_1, j = 0,1,...,n_2\}$.

This p.d.f. is in a multiparameter exponential family with $\theta = \log \frac{p_2(1-p_1)}{p_1(1-p_2)}$, $Y = X_2$, and $U = X_1 + X_2$.

Thus, Theorem 6.4 applies.

Note that hypotheses such as $p_1 = p_2$ and $p_1 \geq p_2$ are equivalent to $\theta = 0$ and $\theta \leq 0$, respectively.

Using the joint distribution of (X_1, X_2), one can show (exercise) that

$$P(Y = y|U = u) = K_u(\theta)\binom{n_1}{u-y}\binom{n_2}{y}e^{\theta y}I_A(y), \quad u = 0,1,...,n_1+n_2,$$

where

$$A = \{y : y = 0,1,...,\min\{u,n_2\}, u-y \leq n_1\}$$

and

$$K_u(\theta) = \left[\sum_{y\in A}\binom{n_1}{u-y}\binom{n_2}{y}e^{\theta y}\right]^{-1}.$$

If $\theta = 0$, this distribution reduces to a known distribution: the hypergeometric distribution $HG(u, n_2, n_1)$ (Table 1.1, page 18).

The following lemma is useful especially when X is from a population in an exponential family with continuous p.d.f.'s.

Lemma 6.7

Suppose that X has the following p.d.f. w.r.t. a σ-finite measure:

$$f_{\theta,\varphi}(x) = \exp\{\theta Y(x) + \varphi^\tau U(x) - \zeta(\theta,\varphi)\},$$

where θ is a real-valued parameter, φ is a vector-valued parameter, and Y (real-valued) and U (vector-valued) are statistics.
Let $V(Y,U)$ be a statistic independent of U when $\theta = \theta_j$, where θ_j's are known values given in the hypotheses in (i)-(iv) of Theorem 6.4.
(i) If $V(y,u)$ is increasing in y for each u, then the UMPU tests in (i)-(iii) of Theorem 6.4 are equivalent to those with Y and (Y,U) replaced by V and with $c_i(U)$ and $\gamma_i(U)$ replaced by constants c_i and γ_i, respectively.
(ii) If there are Borel functions $a(u) > 0$ and $b(u)$ such that $V(y,u) = a(u)y + b(u)$, then the UMPU test in Theorem 6.4(iv) is equivalent to that with Y and (Y,U) replaced by V and with $c_i(U)$ and $\gamma_i(U)$ replaced by constants c_i and γ_i, respectively.

Proof

(i) Since V is increasing in y, $Y > c_i(u)$ is equivalent to $V > d_i(u)$ for some d_i.
The result follows from the fact that V is independent of U so that d_i's and γ_i's do not depend on u when Y is replaced by V.
(ii) Since $V = a(U)Y + b(U)$, the UMPU test in Theorem 6.4(iv) is the same as

$$T_*(V,U) = \begin{cases} 1 & V < c_1(U) \text{ or } V > c_2(U) \\ \gamma_i(U) & V = c_i(U),\ i = 1,2, \\ 0 & c_1(U) < V < c_2(U), \end{cases}$$

subject to $E_{\theta_0}[T_*(V,U)|U = u] = \alpha$ and

$$E_{\theta_0}\left[T_*(V,U)\frac{V - b(U)}{a(U)}\middle| U\right] = \alpha E_{\theta_0}\left[\frac{V - b(U)}{a(U)}\middle| U\right]. \qquad (5)$$

Under $E_{\theta_0}[T_*(V,U)|U = u] = \alpha$, (5) is the same as
$E_{\theta_0}[T_*(V,U)V|U] = \alpha E_{\theta_0}(V|U)$.
Since V and U are independent when $\theta = \theta_0$, $c_i(u)$'s and $\gamma_i(u)$'s do not depend on u and, therefore, T_* does not depend on U.

- If the conditions of Lemma 6.7 are satisfied, then UMPU tests can be derived by working with the distribution of V instead of $P_{Y|U=u}$.
- In exponential families, a $V(Y,U)$ independent of U can often be found by applying Basu's theorem (Theorem 2.4).
- An important application of Theorem 6.4 and Lemma 6.7 is the derivation of UMPU tests in normal families.
- The results presented here are the basic justifications for tests in elementary textbooks concerning parameters in normal families.
- When we consider normal families, γ_i's can be chosen to be 0 since the c.d.f. of Y given $U = u$ or the c.d.f. of V is continuous.

