Neyman-Pearson's Lemma

Neyman-Pearson's Lemma gives us a way to construct the most powerful possible tests for a given level.

The lemma is easier in the special case of simple-simple testing:

Definition A hypothesis is <u>simple</u> if it fully determines a single probability distribution.

The setting we work in then is $H_0: \vartheta = \vartheta_0$ vs. $H_A: \vartheta = \vartheta_A$.
Creating a test we would fix the ~~level~~ power(ϑ_0) and try to maximize the ~~power~~ power(ϑ_A).

The Neyman-Pearson Lemma tells us how to do that:

Theorem For a given $\alpha = \text{power}(\vartheta_0)$, the test that maximizes power(ϑ_A) has a rejection region determined by

$$\frac{\mathcal{L}(Y|\vartheta_0)}{\mathcal{L}(Y|\vartheta_A)} < k$$

where k is chosen to achieve the level α.
Such a test is the most powerful test for H_0 vs. H_A.

Example

$X \sim \text{Exponential}(\vartheta)$, i.e. density is $f(x) = \begin{cases} \vartheta e^{-\vartheta x} & x > 0 \\ 0 & \text{otherwise} \end{cases}$

$H_0 : \vartheta = 1$ $\quad\quad$ $H_A : \vartheta = \vartheta_1$ for some $\vartheta_1 > 1$.

Neyman-Pearson's lemma suggests we use

$\frac{L(1)}{L(\vartheta_1)} = \frac{e^{-x}}{\vartheta_1 e^{-\vartheta_1 x}} < k$ as our rejection region

Solving for x: $\frac{e^{-x}}{e^{-\vartheta_1 x}} < k\vartheta_1$

$$-x - (-\vartheta_1 x) < \log(k\vartheta_1)$$

$$x < \log(k\vartheta_1) / (\vartheta_1 - 1) = k'.$$

To pick k', we calculate the level:

$$\alpha = \mathbb{P}(X < k' \mid \vartheta = 1) = \int_0^{k'} e^{-x}\, dx = 1 - e^{-k'}.$$

Solving we get $\quad k' = -\log(1-\alpha)$

NOTE The resulting rejection region does not depend on H_A!

This is an example of a more general phenomenon:

One-sided Most Powerful Tests

Let our setup be $H_0 : \vartheta \le \vartheta_0$ vs. $H_A : \vartheta > \vartheta_0$.
If the underlying probability densities are nice enough, then the likelihood ratio construction carries over:

Definition A family $p_\vartheta(x)$ of densities has monotone likelihood ratios if there exists some statistic $T = T(x)$ such that if $\vartheta_1 < \vartheta_2$ then $p_{\vartheta_2}(x) / p_{\vartheta_1}(x)$ is a non-decreasing function of T.

Example Exponential families have monotone likelihood ratios if $\eta(\vartheta)$ is increasing.

Definition A uniformly most powerful test maximizes power(ϑ_A) at every ϑ_A.

Theorem The rejection region $\{T > c\}$ for some constant c produces the uniformly most powerful test for a family of densities with monotone likelihood ratios.

Example $\mathcal{N}(\mu, \sigma^2)$ with known σ^2 is an exponential family with $\eta(\mu) = \mu/\sigma^2$ and $T(x) = \sum x_i$.
Since $\eta(\mu)$ is increasing, this has monotone likelihood, and therefore a rejection region on the form $\{\sum x_i > c\}$ for some c.
This agrees with the z-score test from earlier!
So the z-score test is uniformly most powerful!

Example (Exercise 10.94)

$Y_1, \dots, Y_n \sim N(\mu, \sigma^2)$ iid, μ known.

$H_0: \sigma^2 = \sigma_0^2$ vs. $H_A: \sigma^2 = \sigma_1^2$.

Neyman-Pearson: ump test is $\left\{ \frac{L(\sigma_0^2)}{L(\sigma_1^2)} < k \right\}$ for some k.

Ratio is

$$\frac{\exp\left[-\frac{n}{2}\log 2\pi - \frac{n}{2}\log\sigma_0^2 - \frac{1}{2\sigma_0^2}\Sigma(Y_i-\mu)^2\right]}{\exp\left[-\frac{n}{2}\log 2\pi - \frac{n}{2}\log\sigma_1^2 - \frac{1}{2\sigma_1^2}\Sigma(Y_i-\mu)^2\right]} < k$$

Apply log and clear out common terms:

$$\Rightarrow -\frac{n}{2}\log\frac{\sigma_0^2}{\sigma_1^2} - \frac{1}{2}\Sigma\left(\frac{Y_i-\mu}{\sigma_0}\right)^2 + \frac{1}{2}\Sigma\left(\frac{Y_i-\mu}{\sigma_1}\right)^2 < k$$

Rewrite to

$$\left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_0^2}\right)\Sigma(Y_i-\mu)^2 < 2\left(k + n\log\frac{\sigma_0^2}{\sigma_1^2}\right)$$

$$\Sigma(Y_i-\mu)^2 < 2\left(k + n\log\frac{\sigma_0^2}{\sigma_1^2}\right) \Big/ \left(\frac{1}{\sigma_1^2} - \frac{1}{\sigma_0^2}\right) = k'$$

So the rejection region is $\{\Sigma(Y_i-\mu)^2 < k'\}$ for some k'.

$$\alpha = P(\Sigma(Y_i-\mu)^2 < k' \mid \sigma^2 = \sigma_0^2) = P\left(\Sigma\left(\frac{Y_i-\mu}{\sigma_0}\right)^2 < k'\sigma_0^2 \mid \sigma^2 = \sigma_0^2\right)$$
$$= P(W < k'\sigma_0^2 \mid \sigma^2 = \sigma_0^2) = F_{\chi^2(n)}(k'\sigma_0^2)$$

So $k' = F^{-1}_{\chi^2(n)}(\alpha)/\sigma_0^2$ is an appropriate cutoff for the rejection region. It is independent of σ_1^2.

