STAT 801 Lecture 21

Reading for Today's Lecture:
Goals of Today's Lecture:

- Prove the NP lemma.
- Use the NP lemma to find UMP tests.
- Introduce unbiasedness; extend NP lemma to find UMPU tests.

Today's notes

The Neyman Pearson Lemma

Theorem: In testing f_0 against f_1 the probability β of a type II error is minimized, subject to $\alpha \le \alpha_0$ by the test function:

$$\phi(x) = \begin{cases} 1 & \frac{f_1(x)}{f_0(x)} > \lambda \\ \gamma & \frac{f_1(x)}{f_0(x)} = \lambda \\ 0 & \frac{f_1(x)}{f_0(x)} < \lambda \end{cases}$$

where λ is chosen so that for each $y > \lambda$

$$P_0\left(\frac{f_1(X)}{f_0(X)} \ge y\right) < \alpha_0$$

and

$$P_0\left(\frac{f_1(x)}{f_0(x)} \ge \lambda\right) \ge \alpha_0$$

and where γ is any number chosen so that

$$E_0(\phi(X)) = P_0\left(\frac{f_1(X)}{f_0(X)} > \lambda\right) + \gamma P_0\left(\frac{f_1(X)}{f_0(X)} = \lambda\right) = \alpha_0$$

The value of γ is unique if $P_0(\frac{f_1(X)}{f_0(X)} = \lambda) > 0$.
Definition: In the general problem of testing Θ_0 against Θ_1 the level of a test function ϕ is

$$\alpha = \sup_{\theta \in \Theta_0} E_\theta(\phi(X))$$

The power function is

$$\pi(\theta) = E_\theta(\phi(X))$$

A test ϕ^* is a Uniformly Most Powerful level α_0 test if

1. ϕ^* has level $\alpha \le \alpha_o$

2. If ϕ has level $\alpha \le \alpha_0$ then for every $\theta \in \Theta_1$ we have

$$E_\theta(\phi(X)) \le E_\theta(\phi^*(X))$$

Example: In the Binomial(n,p) to test $p = p_0$ versus p_1 for a $p_1 > p_0$ the NP test is of the form

$$\phi(x) = 1(X > k) + \gamma 1(X = k)$$

where we choose k so that

$$P_{p_0}(X > k) \le \alpha_0 < P_{p_0}(X \ge k)$$

and $\gamma \in [0,1)$ so that

$$\alpha_0 = P_{p_0}(X > k) + \gamma P_{p_0}(X = k)$$

This rejection region depends only on p_0 and not on p_1 so that this test is UMP for $p = p_0$ against $p > p_0$. Since this test has level α_0 even for the larger null hypothesis it is also UMP for $p \le p_0$ against $p > p_0$.
Example: In the $N(\mu,1)$ example to test $\mu \le \mu_0$ against $\mu = \mu_1$ for some $\mu_1 > \mu_0$ the NP test is of the form

$$\phi(x) = 1(n^{1/2}(\bar{x} - \mu_0) > z_\alpha)\,.$$

Since $\bar{X}$ has a continuous distribution there is no need to think about γ. Since the rejection region does not depend on μ_1 this test is UMP for $\mu = \mu_0$ against $\mu > \mu_0$. Since it has level α (you can check that the power function is a strictly increasing function of μ) even for the larger null $\mu \le \mu_0$ it is also UMP level α for $\mu \le \mu_0$ against $\mu > \mu_0$.
Proof of the Neyman Pearson lemma: Given a test ϕ with level strictly less than α_0 we can define the test

$$\phi^*(x) = \frac{1-\alpha_0}{1-\alpha}\phi(x) + \frac{\alpha_0 - \alpha}{1-\alpha}$$

has level α_0 and β smaller than that of ϕ. Hence we may assume without loss that $\alpha = \alpha_0$ and minimize β subject to $\alpha = \alpha_0$. However, the argument which follows doesn't actually need this.

For each $\lambda > 0$ we have seen that ϕ_λ minimizes $\lambda\alpha + \beta$ where $\phi_\lambda = 1(f_1(x)/f_0(x) \ge \lambda)$.

As λ increases the level of ϕ_λ decreases from 1 when $\lambda = 0$ to 0 when $\lambda = \infty$. There is thus a value λ_0 where for $\lambda < \lambda_0$ the level is less than α_0 while for $\lambda > \lambda_0$ the level is at least α_0. Temporarily let $\delta = P_0(f_1(X)/f_0(X) = \lambda_0)$. If $\delta = 0$ define $\phi = \phi_\lambda$. If $\delta > 0$ define

$$\phi(x) = \begin{cases} 1 & \frac{f_1(x)}{f_0(x)} > \lambda_0 \\ \gamma & \frac{f_1(x)}{f_0(x)} = \lambda_0 \\ 0 & \frac{f_1(x)}{f_0(x)} < \lambda_0 \end{cases}$$

where $P_0(f_1(X)/f_0(X) < \lambda_0) + \gamma\delta = \alpha_0$. You can check that $\gamma \in [0, 1]$.

