Outline

1) Testing with nuisance parameters

2) Conditional tests

3) Multi-parameter exponential families

Nuisance Parameters

Common setup: Extra unknown parameters which are not of direct interest

$\mathcal{P} = \{P_{\theta,\lambda} : (\theta,\lambda) \in \Omega\}$, $H_0 : \theta \in \Theta_0$ vs $H_1 : \theta \in \Theta_1$

θ parameter of interest

λ nuisance parameter

Issue: λ unknown but might affect type I error or power of a given test

Ex $X_1, \dots, X_n \overset{iid}{\sim} N(\mu, \sigma^2)$ $Y_1, \dots, Y_m \overset{iid}{\sim} N(\nu, \sigma^2)$

μ, ν, σ^2 unknown

$H_0 : \mu = \nu$ vs $H_1 : \mu \neq \nu$

$\theta = \mu - \nu$ $\lambda = (\mu + \nu, \sigma^2)$ or (μ, σ^2)

$H_0 : \theta = 0$ vs $H_1 : \theta \neq 0$

Ex $X \sim \text{Pois}(\mu)$ $Y \sim \text{Pois}(\nu)$ $X \perp\!\!\!\perp Y$

$H_0 : \mu \leq \nu$ vs $H_1 : \mu > \nu$

$\theta = \frac{\mu}{\mu + \nu}$ $\lambda = \mu + \nu$

$H_0 : \theta \leq 1/2$ $H_1 : \theta > 1/2$

Conditional testing

Ex Binomial with random sample size

1) sample $N \sim \text{Pois}(10)$ (e.g., number of survey respondents)
2) sample $X \mid N = n \sim \text{Binomial}(n, \theta)$

Observe N, X, test $H_0: \theta \le 1/2$ vs $H_1: \theta > 1/2$

Should we:

1) Use marginal dist: $X \sim \text{Pois}(10 \cdot \theta)$

   MLR $\Rightarrow$ Reject if $X > c_\alpha^{(1)}$, upper-α quantile of $\text{Pois}(5)$

   $\mathbb{E}_\theta \phi_1(X) \le \alpha$, for $\theta \le 1/2$

2) Use conditional dist: $X \mid N = n \sim \text{Binom}(n, 1/2)$

   MLR $\Rightarrow$ Reject if $X > c_\alpha^{(2)}(n)$, upper-α quant. of $\text{Binom}(n, 1/2)$

   $p(x) = \mathbb{P}_{\theta = 1/2}(X \ge x \mid N = n)$

   $\mathbb{E}_\theta[\phi_2(X \mid N) \mid N] \overset{\text{a.s.}}{\le} \alpha \Rightarrow \mathbb{E}_\theta \phi_2(X \mid N) \le \alpha, \quad \theta \le 1/2$

Neither test is UMP ($\frac{p_{\theta_1}}{p_{\theta_0}}(x, n)$ depends on n, x)

(But ϕ_2 is UMP in conditional problem)

Conditionality Principle

Recall conditionality principle:

N is ancillary, so we should condition on it.

Basic logic: why should we care N is not fixed?

Condition on $N \Longleftrightarrow$ treat as fixed

Conditional model: Candidate cond. dist.s induced by $\mathcal{P}$:

$$\mathcal{Q}_n = \{ q_\theta(x \mid n) : \theta \in \Theta \}$$

$$X \mid N = n \sim \text{Binom}(n, \theta)$$

MLR in $X \Rightarrow \phi_2(X \mid N)$ is optimal conditional test

Prop: If $\phi(X)$ is level α (and unbiased) given $U(X)$ then $\phi(X)$ is level α (and unbiased) marginally

Proof $\mathbb{E}_\theta \phi(X) = \mathbb{E}_\theta [\mathbb{E}_\theta [\phi(X) \mid U(X)]]$ $\boxtimes$

Conditioning and nuisance param.s

Ex (Random sample size, unknown dist.)

$N \sim P^N$ $\qquad$ P^N unknown

$X \mid N = n \sim \text{Binom}(n, \frac{1}{2})$

Infinite-dimensional nuisance parameter P^N

But: Can still use conditional test!

$$\mathbb{E}_{\theta, P^N}[\phi_2(X)] = \mathbb{E}_{\theta, P^N}\left[\mathbb{E}_\theta[\phi_2(X) \mid N]\right] \leq \alpha$$

N is sufficient stat for nuisance param. P^N

(i.e., N sufficient for model with θ known)

$\Leftrightarrow$ Conditioning on N removes P^N from problem

In general: If $U(x)$ suff for λ (suff. when θ known)

then θ is the only parameter of

$$\mathcal{Q}_u = \{Q_\theta(X \mid U(x) = u) : \theta \in \Theta\}$$

Ex (Comparing Poissons)

$X \sim \text{Pois}(\mu) \quad Y \sim \text{Pois}(\nu), \quad X \perp\!\!\!\perp Y$

Test $H_0: \mu \leq \nu$ vs $H_1: \mu > \nu$

Let $N = X + Y \sim \text{Pois}(\mu + \nu)$

$X \mid N = n \sim \text{Binom}(n, \theta) \qquad \theta = \frac{\mu}{\mu + \nu} \qquad (Y = N - X)$

$H_0: \mu \leq \nu \iff \theta \leq \frac{1}{2} \qquad H_1: \mu > \nu \iff \theta > 1/2$

Multiparameter Exp. Families

Assume $X \sim p_{\theta,\lambda}(x) = e^{\theta' T(x) + \lambda' U(x) - A(\theta,\lambda)} h(x)$

$\theta \in \mathbb{R}^s$, $\lambda \in \mathbb{R}^r$, both unknown.

