STATS 300A: Theory of Statistics **Fall 2015**

Lecture 16 — November 17

Lecturer: Lester Mackey *Scribe: Reginald Long, Colin Wei*

Warning: *These notes may contain factual and/or typographic errors.*

16.1 Recap

Last time we discussed UMP unbiased tests for the null hypothesis $H_0 : \theta \in \Omega_0$ versus the alternative given by $H_1 : \theta \in \Omega_1$ and introduced the following key concepts:

Definition 1. ϕ is an unbiased test at level α if $\beta_\phi(\theta_0) \leq \alpha \ \ \forall \theta_0 \in \Omega_0$ and $\beta_\phi(\theta_1) \geq \alpha \ \ \forall \theta_1 \in \Omega_1$.

Definition 2. ϕ is α-similar if $\beta_\phi(\theta) = \alpha \ \ \forall \theta \in \omega$, where $\omega = \overline{\Omega}_0 \cap \overline{\Omega}_1$.

Lemma 1. *(TSH 4.1.1)* If $\beta_\phi(\theta)$ is continuous on Ω for all ϕ, and ϕ_0 is UMP amongst all α-similar level α tests, then ϕ_0 is UMPU at level α.

Last time, we focused on two-sided UMPU tests for one parameter exponential families. Today, we will develop UMPU tests for multiparameter exponential families with nuisance parameters.

16.2 Application of MoUM to our 2-sided testing problem

We continue our discussion of MoUM from last lecture, where our goal is to find an UMPU test. In this setting, $H_0 : \theta = \theta_0$. We will fix a simple alternative $\theta = \theta' \neq \theta_0$ and hope that our best test has no θ' dependence. We would like to maximize power $\int \phi(x) p_{\theta'}(x) d\mu(x)$ subject to

$$\int \phi(x) p_{\theta_0}(x) d\mu(x) = \alpha \tag{16.1}$$

$$\int \phi(x) \frac{d}{d\theta} p_{\theta_0}(x) d\mu(x) = 0. \tag{16.2}$$

For a 1-parameter exponential family, we have

$$p_\theta(x) = h(x) e^{\theta T(x) - A(\theta)} \quad \text{and} \tag{16.3}$$

$$\frac{d}{d\theta} p_\theta(x) = h(x) e^{\theta T(x) - A(\theta)} \left(T(x) - A'(\theta) \right) = p_\theta(x) \left(T(x) - \mathbb{E}_\theta[T(X)] \right). \tag{16.4}$$

By our discussion of MoUM from last lecture, we find that a most powerful test has rejection region defined by

$$p_{\theta'}(x) > k_1 p_{\theta_0}(x) + k_2 \frac{d}{d\theta} p_{\theta_0}(x)$$

for some values of k_1 and k_2, which is equivalent to

$$\frac{e^{(\theta'-\theta_0)T(x)}}{k_1' + k_2' T(x)} > \text{const}$$

with some rearranging.

Now consider the set of values of $T(x)$ satisfying this constraint. Because the constraint is that an exponential function exceeds a linear function, the set of values of $T(x)$ satisfying this constraint is either a one-sided interval

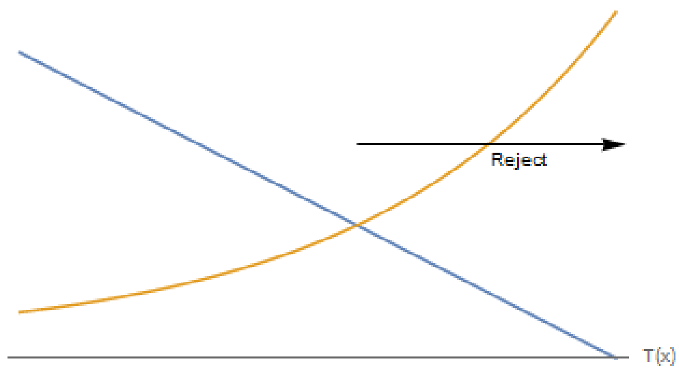

or all points outside a closed interval

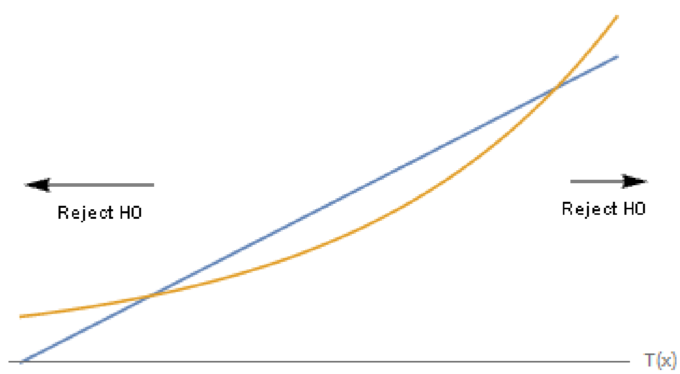

The first possibility will not give rise to an unbiased test, because the result would be a one-sided test with monotone power functions. Therefore any optimal ϕ is of the form