This is the same rejection criterion as in the 1-sample variance test.

Example (10.101)

$Y_1, \dots, Y_n \sim \text{Exponential}(\vartheta)$

So $\mathcal{L}(Y|\vartheta) = \frac{1}{\vartheta^n} \exp\left[-\sum Y_i / \vartheta\right] = \exp\left[-n \log\vartheta - n\bar{Y}/\vartheta\right]$

Exponential family. $-1/\vartheta$ is increasing, so it has monotone likelihoods.

Hence $RR = \{\bar{Y} < c\}$ is UMP for $H_0: \vartheta = \vartheta_0$ vs. $H_A: \vartheta < \vartheta_0$.

To pick c: $\alpha = \mathbb{P}(\bar{Y} < c \mid \vartheta_0) = \mathbb{P}(\sum Y_i < nc \mid \vartheta_0)$.

$\sum Y_i \sim \Gamma(n, \vartheta_0)$ (See Wikipedia page for Exponential Distribution)

So $c = F^{-1}_{\Gamma(n,\vartheta_0)}(\alpha)/n$ works as a rejection cutoff.

Example (10.102)

$Y_1, \dots, Y_n \sim \text{Bernoulli}(p)$.

$H_0: p = p_0$ vs. $H_A: p > p_0$.

$$\mathcal{L}(Y/p) = p^{\sum Y_i}(1-p)^{n-\sum Y_i} = \left(\frac{p}{1-p}\right)^{\sum Y_i}(1-p)^n =$$

$$= \exp\left[\sum Y_i \log\frac{p}{1-p} + n \log 1-p\right].$$

Since $\log\frac{p}{1-p}$ is increasing, this is an exponential family with monotone likelihood — so the $RR = \{\sum Y_i > c\}$ for some c.

$\sum Y_i \sim \text{Binomial}(n, p)$. Therefore, to determine c:

$$\alpha = \mathbb{P}\left(\sum Y_i > c \,/\, p_0\right) = 1 - F_{\text{Binomial}(n,p_0)}(c)$$

So $c = F^{-1}_{\text{Binomial}(n,p_0)}(1-\alpha)$ works as a rejection cutoff for a UMP test.

Example (10.103)

$Y_1, \ldots, Y_n \sim \text{Uniform}(0, \vartheta)$.

$H_0 : \vartheta = \vartheta_0$ vs. $H_A : \vartheta = \vartheta_a$ with $\vartheta_a < \vartheta_0$.

$$\mathcal{L}(Y|\vartheta) = \frac{1}{\vartheta^n} \cdot \mathbb{1}_{[0,\infty)}(\min Y_i) \cdot \mathbb{1}_{(-\infty, \vartheta]}(\max Y_i).$$

$$\frac{\mathcal{L}(\vartheta_0)}{\mathcal{L}(\vartheta_a)} = \left(\frac{\vartheta_a}{\vartheta_0}\right)^n \cdot \mathbb{1}_{[0,\infty)}(\min Y_i) \cdot \mathbb{1}_{(-\infty, \vartheta_a)}(\max Y_i)$$

If $\max Y_i > \vartheta_a$, then automatically favour H_0.
Otherwise the test depends on $\max Y_i$ — so reject if $\max Y_i < k$ for some k.

$\max Y_i$ has PDF $n y^{n-1} \vartheta^{-n}$.
We pick c such that

$$\alpha = \mathbb{P}(\max Y_i < c | \vartheta_0) = \int_0^c n y^{n-1} \vartheta^{-n} dy = \frac{c^n}{\vartheta_0^n}$$

So $c = \vartheta_0 \sqrt[n]{\alpha}$ produces a most powerful test.
Since it does not depend on ϑ_a, it is UMP.

A generic likelihood ratio test

Recall that hypotheses are "just" arbitrary subsets of parameter values.
So we can talk about $H_0: \vartheta \in \Omega_0$ vs $H_A: \vartheta \in \Omega_A$ in high generality — vector valued ϑ; weird shapes of Ω_0 and Ω_A etc.

Write $\Omega = \Omega_0 \cup \Omega_a$ for <u>all</u> parameter values in play.

Then a generic test statistic is

$$\lambda = \frac{\mathcal{L}(\Omega_0)}{\mathcal{L}(\Omega)} = \frac{\max_{\vartheta \in \Omega_0} \mathcal{L}(\vartheta)}{\max_{\vartheta \in \Omega} \mathcal{L}(\vartheta)}$$

If H_0 is true, then $\lambda = 1$. Smaller λ indicates larger difference between $\mathcal{L}(\Omega_0)$ and $\mathcal{L}(\Omega_a)$ — so smaller λ favors rejection.
We get rejection region $RR = \{\lambda < k\}$.

It can be rare that the distribution of λ is known, which makes picking k more difficult.
However — there is a useful theorem for large sample sizes:

Theorem Let r_0 be # free parameters specified by Ω_0 and $r =$ # free in Ω.
Then, for large n,

$$-2 \ln \lambda \sim \chi^2(r_0 - r)$$

Example (10.112–114)

See slides.