[Monday Feb. 15, 2021 Lecture 12] (No class Wed.).

3 Uniformly Most Powerful Tests

The case considered here is that θ is modeled as a deterministic but unknown constant (or as random, but with unknown statistics). Then a Bayes test is not meaningful so we look at Neyman-Pearson methods.

- Say parameter set is given as a disjoint union: $\Lambda = \Lambda_0 \cup \Lambda_1$. Hypothesis H_0 corresponds to a state of nature P_θ where $\theta \in \Lambda_0$. Similarly, for H_1.
- Let $\delta(y)$ be a randomized decision rule for H_0 vs. H_1. Define $\delta : \Gamma \to [0,1]$
  - False-alarm probabilities.
  $$P_F(\delta;\theta) = \mathrm{E}_\theta\{\delta(Y)\}, \quad \text{for } \theta \in \Lambda_0.$$
  - Detection probabilities.
  $$P_D(\delta;\theta) = \mathrm{E}_\theta\{\delta(Y)\}, \quad \text{for } \theta \in \Lambda_1.$$
- A Uniformly Most Powerful (UMP) test of level α is one that maximizes
  $$P_D(\delta;\theta)$$
  for every $\theta \in \Lambda_1$ subject to
  $$P_F(\delta;\theta) \le \alpha$$
  for all $\theta \in \Lambda_0$.

- UMP tests do not always exist.

3.1 Example: UMP Tests Don't Always Exist

- $\Lambda = \Lambda_0 \cup \Lambda_1$
- Suppose H_0 is simple, i.e., $\Lambda_0 = \{\theta_0\}$
- Suppose that P_θ has a density $f_\theta(\cdot)$ for each $\theta \in \Lambda$ and consider the Neyman-Pearson problem for testing
$$\begin{aligned} H_0 &: Y \sim P_{\theta_0} \\ &\text{vs.} \\ H_\theta &: Y \sim P_\theta \end{aligned}$$
for some fixed $\theta \in \Lambda_1$. To be clear, at the moment we are considering a simple hypothesis test.
- We know from the NPL that there exists a most powerful α-level test for this problem with a critical region of the form
$$\Gamma_\theta = \{y \in \Gamma : f_\theta(y) > \tau f_{\theta_0}(y)\}$$
where τ and a possible randomization are chosen to give a size α test. Also from the NPL we know that the test is essentially unique and that any other α-level test will have smaller power.

- So for two distinct parameter values θ', $\theta'' \in \Lambda_1$ the test with critical region $\Gamma_{\theta'}$ will have a smaller power for testing

$$H_0 \text{ vs. } H_{\theta''}$$

than the test with $\Gamma_{\theta''}$ (and vice versa) *unless* $\Gamma_{\theta'}$ and $\Gamma_{\theta''}$ are essentially identical.

A UMP test for

$$H_0 \textit{ vs. } H_\theta : Y \sim P_\theta,\ \theta \in \Lambda_1$$

exists if and only if the critical region Γ_θ *is (essentially) the same for all* $\theta \in \Lambda_1$.

Synonymously, we can say a UMP test exists if and only if the LRT for every $\theta \in \Lambda_1$ can be completely defined (including threshold) without knowledge of θ.

3.2 Example: UMP Testing of Location

- Consider the family of distributions $\{P_\theta : \theta \in \Lambda\}$ where Λ is a subset of $\mathcal{R}$ and P_θ is $\mathcal{N}(\theta, \sigma^2)$. σ^2 is Known.

- Consider the hypothesis pair

$$\begin{aligned} H_0 : &\quad \theta = \mu_0 \\ &\text{vs.} \\ H_1 : &\quad \theta > \mu_0 \end{aligned}$$

where μ_0 is a fixed real number. Therefore, we have a simple null hypothesis $\Lambda_0 = \{\mu_0\}$ and a composite alternative $\Lambda_1 = (\mu_0, \infty)$.

- From the previous example (i.e., location testing with Gaussian error) we know that for each fixed $\theta \in \Lambda_1$ the most powerful α-level test for H_0 versus $Y \sim \mathcal{N}(\theta, \sigma^2)$ has a critical region

$$\Gamma_\theta = \{y \in \Gamma : y > \underline{\sigma\Phi^{-1}(1-\alpha) + \mu_0}\}.$$

This region does not depend upon θ (note that θ is restricted to be $> \mu_0$) and thus it gives a UMP test for $H_0 : \theta = \mu_0$ vs. $H_1 : \theta > \mu_0$.