How to test $H_0 : \theta \in \Theta_0$ vs $H_1 : \theta \in \Theta_1$?

<u>Idea</u>: Condition on $U(X)$ to eliminate dep. on λ

1) <u>Sufficiency reduction</u> to $(T(X), U(X))$

$$(T, U) \sim p_{\theta,\lambda}(t,u) = e^{\theta' t + \lambda' u - A(\theta,\lambda)} g(t,u)$$

density wrt e.g. Lebesgue on $\mathbb{R}^{s+r}$

[$g(t,u)\,dt\,du$ = push-forward of $h(x)\,d\mu(x)$]

2) <u>Condition on U:</u>

$$q_\theta(t \mid u) = \frac{p_{\theta,\lambda}(t,u)}{\int p_{\theta,\lambda}(z,u)\,dz}$$

$$= \frac{e^{\theta' t + \cancel{\lambda' u} - \cancel{A(\theta,\lambda)}} g(t,u)}{\underbrace{\int e^{\theta' z + \cancel{\lambda' u} - \cancel{A(\theta,\lambda)}} g(z,u)\,dz}_{e^{B_u(\theta)} =}}$$

$$= e^{\theta' t - B_u(\theta)} g(t,u)$$

3) Conditional test:

Test $H_0: \theta \in \Theta_0$ vs. $H_1: \theta \in \Theta_1$ in

s-parameter model $\mathcal{Q}_u = \{q_\theta(t|u) : \theta \in \Theta\}$

Note if $s=1$, this family has MLR in T

Even if $s>1$, we still have gotten rid of λ

Ex (Comparing Poissons, Cont'd)

$X \sim \text{Pois}(\mu)$ $\quad Y \sim \text{Pois}(\nu)$, $\quad X \perp\!\!\!\perp Y$

$H_0: \mu \le \nu$ vs $H_1: \mu > \nu$

$$p_{\mu,\nu}(x,y) = \frac{\mu^x e^{-\mu}}{x!} \cdot \frac{\nu^y e^{-\nu}}{y!}$$

$$= e^{x \log \mu + y \log \nu - (\mu+\nu)} \cdot \frac{1}{x!y!}$$

want $T = X$, $u = X+Y$

$$= e^{x(\log \mu - \log \nu) + (x+y)\log \nu - (\mu+\nu)} \frac{1}{x!y!}$$

$H_0: \mu \le \nu \iff \log \frac{\mu}{\nu} \le 0 \quad \left(\iff \frac{\mu}{\mu+\nu} \le \frac{1}{2}\right)$

$\rightsquigarrow$ Condition on $U = X+Y$, reject for large X

<u>Theorem</u> Let $\mathcal{P}$ be full rank exp. fam. with densities $$p_{\theta,\lambda}(x) = e^{\theta T(x) + \lambda' U(x) - A(\theta,\lambda)} h(x)$$

$\theta \in \mathbb{R}$, $\lambda \in \mathbb{R}^r$, $(\theta,\lambda) \in \Omega$ open, θ_0 possible

a) To test $H_0: \theta \le \theta_0$ vs. $H_1: \theta > \theta_0$, there is a UMPU test $\phi^*(x) = \psi(T(x); U(x))$ where

$$\psi(t;u) = \begin{cases} 1 & t > c(u) \\ \gamma(u) & t = c(u) \\ 0 & t < c(u) \end{cases}$$

With $c(u)$, $\gamma(u)$ chosen to make

$$\mathbb{E}_{\theta_0}[\phi^*(X) \mid U(X) = u] = \alpha$$

b) To test $H_0: \theta = \theta_0$ vs. $H_1: \theta \ne \theta_0$ there is a UMPU test $\phi^*(x) = \psi(T(x); U(x))$ where

$$\psi(t;u) = \begin{cases} 1 & t < c_1(u) \text{ or } t > c_2(u) \\ \gamma_i(u) & t = c_i(u) \\ 0 & t \in (c_1(u), c_2(u)) \end{cases}$$

with $c_i(u)$, $\gamma_i(u)$ chosen to make

$$\mathbb{E}_{\theta_0}[\phi^*(X) \mid U(X) = u] = \alpha$$

$$\mathbb{E}_{\theta_0}[T(X)(\phi^*(X) - \alpha) \mid U(X) = u] = 0$$

[Note λ has disappeared from the problem.]