One-sample problems

Let $X_1,...,X_n$ be i.i.d. from $N(\mu,\sigma^2)$ with unknown $\mu \in \mathscr{R}$ and $\sigma^2 > 0$, where $n \geq 2$.
The joint p.d.f. of $X = (X_1,...,X_n)$ is

$$\frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left\{ -\frac{1}{2\sigma^2}\sum_{i=1}^n x_i^2 + \frac{\mu}{\sigma^2}\sum_{i=1}^n x_i - \frac{n\mu^2}{2\sigma^2} \right\}.$$

Tests concerning σ^2

Consider first hypotheses concerning σ^2.
The p.d.f. of X is in a multiparameter exponential family with $\theta = -(2\sigma^2)^{-1}$, $\varphi = n\mu/\sigma^2$, $Y = \sum_{i=1}^n X_i^2$, and $U = \bar{X}$.
By Basu's theorem, $V = (n-1)S^2$ is independent of $U = \bar{X}$ (Example 2.18), where S^2 is the sample variance.
Also,

$$\sum_{i=1}^n X_i^2 = (n-1)S^2 + n\bar{X}^2,$$

i.e., $V = Y - nU^2$.
Hence the conditions of Lemma 6.7 are satisfied.
Since V/σ^2 has the chi-square distribution χ^2_{n-1} (Example 2.18), values of c_i's for hypotheses in (i)-(iii) of Theorem 6.4 are related to quantiles of χ^2_{n-1}.
For testing $H_0 : \theta = \theta_0$ vs $H_1 : \theta \neq \theta_0$ (which is equivalent to testing $H_0 : \sigma^2 = \sigma_0^2$ vs $H_1 : \sigma^2 \neq \sigma_0^2$), $d_i = c_i/\sigma_0^2$, $i = 1, 2$, are determined by

$$\int_{d_1}^{d_2} f_{n-1}(v)dv = 1 - \alpha \qquad \text{and} \qquad \int_{d_1}^{d_2} v f_{n-1}(v)dv = (n-1)(1-\alpha),$$

where f_m is the Lebesgue p.d.f. of the chi-square distribution χ^2_m.
Since $vf_{n-1}(v) = (n-1)f_{n+1}(v)$, d_1 and d_2 are determined by

$$\int_{d_1}^{d_2} f_{n-1}(v)dv = \int_{d_1}^{d_2} f_{n+1}(v)dv = 1-\alpha.$$

If $n-1 \approx n+1$, then d_1 and d_2 are nearly the $(\alpha/2)$th and $(1-\alpha/2)$th quantiles of χ^2_{n-1}, respectively, in which case the UMPU test in Theorem 6.4(iv) is the same as the "equal-tailed" chi-square test for H_0 in elementary textbooks.

Tests concerning μ

Consider next hypotheses concerning μ.
The p.d.f. of X has is in a multiparameter exponential family with $Y = \bar{X}$, $U = \sum_{i=1}^n (X_i - \mu_0)^2$, $\theta = n(\mu - \mu_0)/\sigma^2$, and $\varphi = -(2\sigma^2)^{-1}$.
For testing hypotheses $H_0 : \mu \leq \mu_0$ versus $H_1 : \mu > \mu_0$, we take V to be $t(X) = \sqrt{n}(\bar{X} - \mu_0)/S$.
By Basu's theorem, $t(X)$ is independent of U when $\mu = \mu_0$.
Hence it satisfies the conditions in Lemma 6.7(i).

From Examples 1.16 and 2.18, $t(X)$ has the t-distribution t_{n-1} when $\mu = \mu_0$.
Thus, $c(U)$ in Theorem 6.4(i) is the $(1-\alpha)$th quantile of t_{n-1}.
For the two-sided hypotheses $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$, the statistic $V = (\bar{X} - \mu_0)/\sqrt{U}$ satisfies the conditions in Lemma 6.7(ii) and has a distribution symmetric about 0 when $\mu = \mu_0$.
Then the UMPU test in Theorem 6.4(iv) rejects H_0 when $|V| > d$, where d satisfies $P(|V| > d) = \alpha$ when $\mu = \mu_0$.
Since

$$t(X) = \sqrt{(n-1)n}V(X)/\sqrt{1 - n[V(X)]^2},$$

the UMPU test rejects H_0 if and only if $|t(X)| > t_{n-1,\alpha/2}$, where $t_{n-1,\alpha}$ is the $(1-\alpha)$th quantile of the t-distribution t_{n-1}.
The UMPU tests derived here are the so-called one-sample t-tests in elementary textbooks.
The power function of a one-sample t-test is related to the noncentral t-distribution introduced in §1.3.1 (see Exercise 36).