

Lecture 16 – Most powerful tests for one-sided hypotheses

Last Modified: October 23, 2026

Outline:

16.1 From most powerful to uniformly most powerful	2
16.2 Monotone likelihood ratios	2
16.3 UMP tests for one-sided hypotheses	4

Finishing example from Lecture 15:

Example 1 (Exponential rate). Suppose $X_1, \dots, X_n \stackrel{\text{iid}}{\sim} \text{Exp}(\theta)$, where $\theta > 0$ is the rate. Find the most powerful level α test for

$$H_0 : \theta = \theta_0 \quad \text{versus} \quad H_1 : \theta = \theta_1,$$

where $\theta_1 > \theta_0$ and $0 < \alpha < 1$.

The likelihood function is

$$L_{\underline{x}}(\theta) = \theta^n \exp \left\{ -\theta \sum_{i=1}^n x_i \right\} \mathbb{1}_{\{x_1 > 0, \dots, x_n > 0\}}.$$

So the likelihood ratio is

$$\frac{L_{\underline{x}}(\theta_1)}{L_{\underline{x}}(\theta_0)} = \left(\frac{\theta_1}{\theta_0} \right)^n \exp \left\{ -(\theta_1 - \theta_0) \sum_{i=1}^n x_i \right\}.$$

Since $\theta_1 > \theta_0$, this ratio is decreasing in $\sum_{i=1}^n x_i$. By the Neyman-Pearson lemma, the most powerful test therefore rejects when $\sum_{i=1}^n X_i < c$, with c chosen to give size α . This makes sense because θ_1 being larger than θ_0 means the mean of $\text{Exp}(\theta_1)$ is *smaller* than the mean of $\text{Exp}(\theta_0)$.

Let us now choose c . Under H_0 , we have that $\theta_0 X_i \stackrel{\text{iid}}{\sim} \text{Exp}(1)$, so

$$\theta_0 \sum_{i=1}^n X_i \sim \text{Gamma}(n, 1).$$

Let $F_{n,1}$ be the cdf of $\text{Gamma}(n, 1)$ and let $q_\alpha = F_{n,1}^{-1}(\alpha)$ be its *lower-tail* quantile: $\mathbb{P}(Y < q_\alpha) = \alpha$ for $Y \sim \text{Gamma}(n, 1)$. Then we should choose c so that

$$\alpha = \mathbb{P}_{\theta_0} \left(\sum_{i=1}^n X_i < c \right) = \mathbb{P}_{\theta_0} \left(\theta_0 \sum_{i=1}^n X_i < \theta_0 c \right) = \mathbb{P}(\text{Gamma}(n, 1) \leq \theta_0 c).$$

Thus we should set $\theta_0 c = q_\alpha$, or $c = q_\alpha / \theta_0$. We conclude that the most powerful level α test rejects when

$$\sum_{i=1}^n X_i < \frac{q_\alpha}{\theta_0}.$$

16.1 From most powerful to uniformly most powerful

The Neyman–Pearson lemma gives a most powerful (MP) test when both hypotheses are simple. For a composite alternative (in which Θ_1 consists of more than one parameter value), we want a single test that is most powerful at every alternative parameter value. The key question is therefore: *do the different simple alternatives all lead to the same test?*

Suppose first that we test

$$H_0 : \theta = \theta_0 \quad \text{versus} \quad H_1 : \theta \in \Theta_1.$$

For each fixed $\theta_1 \in \Theta_1$, use Neyman–Pearson to find an MP level- α test against θ_1 . Suppose the *same* test ψ^* works for every $\theta_1 \in \Theta_1$. We claim it is UMP level- α for the original hypotheses. To see why, fix any other level- α ψ for testing $H_0 : \theta = \theta_0$ vs $H_1 : \theta \in \Theta_1$, and fix any $\theta_1 \in \Theta_1$. Then ψ is also a valid level- α test for $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_1$. But ψ^* is MP for this test, so $\beta_{\psi^*}(\theta_1) \geq \beta_{\psi}(\theta_1)$. This is the definition of ψ^* being level- α for $H_0 : \theta = \theta_0$ vs $H_1 : \theta \in \Theta_1$.

What if the null is composite too? As we now show, it suffices to solve the simple problems using one reference point θ_0 in the null, and then check that the resulting test has level α over the entire null.