Proof Sketch

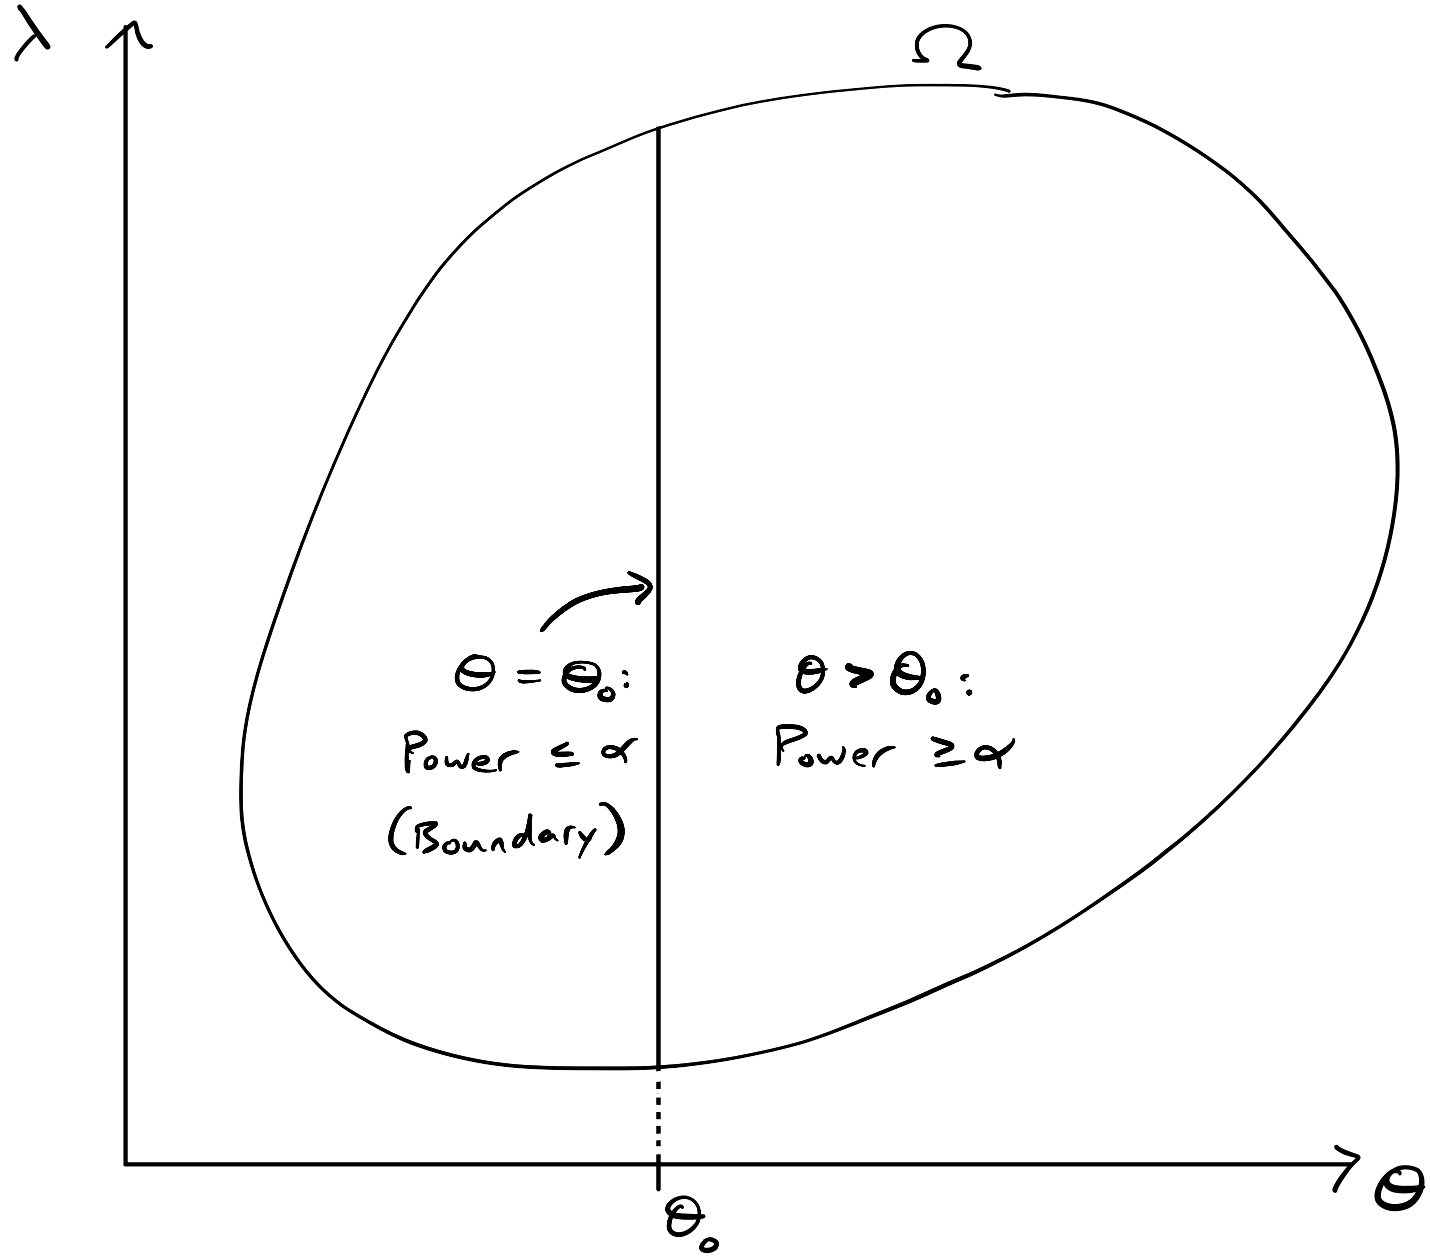

1) Any unbiased test has $\beta(\theta_0, \lambda) = \alpha \quad \forall \lambda$

(continuity of $\beta(\theta, \lambda)$)

2) Power $\equiv \alpha$ on boundary $\Rightarrow$ $\mathbb{E}_{\theta_0}[\phi \mid U] \overset{a.s.}{=} \alpha$

($U(X)$ complete sufficient on boundary submodel)

3) ϕ^* optimal among all tests with conditional level α

(by reduction to univariate model)

Proof Assume ϕ any unbiased test

Step 1: $\mathbb{E}_{\theta,\lambda}|\phi(X)| \le 1 < \infty \quad \forall(\theta,\lambda)\in\Omega$

(Keener Thm 2.4)

$\Rightarrow \mathbb{E}_{\theta,\lambda}\phi(X)$ infinitely diff. on Ω, can diff. under $\int$

ϕ unbiased $\Rightarrow \mathbb{E}_{\theta_0,\lambda}[\phi(X)] = \alpha \quad \forall (\theta_0,\lambda)\in\Omega$

Step 2: Boundary submodel: $\mathcal{P}_{\theta_0} = \{P_{\theta_0,\lambda} : (\theta_0,\lambda)\in\Omega\}$

$$p_{\theta_0,\lambda}(x) = e^{\lambda' U(x) - A(\theta_0,\lambda)} \cdot \frac{e^{\theta_0 T(x)}}{h(x)}$$

$\mathcal{P}_{\theta_0}$ is full-rank, s-param exp. fam, $U(X)$ comp. suff.

Let $f(u) = \mathbb{E}_{\theta_0}[\phi(X) \mid U(X)=u] - \alpha$

$$\mathbb{E}_{\theta_0,\lambda}[f(U(X))] = \mathbb{E}_{\theta_0,\lambda}[\phi(X)] - \alpha = 0 \quad \forall\lambda$$

$$\Rightarrow f(u) \overset{a.s.}{=} 0$$