$$\phi(x) = \begin{cases} 1 & \text{if } T(x) > C_1 \text{ or } T(x) < C_2 \\ \gamma_i & \text{if } T(x) = C_i \\ 0 & \text{otherwise} \end{cases} .$$

A simplification is possible if $T(x)$ is symmetrically distributed under θ_0. Then the optimal test rejects whenever $|T(x) - \mu| > \text{const}$. Such tests are called **equitailed tests**.

Example 1. Suppose $X_1, ..., X_n \overset{\text{iid}}{\sim} N(0, \sigma^2)$ with $H_0 : \sigma = \sigma_0$ and $H_1 : \sigma \neq \sigma_0$.

The optimal test has an acceptance region of the form

$$C_1 \leq \frac{\sum_i X_i^2}{\sigma_0^2} \leq C_2$$

The middle expression is a sufficient statistic that is χ_n^2 distributed under H_0. How do we choose C_1 and C_2? Let f_n be density χ_n^2. The level constraint is:

$$P(C_1 \leq T(X) \leq C_2) = \int_{C_1}^{C_2} f_n(y)dy = 1 - \alpha,$$

and the derivative constraint (16.2) with substitution (16.4) gives

$$\mathbb{E}_{\theta_0}[T(X)\phi(X)] = \mathbb{E}_{\theta_0}[T(x)]\mathbb{E}_{\theta_0}[\phi(X)].$$

On the left side $\phi(x) = 1$ on the complement of (C_1, C_2), and on the right, the mean of a χ_n^2 is n, and the level is α. Thus,

$$\int_{(C_1,C_2)^c} y f_n(y) = n\alpha.$$

For χ_n^2 distributions, $y f_n(y) = n f_{n+1}(y)$, and the derivative constraint ultimately becomes

$$\int_{C_1}^{C_2} f_{n+2}(y) = 1 - \alpha.$$

We have two integral equations, and we can solve them for the unknown boundaries C_1, C_2.

16.3 Optimal Unbiased Testing in Multiparameter Exponential Families

Let the density of an exponential family with the natural parameters $(\theta, \lambda_1, \cdots, \lambda_k) \in \mathbb{R}^{k+1}$ be

$$\mathbb{P}_{\theta,\lambda}(x) = \exp\left(\theta U(x) + \sum_{i=1}^{k} \lambda_i T_i(x) - A(\theta, \lambda)\right) h(x). \tag{16.5}$$

We want to frame a test for the null hypothesis given by $H_0 : \theta \leq \theta_0$ against the alternative given by $H_1 : \theta > \theta_0$ in the presence of nuisance parameters λ. On the boundary $\omega = \{(\theta, \lambda) : \theta = \theta_0\}$, θ is known which would imply $T = (T_1, T_2, \cdots, T_k)$ is sufficient for λ. Hence the conditional distribution of X and subsequently that of $U(X)|T(X)$ has no λ dependence on ω. This point is worth repeating: **conditioning can eliminate the influence of nuisance parameters!** What is more, for each t, $U(X)|T(X) = t$ forms a one parameter exponential family. Here is a proof of this in the discrete case.

Proof.

$$
\begin{aligned}
\mathbb{P}_{\theta,\lambda}(U(X) = u | T(X) = t) &= \frac{\exp\left(\theta u + \sum_{i=1}^{k} \lambda_i t_i\right) \sum_{x:T(x)=t,U(x)=u} h(x)}{\sum_{x:T(x)=t} \exp\left(\theta u(x) + \sum_{i=1}^{k} \lambda_i t_i\right) h(x)} \\
&= c_t(\theta) \exp(\theta u)\, g(t,u)
\end{aligned}
$$

□

These facts suggest the following strategy:

1. Condition on $T(X) = t$.
2. For each value of t, construct a "best conditional test" $\phi(u,t)$ which maximizes conditional power $\mathbb{E}_\theta[\phi(u,t)|T = t]\ \ \forall\theta > \theta_0$ at conditional level α, i.e., $\mathbb{E}[\phi(u,t)|T = t] \leq \alpha\ \ \forall\theta \leq \theta_0$.
3. Check whether this test is UMPU at level α.