Proposition 1 (From simple comparisons to a composite problem). Consider $H_0 : \theta \in \Theta_0$ versus $H_1 : \theta \in \Theta_1$. Suppose that there is some $\theta_0 \in \Theta_0$ and there is a test ψ^* that is MP of level α for $\theta = \theta_0$ versus $\theta = \theta_1$ for every $\theta_1 \in \Theta_1$. If also

$$\sup_{\theta \in \Theta_0} \mathbb{P}_{\theta}(\psi^*(\underline{X}) = 1) \leq \alpha,$$

then ψ^* is UMP of level α for the composite problem.

Remark 1 (Which θ_0 should we try?). Recall from the Neyman–Pearson lemma that the MP test ψ^* we know how to construct for simple hypotheses, θ_0 vs θ_1 , has size exactly α . This means $\mathbb{P}_{\theta_0}(\psi^*(\underline{X}) = 1) = \alpha$ exactly. Since this same ψ^* is now required to satisfy $\mathbb{P}_{\theta}(\psi^*(\underline{X}) = 1) \leq \alpha$ for every θ in the null, the θ_0 in the null for which equality holds should be the one for which rejecting is most probable.

Proof. Let ψ be any competing level- α test for the composite problem. Because θ_0 belongs to the null, $\mathbb{P}_{\theta_0}(\psi(\underline{X}) = 1) \leq \alpha$. Thus ψ is also an allowable competitor in each simple problem. By the MP property,

$$\mathbb{P}_{\theta_1}(\psi^*(\underline{X}) = 1) \geq \mathbb{P}_{\theta_1}(\psi(\underline{X}) = 1) \quad \text{for every } \theta_1 \in \Theta_1.$$

This is exactly the power comparison required for UMP. The extra size condition ensures that ψ^* itself is allowable for the composite problem. $\square$

Remark 2. Enlarging the null makes the level constraint more demanding: a test must now control its rejection probability at more parameter values. It also leaves fewer competing tests. Thus an MP comparison proved using only θ_0 is enough, *provided our proposed test still satisfies the stronger level constraint.*

16.2 Monotone likelihood ratios

We now specify a property of the statistical model f_{θ} which guarantees that the same ψ_* is the level- α MP for θ_0 vs θ_1 , regardless of the θ_1 in the alternative.

Definition 1. A family of pdfs or pmfs with scalar parameter θ has an **increasing monotone likelihood ratio (MLR)** in a statistic $T(\underline{x})$ if, for every $\theta_2 > \theta_1$,

$$\frac{L_{\underline{x}}(\theta_2)}{L_{\underline{x}}(\theta_1)}$$

depends on $\underline{x}$ only through $T(\underline{x})$ and is increasing (technically, non-decreasing) in $T(\underline{x})$. A **decreasing MLR** is defined in the same way, with increasing replaced by decreasing (technically, by non-increasing). When only the denominator is zero, interpret the ratio as $+\infty$. Sample points $\underline{x}$ such that $L_{\underline{x}}(\theta_1) = L_{\underline{x}}(\theta_2) = 0$ can be ignored.

An increasing MLR means that larger values of T provide stronger evidence for the larger parameter value. Crucially, *every pair* $\theta_2 > \theta_1$ orders the observations in the same way. This is what lets the simple Neyman–Pearson tests share a rejection rule.

Example 2 (Exponential families). Suppose the likelihood has the form

$$L_{\underline{x}}(\theta) = h(\underline{x}) \exp\{b(\theta)T(\underline{x}) - c(\theta)\}$$

on a support that does not depend on θ . Then

$$\frac{L_{\underline{x}}(\theta_2)}{L_{\underline{x}}(\theta_1)} = \exp\{[b(\theta_2) - b(\theta_1)]T(\underline{x}) - [c(\theta_2) - c(\theta_1)]\}.$$

The last term is constant as the data vary. Thus the family has an increasing MLR if b is non-decreasing, and a decreasing MLR if b is non-increasing.

Example 3 (Exponential observations). Suppose $X_i \stackrel{\text{iid}}{\sim} \text{Exp}(\theta)$, where $\theta > 0$ is the *rate*. On the positive sample space,

$$L_{\underline{x}}(\theta) = \theta^n \exp\left\{-\theta \sum_{i=1}^n x_i\right\}.$$

For $\theta_2 > \theta_1$,

$$\frac{L_{\underline{x}}(\theta_2)}{L_{\underline{x}}(\theta_1)} = \left(\frac{\theta_2}{\theta_1}\right)^n \exp\left\{-(\theta_2 - \theta_1) \sum_{i=1}^n x_i\right\}.$$

This is decreasing in $T(\underline{x}) = \sum_{i=1}^n x_i$. A larger rate means shorter waiting times, so a small sum is evidence for a larger θ .