$$\Rightarrow \mathbb{E}_{\theta_0}[\phi(X) \mid U(X)=u] = 0 \quad \forall u$$

Two-sided case: $g(u) = \frac{d}{d\theta}\mathbb{E}_{\theta_0}[\phi \mid U=u]$

$$= \mathbb{E}_{\theta_0}[(T - \mathbb{E}_{\theta_0}[T\mid u])\phi \mid u]$$

$$= \mathbb{E}_{\theta_0}[T(\phi-\alpha) \mid u]$$

$$\mathbb{E}_{\theta_0,\lambda}\, g(u) = \mathbb{E}_{\theta_0,\lambda}[T(\phi-\alpha)] = \frac{\partial}{\partial\theta}\beta_\phi(\theta_0) = 0 \;\forall\lambda$$

$$\Rightarrow \frac{d}{d\theta}\mathbb{E}_{\theta_0}[\phi \mid u] \overset{a.s.}{=} 0$$ (Cond'l power has derivative 0 at θ_0)

Step 3: For any value u, the conditional model is

$$q_\theta(t \mid u) = e^{\theta t - B_u(\theta)} g(t,u), \quad \text{1-param. exp. fam}$$

In one- / two-sided case, we have shown

$\psi(t;u)$ is UMP / UMPU in Q_u

Let $\bar{\phi}(t;u) = \mathbb{E}\left[\phi(X) \mid T(X)=t, U(X)=u\right]$

$$\mathbb{E}_\theta\left[\bar{\phi}(T;u) \mid U=u\right] = \mathbb{E}_\theta\left[\phi(X) \mid U(X)=u\right]$$
$$= \alpha \text{ if } \theta = \theta_0$$

$\Rightarrow \bar{\phi}(\cdot\,;u)$ is a (cond'l) test of H_0 vs. H_1 in Q_u with power $= \alpha$ at boundary

One-sided case:

(or $\theta \le \theta_0$)

$\psi(t;u)$ is the UMP test of $\theta = \theta_0$ vs $\theta > \theta_0$ in Q_u, which is a 1-param. exp. fam.