- Let δ_1 denote the decision rule for the critical region Γ_θ (as seen previously, randomization is not required). The detection probabilities are:

$$P_D(\delta_1; \theta) = 1 - \Phi\left(\Phi^{-1}(1-\alpha) - \frac{\theta - \mu_0}{\sigma}\right)$$

for $\theta > \mu_0$.

- Now for the same family of distributions consider the hypothesis testing problem

$$\begin{aligned} H_0 :& \quad \theta = \mu_0 \\ \text{vs.}& \\ H_1 :& \quad \theta \neq \mu_0 \end{aligned}$$

where μ_0 is a fixed real number. Therefore, we have the same simple null hypothesis $\Lambda_0 = \{\mu_0\}$ and a new composite alternative $\Lambda_1 = (-\infty, \mu_0) \cup (\mu_0, \infty)$.

- For $\theta > \mu_0$ the critical region of the most powerful α-level test is as before, but for $\theta < \mu_0$ it is different. With the two cases included in the same formula:

$$\Gamma_\theta = \begin{cases} \{y \in \Gamma : y > \sigma\Phi^{-1}(1-\alpha) + \mu_0\} & \text{for} \quad \theta > \mu_0 \\ \{y \in \Gamma : y < \sigma\Phi^{-1}(\alpha) + \mu_0\} & \text{for} \quad \theta < \mu_0 \end{cases}.$$

  In the sense that the critical region "switches" as a function of θ between these two cases, it depends upon θ and

  No UMP test exists for this two-sided hypothesis testing problem.

- Let δ_2 be the test corresponding to critical region Γ_θ with $\theta < \mu_0$. Then we can show

$$P_D(\delta_2; \theta) = \Phi\left(\Phi^{-1}(\alpha) - \frac{\theta - \mu_0}{\sigma}\right)$$

  for $\theta < \mu_0$.

- We can certainly extend the definitions of $P_D(\delta_1; \theta)$ and $P_D(\delta_2; \theta)$ to $\theta \in \mathcal{R}$ and then plot the two power functions together on the same axis. Note that $P_D(\delta_1; \theta)$ increases as θ increases while $P_D(\delta_2; \theta)$ decreases as θ increases. The curves cross when θ is such that

$$\Phi\left(\Phi^{-1}(\alpha) - \frac{\theta - \mu_0}{\sigma}\right) = 1 - \Phi\left(\Phi^{-1}(1-\alpha) - \frac{\theta - \mu_0}{\sigma}\right)$$

which happens for $\theta = \mu_0$ where both sides equal α.

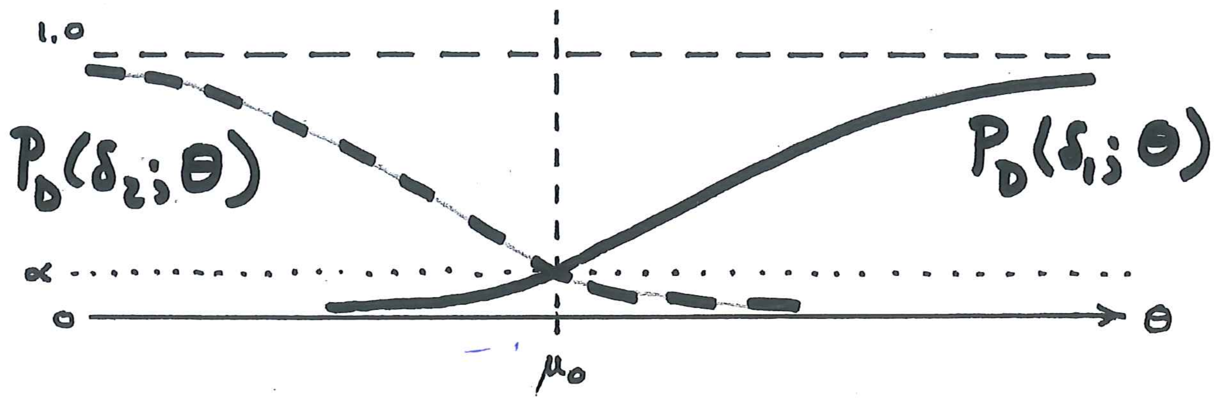

- This shows that neither test performs well outside of its region of optimality. A more reasonable test than either δ_1 or δ_2 would compare $|y - \mu_0|$ to a threshold, but this cannot be UMP for $H_0 : \theta = \mu_0$ vs. $H_1 : \theta \neq \mu_0$ because such does not exist.

3.3 Unbiasedness

- Previous example illustrates how the UMP criterion is too strong for some problems since it is not useful to aim for a criterion for which a test does not exist.