Example 4 (Uniform observations). Suppose $X_i \stackrel{\text{iid}}{\sim} \text{Unif}(0, \theta)$, where $\theta > 0$. On the positive sample space,

$$L_{\underline{x}}(\theta) = \theta^{-n} \mathbb{1}_{\{x_{(n)} < \theta\}}.$$

For $\theta_2 > \theta_1$, the ratio on the union of the two supports is

$$\frac{L_{\underline{x}}(\theta_2)}{L_{\underline{x}}(\theta_1)} = \begin{cases} (\theta_1/\theta_2)^n, & x_{(n)} < \theta_1, \\ +\infty, & \theta_1 \leq x_{(n)} < \theta_2. \end{cases}$$

It is non-decreasing in $T(\underline{x}) = x_{(n)}$. Below θ_1 , the ratio is constant. Once the maximum exceeds θ_1 , the smaller parameter cannot explain the data at all, so we are completely sure the samples must have been drawn from $\text{Unif}(0, \theta_2)$ rather than $\text{Unif}(0, \theta_1)$.

16.3 UMP tests for one-sided hypotheses

Theorem 1 (Karlin–Rubin). Consider

$$H_0 : \theta \leq \theta_0 \quad \text{versus} \quad H_1 : \theta > \theta_0.$$

Suppose the family has an increasing MLR in T . If k is such that

$$\mathbb{P}_{\theta_0}(T > k) = \alpha,$$

then the test with rejection region $\mathcal{R} = \{\underline{x} : T(\underline{x}) > k\}$ is UMP of level α .

If instead the family has a decreasing MLR in T , the corresponding UMP test rejects when $T < k$, with $\mathbb{P}_{\theta_0}(T < k) = \alpha$.

Proof for increasing MLR.

Simple alternatives all give the same test: Fix any $\theta_1 > \theta_0$. By MLR, the likelihood ratio $L_{\underline{x}}(\theta_1)/L_{\underline{x}}(\theta_0)$ increases with T . Neyman–Pearson therefore tells us that the MP test for θ_0 vs θ_1 rejects when T is large. So the test is $\psi^*(\underline{x}) = \mathbb{1}_{\{T(\underline{x}) > k\}}$, where k is chosen so that

$$\mathbb{P}_{\theta_0}(\psi^*(\underline{X}) = 1) = \mathbb{P}_{\theta_0}(T > k) = \alpha.$$

This value of k does not depend on θ_1 ! Therefore ψ^* is level- α MP for each $\theta_1 > \theta_0$.

Rejection probabilities increase as θ approaches θ_0 from below: We will show that

$$\mathbb{P}_{\theta}(T > k) \leq \mathbb{P}_{\theta_0}(T > k) = \alpha \quad \text{for every } \theta \leq \theta_0. \quad (16.1)$$

Here, the equality to α is by design, and the inequality is what is new. We prove (16.1) below using the MLR property.

Conclude: We conclude that the test ψ^* is level α , since (16.1) gives

$$\sup_{\theta \leq \theta_0} \mathbb{P}_{\theta}(T > k) = \mathbb{P}_{\theta_0}(T > k) = \alpha.$$

We have now checked both ingredients of Proposition 1: we found a test ψ^* which is MP for θ_0 against any $\theta_1 \in \Theta_1$, and we showed ψ^* has level α over the entire null. Hence it is UMP. For decreasing MLR, apply the same argument to $-T$. $\square$

Proof of (16.1) (optional). First, we parse the likelihood ratio. The MLR assumption implies in particular that T is sufficient. Using the definition of sufficiency (recall Lec 3), we can write $f_{\theta}(\underline{x}) = h(\underline{x})g_{\theta}(T(\underline{x}))$, where g_{θ} is the pdf/pmf of T . Therefore, for any two parameter values θ_a, θ_b , we have

$$\frac{L_{\underline{x}}(\theta_b)}{L_{\underline{x}}(\theta_a)} = \frac{g_{\theta_b}(T(\underline{x}))h(\underline{x})}{g_{\theta_a}(T(\underline{x}))h(\underline{x})} = \frac{g_{\theta_b}(T(\underline{x}))}{g_{\theta_a}(T(\underline{x}))}.$$