Two-sided case:

$\psi(t;u)$ is the UMP test of $\theta = \theta_0$ vs. $\theta \ne \theta_0$ among tests with power $= \alpha$, $\frac{d}{d\theta}$power $= 0$ @ θ_0

In either case ψ has higher <u>cond</u>. power than $\bar{\phi}$, a.s.

For $(\theta, \lambda) \in \Omega_1$:

$$\begin{aligned}
\mathbb{E}_{\theta,\lambda}[\phi(X)] &= \mathbb{E}_{\theta,\lambda}\left[\mathbb{E}_{\theta}\left[\bar{\phi}(T; u) \mid u \right] \right] \\
&\le \mathbb{E}_{\theta,\lambda}\left[\mathbb{E}_{\theta}\left[\psi(T; u) \mid u \right] \right] \\
&= \mathbb{E}_{\theta,\lambda}\left[\phi^*(X) \right]
\end{aligned}$$

Permutation Tests

Even if we don't get a UMPU test at the end, conditioning on null suff. stat. still helps.

Ex. $X_1, \dots, X_n \overset{iid}{\sim} P \quad Y_1, \dots, Y_m \overset{iid}{\sim} Q \quad H_0: P = Q \quad H_1: P \neq Q$

Under H_0, $P = Q$, $X_1, \dots, X_n, Y_1, \dots, Y_m \overset{iid}{\sim} P$

Let $(Z_1, \dots, Z_{n+m}) = (X_1, \dots, X_n, Y_1, \dots, Y_m)$

Under H_0, $U(Z) = (Z_{(1)}, \dots, Z_{(n+m)})$ compl. suff

Let $S_{n+m} = \{\text{Permutations on } n+m \text{ elements}\}$

$$(X, Y) \mid U \overset{H_0}{\sim} \text{Unif}\left(\{\pi U : \pi \in S_{n+m}\}\right)$$

Thus, for any test stat T, if $P = Q$,

$$\mathbb{P}_{P,Q}\left(T(Z) \geq t \mid U\right) = \frac{1}{(n+m)!} \sum_{\pi \in S_{n+m}} 1\{T(\pi Z) \geq t\}$$

Monte Carlo test: In practice, we sample

$\pi_1, \dots, \pi_B \overset{iid}{\sim} S_{n+m}$, e.g. $B = 1000$

Then $Z, \pi_1 Z, \dots, \pi_B Z \overset{iid}{\sim} \text{Unif}(S_{n+m} U)$ under H_0

MC p-value $$p = \frac{1}{1+B}\left(1 + \sum_{b=1}^{B} 1\{T(Z) \leq T(\pi_b Z)\}\right)$$

$$\overset{H_0}{\sim} \text{Unif}\left(\left\{\frac{1}{1+B}, \dots, \frac{B-1}{1+B}, 1\right\}\right) \quad (\text{if no ties})$$

$(p \geq \text{Unif}(\cdot)$ if there are ties $)$

One-sample t-test $(n=2)$

Ex $\quad X_1, X_2 \overset{iid}{\sim} N(\mu, \sigma^2) \qquad \sigma^2 > 0$ unknown

$H_0: \mu = 0$ vs. $H_1: \mu \neq 0$

$$p_{\mu,\sigma^2}(x) = \frac{1}{2\pi\sigma^2} e^{-\frac{1}{2\sigma^2}\sum_{i=1}^{2}(x_i-\mu)^2}$$

$$= e^{\underbrace{\frac{\mu}{\sigma^2}}_{\theta} \underbrace{\sum X_i}_{T=\bar{x}} \underbrace{-\frac{1}{2\sigma^2}}_{\lambda} \underbrace{\sum X_i^2}_{u=\|x\|^2} - \frac{\mu}{\sigma^2}} \cdot \frac{1}{2\pi\sigma^2}$$

Optimal test rejects when $T = X_1 + X_2$

is extreme given $T = \|X\|^2$

If $\mu = 0$, p is rotationally symmetric

$\Rightarrow X$ rotationally symmetric, $\frac{X}{\|X\|} \perp\!\!\!\perp \|X\|^2$

$\Leftrightarrow X \,/\, \|X\|^2 = u \overset{H_0}{\sim} \text{Unif}(\sqrt{u} \cdot S^1)$