- Sometimes can overcome the difficulty by applying more constraints to eliminate unreasonable tests.
- One such condition is to require <u>unbiasedness</u> meaning we require[4]

$$P_D(\delta;\theta) \geq \alpha$$

for all $\theta \in \Lambda_1$ in addition to the constraint $P_F(\delta;\theta) \leq \alpha$ for all $\theta \in \Lambda_0$. This would have eliminated both δ_1 and δ_2 from consideration in the previous example.

4 Locally Most Powerful Tests

- Consider the case where Λ is of the form $[\theta_0, \infty)$ with $\Lambda_0 = \{\theta_0\}$ and $\Lambda_1 = (\theta_0, \infty)$.
- Such comes up in many signal detection problems in which $\theta_0 = 0$ and θ is a signal amplitude parameter.
- Often we are primarily interested in the case where, under H_1, θ is close to θ_0. When θ is a signal amplitude parameter this would correspond to small signal strength (i.e.,

[4] Recalling the actual definition of detection and false alarm probabilties we can give a more symmetric definition for unbiasedness. Namely δ is unbiased in this composite binary hypothesis testing problem if

$$E_\theta\{\delta(Y)\} \text{ is } \begin{cases} \geq \alpha & \text{for all } \theta \in \Lambda_1 \\ \leq \alpha & \text{for all } \theta \in \Lambda_0 \end{cases}.$$

the low signal-to-noise ratio regime).

- Consider a decision rule δ. Then subject to "regularity conditions" we may expand $P_D(\delta;\theta)$ in a Taylor series about θ_0: $P_F(\delta)$

$$P_D(\delta;\theta) = P_D(\delta;\theta_0) + (\theta-\theta_0)P_D'(\delta;\theta_0) + O((\theta-\theta_0)^2)$$

where $P_D'(\delta;\theta_0) = \frac{\partial}{\partial\theta}P_D(\delta;\theta)$.

- Note that $P_D(\delta;\theta_0) = P_F(\delta)$ so for all α sized tests we see that for θ near θ_0

$$P_D(\delta;\theta) \approx \alpha + (\theta-\theta_0)P_D'(\delta;\theta_0).$$

Conclude that for θ near θ_0 we can achieve an approximate maximum power with size α by choosing δ to maximize $P_D'(\delta;\theta_0)$.

- A test which maximizes $P_D'(\delta;\theta_0)$ subject to a false alarm constraint $P_F(\delta) \leq \alpha$ is called an α-level locally most powerful (LMP) test or simply a locally optimum test.

4.1 The general structure of LMP

- Assume that P_θ has density f_θ for each $\theta \in \Lambda_1$. Then we can write

$$P_D(\delta;\theta) = \mathrm{E}_\theta\{\delta(Y)\} = \int_\Gamma \delta(y)f_\theta(y)\mathrm{d}y.$$

- If the family $\{f_\theta(y) : \theta \in \Lambda_1\}$ is sufficiently regular that differentiation wrt θ and integration wrt y may be interchanged then

$$P_D'(\delta;\theta) = \int_\Gamma \delta(y) \left.\frac{\partial}{\partial\theta} f_\theta(y)\right|_{\theta=\theta_0} dy.$$

- Comparison of this expression with our previous work on NP testing for simple hypotheses shows that the α-level LMP problem is the same as the α-level NP design problem where we replace $f_1(y)$ with

$$\left.\frac{\partial}{\partial\theta} f_\theta(y)\right|_{\theta=\theta_0}.$$

- From this analogy (within regularity) an α-level LMP test for $H_0 : \theta = \theta_0$ vs. $H_1 : \theta > \theta_0$ is given by

$$\delta_{lo}(y) = \begin{cases} 1 & & > \\ \gamma & \text{if} \quad \left.\frac{\partial}{\partial\theta} f_\theta(y)\right|_{\theta=\theta_0} & = \; \eta f_{\theta_0}(y) \\ 0 & & > \end{cases}$$

where η and γ are chosen st $P_F(\delta_{lo}) = \alpha$.

5 Generalized Likelihood Ratio Tests

- In the absense of the applicability of any of the above a test for composite hypothesis which is often used is to compare

$$\frac{\max_{\theta \in \Lambda_1} f_\theta(y)}{\max_{\theta \in \Lambda_0} f_\theta(y)}$$

to a threshold.

- Called a generalized likelihood ratio test (GLRT) or maximum likelihood test.

Maximum Likelihood.

$V(\pi)$ = min. Bayes risk for a Bayes rule as a function of $\pi = P(H_0 \text{ true})$.