Now assume $\theta_a < \theta_b$ and use g_a, g_b as shorthand for $g_{\theta_a}, g_{\theta_b}$. Increasing MLR gives that $t \mapsto g_b(t)/g_a(t)$ is nondecreasing in t . Therefore, if we take any two values $u < t$, we have

$$\frac{g_b(t)}{g_a(t)} \geq \frac{g_b(u)}{g_a(u)}, \quad \text{so} \quad g_b(t)g_a(u) \geq g_a(t)g_b(u).$$

We now take this inequality and integrate it over all $t > k$. This gives

$$\mathbb{P}_{\theta_b}(T > k)g_a(u) = \int_k^{\infty} g_b(t)g_a(u)dt \geq \int_k^{\infty} g_a(t)g_b(u)dt = \mathbb{P}_{\theta_a}(T > k)g_b(u).$$

Next, integrate over $u \leq k$ to get

$$\begin{aligned}\mathbb{P}_{\theta_b}(T > k)\mathbb{P}_{\theta_a}(T \leq k) &= \mathbb{P}_{\theta_b}(T > k) \int_{-\infty}^k g_a(u) du \\ &\geq \mathbb{P}_{\theta_a}(T > k) \int_{-\infty}^k g_b(u) du = \mathbb{P}_{\theta_a}(T > k)\mathbb{P}_{\theta_b}(T \leq k)\end{aligned}$$

Write $p_a = \mathbb{P}_{\theta_a}(T > k)$ and $p_b = \mathbb{P}_{\theta_b}(T > k)$. The inequality is $p_b(1 - p_a) \geq p_a(1 - p_b)$; canceling the product terms gives $p_b \geq p_a$. Recalling what p_b and p_a are, we conclude that

$$\mathbb{P}_{\theta_b}(T > k) \geq \mathbb{P}_{\theta_a}(T > k) \quad \text{for every } k, \quad \text{for every } \theta_b > \theta_a.$$

Thus larger parameter values give larger upper-tail probabilities. This proves (16.1) $\square$

Remark 3 (Other one-sided hypotheses). The same test is also UMP when the null is just $\theta = \theta_0$. For the reversed hypotheses $H_0 : \theta \geq \theta_0$ versus $H_1 : \theta < \theta_0$, reverse the rejection direction: an increasing MLR gives a lower-tail test, and a decreasing MLR gives an upper-tail test. In each case, calibrate at θ_0 .

Example 5 (Testing an exponential rate). Suppose $X_i \stackrel{\text{iid}}{\sim} \text{Exp}(\theta)$ and consider

$$H_0 : \theta \leq \theta_0 \quad \text{versus} \quad H_1 : \theta > \theta_0.$$

The decreasing MLR in $T = \sum_{i=1}^n X_i$ gives the UMP rejection region

$$\mathcal{R} = \left\{ \underline{x} : \sum_{i=1}^n x_i < k \right\}, \quad \mathbb{P}_{\theta_0} \left(\sum_{i=1}^n X_i < k \right) = \alpha.$$

This is the same test as for the simple null in the previous lecture. Smaller rates make waiting times longer, so rejecting for a small sum is less likely anywhere else in the null. Consequently, calibration at the boundary controls the whole null.

Example 6 (Testing a uniform endpoint). Suppose $X_i \stackrel{\text{iid}}{\sim} \text{Unif}(0, \theta)$ and consider

$$H_0 : \theta \geq \theta_0 \quad \text{versus} \quad H_1 : \theta < \theta_0.$$

The increasing MLR in $X_{(n)}$ tells us to reject for a small maximum. For $0 < k < \theta_0$, independence gives

$$\mathbb{P}_{\theta_0}(X_{(n)} < k) = \mathbb{P}_{\theta_0}(X_1 < k, \dots, X_n < k) = \left(\frac{k}{\theta_0} \right)^n.$$

Equating this to α gives $k = \theta_0 \alpha^{1/n}$, so the UMP test rejects when

$$X_{(n)} < \theta_0 \alpha^{1/n}.$$

We can check the composite-null constraint directly: for every $\theta \geq \theta_0$,

$$\mathbb{P}_{\theta}(X_{(n)} < \theta_0 \alpha^{1/n}) = \left(\frac{\theta_0 \alpha^{1/n}}{\theta} \right)^n = \alpha \left(\frac{\theta_0}{\theta} \right)^n \leq \alpha.$$

The same test is also UMP for the simple null $\theta = \theta_0$.