8: UNIFORMLY MOST POWERFUL TESTS

The Neyman-Pearson Lemma can be used in certain cases to derive optimal tests of a simple null versus a *composite* alternative. First, we must define a concept of optimality for this case.

Definition: The region C is a **uniformly most powerful critical region** of size α for testing the simple hypothesis H_0 against a composite alternative hypothesis H_1 if C is a best critical region of size α for testing H_0 against each simple hypothesis in H_1. The resulting test is said to be **uniformly most powerful.**

Eg: Suppose $X_1, \ldots, X_n$ are *iid* $N(\theta, 1)$ and we want to test $H_0: \theta = \theta'$ versus $H_1: \theta > \theta'$. Let's consider each simple hypothesis in H_1 separately. We know from the Neyman-Pearson Lemma that a best critical region for testing H_0 versus the simple alternative $A: \theta = \theta''$ (where $\theta'' > \theta'$) is the region $(x_1, \ldots, x_n)$ such that $\bar{x} \geq k$. But k does not depend on θ''; it only depends on θ', α and n. To see this explicitly, note that the z-statistic is $z = \sqrt{n}\,(\bar{x} - \theta')$, so that a level-$\alpha$ test rejects H_0 when $z \geq z_\alpha$, or equivalently, when $\bar{x} \geq z_\alpha / \sqrt{n} + \theta'$. So the critical region defined by $\bar{x} \geq k$, where $k = z_\alpha / \sqrt{n} + \theta'$, is a best critical region of size α

for testing H_0 against *each* simple hypothesis in H_1. (It is important to realize that this critical region remains fixed as θ'' ranges through the values in H_1). Therefore the test which rejects H_0 whenever $\bar{x} \geq k$ is uniformly most powerful.

Equivalently, we can say that the power function $K_1(\theta)$ for this test is at least as large as the power function for any other level α test, at *all* values of θ which exceed θ'.

- Uniformly most powerful tests do not always exist. In Example 3, page 407, Hogg and Craig show that there is no uniformly most powerful test of $H_0 : \theta = \theta'$ versus the two-sided alternative $H_1 : \theta \neq \theta'$ when $X_1, \ldots, X_n$ are *iid* $N(\theta, 1)$. (There is a flaw in their proof, but it can be fixed.)

- Another difficulty with the kind of optimality theory covered by the Neyman-Pearson Lemma is that it must be assumed that the p.d.f. of the population is known, except for a finite number of parameters. This assumption certainly makes the problems of statistical inference easier to think about, but in practice it will almost never be true. How, for example, can we be sure that our data came from an *exact* normal distribution?