where $S^1 =$ unit circle

Geometric Picture

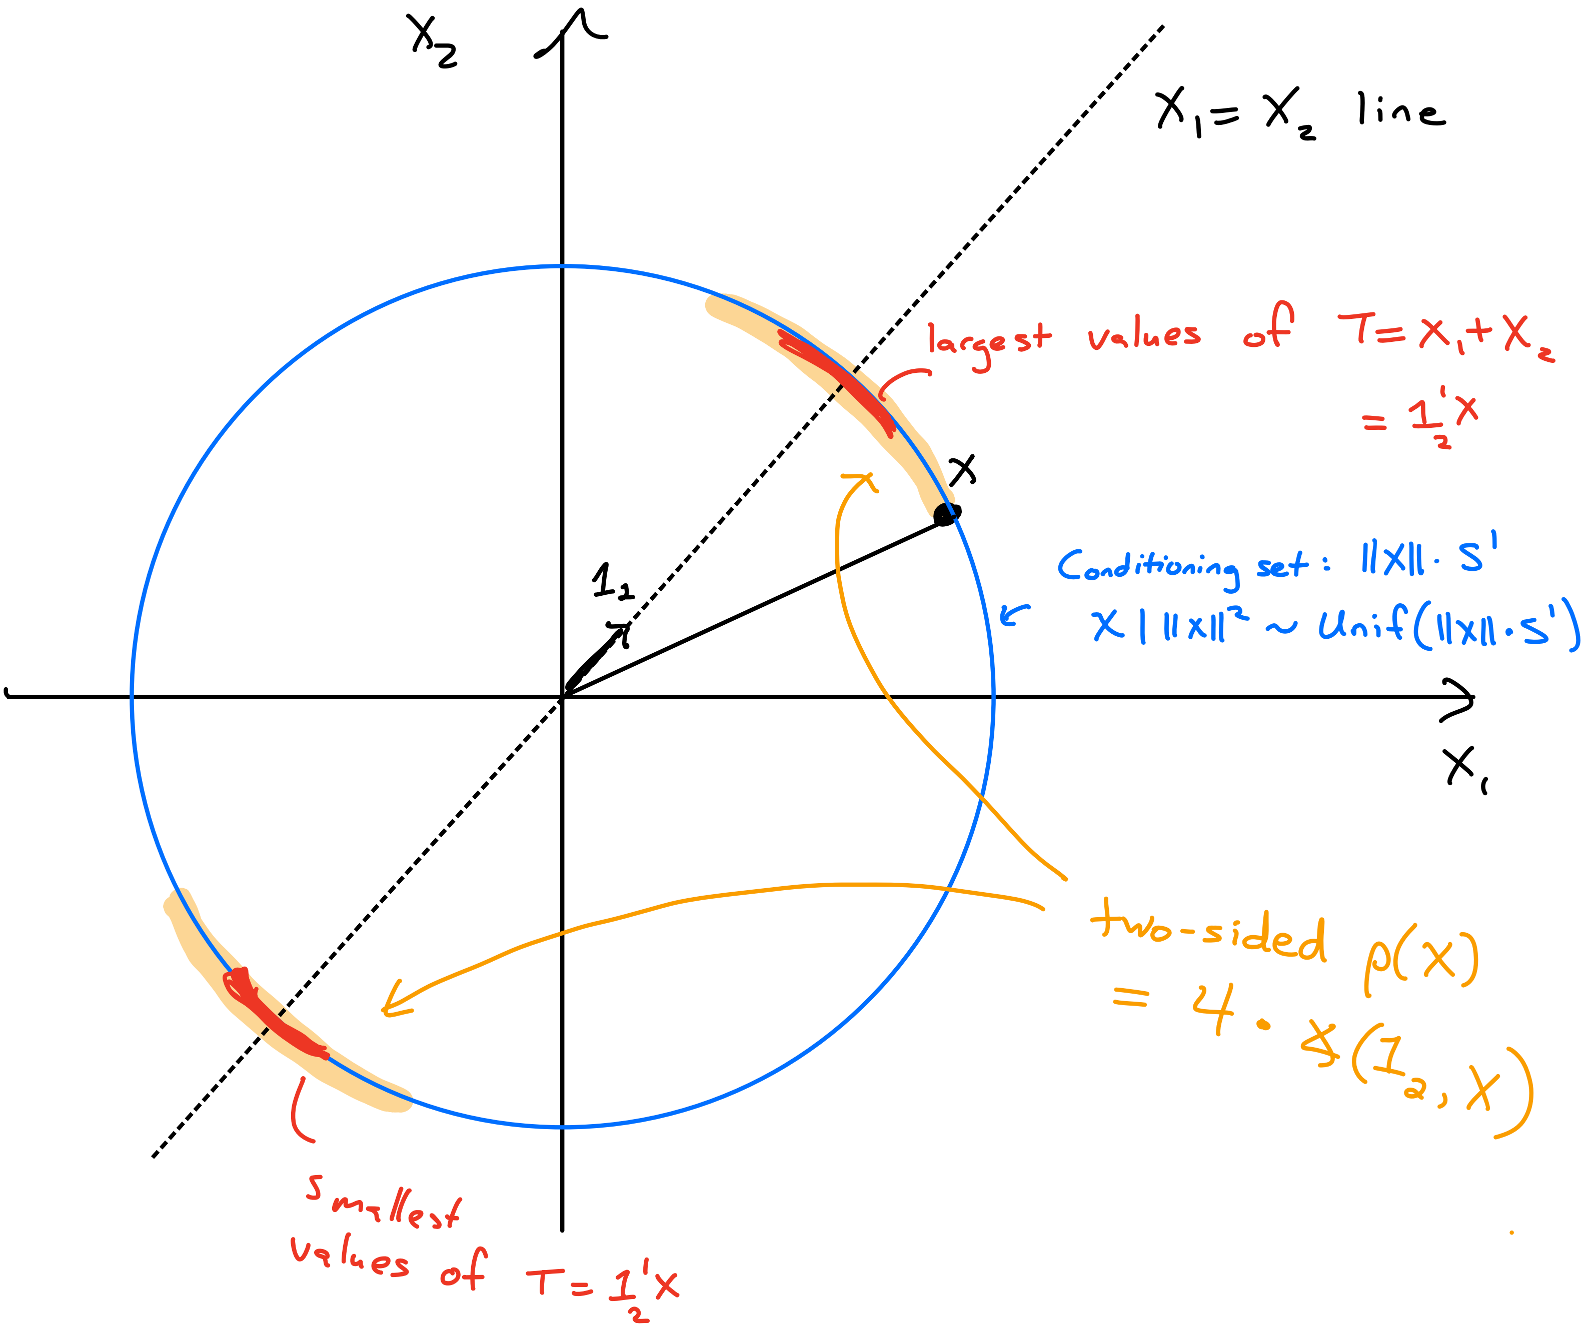

t - statistic

Above test rejects for

- conditionally extrene $1_2'X$ given $\|X\|^2$

(equiv.) • marginally extreme $R = \frac{1_2'X}{\sqrt{2}\cdot\|X\|} = \cos\measuredangle(1_2, X)$

(since $R \perp\!\!\!\perp \|X\|^2$)

(equiv.) • marginally extreme $T = \frac{R}{\sqrt{1-R^2}} \overset{H_0}{\sim} t_1$

$$= \cot\measuredangle(1_2, X)$$

Geometric Picture

$$T = \frac{R}{\sqrt{1-R^2}} = \frac{\| Proj_{1_2} X \|}{\| Proj^{\perp}_{1_2} X \|} \cdot sgn(\bar{x})$$