Since, for each t, the one-parameter exponential family distribution of $U(X)|T(X) = t$ has MLR in $U(X)$, our "best test" for a fixed t will be a one-sided test of the form

$$
\phi(u,t) = \begin{cases} 1 & \text{if } u > c(t) \\ \gamma(t) & \text{if } u = c(t) \\ 0 & \text{if } u < c(t) \end{cases}.
$$

By the proof of the MLR theorem (TSH Thm 3.4.1), if $c(t)$ and $\gamma(t)$ are chosen to satisfy the conditional similarity constraint

$$
\mathbb{E}_{\theta_0}(\phi(U,T)|T = t) = \alpha \tag{16.6}
$$

then it will also satisfy the conditional level constraint

$$
\mathbb{E}_\theta(\phi(U,T)|T = t) \leq \alpha\ \ \forall\theta \leq \theta_0 \tag{16.7}
$$

and the maximum conditional power property, that

$$
\mathbb{E}_\theta(\phi(U,T)|T = t) \quad \text{is maximized} \quad \forall\,\theta > \theta_0, \tag{16.8}
$$

among all tests that satisfy (16.6). Since (16.7) and (16.8) hold for every t, we may take expectations and deduce that $\phi(U,T)$ is level α and UMP amongst level α tests satisfying (16.6). However, (16.6) is a more stringent requirement than α-similarity, so it remains to show that $\phi(U,T)$ is also UMP amongst level α, α-similar tests. This would imply that ϕ is UMPU unbiased at level α.

16.4 Characterization of Similar Tests

We will achieve our goal by developing a new characterization of similar tests based on conditioning on a complete sufficient statistic. We begin by assigning a name to property (16.6).

Definition 3. A test ϕ has α-**Neyman structure** if T is sufficient for $\{\mathbb{P}_\gamma, \gamma \in \omega\}$ and ϕ satisfies $\mathbb{E}(\phi(X)|T(X)) = \alpha$ a.s. $\mathbb{P}_\gamma$ for all $\gamma \in \omega$, $\omega = \bar{\Omega}_0 \cap \bar{\Omega}_1$.

Notice that if ϕ has α-Neyman structure, then ϕ is also α-similar, since $\mathbb{E}_\gamma \phi(X) = \mathbb{E}_\gamma(\mathbb{E}(\phi(X)|T(X))) = \alpha$ for all $\gamma \in \omega$. Moreover, if T is *complete* and sufficient for $\{\mathbb{P}_\gamma : \gamma \in \omega\}$, then every α-similar test has α-Neyman structure with respect to T.

Proof. Suppose ϕ is α-similar, and let $\Psi(T) = \mathbb{E}(\phi(X) - \alpha|T)$, which has no γ dependence as T is sufficient. Then $\mathbb{E}_\gamma(\Psi(T)) = \mathbb{E}_\gamma(\phi(X) - \alpha) = 0, \ \forall \gamma \in \omega$, which in turn implies that $\Psi(T) = 0$ *a.s.* by completeness. □

Therefore, Neyman structure and similarity are equivalent whenever a complete sufficient statistic T exists. We now focus on exponential families, as they typically yield complete sufficient statistics on the boundary. For exponential family with density $f(x) \propto \exp\left(\theta U(x) + \sum_i \lambda_i T_i(x)\right) h(x)$. If T is complete on ω, then there exists an UMPU test for $\theta \le \theta_0$ vs. $\theta > \theta_0$ with form

$$\phi(U,T) = \begin{cases} 1 & \text{if } U > c(T) \\ \gamma(T) & \text{if } U = c(T) \\ 0 & \text{if } U < c(T) \end{cases} \tag{16.9}$$

with $c(T)$ and $\gamma(T)$ chosen such that $E_{\theta_0}[\phi(U,T)|T] = \alpha$.

Why is this true? By the previous section, ϕ is UMP amongst level α tests satisfying (16.2), but (16.2) is equivalent to α-Neyman structure for this testing scenario. This implies that ϕ is UMP amongst tests with α-Neyman structure. T is complete and sufficient for $\mathbb{P}_{\theta_0}$, so every α-similar test has α-Neyman structure w.r.t T. Hence, ϕ is UMP amongst α-similar tests. Finally, since an exponential family has continuous power functions, ϕ is UMPU at level α.