Now ϕ has level α_0 and according to the theorem above minimizes $\alpha + \lambda_0\beta$. Suppose ϕ^* is some other test with level $\alpha^* \le \alpha_0$. Then

$$\lambda_0\alpha_\phi + \beta_\phi \le \lambda_0\alpha_{\phi^*} + \beta_{\phi^*}$$

We can rearrange this as

$$\beta_{\phi^*} \ge \beta_\phi + (\alpha_\phi - \alpha_{\phi^*})\lambda_0$$

Since

$$\alpha_{\phi^*} \le \alpha_0 = \alpha_\phi$$

the second term is non-negative and

$$\beta_{\phi^*} \ge \beta_\phi$$

which proves the Neyman Pearson Lemma.

The examples show a phenomenon which is somewhat general. What happened was this. For any $\mu > \mu_0$ the likelihood ratio f_μ/f_0 is an increasing function of $\sum X_i$. The rejection region of the NP test is thus always a region of the form $\sum X_i > k$. The value of the constant k is determined by the requirement that the test have level α_0 and this depends on μ_0 not on μ_1.

Definition: The family $f_\theta; \theta \in \Theta \subset R$ has monotone likelikelood ratio with respect to a statistic $T(X)$ if for each $\theta_1 > \theta_0$ the likelihood ratio $f_{\theta_1}(X)/f_{\theta_0}(X)$ is a monotone increasing function of $T(X)$.

Theorem: For a monotone likelihood ratio family the Uniformly Most Powerful level α test of $\theta \le \theta_0$ (or of $\theta = \theta_0$) against the alternative $\theta > \theta_0$

is

$$\phi(x) = \begin{cases} 1 & T(x) > t_\alpha \\ \gamma & T(X) = t_\alpha \\ 0 & T(x) < t_\alpha \end{cases}$$

where $P_0(T(X) > t_\alpha) + \gamma P_0(T(X) = t_\alpha) = \alpha_0$.

A typical family where this will work is a one parameter exponential family. In almost any other problem the method doesn't work and there is no uniformly most powerful test. For instance to test $\mu = \mu_0$ against the two sided alternative $\mu \neq \mu_0$ there is no UMP level α test. If there were its power at $\mu > \mu_0$ would have to be as high as that of the one sided level α test and so its rejection region would have to be the same as that test, rejecting for large positive values of $\bar{X} - \mu_0$. But it also has to have power as good as the one sided test for the alternative $\mu < \mu_0$ and so would have to reject for large negative values of $\bar{X} - \mu_0$. This would make its level too large.

The favourite test is the usual 2 sided test which rejects for large values of $|\bar{X} - \mu_0|$ with the critical value chosen appropriately. This test maximizes the power subject to two constraints: first, that the level be α and second that the test have power which is minimized at $\mu = \mu_0$. This second condition is really that the power on the alternative be larger than it is on the null.

Definition: A test ϕ of Θ_0 against Θ_1 is unbiased level α if it has level α and, for every $\theta \in \Theta_1$ we have

$$\pi(\theta) \geq \alpha\,.$$

When testing a point null hypothesis like $\mu = \mu_0$ this requires that the power function be minimized at μ_0 which will mean that if π is differentiable then

$$\pi'(\mu_0) = 0$$

We now apply that condition to the $N(\mu, 1)$ problem. If ϕ is any test function then

$$\pi'(\mu) = \frac{\partial}{\partial \mu} \int \phi(x) f(x, \mu) dx$$

We can differentiate this under the integral and use

$$\frac{\partial f(x, \mu)}{\partial \mu} = \sum (x_i - \mu) f(x, \mu)$$

to get the condition

$$\int \phi(x) \bar{x} f(x, \mu_0) dx = \mu_0 \alpha_0$$

Consider the problem of minimizing $\beta(\mu)$ subject to the two constraints $E_{\mu_0}(\phi(X)) = \alpha_0$ and $E_{\mu_0}(\bar{X}\phi(X)) = \mu_0\alpha_0$. Now fix two values $\lambda_1 > 0$ and λ_2 and minimize

$$\lambda_1\alpha + \lambda_2 E_{\mu_0}[(\bar{X} - \mu_0)\phi(X)] + \beta$$

The quantity in question is just

$$\int [\phi(x)f_0(x)(\lambda_1 + \lambda_2(\bar{X} - \mu_0)) + (1 - \phi(x))f_1(x)]dx$$

As before this is minimized by

$$\phi(x) = \begin{cases} 1 & \frac{f_1(x)}{f_0(x)} > \lambda_1 + \lambda_2(\bar{X} - \mu_0) \\ 0 & \frac{f_1(x)}{f_0(x)} < \lambda_1 + \lambda_2(\bar{X} - \mu_0) \end{cases}$$