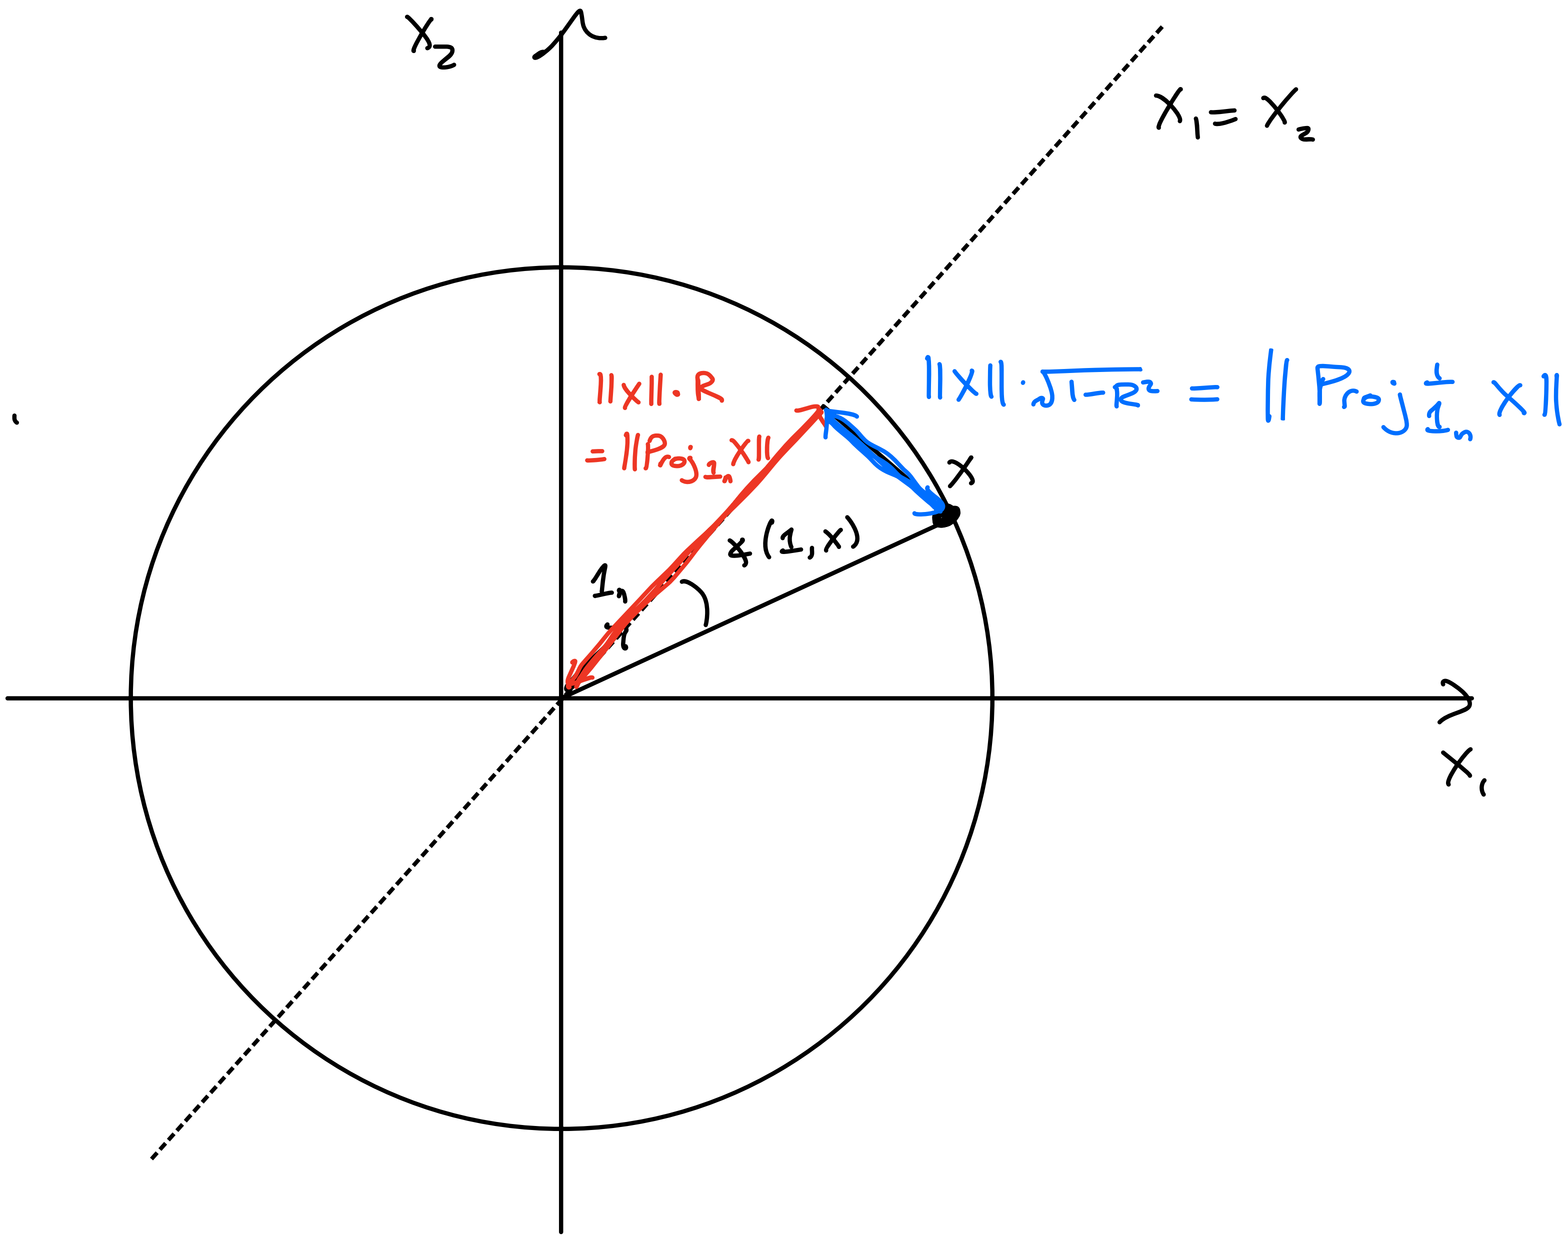

Next major theme: ratios of projections

One-sample t-test

Ex $X_1, \ldots, X_n \overset{iid}{\sim} N(\mu, \sigma^2)$ $\quad \mu, \sigma^2$ unknown

$H_0: \mu = \mu_0$ vs $H_1: \mu \neq \mu_0$

Same logic: Reject for extreme $\bar{X}$ given $\|X\|^2$

$\iff$ Reject for extreme

$$R = \frac{1_n' X}{\sqrt{n} \cdot \|X\|} = \frac{\sqrt{n}\, \bar{X}}{\|X\|}$$

$\iff$ Reject for extreme

$$T = \sqrt{n-1} \cdot \frac{R}{\sqrt{1-R^2}}$$

$$= \sqrt{n-1} \cdot \frac{\sqrt{n}\, \bar{X}}{\sqrt{\|X\|^2 - n\bar{X}^2}} = \frac{\sqrt{n}\, \bar{X}}{\sqrt{S^2}}$$

for sample variance

$$S^2 = \frac{1}{n-1} \sum_{i=1}^{n} (X_i - \bar{X})^2$$

$$= \frac{1}{n-1} \| X - \bar{X} 1_n \|^2 \quad \leftarrow \|\mathrm{Proj}_{1_n}^{\perp} X\|^2$$

$$= \frac{1}{n-1} \left(\|X\|^2 - n\bar{X}^2 \right)$$

$$\sim \frac{\sigma^2}{n-1} \chi^2_{n-1}$$

Next major theme: ratios of projections