Example 2. Suppose that $X \sim \text{Poi}(\nu)$ and $Y \sim \text{Poi}(\mu)$ for X and Y independent. You might view X and Y as the number of successful recoveries from a disease under two different treatments. Our goal is to test the null hypothesis $H_0 : \mu \le \nu$ against the alternative $H_1 : \mu > \nu$. Equivalently, we test $H_0 : \log(\mu/\nu) \le 0$ against $H_1 : \log(\mu/\nu) > 0$. This rewriting emphasizes our singular interest in the ratio μ/ν. Any additional information in (μ, ν) can be regarded as nuisance.

The joint density of (X, Y) is given by $\exp(-\nu - \mu) \exp(x \log \nu + y \log \mu) / x!y!$ which is proportional to:

$$\frac{1}{x!y!} \exp\left(y \log(\mu/\nu) + (x + y) \log \nu\right)$$

The above is a 2-dimensional exponential family with sufficient statistics $U = Y$ and $T = X + Y$. The natural parameters are $\theta = \log(\mu/\nu)$ and $\lambda = \log(\nu)$ where λ acts as the nuisance parameter.

On the boundary, $\omega = \{(\theta,\lambda) : \theta = 0\}$, (U,T) is a one-dimensional full-rank exponential family with density $\propto h(u,t)\exp(t\lambda)$. This implies that T is complete sufficient on ω.

The conditional distribution of $U|T$ on ω is Bin($p = \frac{\mu}{\nu+\mu}, n = T$), so T determines the number of coin flips, and the probability of each coin coming up "heads" is determined by μ and ν. The critical value $c(T)$ is going to be determined at the boundary point, Bin($p = \frac{1}{2}, n = T$) (since $\theta = \log(\frac{\mu}{\nu}) = 0$).

Example 3. $X_1, ..., X_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$, where both μ and σ^2 are unknown. The joint density of $X_1, ..., X_n$ is

$$\propto \exp\left(-\frac{1}{2\sigma^2}\sum_i x_i^2 + \frac{\mu}{\sigma^2}\sum_i x_i\right) \tag{16.10}$$

Testing the variance We will consider testing $H_0 : \sigma \geq \sigma_0$ vs. $H_1 : \sigma < \sigma_0$, which is equivalent to testing $\theta \geq \theta_0$ vs. $\theta < \theta_0$, under the choices $\theta = -\frac{1}{2\sigma^2}$, $\lambda = \frac{\mu}{\sigma^2}$, $U = \sum_i {x_i}^2$, and $T = \sum_i x_i$.

The UMPU test rejects for small U, i.e., $\phi(X) = \mathbb{I}\left(\sum_{i=1}^n X_i^2 \leq C_\alpha(\bar{X})\right)$. To determine the critical value $C_\alpha(\bar{X})$, we notice that

$$\begin{aligned}
\alpha &= \mathbb{P}_{\sigma_0}\left[\sum_{i=1}^n X_i^2 \leq C_\alpha(\bar{X})|\bar{X}\right] = \mathbb{P}_{\sigma_0}\left[\sum_{i=1}^n X_i^2 - n\bar{X}^2 \leq C_\alpha(\bar{X}) - n\bar{X}^2|\bar{X}\right] \\
&= \mathbb{P}_{\sigma_0}\left[\sum_{i=1}^n (X_i - \bar{X})^2 \leq C'_\alpha(\bar{X})|\bar{X}\right] \\
&= \mathbb{P}_{\sigma_0}\left[\sum_{i=1}^n (X_i - \bar{X})^2 \leq C''_\alpha\right] \\
&\qquad \left(\text{Basu's theorem} \Rightarrow \sum_{i=1}^n (X_i - \bar{X})^2 \perp \bar{X} \Rightarrow \text{ we can choose } C'_\alpha(\bar{X}) \text{ constant w.r.t } \bar{X}\right) \\
&= \mathbb{P}_{\sigma_0}\left[\underbrace{\frac{\sum_{i=1}^n (X_i - \bar{X})^2}{\sigma_0^2}}_{\chi^2_{n-1}} \leq \frac{C''_\alpha}{\sigma_0^2}\right]
\end{aligned}$$

implying that $C''_\alpha = \sigma_0^2 z_{n-1,\alpha}$ where $z_{n-1,\alpha}$ is the α quantile of χ^2_{n-1}. Therefore the UMPU test rejects if $\sum_{i=1}^n (X_i - \bar{X})^2 \leq \sigma_0^2 z_{n-1,\alpha}$.