The likelihood ratio f_1/f_0 is simply

$$\exp\{n(\mu_1 - \mu_0)\bar{X} + n(\mu_0^2 - \mu_1^2)/2\}$$

and this exceeds the linear function

$$\lambda_1 + \lambda_2(\bar{X} - \mu_0)$$

for all $\bar{X}$ sufficiently large or small. That is, the quantity

$$\lambda_1\alpha + \lambda_2 E_{\mu_0}[(\bar{X} - \mu_0)\phi(X)] + \beta$$

is minimized by a rejection region of the form

$$\{\bar{X} > K_U\} \cup \{\bar{X} < K_L\}$$

To satisfy the constraints we adjust K_U and K_L to get level α and $\pi'(\mu_0) = 0$. The second condition shows that the rejection region is symmetric about μ_0 and then we discover that the test rejects for

$$\sqrt{n}|\bar{X} - \mu_0| > z_{\alpha/2}$$

Now you have to mimic the Neyman Pearson lemma proof to check that if λ_1 and λ_2 are adjusted so that the unconstrained problem has the rejection region given then the resulting test minimizes β subject to the two constraints.

A test ϕ^* is a Uniformly Most Powerful Unbiased level α_0 test if

1. ϕ^* has level $\alpha \le \alpha_0$.

2. ϕ^* is unbiased.

3. If ϕ has level $\alpha \le \alpha_0$ then for every $\theta \in \Theta_1$ we have

$$E_\theta(\phi(X)) \le E_\theta(\phi^*(X))$$

Conclusion: The two sided z test which rejects if

$$|Z| > z_{\alpha/2}$$

where

$$Z = n^{1/2}(\bar{X} - \mu_0)$$

is the uniformly most powerful unbiased test of $\mu = \mu_0$ against the two sided alternative $\mu \ne \mu_0$.

Nuisance Parameters

What good can be said about the t-test? It's UMPU.

Suppose $X_1, \ldots, X_n$ are iid $N(\mu, \sigma^2)$ and that we want to test $\mu = \mu_0$ or $\mu \le \mu_0$ against $\mu > \mu_0$. Notice that the parameter space is two dimensional and that the boundary between the null and alternatives is

$$\{(\mu, \sigma); \mu = \mu_0, \sigma > 0\}$$

If a test has $\pi(\mu, \sigma) \le \alpha$ for all $\mu \le \mu_0$ and $\pi(\mu, \sigma) \ge \alpha$ for all $\mu > \mu_0$ then we must have $\pi(\mu_0, \sigma) = \alpha$ for all σ because the power function of any test must be conntinuous. (This actually uses the dominated convergence theorem; the power function is an integral.)

Now think of $\{(\mu, \sigma); \mu = \mu_0\}$ as a parameter space for a model. For this parameter space you can check that

$$S = \sum (X_i - \mu_0)^2$$

is a complete sufficient statistic. Remember that the definitions of both completeness and sufficiency depend on the parameter space. Suppose $\phi(\sum X_i, S)$ is an unbiased level α test. Then we have

$$E_{\mu_0,\sigma}(\phi(\sum X_i, S)) = \alpha$$

for all σ. Condition on S and get

$$E_{\mu_0,\sigma}[E(\phi(\sum X_i, S)|S)] = \alpha$$

for all σ. Sufficiency guarantees that

$$g(S) = E(\phi(\sum X_i, S)|S)$$

is a statistic and completeness that

$$g(S) \equiv \alpha$$

Now let us fix a single value of σ and a $\mu_1 > \mu_0$. To make our notation simpler I take $\mu_0 = 0$. Our observations above permit us to condition on $S = s$. Given $S = s$ we have a level α test which is a function of $\bar{X}$.

If we maximize the conditional power of this test for each s then we will maximize its power. What is the conditional model given $S = s$? That is, what is the conditional distribution of $\bar{X}$ given $S = s$? The answer is that the joint density of $\bar{X}, S$ is of the form

$$f_{\bar{X},S}(t,s) = h(s,t)\exp\{\theta_1 t + \theta_2 s + c(\theta_1,\theta_2)\}$$

where $\theta_1 = n\mu/\sigma^2$ and $\theta_2 = -1/\sigma^2$.

This makes the conditional density of $\bar{X}$ given $S = s$ of the form

$$f_{\bar{X}|s}(t|s) = h(s,t)\exp\{\theta_1 t + c^*(\theta_1, s)\}$$

Notice the disappearance of θ_2. Notice that the null hypothesis is actually $\theta_1 = 0$. This permits the application of the NP lemma to the conditional family to prove that the UMP unbiased test must have the form

$$\phi(\bar{X}, S) = 1(\bar{X} > K(S))$$

where $K(S)$ is chosen to make the conditional level α. The function $x \mapsto x/\sqrt{a - x^2}$ is increasing in x for each a so that we can rewrite ϕ in the form

$$\phi(\bar{X}, S) = 1(n^{1/2}\bar{X}/\sqrt{n[S/n - \bar{X}^2]/(n-1)} > K^*(S))$$

for some K^*. The quantity

$$T = \frac{n^{1/2}\bar{X}}{\sqrt{n[S/n - \bar{X}^2]/(n-1)}}$$

is the usual t statistic and is exactly independent of S (see Theorem 6.1.5 on page 262 in Casella and Berger). This guarantees that

$$K^*(S) = t_{n-1,\alpha}$$

and makes our UMPU test the usual t test.